
TOMORROW starts here.
cisco
Cisco live!

# VxLAN Deployment – Use Cases and Best Practices

BRKDCT-2643



Azeem Suleman
Solutions Architect – Cisco Advanced Services



Cisco live!

# Contributions

Thanks to the team:

- Abhishek Saxena
- Mehak Mahajan
- Lilian Quan
- Bradley Wong
- Mike Herbert
- Sandeep Subramaniam

# Abstract

With growing adoption of virtualisation in customer environment and large number of workload mobility requirements in data centre; overlays are becoming key technology.  VXLAN is one the overlay technology.

This session will discuss about VXLAN design, deployment and best practices in Data Centre environment. It will also cover:

1) Explain VXLAN deployment steps on top design use cases
2) Configuration and key troubleshooting tips needed during deployment

# Agenda

- Overlays
- Introduction to VXLAN
- VXLAN Design
- Deployment Steps
- Troubleshooting
- Key Takeaways
- References




# Overlays

# Underlay vs Overlays

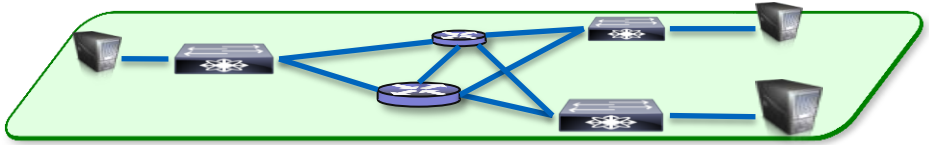

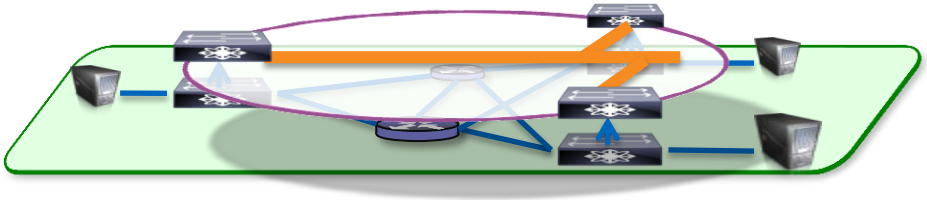

## Robust Underlay/Fabric

- High Capacity Resilient Fabric
- Intelligent Packet Handling
- Programmable & Manageable

## Flexible Overlay Virtual Network

- Mobility – Track end-point attach at edges
- Scale – Reduce core state
  - Distribute and partition state to network edge
- Multi-tenancy – Share Network resources
- Flexibility/Programmability
  - Reduced number of touch points

# Overlay Taxonomy

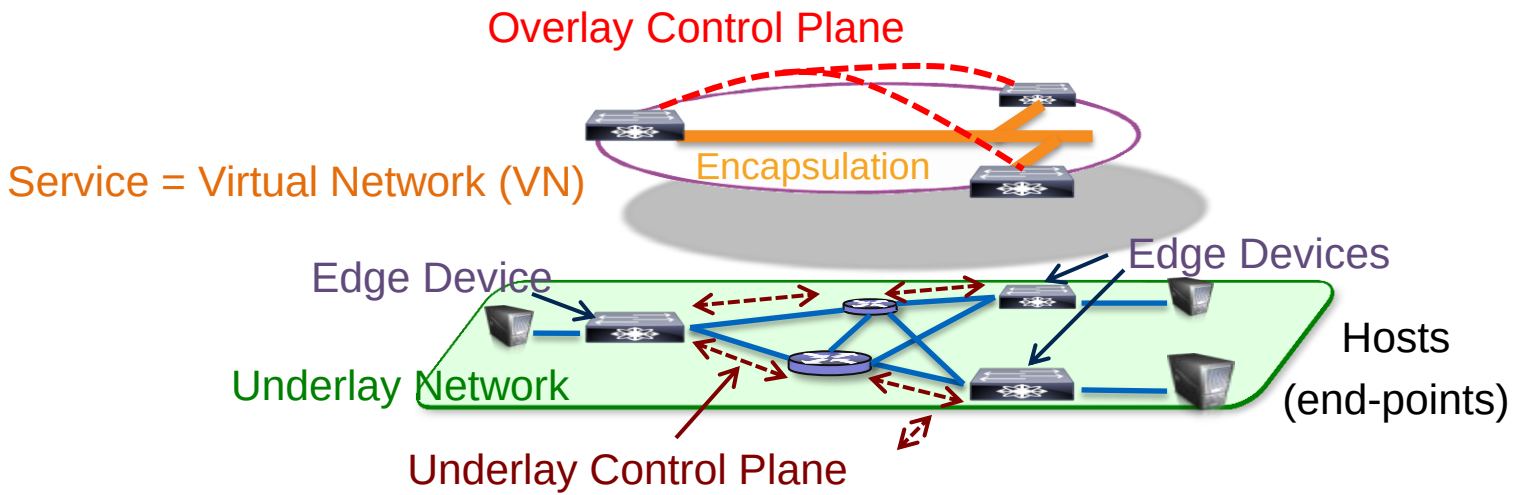

# Why Overlays

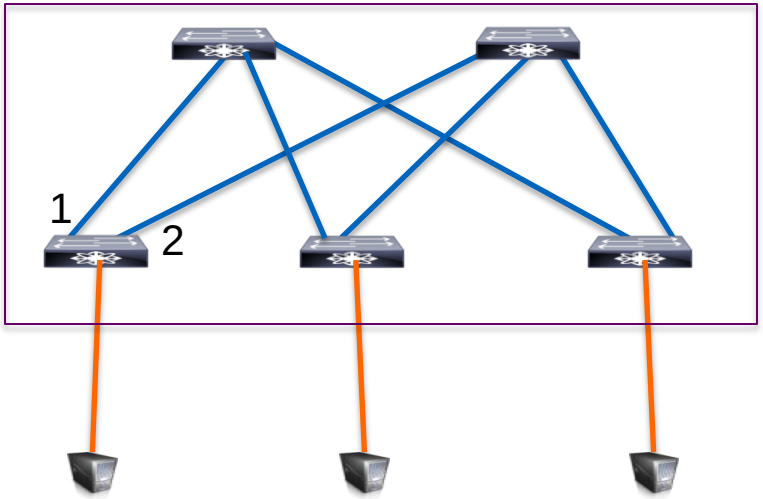


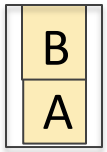


- Split Location from Identity
- Virtualisation
- Logical Networks

# Types of Overlay Service

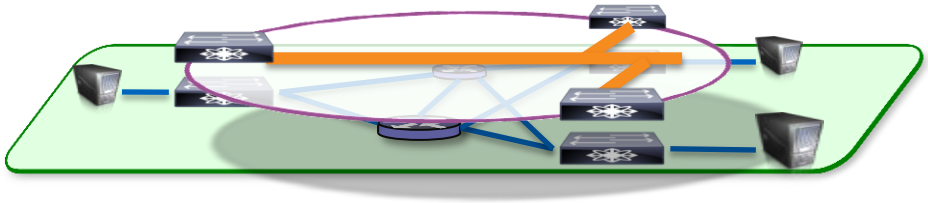

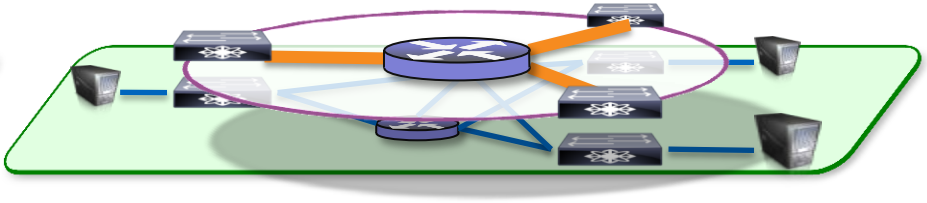

## Layer 2 Overlays

- Emulate a LAN segment
- Transport Ethernet Frames (IP and non-IP)
- Single subnet mobility (L2 domain)
- Exposure to open L2 flooding
- Useful in emulating physical topologies

## Layer 3 Overlays

- Abstract IP based connectivity
- Transport IP Packets
- Full mobility regardless of subnets
- Contain network related failures (floods)
- Useful in abstracting connectivity and policy

# Types of Overlay Edge Devices

## Network Overlays

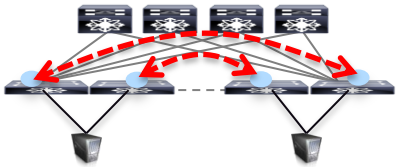

Physical Physical

- Router/switch end-points
- Protocols for resiliency/loops
- Traditional VPNs
- OTV, VXLAN, VPLS, LISP

## Host Overlays

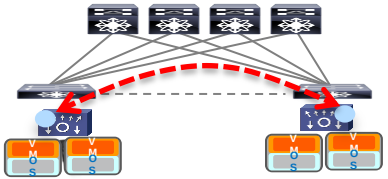


Virtual Virtual

- Virtual end-points only
- Single admin domain
- VXLAN, NVGRE, STT

## Integrated Overlays

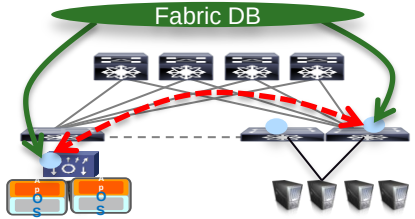


Virtual Physical

- Physical and Virtual
- Resiliency + Scale
- x-organisations/federation
- Open Standards

Tunnel End-points

# Introduction to VXLAN

# VXLAN Overview

Challenges VXLAN addresses:

- VLAN Scalability (4K) – VXLAN extends the L2 Segment ID field to 24-bits, potentially allowing up to 16 million unique L2 Segments over the same network
- VM mobility restricted within a VLAN – VXLAN encapsulates L2 frame in IP-UDP header allowing L2 adjacency across router boundaries

VXLAN Technology Overview:

- MAC-in-UDP encapsulation
- Leverages multicast in the transport network to simulate flooding behaviour for broadcast, unknown unicast and multicast in the layer 2 segment
- Leverage ECMP to achieve optimal path usage over the transport network

# VXLAN Overview

VXLAN can be implemented on both Hypervisor-based Virtual Switches to allow for scalable VM deployments, as well as on Physical switches, which provides the ability to bridge VXLAN segments back into VLAN segments. In these cases, the Physical Switch instantiates a VTEP, and function as a VXLAN Gateway…

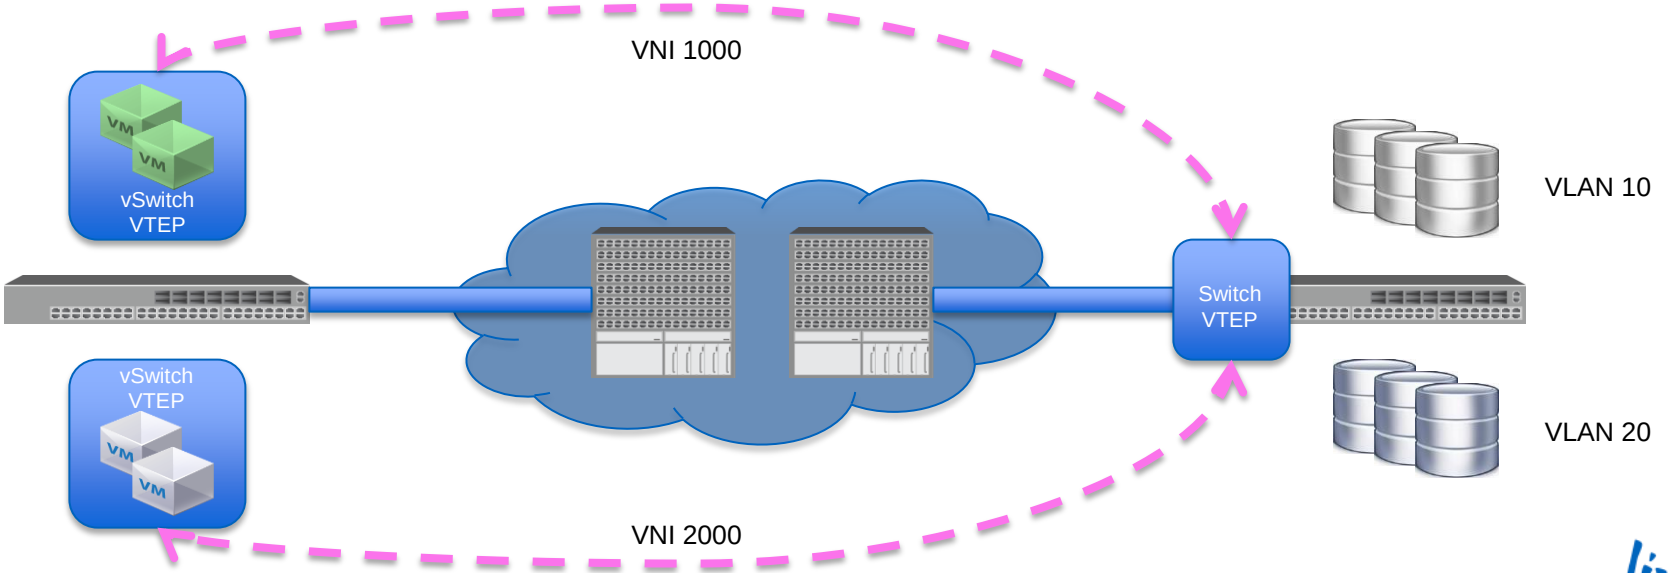

# Virtual eXtensible LAN (VXLAN)

- Virtual eXtensible LAN (VXLAN) is a Layer 2 overlay scheme over a Layer 3 network.
- A 24-bit VXLAN Segment ID or VXLAN Network Identifier (VNI) is included in the encapsulation to provide up to 16M VXLAN segments for traffic isolation / segmentation, in contrast to 4K segments achievable with VLANs.
  - Each of these segments represents a unique Layer 2 broadcast domain, and can be administered in such a way that it can uniquely identify a given tenants address space or subnet.

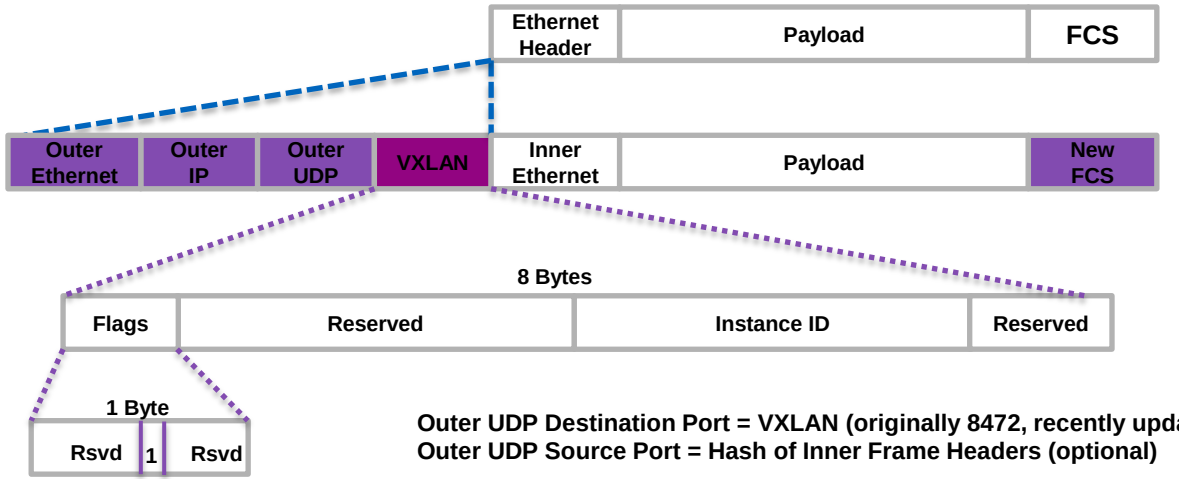


**Outer UDP Destination Port = VXLAN (originally 8472, recently updated to 4789)**
**Outer UDP Source Port = Hash of Inner Frame Headers (optional)**

