Cisco live!

January 29 – February 2, 2018 · Barcelona

BRKSPG-2202

Deploying Carrier Ethernet Services

Kashif Islam, Solutions Architect

Cisco live!

Cisco Spark

Questions?

Use Cisco Spark to communicate with the speaker after the session

How

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

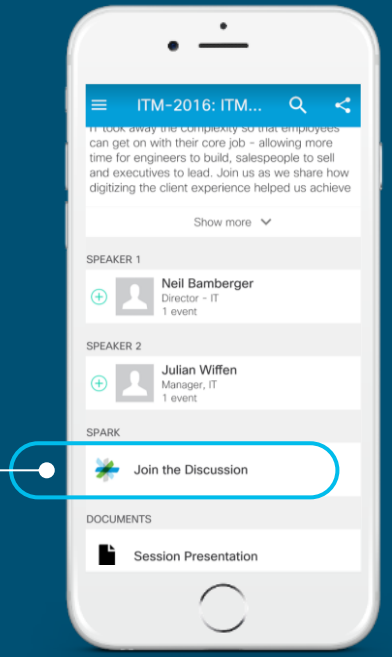

cs.co/ciscolivebot#BRKSPG-2202

Agenda

- Cisco Metro Fabric Overview
- Metro Services and Use Cases
- Metro Underlay Configuration
- IOS-XR Services Configuration Toolset
- Metro Services Configuration
- Orchestration and Management
- Summary

Cisco Metro Fabric Overview

Cisco SP Fabric Designs principals:

Simple, Scalable, Automatable

Network Location

Metro & Access

Core

Peering

Designs

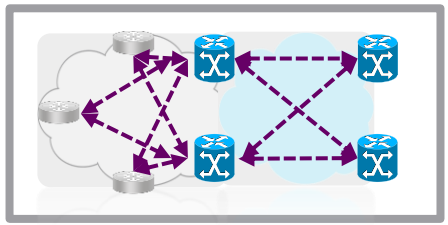

Metro Fabric

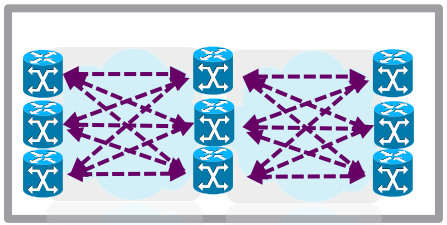

Core Fabric

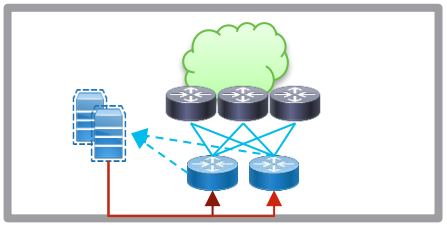

Peering Fabric

Building Blocks

CLOS Fabric

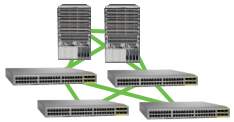

SR/EVPN

Automation

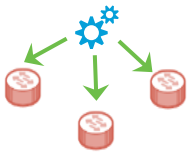

Telemetry and Analytics

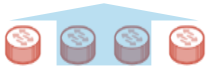

YANG data models

Cisco Metro Fabric Building Blocks

CLOS Fabric

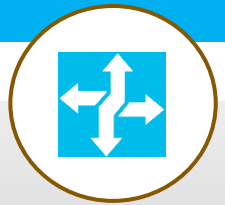

Industry leader: ASR9K

Dense, Scalable: NCS 5500

Segment Routing

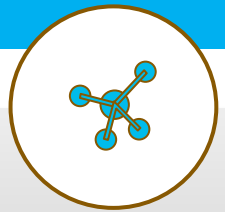

Unified Forwarding Plane with Explicit Path Control and Traffic Engineering

BGP Based VPN

Common Control Plane for L2, L3 and IRB

Automation

Programmability and analytics with YANG data models and telemetry

Compass Metro Fabric - High-Level Domain View

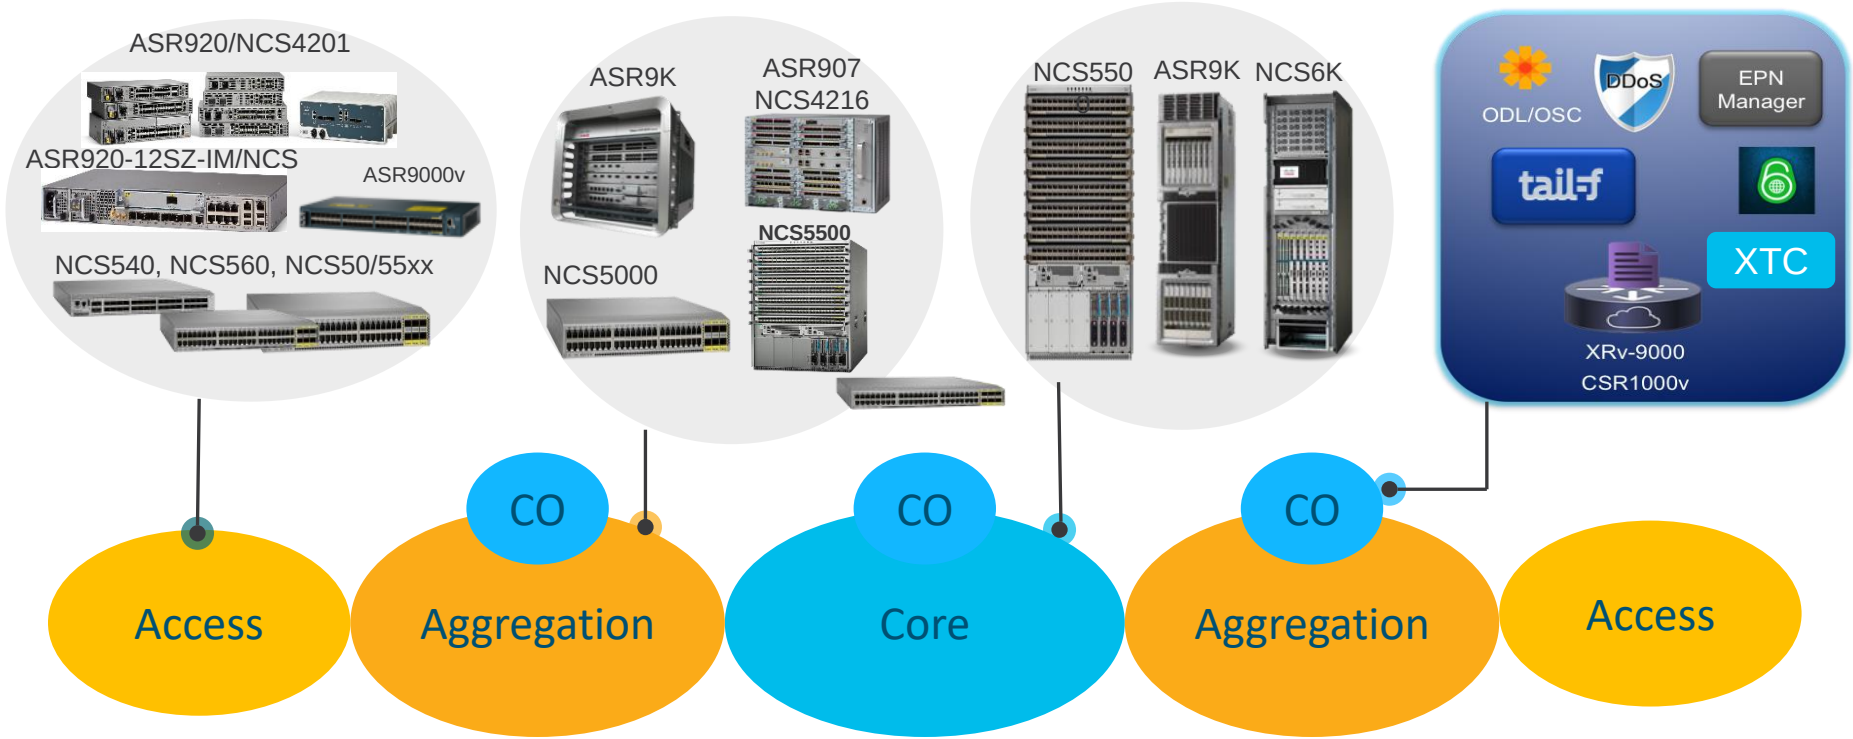

Metro Services and Use Cases

Metro Fabric Services - Highlights

- Transport Independent => interop. with existing services
- Seamless Integration with existing L2VPN
- End-To-End and Hierarchical Services