# VXLAN Frame Format

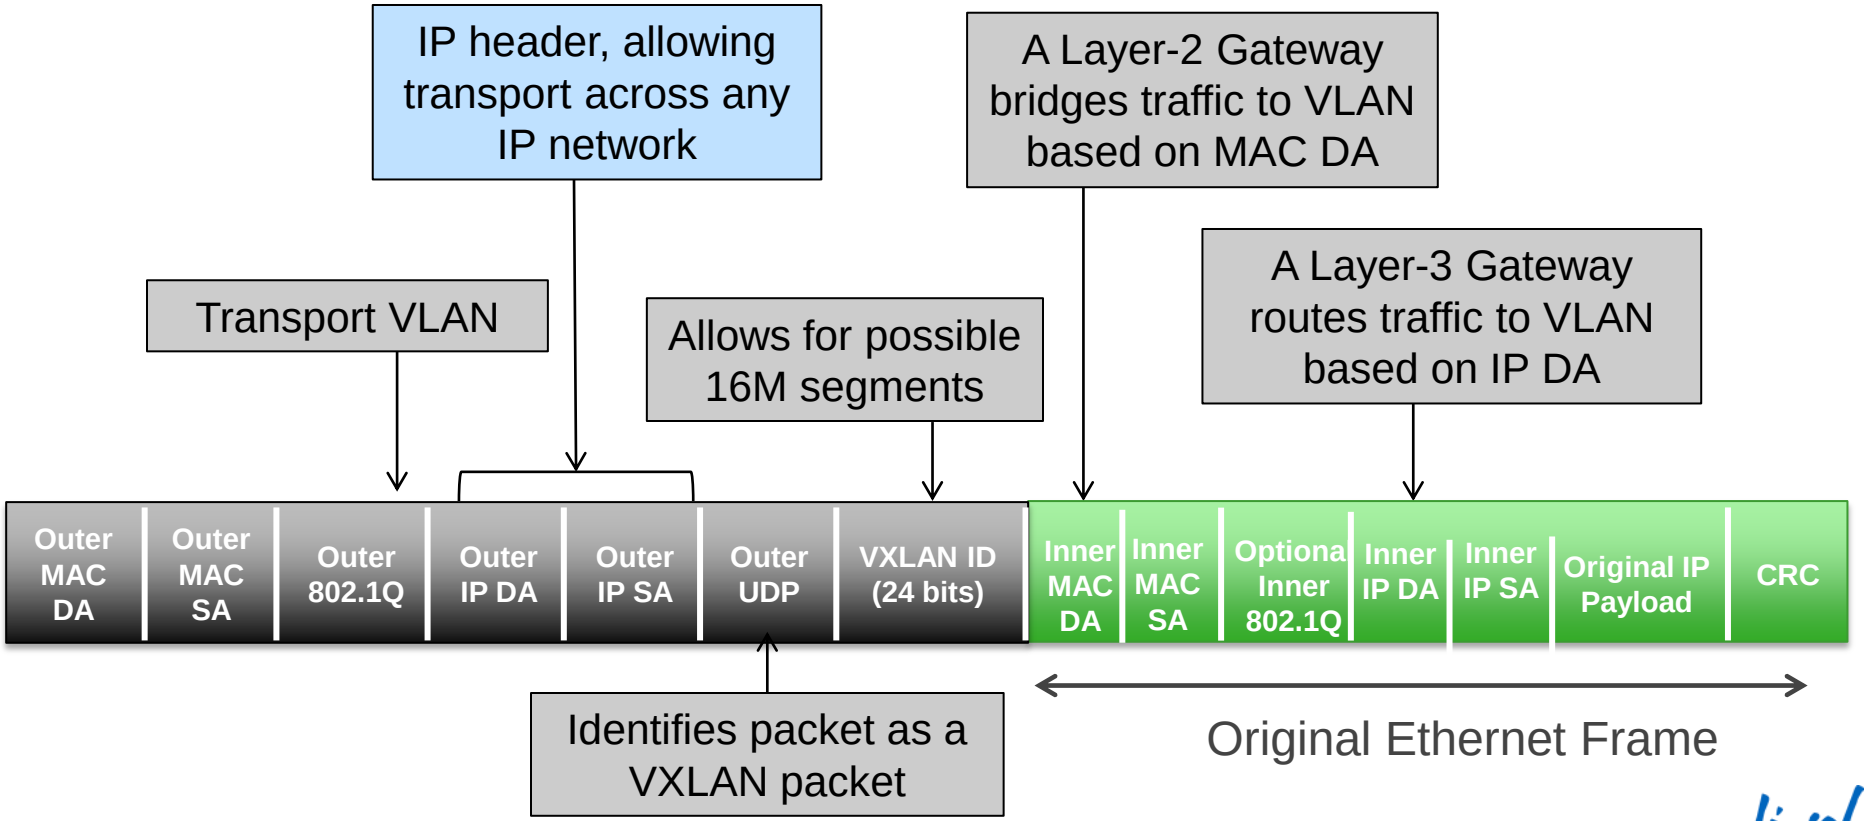

# VXLAN Frame Format (Cont.)

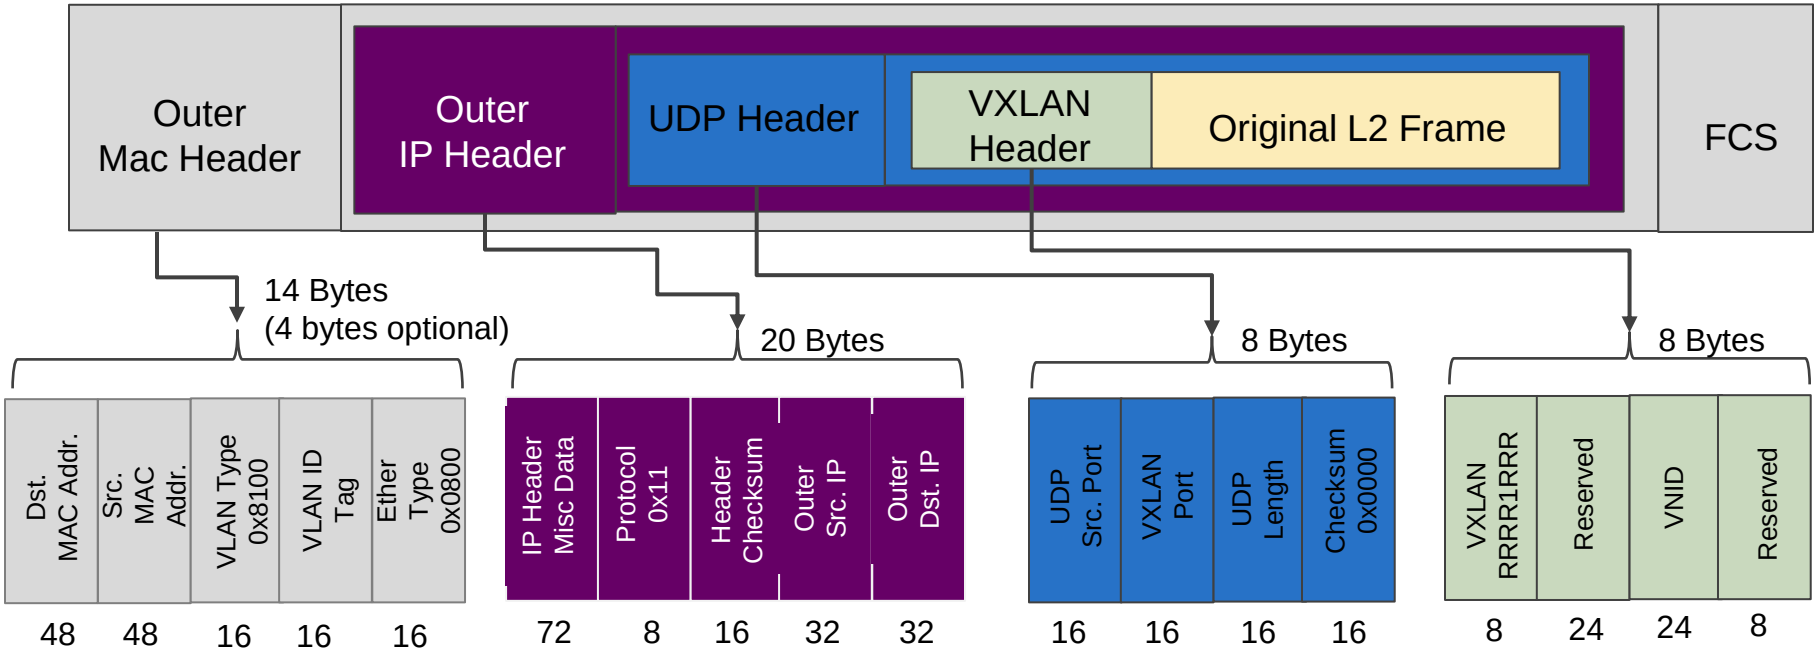

# VXLAN Key Terminology

- VTEP (Virtual Tunnel Endpoint)
- NVE (Network Virtualisation Endpoint)
- VNI (VXLAN Network Identifier or VXLAN Segment ID)
- VXLAN Gateway
- Transit
- Remote VTEP
- Delivery Group (DG)
- BUM

# VXLAN VTEP

VXLAN terminates its tunnels on VTEPs (Virtual Tunnel End Point).

Each VTEP has two interfaces - one to provide bridging function for local hosts, the other has an IP identification in the core network for VxLAN encapsulation/de-encapsulation.

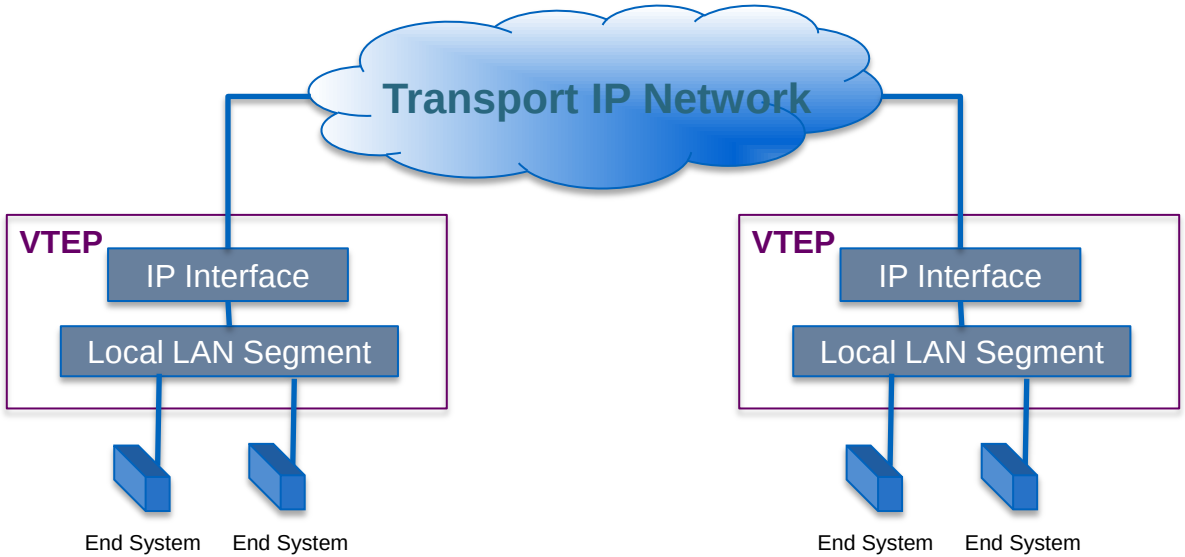

# Handling of Multi-Destination Traffic

Since a control/signalling protocol has not been defined, emulation of Multi-Destination traffic (Broadcast, Multicast, Unknown Unicast) is handled through the VXLAN IP underlay through the use of segment control multicast groups…

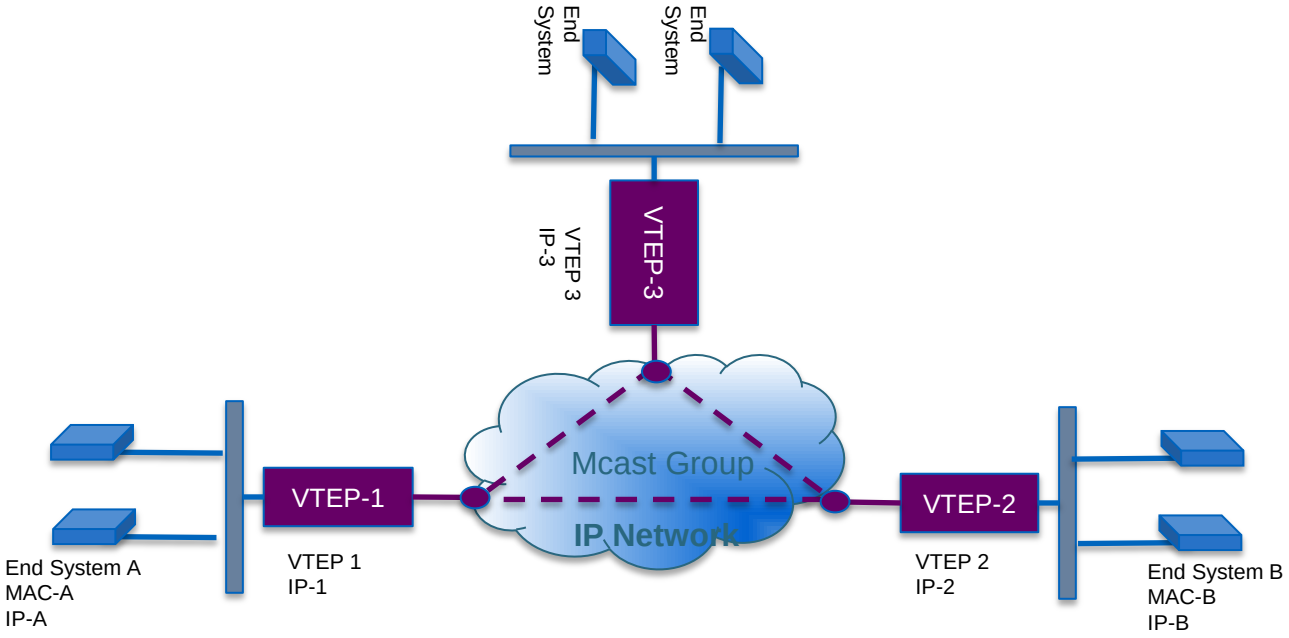

# VXLAN Overview
## VTEP Discovery & Address Learning

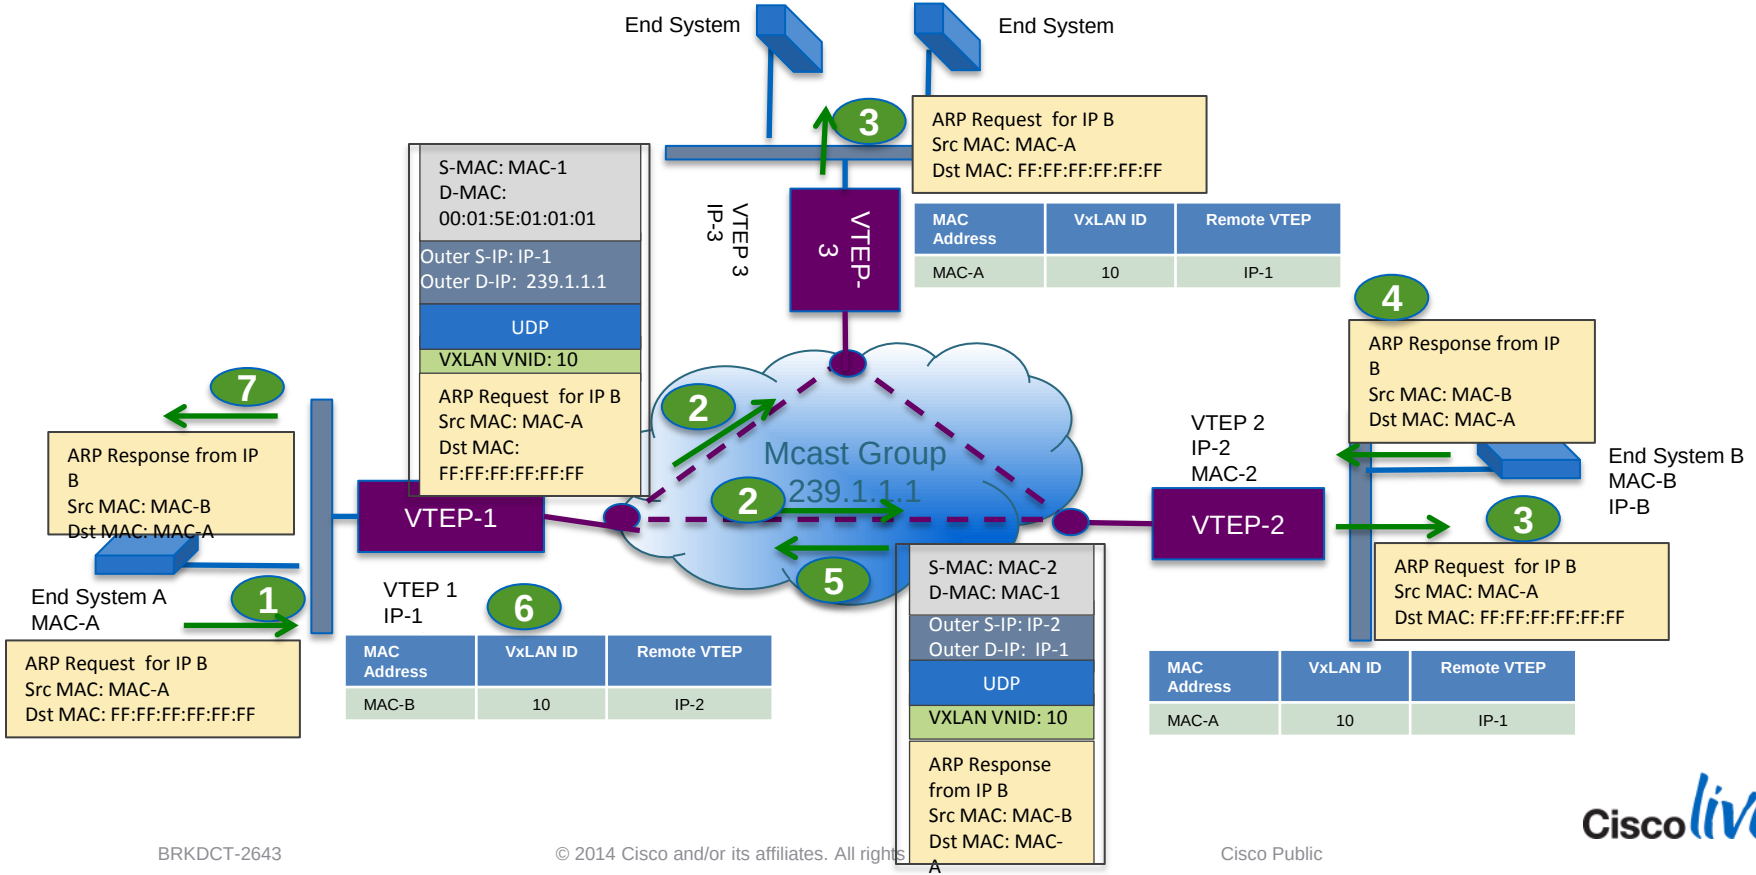

# VXLAN Overview

## Unicast Forwarding Packet Flow

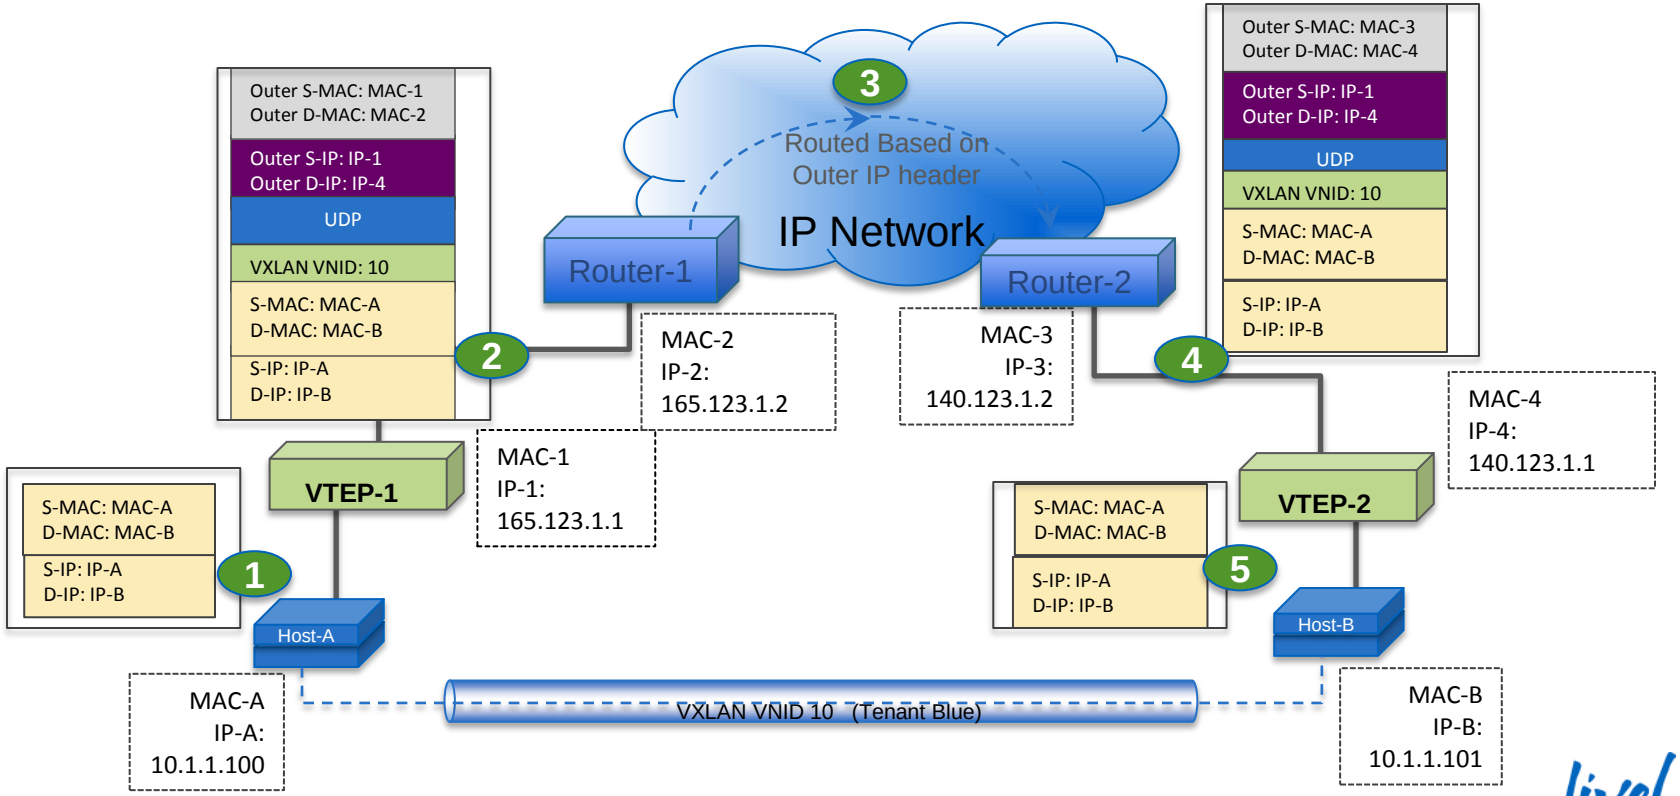

# Nexus 9000 Series
## VXLAN Support

VXLAN is supported across the Nexus 9000 series platforms. The VXLAN Gateway functionality is supported across all form factors and line cards. Integrated routing functionality is only supported on ACI-enabled Modules…

**Nexus 9300 Series**

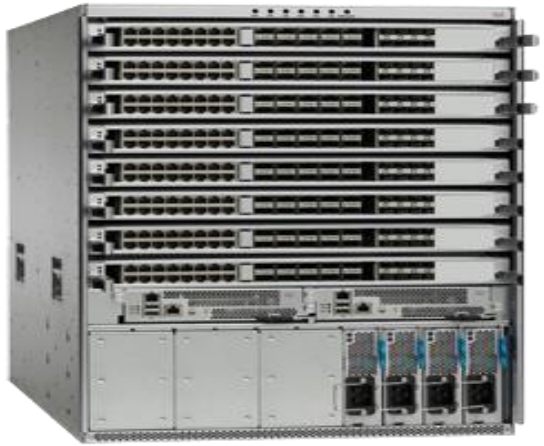

**Nexus 9500 Series**

# Modular Switch Platform – Nexus 9500

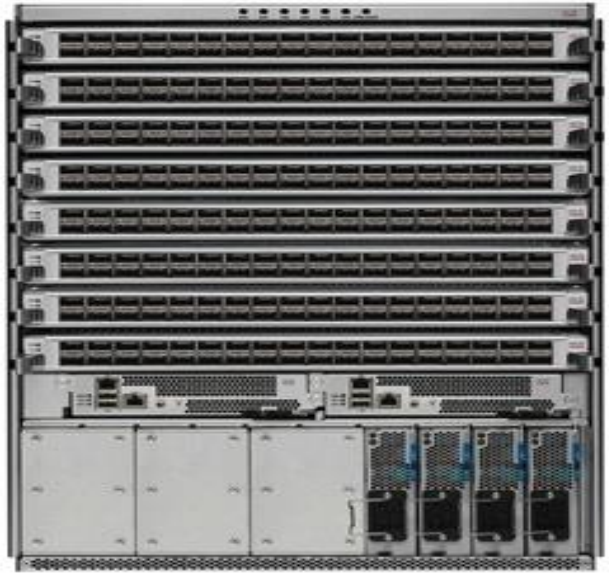

**Nexus 9508**

- 13 RU high
- 30Tbps fabric today
- Up to 288p 40G & 1,152p 10G
- Headroom for 100G densities (connectors, power)
- Supervisors w/ quad core CPU and default 64GB SSD

## ACI Ready Linecards

48 ports 10G SFP+ & 4 ports 40G QSFP+

48 ports 1/10G-T & 4 ports 40G QSFP+
(non blocking)

36 ports 40G QSFP+
(1.5:1 oversubscribed)

## Unified Fabric Linecards

36 ports 40G QSFP+

# Fixed Switch Platform – Nexus 9300

## Nexus 9396PQ

- 48 port 10G SFP+ & 12 port 40G QSFP+
- 2 RU

## Nexus 93128TX

- 96 port 1/10G-T & 8 port 40G QSFP+
- 3 RU

## ACI Ready Uplink Module

- 12 port 40G QSFP+
- Additional 40MB buffer
- Full VXLAN Bridging & Routing Capability

## Nexus 9300 - Common

- Redundant FAN and Power Supply
- Front-to-back and Back-to-Front airflow
- Dual or Quad Core CPU with default 64GB SDD

# Enhanced NX-OS
## Purpose-Built Data Centre OS

- **Modern:** 64-bit Linux 3.4.10 Kernel
- **Comprehensive:** Purpose-built Data Centre Feature Set, including L2/L3/VXLAN
- **Modular:** Code runs in DRAM only when invoked
- **Fault Containment:** Complete process isolation for both features & services
- **Resiliency:** Restartable user-space network stack & drivers
- **Management Infrastructure:** CLI, SNMP, NetConf/XML, OnePK, Open Containers, JSON

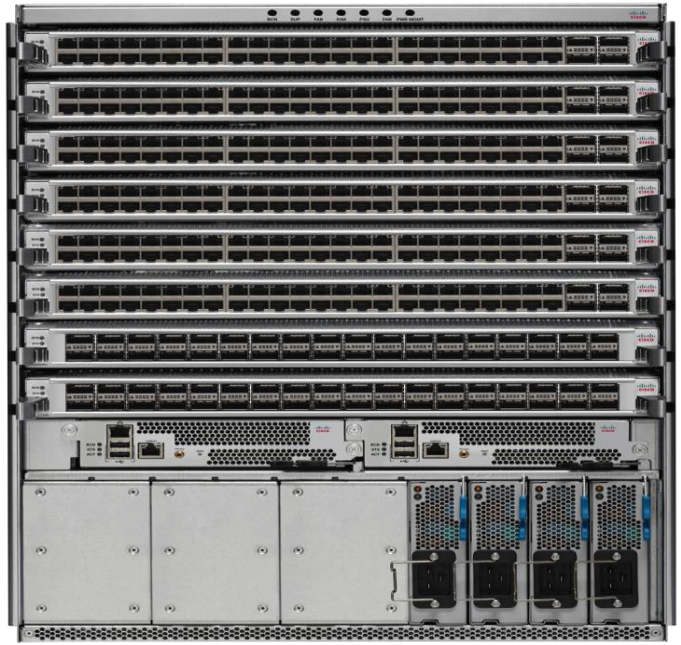

# Supported Platforms

| Platform | NX-OS Version | |
|---|---|---|
| | Minimum | Recommended |
| Nexus 9500 | 6.1.2I3.1.x | 6.1.2.I3.1.x |
| Nexus 9300 | 6.1.2I2.1.x | 6.1.2.I2.1.x |
| Nexus 3100 (3132/3172) | 6.x | Q2 CY14 |
| Nexus 6000 | 7.0(0)N1x | Q3 CY14 |
| Nexus 5600 | 7.0(0)N1x | Q4 CY14 |
| Nexus 7000 with F3 | 7.0.x | Q4 CY14 |

* There is no licensing cost for VXLAN – Enhance Layer 3

VXLAN gateway bridges traffic between VXLAN segment and another physical / logical layer 2 domain (such as a VLAN)…

| VLAN ID | VXLAN ID |
|---|---|
| 10 | 1010 |
| 20 | 1020 |

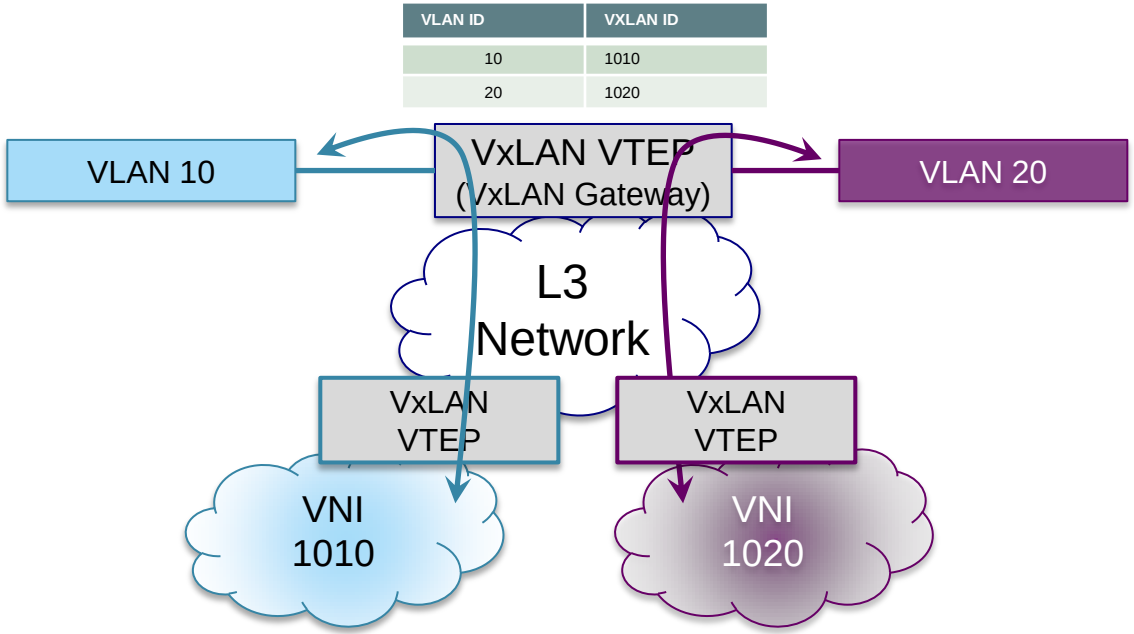

# Nexus 9000 Series
## VXLAN Gateway

The Nexus 9000 series supports VXLAN Gateway function, allowing VLANs to be bridged/mapped to VXLAN Segments and vice versa…

```
feature nv overlay
feature vn-segment-vlan-based

interface et4/13
  switchport
  switchport access vlan 10
  no shut
interface nve1
  no shutdown
  source-interface loopback0
  overlay-encapsulation vxlan
  member vni 1010 mcast-group 230.1.1.1

vlan 10
  vn-segment 1010
```

```
switch# show nve vni
Interface          VNI        Multicast-group   VNI State
-----------------  --------   ---------------   ---------
nve1               1010        230.1.1.1            up
switch# show nve peers
Interface           Peer-IP              VNI         Up Time
-------------       ----------------     --------    -------
nve1                10.1.1.2             1010       00:52:24
switch#
```

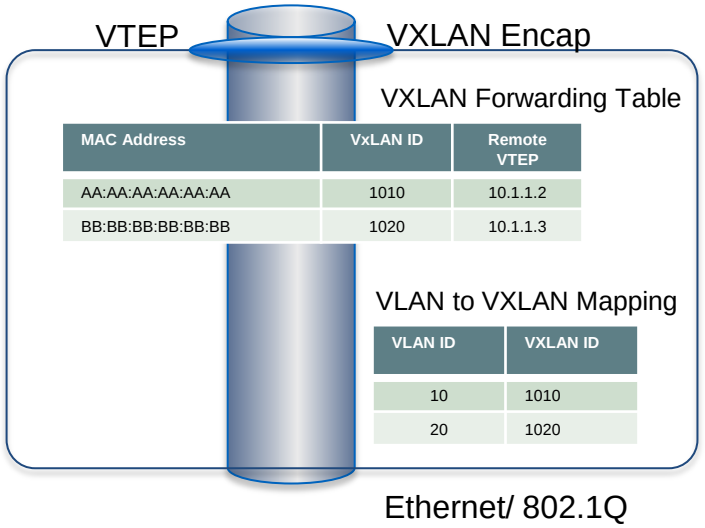

VXLAN “Bridging” bridges traffic between VXLAN segments

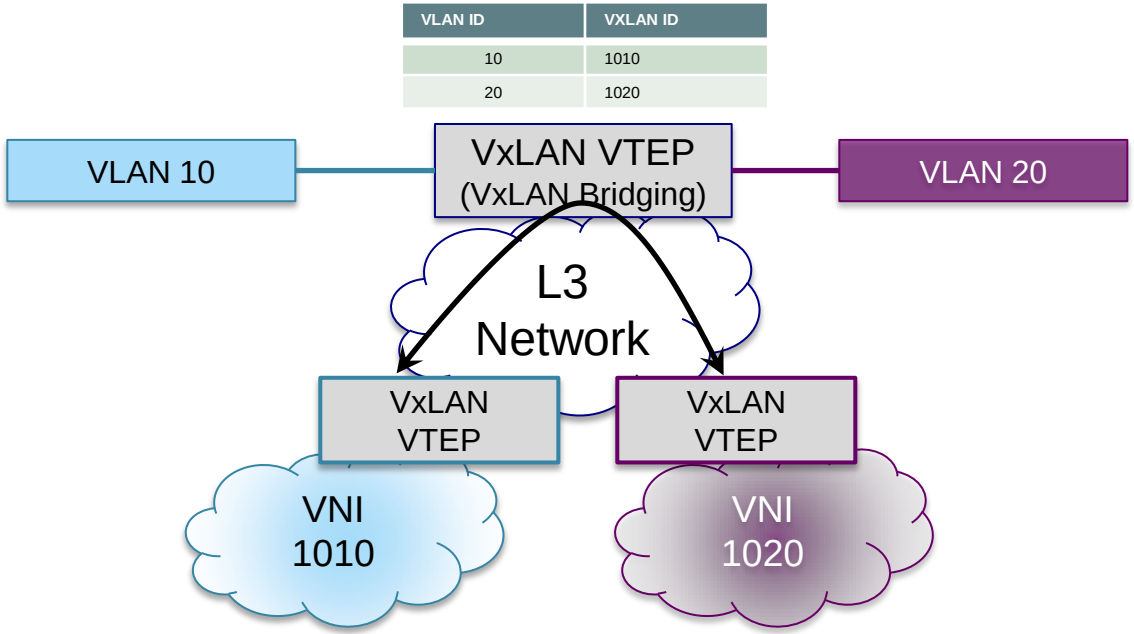

# VXLAN Forwarding in 9300

## VXLAN Bridging and Gateway

- VXLAN Encapsulation and De-encapsulation occur on NFE
  - VLAN -> VXLAN (gateway)
  - VXLAN - VLAN (gateway)
  - VXLAN – VXLAN (bridging)
- Bridging and Gateway are independent of the port type (1/10/40G ports)

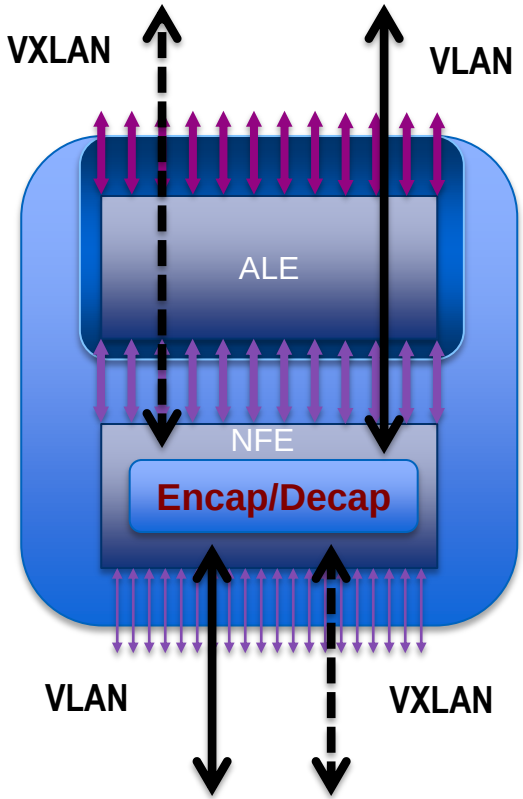

# VXLAN Forwarding – Nexus 9500

## VXLAN Bridging and Gateway

- VXLAN Encapsulation and De-encapsulation occur on NFE
  - VLAN -> VXLAN (gateway)
  - VXLAN - VLAN (gateway)
  - VXLAN – VXLAN (bridging)