  - Provides scale and provisioning simplification
- BGP-based services:
  - L3VPN (VPNv4/VPNv6)
  - EVPN
  - Traditional L2VPN
- Provisioning CLI and Automation

End to End Metro Services

Service	Technology
L3VPN	MP-BGP VPNv4/6
L2 P2P	EVPN-VPWS • Multi/Single-Homed • All/Single-Active Legacy EoMPLS (Static PW)
L2 Multipoint	EVPN • Multi/Single-Homed • All/Single-Active
L2/L3 Multipoint	EVPN • Multi/Single-Homed • All/Single-Active Anycast-IRB

End to End Metro Services

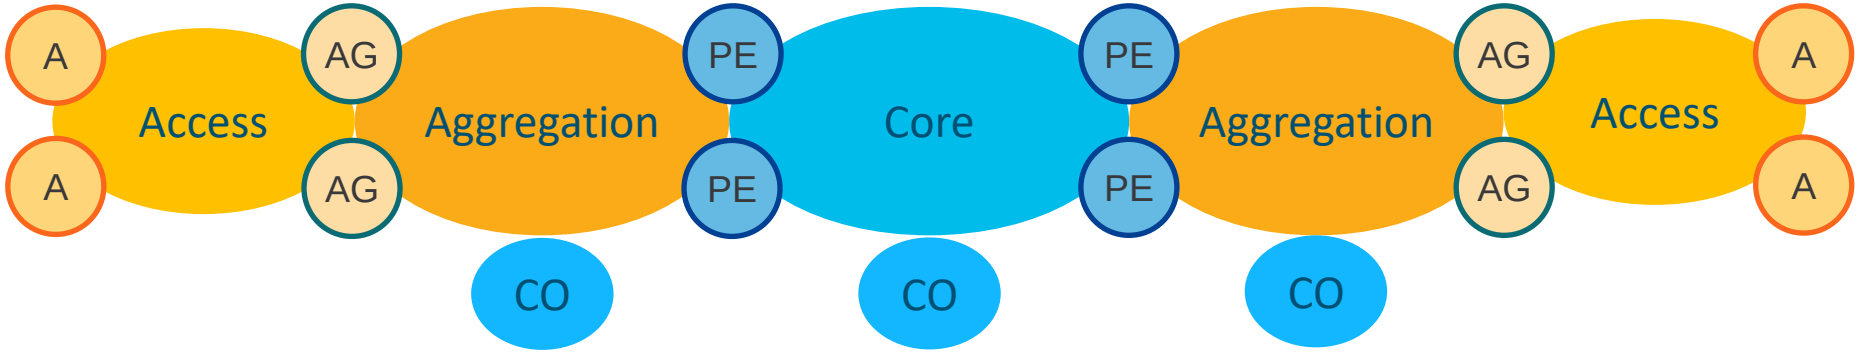

Distributed CO - service termination

Centralized CO - service termination

Access-To-Access - service termination

MP-BGP L3VPN
EVPN
EVPN-VPWS
Legacy EoMPLS PW

Hierarchical Metro Services

Service	Technology in Access/Aggregation	Technology in Core
L3VPN	EVPN-VPWS • Multi/Single-Homed • All/Single-Active Anycast-Static-PW	MP-BGP VPNv4/6 (IRB/PWHE)
L2 P2P	Not-Required	
L2 Multipoint	EVPN-VPWS • Multi/Single-Homed • All/Single-Active Anycast-Static-PW	EVPN • Multi/Single-Homed • All/Single-Active
L2/L3 Multipoint	EVPN-VPWS • Multi/Single-Homed • All/Single-Active Anycast-Static-PW	EVPN • Multi/Single-Homed • All-Active (Anycast-IRB) • Single-Active (PWHE)