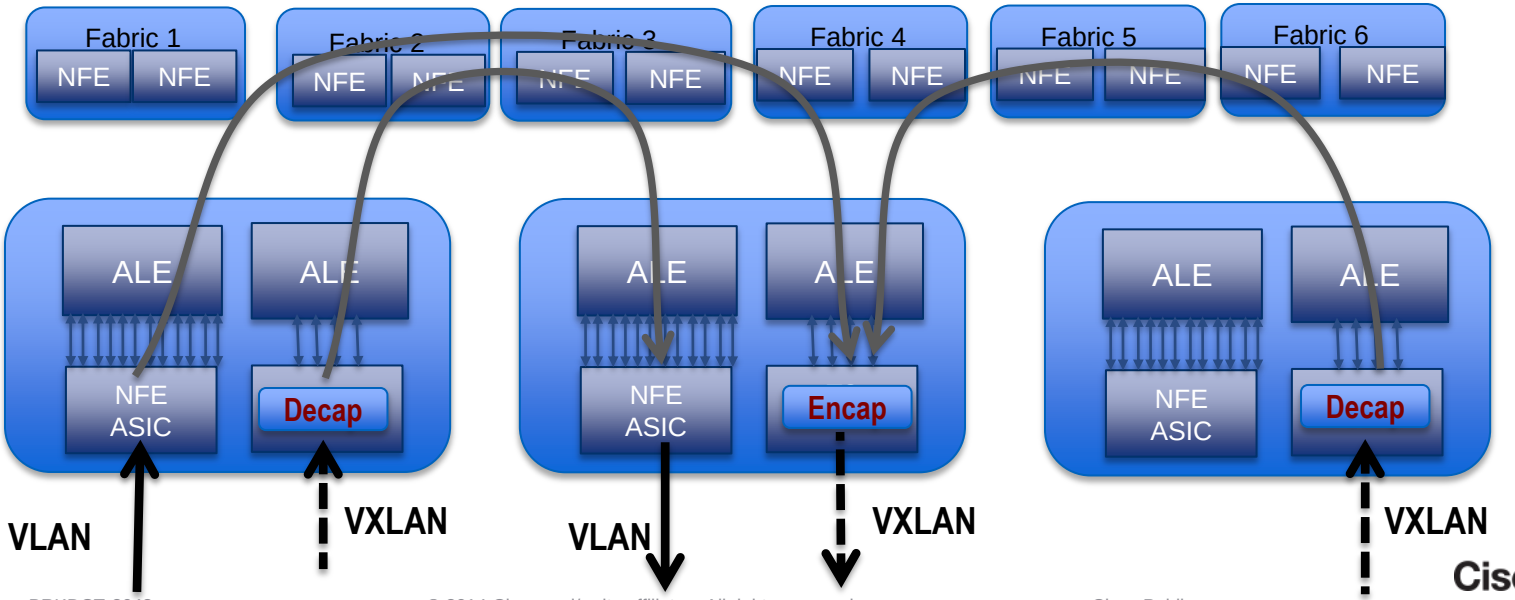

# Nexus 9000 Series
## VXLAN Routed Mode

VXLAN routed mode ‘routes’ traffic between VXLAN segments and between VXLAN another physical / logical layer 2 domain (such as a VLAN)…

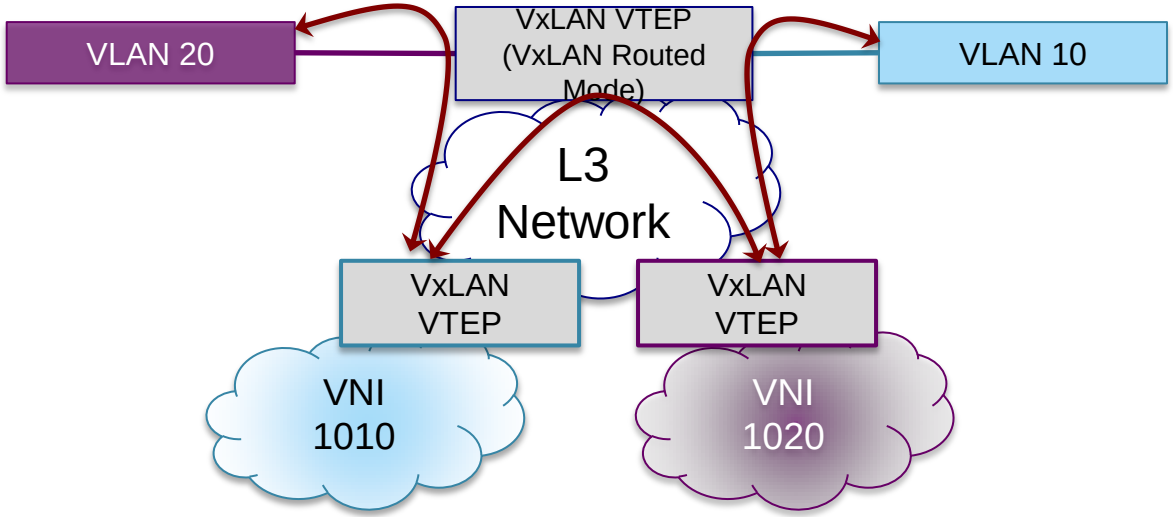

# VXLAN Forwarding – Nexus 9000

## VXLAN Routing

- VXLAN Routing is not supported currently on Packet Processor
- Additional recirculation required for VXLAN routing through Extended Packet Processor

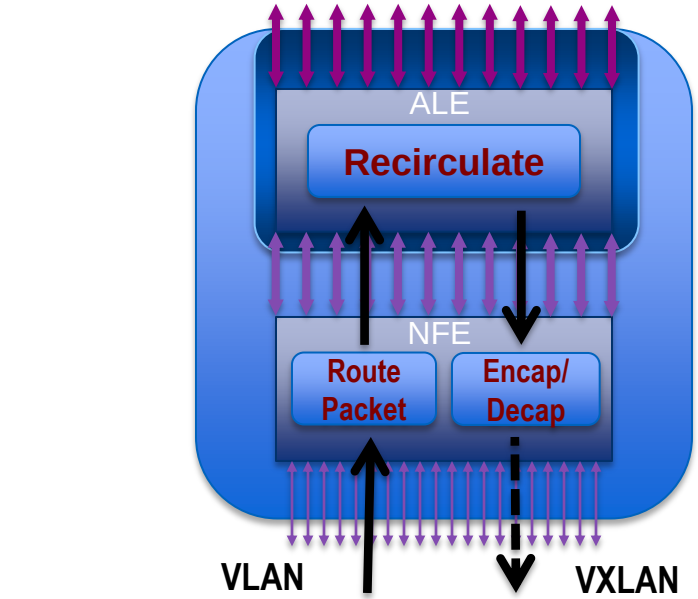


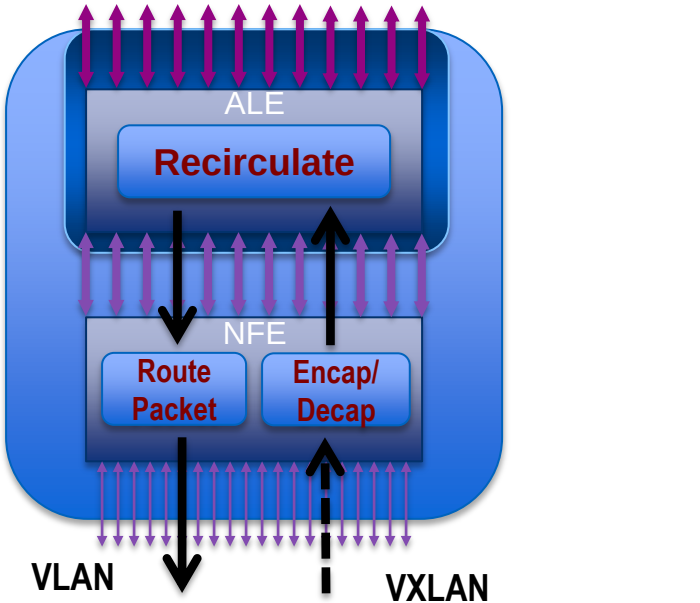

# VXLAN Forwarding – Nexus 9500

## VXLAN Routing

- VXLAN Encapsulation and De-encapsulation occur on NFE
  - Decap on ingress line card, Encap on egress
- Routing occurs on the ingress NFE and 'if' required on fabric NFE
- Routed VXLAN is recirculated by ALE ASIC

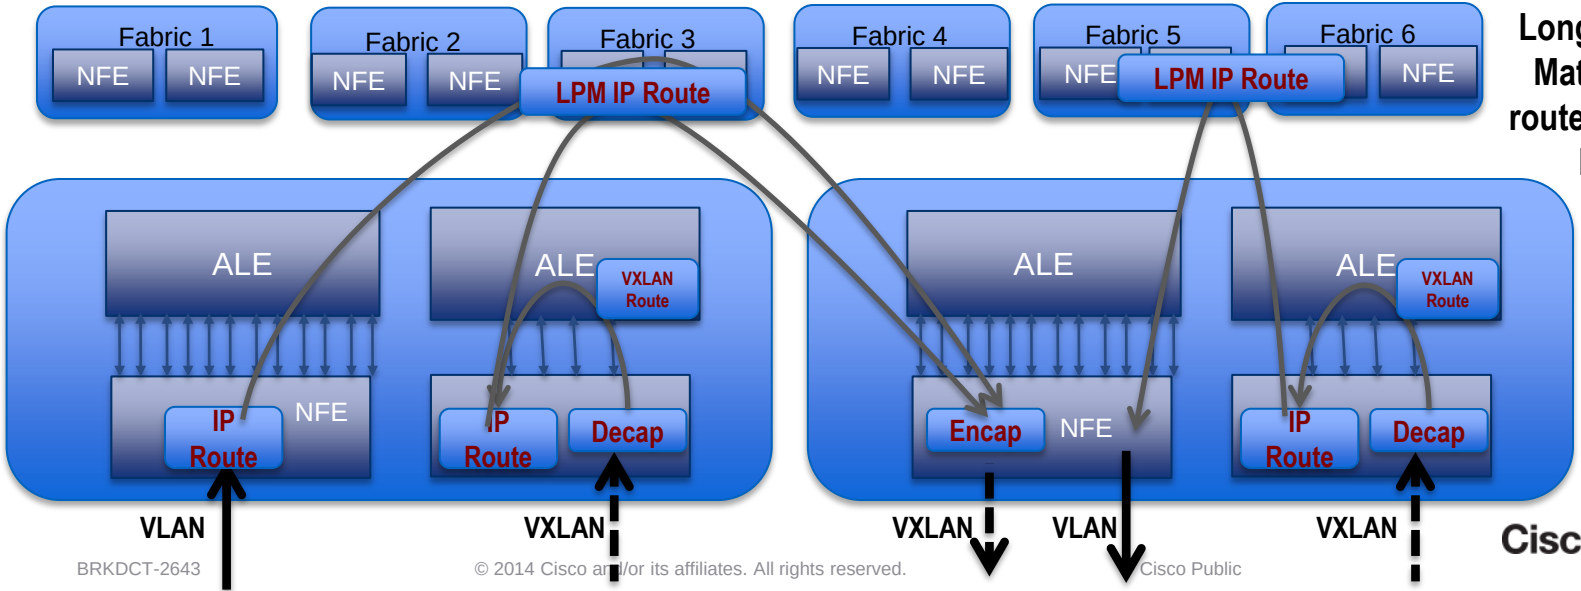




# VXLAN Design

# VXLAN Designs

- Explore a variety of VXLAN designs and evaluate how they meet key design criteria
- Assumption is the choice to go Overlay/VXLAN has already been made
  - Objective is not to debate "Layer 2 vs. Layer 3"  or "why overlay in the Data Centre?"
- Introduce concepts / design building-blocks to help you build a design that meets your requirements

# VXLAN Designs

High-level design options considered in this presentation are in the following areas:

- Routed Access + IP Mobility
- L2 extension across Pod / Multi-tenancy
- Data Centre Interconnect (DCI)

# Classic Design

- Simplest design option
- Extension of traditional aggregation/access designs in spine / leaf

Immediate benefits:

- Simplified configuration
- Extension of VLANS
- Traffic distribution using ECMP
- IP Mobility

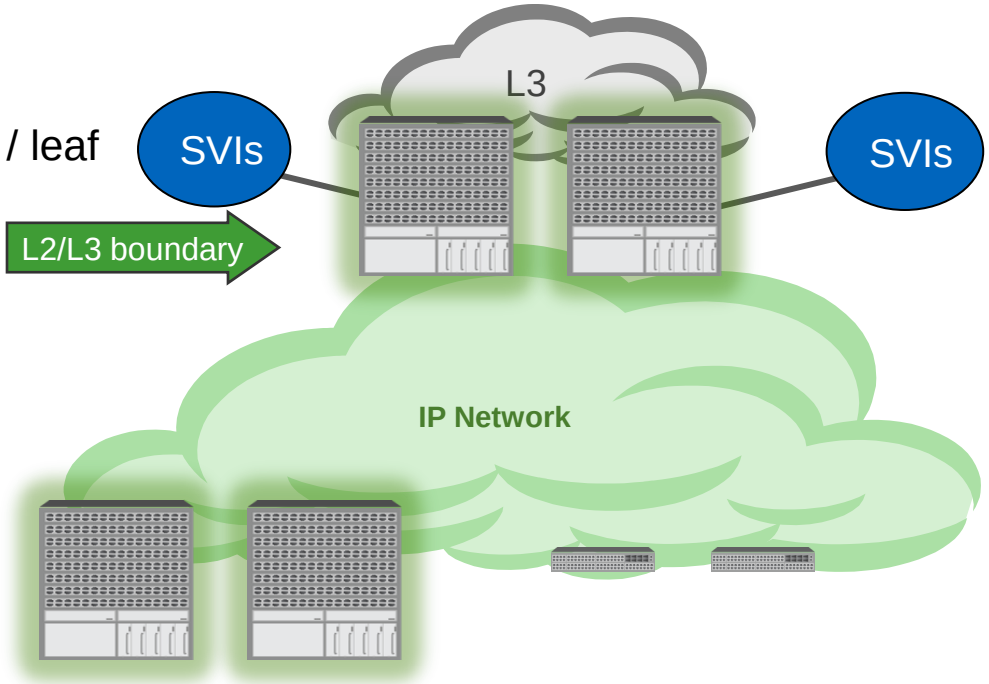

# VXLAN Gateway – Routed Access + IP Mobility

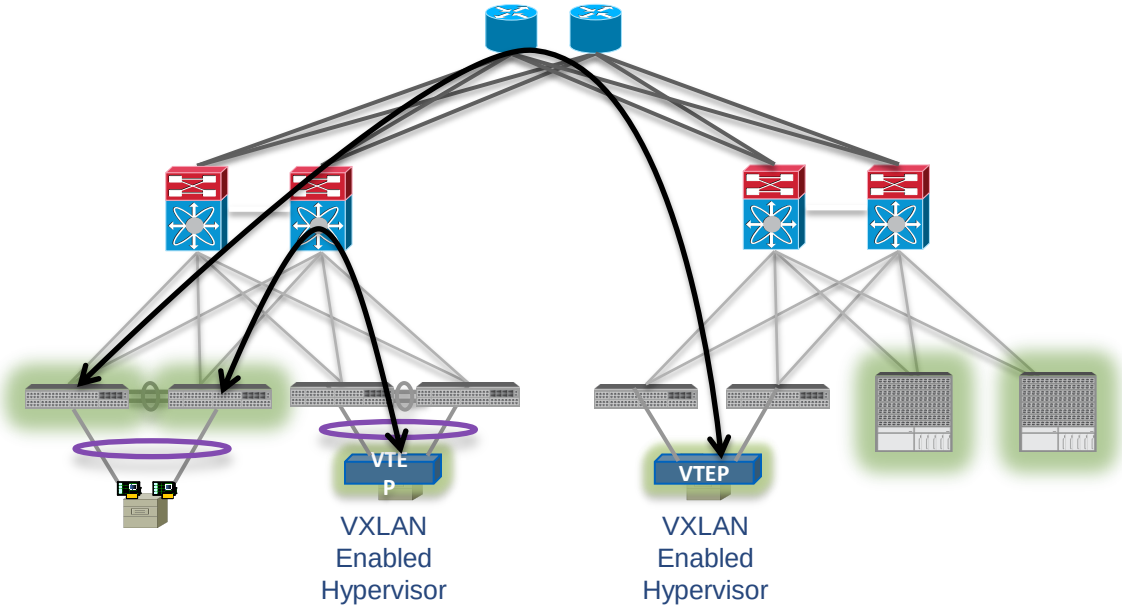


- VXLAN Gateway defined at access layer (leaf) – Nexus 9000
- Multicast needs to be enabled for VXLAN to work on the source interface
- Next hop of VTEP needs to be Layer 3
- vPC needs peer gateway
- Only 1:1 mapping is allowed for VXLAN to VLAN
- Recommended N9K to be configured as STP root switch in each L2 network
- Link discovery protocols like CDP, LLDP will not discover neighbours on the remote VTEPs
- Virtual to physical migration (P2V)

# VXLAN Forwarding
## vPC

- When vPC is enabled an ‘anycast’ VTEP address is programmed on both vPC peers
- Symmetrical forwarding behaviour on both peers provides
- Multicast topology prevents BUM traffic being sent to the same IP address across the L3 network (prevents duplication of flooded packets)
- vPC peer-gateway feature must be enabled on both peers
- VXLAN header is ‘not’ carried on the vPC Peer link (MCT link)

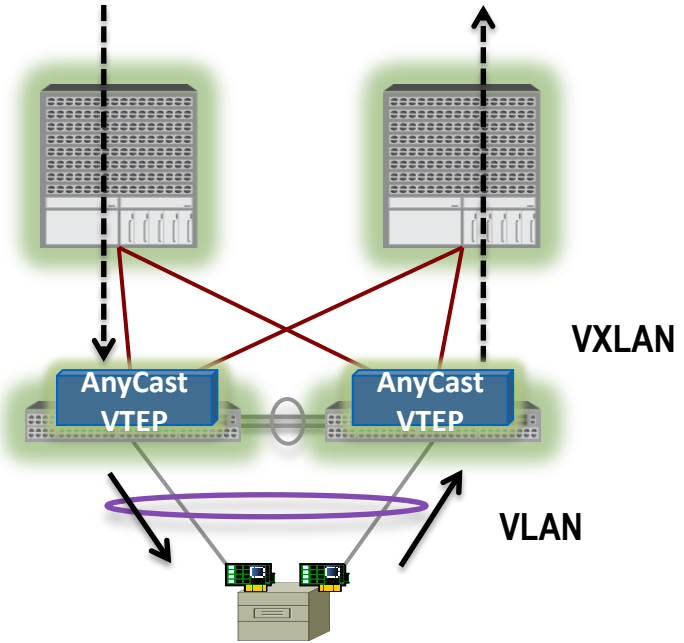

# VXLAN Forwarding
## Design Considerations

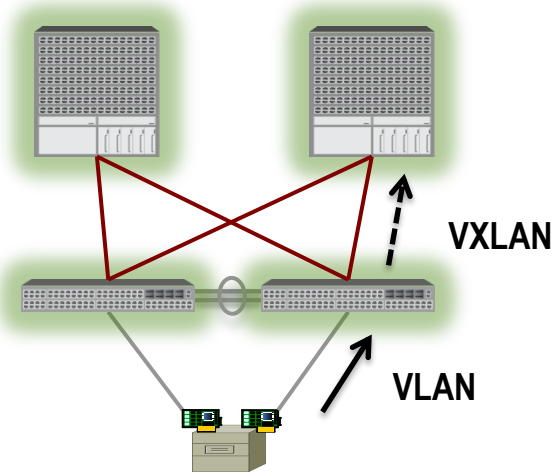


When VXLAN is being routed the next hop for VXLAN encapsulated frames needs to be over an L3 interface

Alternatively, all SVIs from a VXLAN Gateway must point to the same physical next hop

[same VXLAN header MAC DA for all VXLAN encapsulated packets sent from the same physical port]