Hierarchical Metro Services

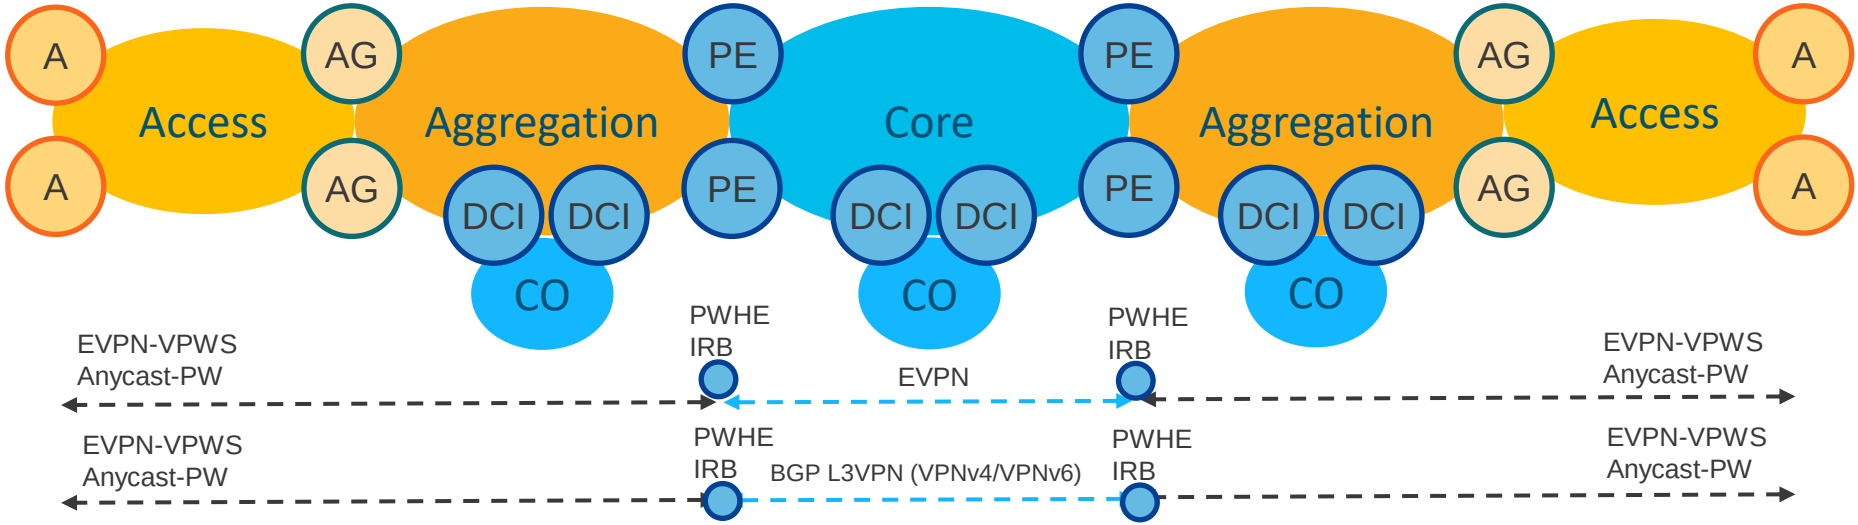

Metro Fabric Underlay Configuration

Cisco Metro Fabric Building Blocks

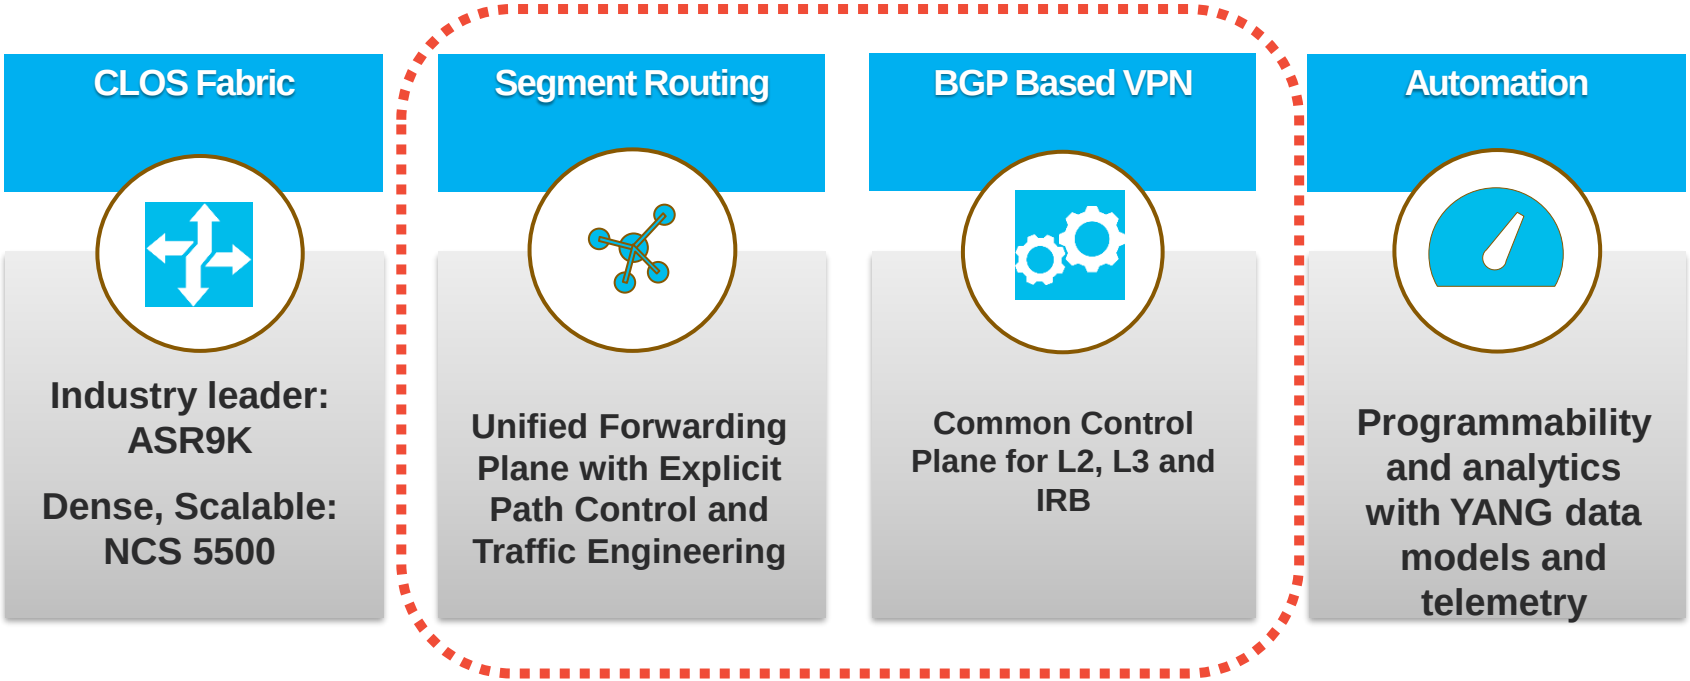

Segment Routing Configuration Basics

- Configured with the IGP Routing Protocols – ISIS and OSPF
- Requires: Enabling SR and configuring Prefix SID
- Prefix SID
  - **Globally Significant** →SR Global Block (SRGB)
  - SRGB advertised with router capabilities TLV
  - Configured as an **absolute value or an index**
  - Advertised as globally unique index

    E.g. index **1** → SID is 16,000 + **1** = 16,001
- Adjacency SID
  - **Locally significant**
  - Automatically allocated for each adjacency
  - Always encoded as an absolute value