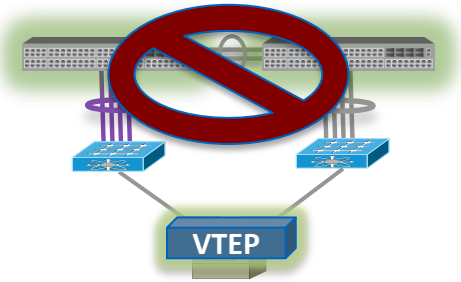


VXLAN VTEP downstream of a Nexus 2000 FEX is not supported

# VXLAN Design with VXLAN Bridging Only

## L2 Extension across Pods

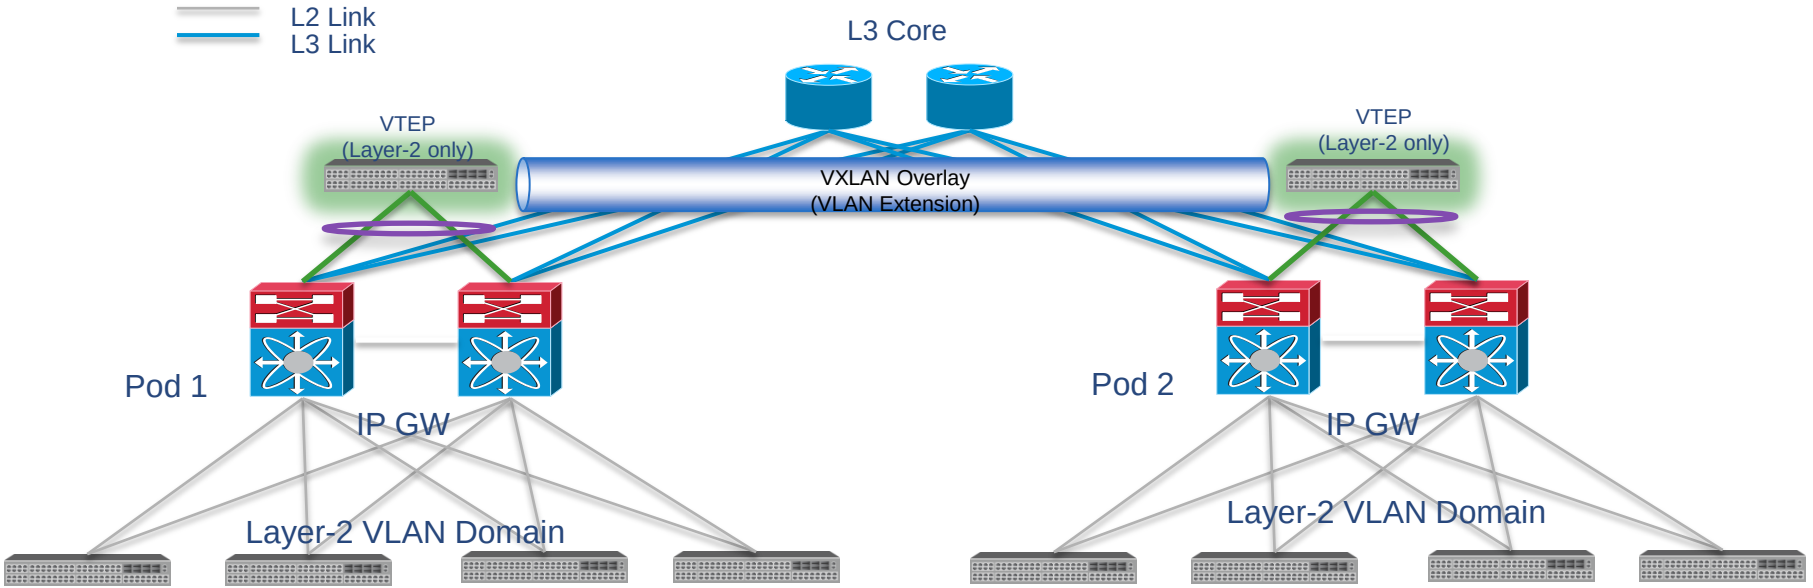

# VXLAN Design with VXLAN Bridging Only

## STP/VPC Replacement - Routing off-box

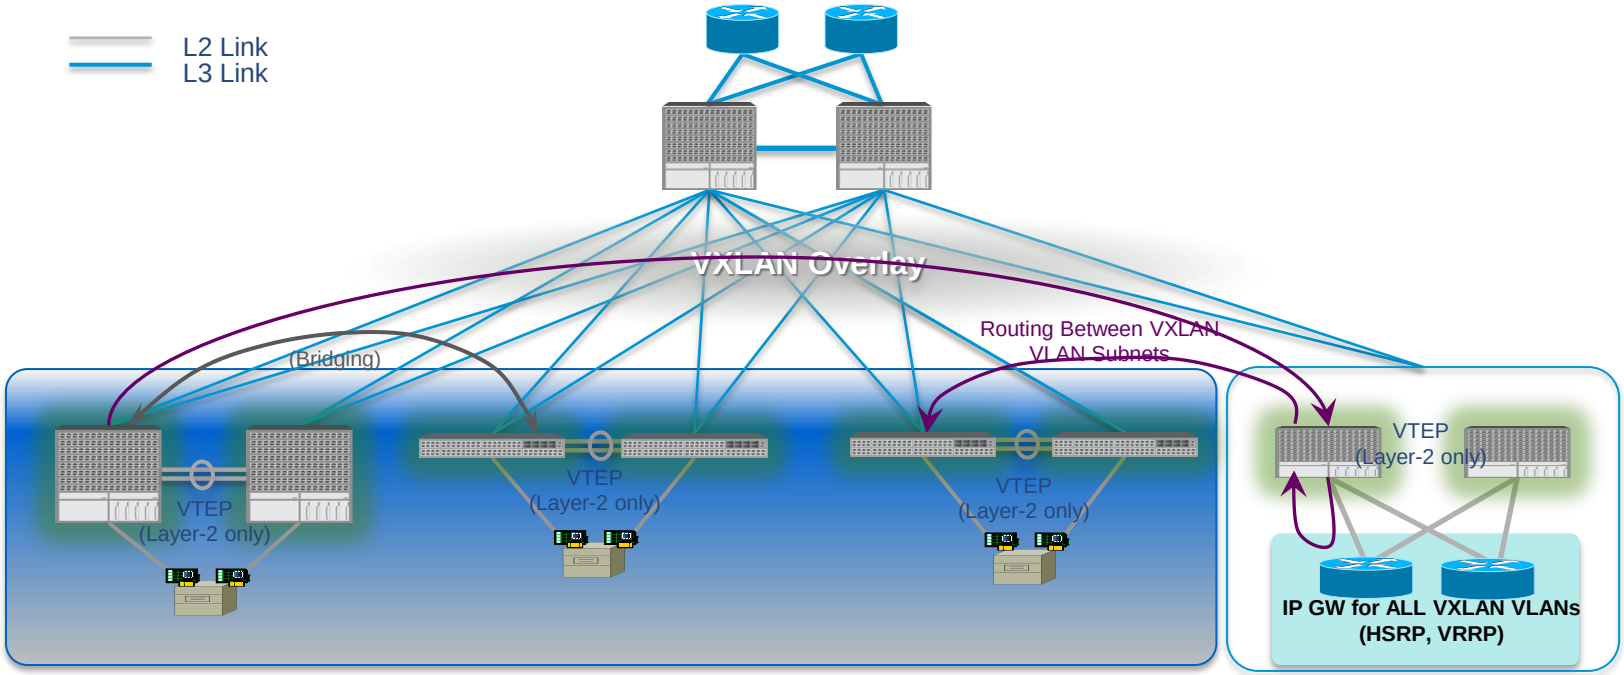

# VXLAN Design with VXLAN Bridging Only

## Spine-Leaf Deployment

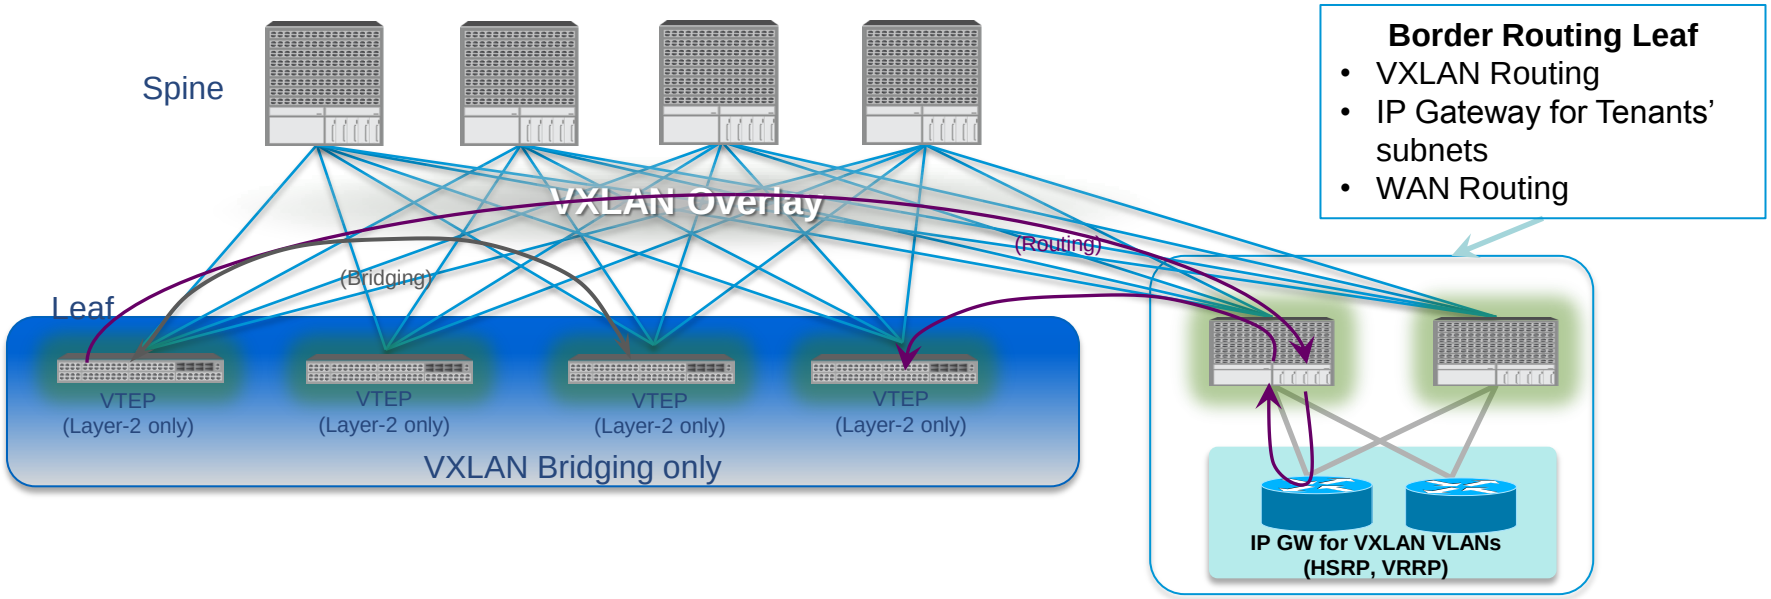

# VXLAN Design with VXLAN Bridging + Routing

## STP/VPC Replacement

### Option A: Centralised IP Gateway and Inter-VXLAN Routing

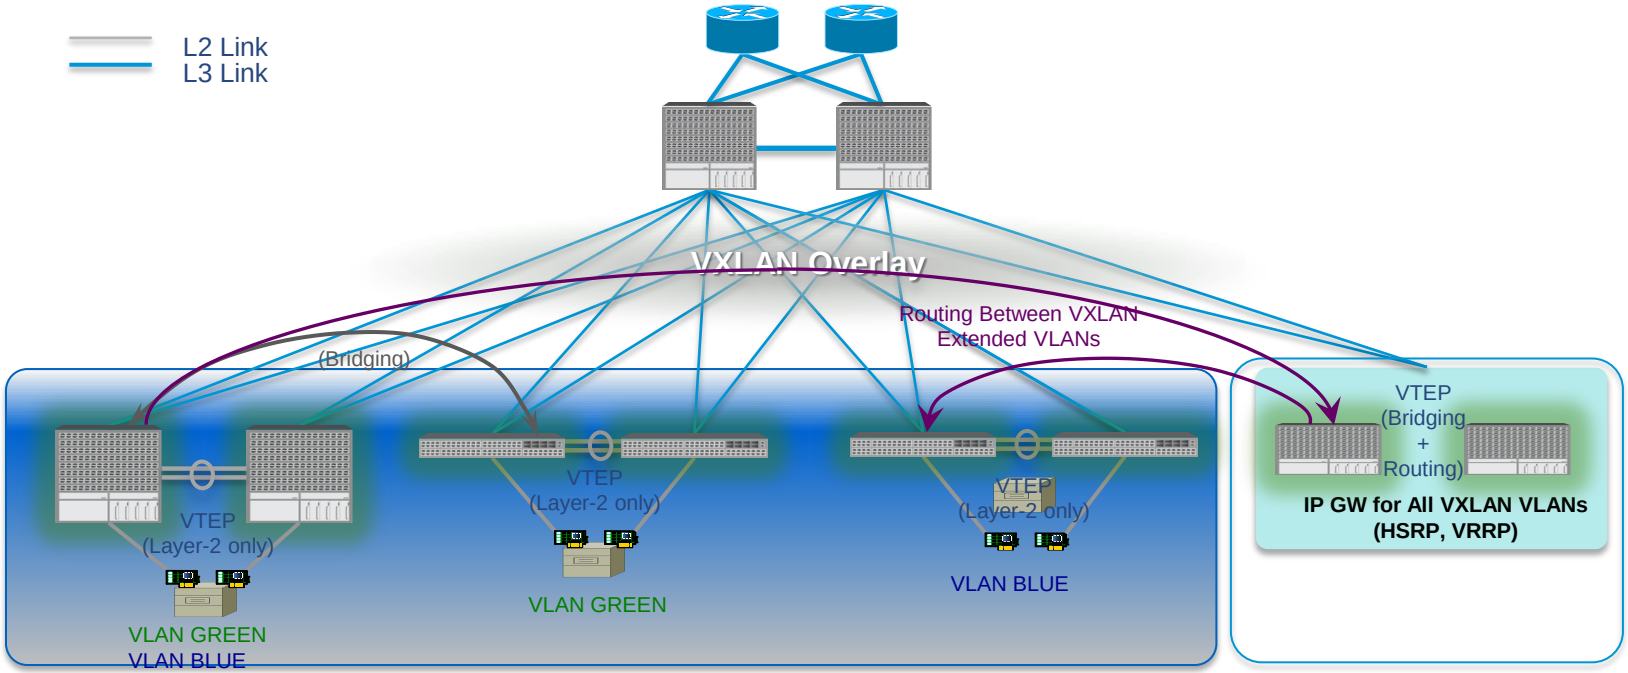

# VXLAN Design with VXLAN Bridging + Routing
## STP/VPC Replacement

### Option B: Distributed IP Gateway and Inter-VXLAN Routing

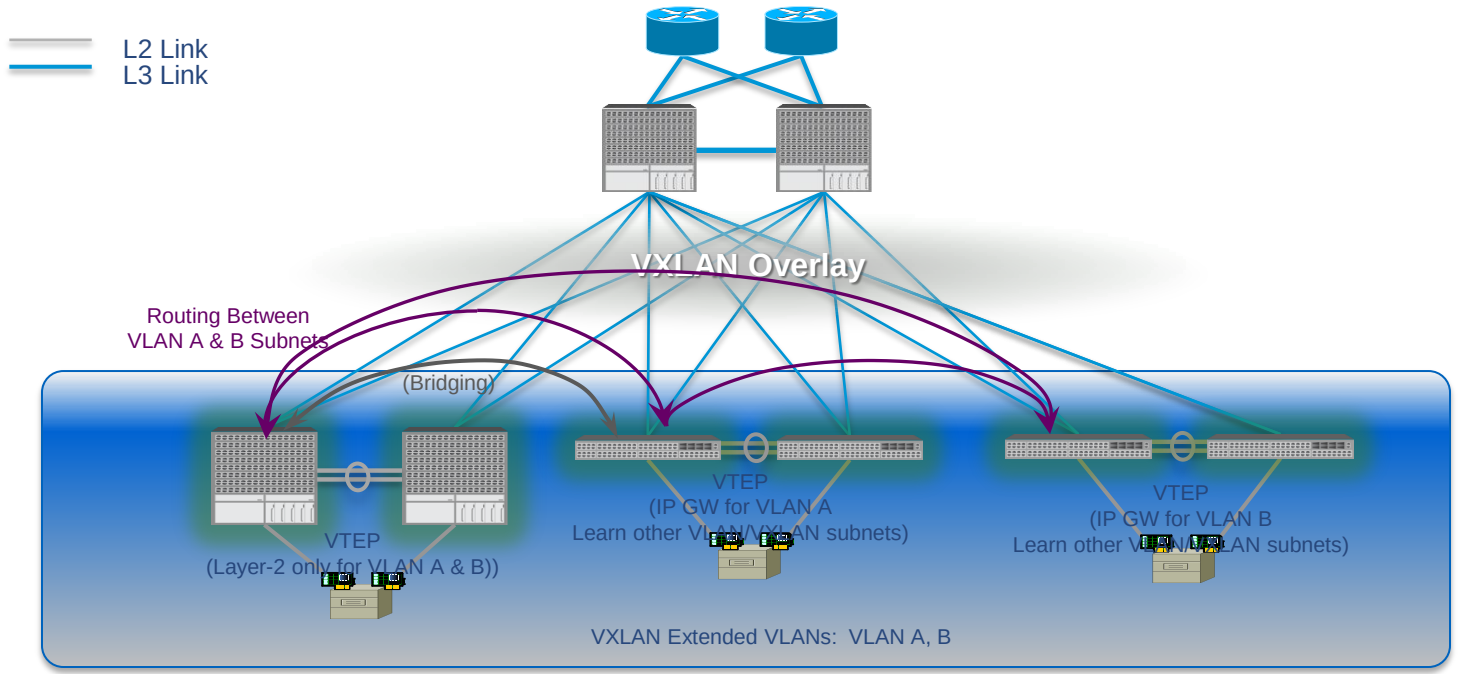

# VXLAN Design with VXLAN Bridging + Routing
## Spine-Leaf Deployment

### Option A: Centralised IP Gateway and Inter-VXLAN Routing

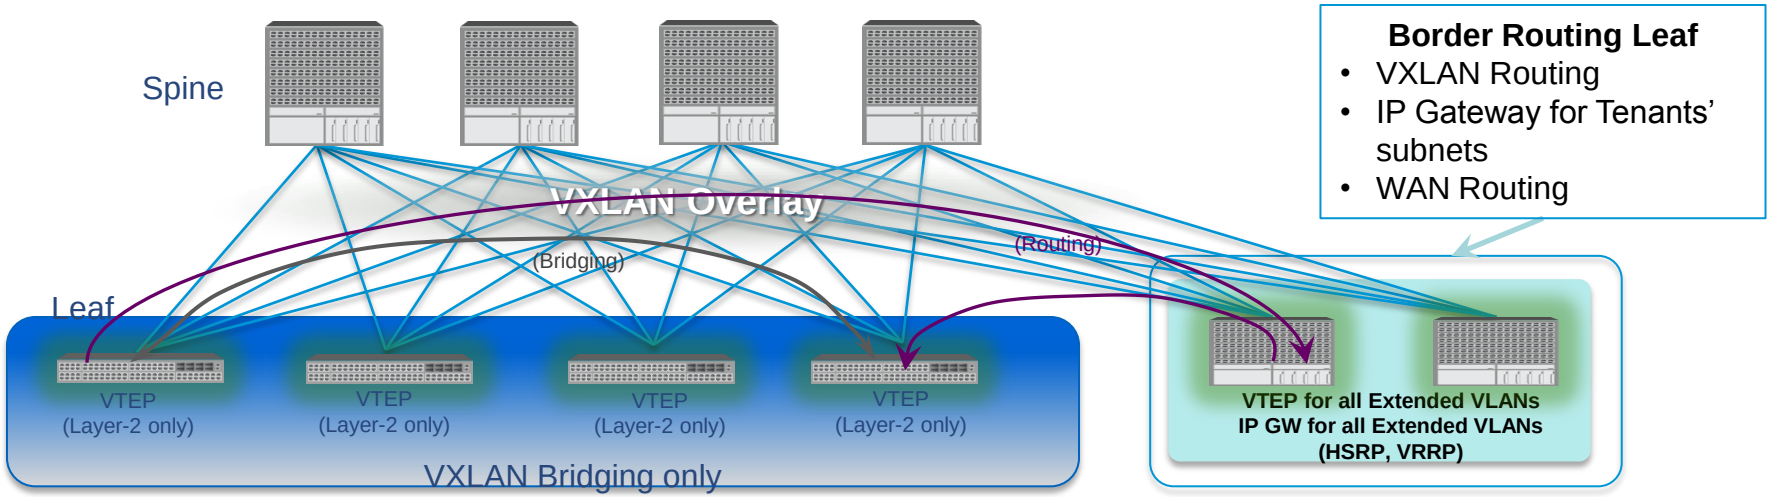

# VXLAN Design with VXLAN Bridging + Routing
## Spine-Leaf Deployment

### Option B: Distributed IP Gateway and Inter-VXLAN Routing

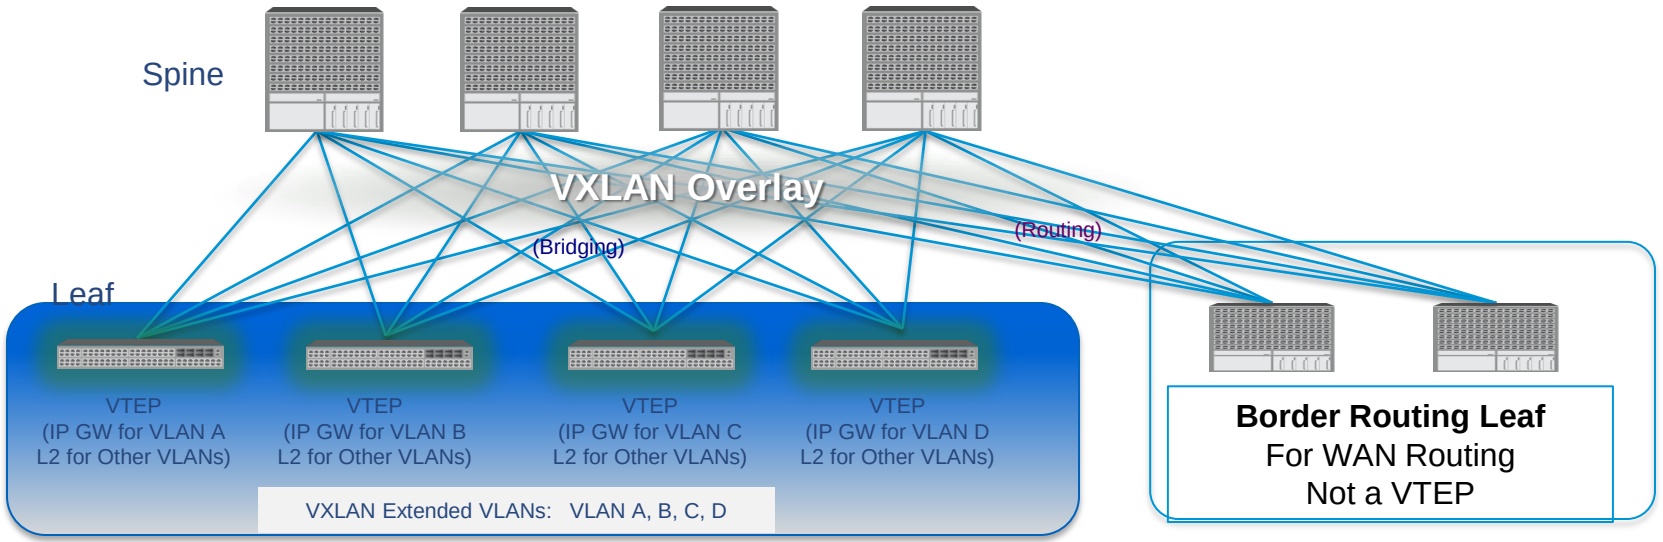

# VXLAN Design with VXLAN Bridging + Routing + Anycast IP GW - STP/VPC Replacement

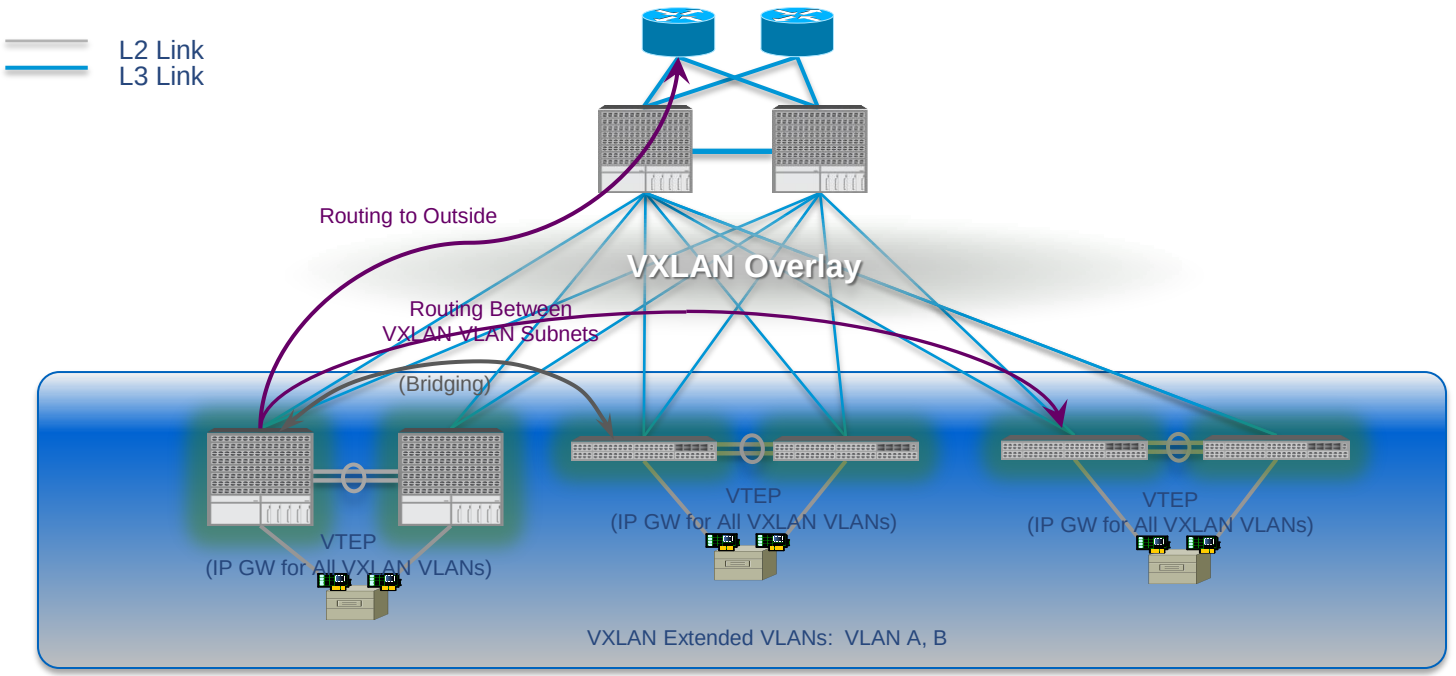

# VXLAN Design with VXLAN Bridging + Routing + Anycast IP Gateway - Spine-Leaf Deployment

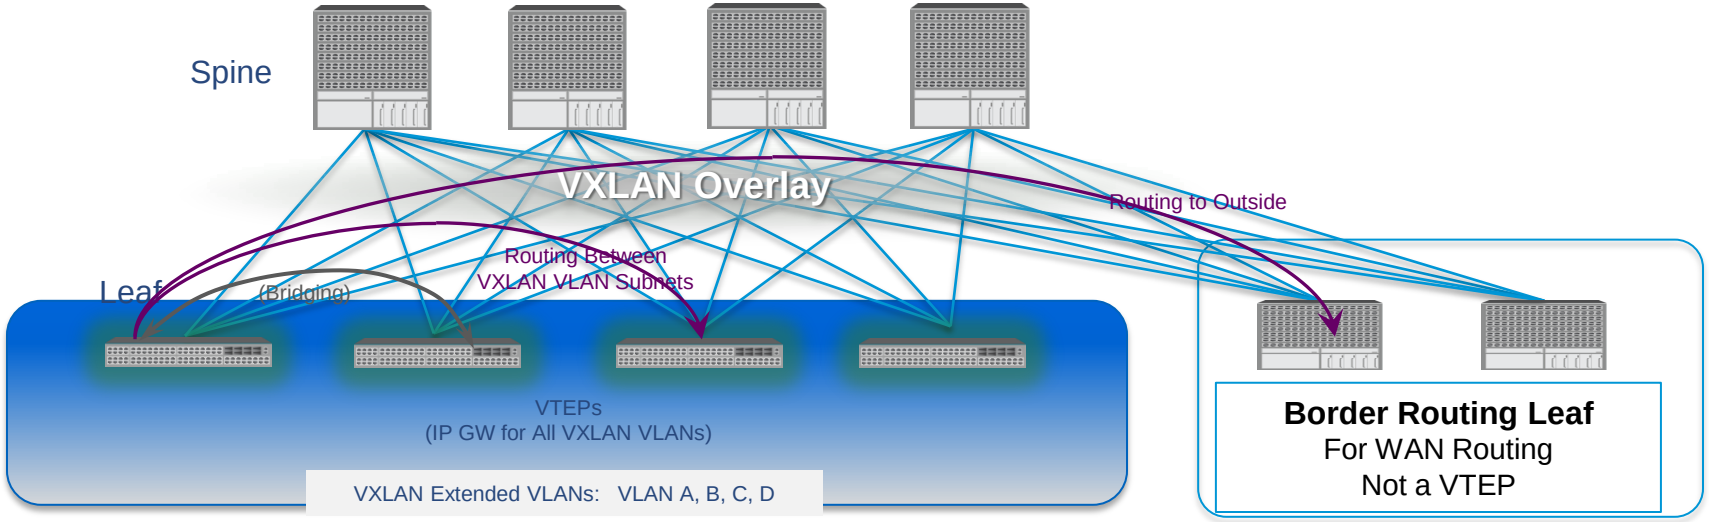

# VXLAN – Unsupported Bud-node Topology

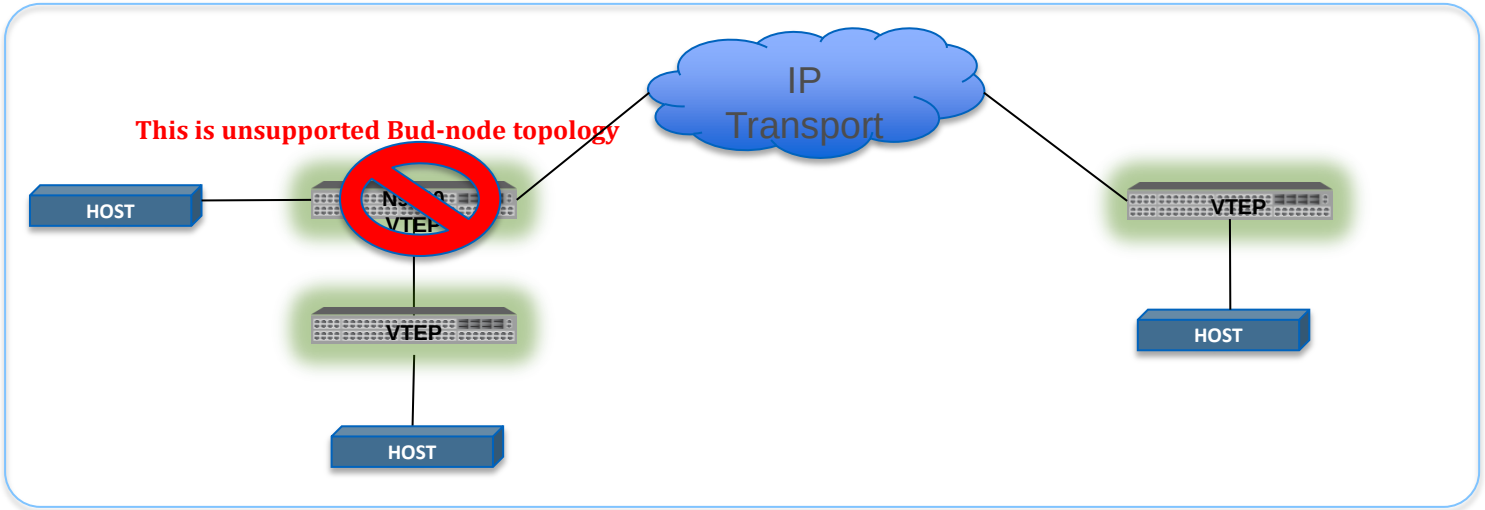

# VXLAN – Supported and Unsupported FEX Topology

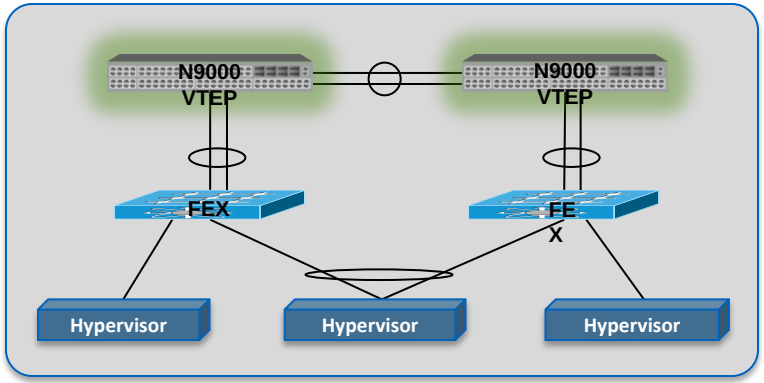


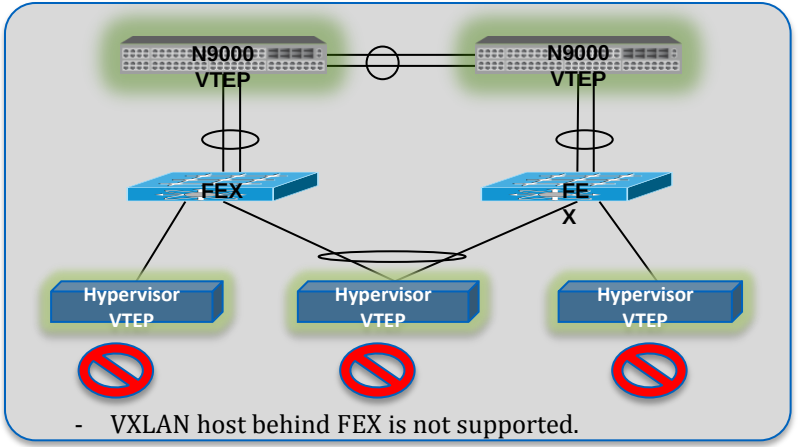


- VXLAN host behind FEX is not supported.

# VXLAN Deployment Steps

# VXLAN Configuration

## Enabling VXLAN feature

```
switch(config)# [no] feature nv overlay
```

## Enable VLAN model of VXLAN bridging

```
switch(config)# [no] feature vn-segment-vlan-based
```

## Enable Multicast

```
switch(config)# [no] feature pim
```

## Make sure source interface and next hop L3 interface should have multicast enabled

```
switch(config)# interface loopback 0
switch(config-if)# ip pim sparse-mode
```

## Creating NVE (Overlay) interface

```
switch(config)# interface nve <x>
switch(config-if-nve)# source-interface <src-if>
switch(config-if-nve)# overlay-encapsulation <vxlan|nvgre>
switch(config-if-nve-si)# member vni [<range>| auto] [mcast-group <start addr> [<end addr>]]
```

# VXLAN Configuration

Default UDP Port is 4789

Configure VXLAN UDP Port

```
switch(config)# [no] vxlan udp port <number>
```

Creating VXLAN VNIs (optional)

```
switch(config)# [no] vni <range>
```

VLAN to VXLAM Mapping configuration

```
switch(config)# vlan <number>
switch(config-vlan)# vn-segment <vnid>
```

# VXLAN vPC Configuration

## Loopback Interface Configuration

```
interface loopback0
 no ip redirects
 ip address 1.1.1.1/24
 ip address 1.1.1.3/24 Secondary
 ip pim sparse-mode
```

The secondary ip address is used by the emulated VTEP for VXLAN

The emulated ip address allows the network to load balance the traffic destined to vPC connected devices without using MCT.

Just make sure all the configs are identical between vPC Primary and vPC Secondary

# VXLAN Sample Topology - VTEP

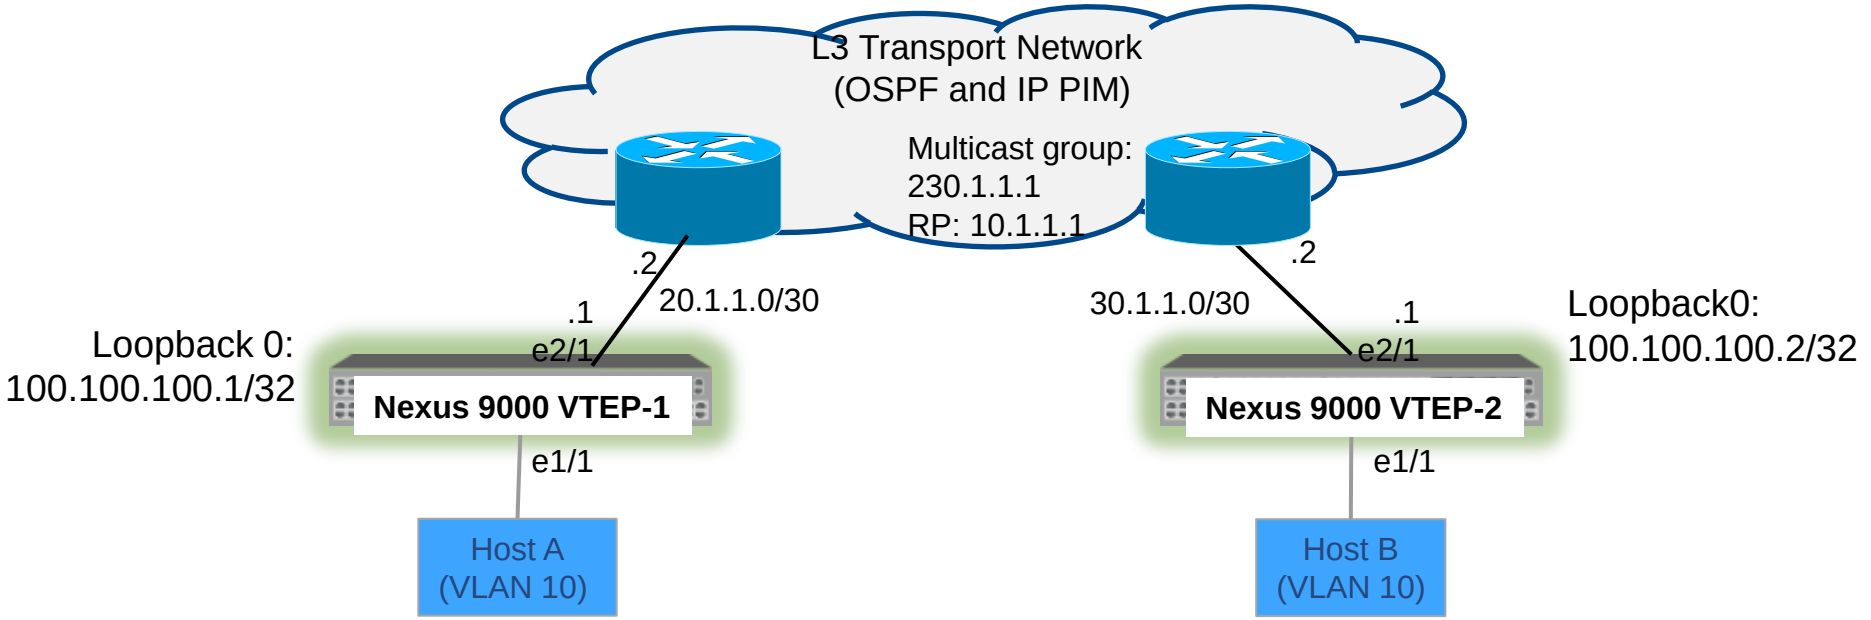

```
n9k-vtep-1(config)# feature ospf
n9k-vtep-1(config)# feature pim
n9k-vtep-1(config)# router ospf 1
n9k-vtep-1(config-router)# router-id 100.100.100.1
n9k-vtep-1(config)# ip pim rp-address 10.1.1.1 group-list 224.0.0.0/4
```

```
n9k-vtep-1(config)# interface loopback0
n9k-vtep-1(config-if)# ip address 100.100.100.1/32
n9k-vtep-1(config-if)# ip router ospf 1 area 0.0.0.0
n9k-vtep-1(config-if)# ip pim sparse-mode
```

```
n9k-vtep-1(config)# interface e2/1
n9k-vtep-1(config-if)# ip address 20.1.1.1/30
n9k-vtep-1(config-if)# ip router ospf 1 area 0.0.0.0
n9k-vtep-1(config-if)# ip pim sparse-mode
```

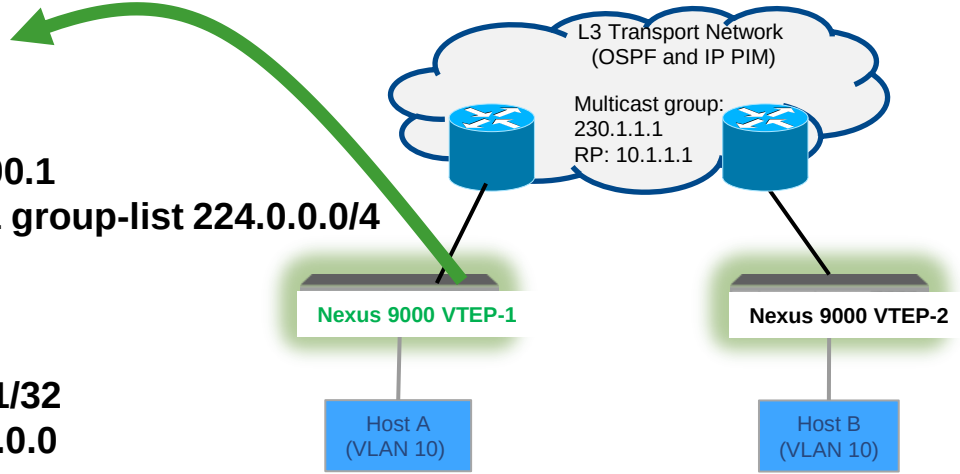

```
n9k-vtep-1(config)# feature nv overlay
n9k-vtep-1(config)# feature vn-segment-vlan-based

n9k-vtep-1(config)# interface e1/1
n9k-vtep-1(config-if)# switchport
n9k-vtep-1(config-if)# switchport access vlan 10
n9k-vtep-1(config-if)# no shutdown


n9k-vtep-1(config)# interface nve1
n9k-vtep-1(config-if)# no shutdown
n9k-vtep-1(config-if)# source-interface loopback0
n9k-vtep-1(config-if)# overlay-encapsulation vxlan
n9k-vtep-1(config-if)# member vni 1010 mcast-group 230.1.1.1

n9k-vtep-1(config)# vlan 10
n9k-vtep-1(config-vlan)# vn-segment 1010
```

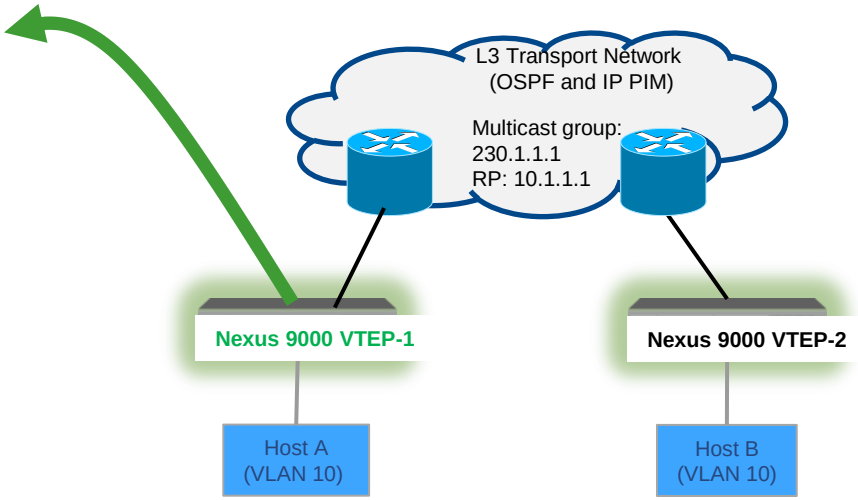

```
n9k-vtep-2(config)# feature ospf
n9k-vtep-2(config)# feature pim
n9k-vtep-2(config)# router ospf 1
n9k-vtep-2(config-router)# router-id 100.100.100.2
n9k-vtep-2(config)# ip pim rp-address 10.1.1.1 group-list 224.0.0.0/4

n9k-vtep-2(config)# interface loopback0
n9k-vtep-2(config-if)# ip address 100.100.100.2/32
n9k-vtep-2(config-if)# ip router ospf 1 area 0.0.0.0
n9k-vtep-2(config-if)# ip pim sparse-mode

n9k-vtep-2(config)# interface e2/1
n9k-vtep-2(config-if)# ip address 30.1.1.1/30
n9k-vtep-2(config-if)# ip router ospf 1 area 0.0.0.0
n9k-vtep-2(config-if)# ip pim sparse-mode
```

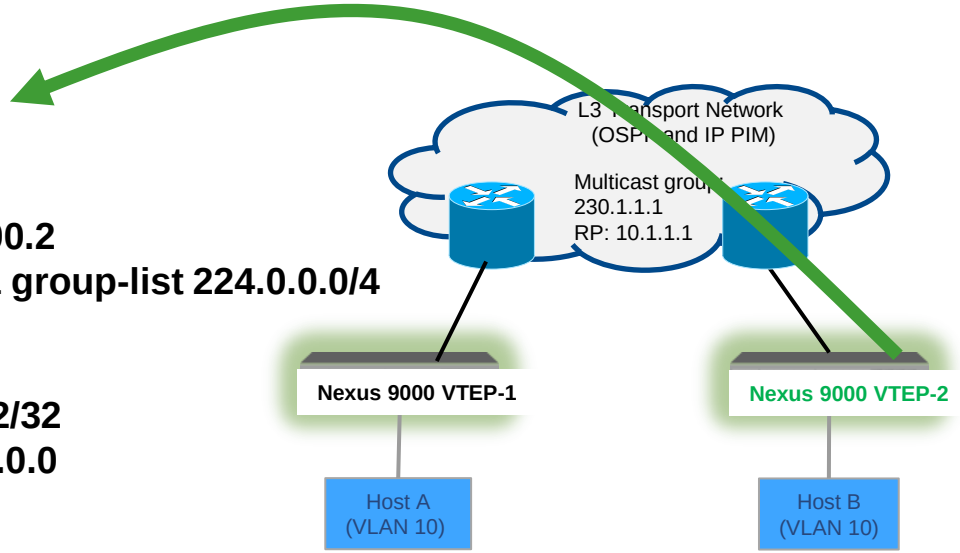

n9k-vtep-2(config)# **feature nv overlay**
n9k-vtep-2(config)# **feature vn-segment-vlan-based**

n9k-vtep-2(config)# **interface e1/1**
n9k-vtep-2(config-if)# **switchport**
n9k-vtep-2(config-if)# **switchport access vlan 10**
n9k-vtep-2(config-if)# **no shutdown**

n9k-vtep-2(config)# **interface nve1**
n9k-vtep-2(config-if)# **no shutdown**
n9k-vtep-2(config-if)# **source-interface loopback0**
n9k-vtep-2(config-if)# **overlay-encapsulation vxlan**
n9k-vtep-2(config-if)# **member vni 1010 mcast-group 230.1.1.1**

n9k-vtep-2(config)# **vlan 10**
n9k-vtep-2(config-vlan)# **vn-segment 1010**

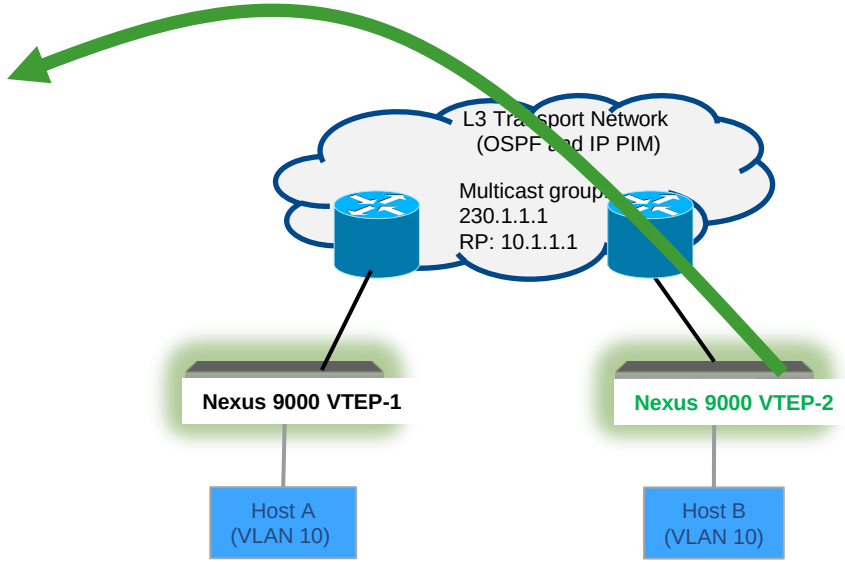

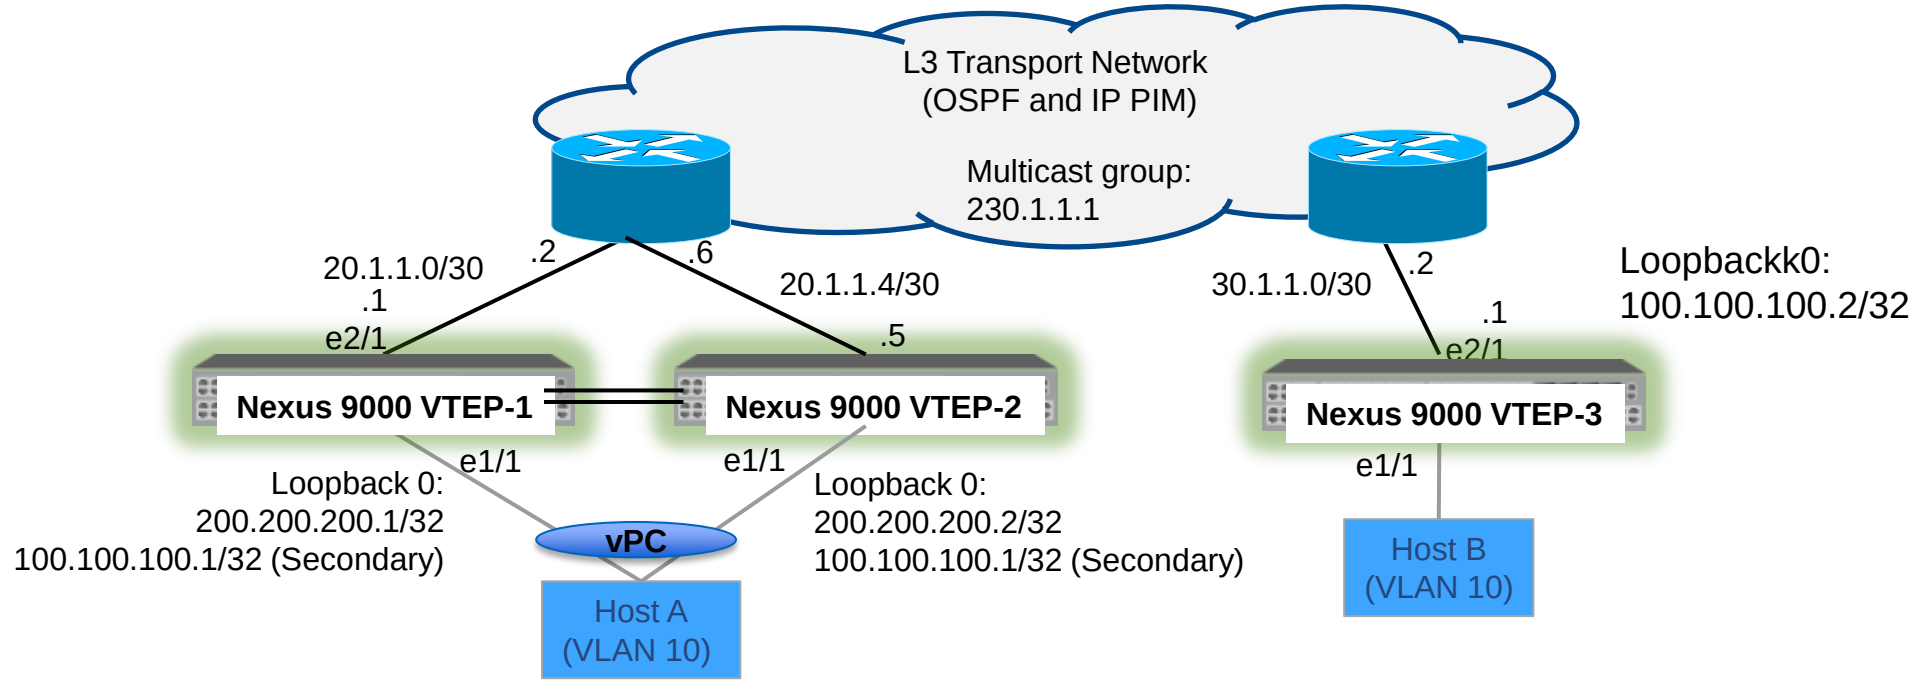

```
n9k-vtep-1(config)# feature ospf
n9k-vtep-1(config)# feature pim
n9k-vtep-1(config)# router ospf 1
n9k-vtep-1(config-router)# router-id 200.200.200.1
n9k-vtep-1(config)# ip pim rp-address 10.1.1.1 group-list 224.0.0.0/4

n9k-vtep-1(config)# interface loopback0
n9k-vtep-1(config-if)# ip address 200.200.200.1/32
n9k-vtep-1(config-if)# ip address 100.100.100.1/32 secondary
n9k-vtep-1(config-if)# ip router ospf 1 area 0.0.0.0
n9k-vtep-1(config-if)# ip pim sparse-mode

n9k-vtep-1(config)# interface e2/1
n9k-vtep-1(config-if)# ip address 20.1.1.1/30
n9k-vtep-1(config-if)# ip router ospf 1 area 0.0.0.0
n9k-vtep-1(config-if)# ip pim sparse-mode

n9k-vtep-1(config)# feature nv overlay
n9k-vtep-1(config)# feature vn-segment-vlan-based
```

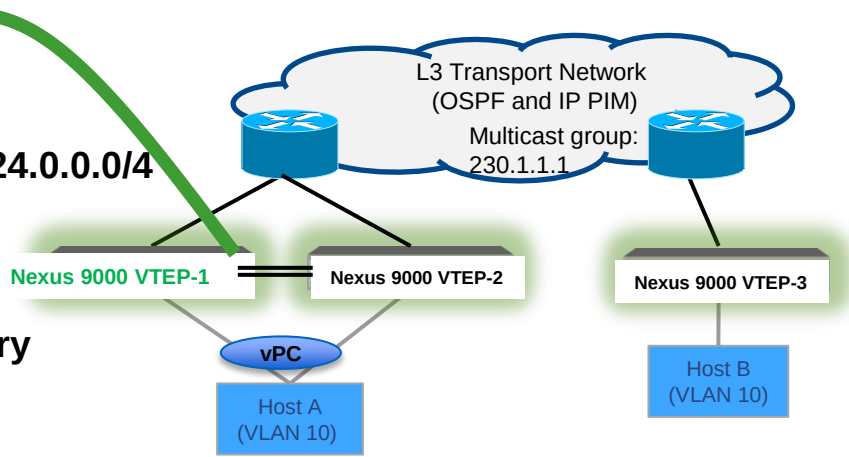

```
n9k-vtep-1(config)# interface port-channel 10
n9k-vtep-1(config-if)# vpc 10
n9k-vtep-1(config-if)# switchport
n9k-vtep-1(config-if)# switchport mode access
n9k-vtep-1(config-if)# switchport access vlan 10
n9k-vtep-1(config-if)# no shutdown

n9k-vtep-1(config)# interface e1/1
n9k-vtep-1(config-if)# channel-group 10 mode active
n9k-vtep-1(config-if)# no shutdown

n9k-vtep-1(config)# interface nve1
n9k-vtep-1(config-if)# no shutdown
n9k-vtep-1(config-if)# source-interface loopback0
n9k-vtep-1(config-if)# overlay-encapsulation vxlan
n9k-vtep-1(config-if)# member vni 1010 mcast-group 230.1.1.1

n9k-vtep-1(config)# vlan 10
n9k-vtep-1(config-vlan)# vn-segment 1010
```

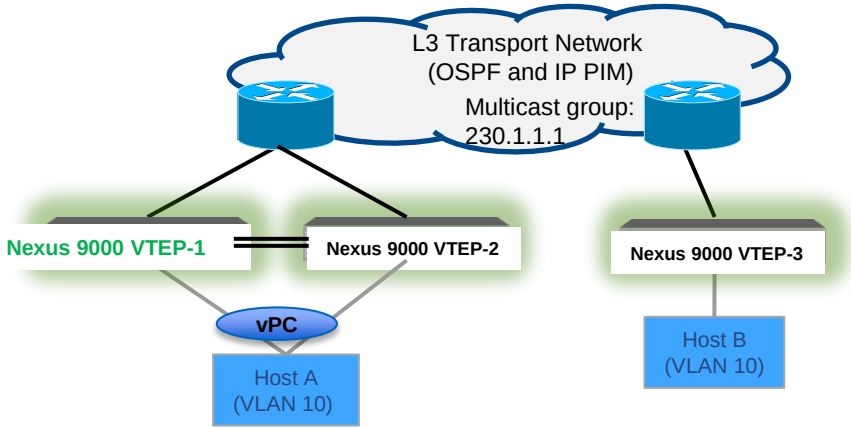

```
n9k-vtep-2(config)# feature ospf
n9k-vtep-2(config)# feature pim
n9k-vtep-2(config)# router ospf 1
n9k-vtep-2(config-router)# router-id 200.200.200.2
n9k-vtep-2(config)# ip pim rp-address 10.1.1.1 group-list 224.0.0.0/4

k-vtep-2(config)# interface loopback0
n9k-vtep-2(config-if)# ip address 200.200.200.2/32
n9k-vtep-2(config-if)# ip address 100.100.100.1/32 secondary
n9k-vtep-2(config-if)# ip router ospf 1 area 0.0.0.0
n9k-vtep-2(config-if)# ip pim sparse-mode

n9k-vtep-2(config)# interface e2/1
n9k-vtep-2(config-if)# ip address 20.1.1.5/30
n9k-vtep-2(config-if)# ip router ospf 1 area 0.0.0.0
n9k-vtep-2(config-if)# ip pim sparse-mode

n9k-vtep-2(config)# feature nv overlay
n9k-vtep-2(config)# feature vn-segment-vlan-based
```

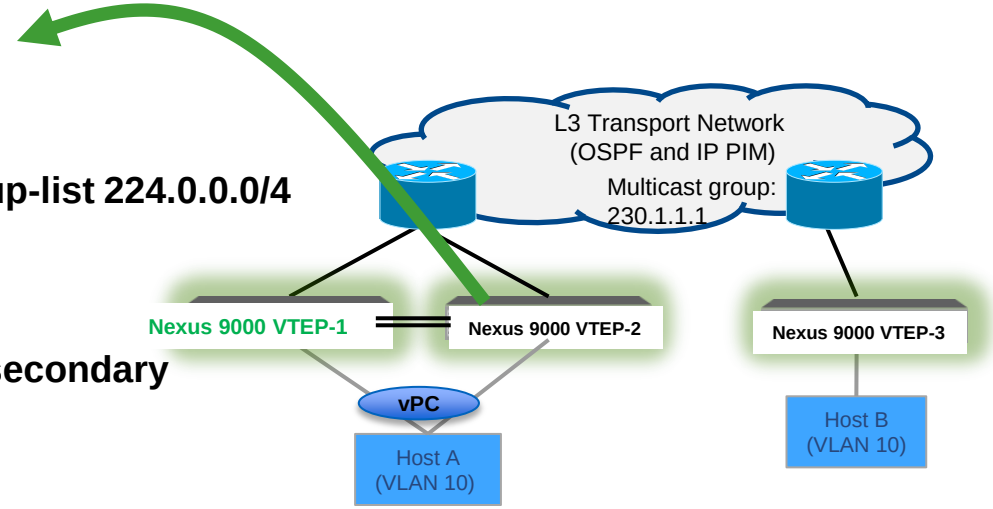

```
n9k-vtep-2(config)# interface port-channel 10
n9k-vtep-2(config-if)# vpc 10
n9k-vtep-2(config-if)# switchport
n9k-vtep-2(config-if)# switchport mode access
n9k-vtep-2(config-if)# switchport access vlan 10
n9k-vtep-2(config-if)# no shutdown

n9k-vtep-2(config)# interface e1/1
n9k-vtep-2(config-if)# channel-group 10 mode active
n9k-vtep-2(config-if)# no shutdown

n9k-vtep-2(config)# interface nve1
n9k-vtep-2(config-if)# no shutdown
n9k-vtep-2(config-if)# source-interface loopback0
n9k-vtep-2(config-if)# overlay-encapsulation vxlan
n9k-vtep-2(config-if)# member vni 1010 mcast-group 230.1.1.1

n9k-vtep-2(config)# vlan 10
n9k-vtep-2(config-vlan)# vn-segment 1010
```

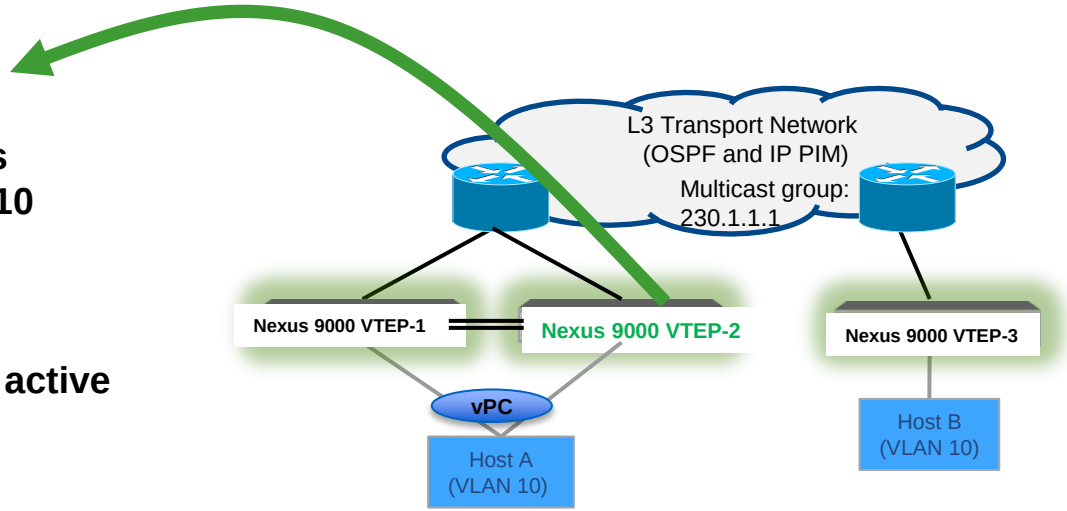

```
n9k-vtep-3(config)# feature nv overlay
n9k-vtep-3(config)# feature vn-segment-vlan-based

n9k-vtep-3(config)# interface e1/1
n9k-vtep-3(config-if)# switchport
n9k-vtep-3(config-if)# switchport access vlan 10
n9k-vtep-3(config-if)# no shutdown

n9k-vtep-3(config)# interface nve1
n9k-vtep-3(config-if)# no shutdown
n9k-vtep-3(config-if)# source-interface loopback0
n9k-vtep-3(config-if)# overlay-encapsulation vxlan
n9k-vtep-3(config-if)# member vni 1010 mcast-group 230.1.1.1

n9k-vtep-3(config)# vlan 10
n9k-vtep-3(config-vlan)# vn-segment 1010
```

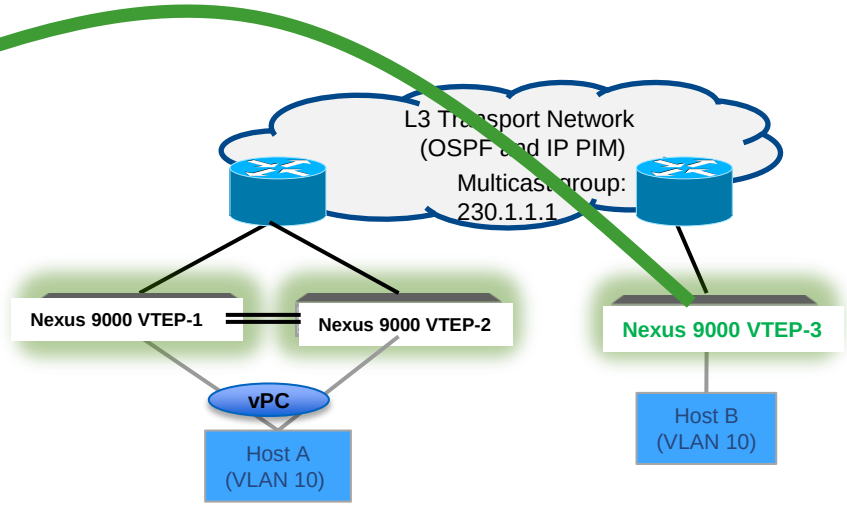

# VXLAN Troubleshooting

# Troubleshooting

- NVE Interface State
  - NVE source interface is configured
  - The source interface has IPv4 address is configured
  - Source interface is OPER up
  - NVE interface is admin up and has multicast group associated to it.

- Restrictions
  - Only 1 NVE local interface is supported
  - VXLAN gateway shouldn’t be connected to multiple core routers via a layer 2 switch
  - Support for downstream multicast receivers

# Troubleshooting

## Verify Overlay Peer

```
switch# show nve peers
Interface           Peer-IP              VNI         Up Time
-------------       -----------------    --------    -------
nve1                10.1.1.2             10000       00:52:24
```

```
switch# show nve peers interface nve1
Interface           Peer-IP              VNI         Up Time
-------------       -----------------    --------    -------
nve1                10.1.1.2             10000       00:52:24
```

```
switch# show nve peers vni 1010
Interface           Peer-IP              VNI         Up Time
-------------       -----------------    --------    -------
nve1                10.1.1.2             10000       00:52:24
```

- It displays VTEP peers and the associated VNI with them. The output displays time since the peer was last detected / learnt.

## Verify Overlay interface

```
switch# show nve vni
Interface           VNI        Multicast-group   VNI State
----------------    --------   ---------------   ---------
nve1                10000      230.1.1.1           up
```

- List of all VNIs that are associated with various nve interfaces and the associated multicast IP address.

# Troubleshooting (Cont.)

## Verify MAC address is learnt

```
switch# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O -
Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC
Peer-Link,
        (T) - True, (F) - False
   VLAN     MAC Address      Type      age     Secure NTFY Ports
---------+-----------------+--------+---------+------+----+-------
-----------
*  100     0000.aa01.0001   dynamic  0         F      F    nve1
*  100     0000.bb01.0001   dynamic  0         F      F    nve1
*  100     0000.bb01.0001   dynamic  0         F      F    Eth1/46
```

## Few Other Commands

```
switch# show logging level nve

switch# show tech-support nve

switch# show runn interface nve 1

switch# show nve interface

switch# show nve vxlan-params
```

- List of all VNIs that are associated with various nve interfaces and the associated multicast IP address.

# Troubleshooting (Cont.)

## Few Debug commands

```
switch# debug nve errors

switch# debug nve events

switch# debug nve all

switch# debug nve pim-library
```

# VXLAN Control Plane Alternatives – Future

- Challenges:
  - Data Centre Core may not support multicast.
  - Scale challenges with multicast flooding.
  - VXLAN Membership / Auto-discovery of VTEP.
  - Support of Gateway multi-homing.
  - VNI to IP multicast group mapping has to be managed.

- Alternatives:
  - BGP Control Plane
  - LISP Control Plane
  - API

# Key Takeaways – VXLAN Technology

- VXLAN is simple
  - Keeps the attractive aspects of Layer 2 – No re-addressing, simple configuration and deployment
  - Integrates stability and scale of Layer 3
- VXLAN is efficient
  - Proper utilisation of ECMP
  - Optimal path between any two nodes
- VXLAN is scalable
  - Can extend a bridged domain without extending the risks generally associated with Layer 2 and beyond 4K VLAN limit
- VXLAN Control Plane (Future)
  - BGP and LISP

# Key Takeaways – VXLAN Design

- You can deploy VXLAN today, with traditional network designs
- VXLAN introduces immediate, tangible benefits to any design:
  - Simple configuration, leverage parallel network paths, extend VLANs safely etc.
- Provides multiple design options to help you build a network that meets your requirements

# References

VXLAN Overview: Cisco Nexus 9000 Series Switches
http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps13386/white-paper-c11-729383_ns1261_Networking_Solutions_White_Paper.html

* Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide

VXLAN: A Framework for Overlaying Virtualised Layer 2 Networks over Layer 3 Networks
http://www.ietf.org/id/draft-mahalingam-dutt-dcops-vxlan-07.txt

* Will be available at the time of FCS of the software code




# Q & A

# Complete Your Online Session Evaluation

## Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm

**Learn online with Cisco Live!**

Visit us online after the conference for full access to session videos and presentations.
www.CiscoLiveAPAC.com


CISCO
TM