IS-IS Configuration – Example

```
router isis 1
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 address-family ipv6 unicast
  metric-style wide
  segment-routing mpls
 !
 interface Loopback0
  passive
  address-family ipv4 unicast
   prefix-sid index 1
  !
  address-family ipv6 unicast
   prefix-sid absolute 20001
  !
 !
```

Wide metrics

enable SR IPv4 control plane and SR MPLS data plane on all ipv4 interfaces in this IS-IS instance

Wide metrics

enable SR IPv6 control plane and SR MPLS data plane on all ipv6 interfaces in this IS-IS instance

Ipv4 Prefix-SID value for loopback0 (Index translate to 16001 absolute vlaue)

Ipv6 Prefix-SID value for loopback0

OSPF Configuration Example

```
router ospf 1
 router-id 1.1.1.1
 segment-routing mpls

area 0
  interface Loopback0
   passive enable
   prefix-sid absolute 16001
  !
 !
!
```

Enable SR on all areas

Prefix-SID for loopback0

EVPN Flavors

- EVPN provides an evolution of Ethernet services
  - BGP control-plane for Ethernet Segment and MAC distribution and learning over MPLS core
  - Same principles and operational experience of IP VPNs
- BGP control plane provides a familiar, consistent and flexible configuration interface
- Multi-vendor solutions for P2P and MP services

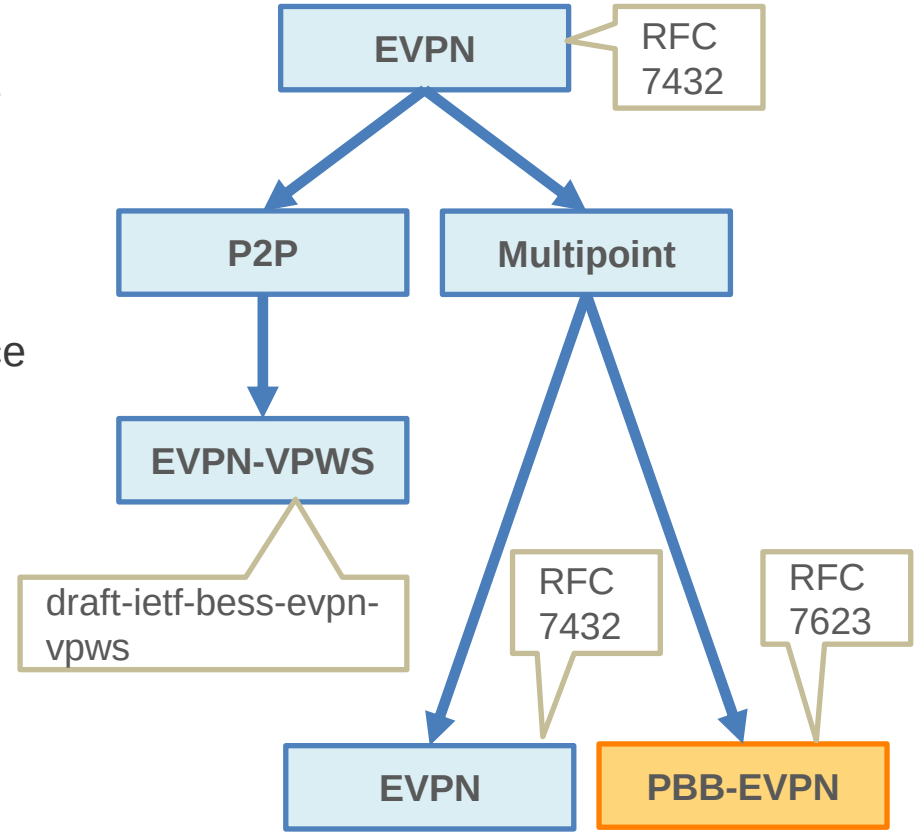

EVPN control plane with BGP

- New BGP NLRI to advertise MACs and IPs for next hop resolution
  - AFI=25 (L2VPN) SAFI=70 (EVPN)
  - IPv4 and IPv6 support
  - Control over MAC learning
  - ECMP for multihomed CEs
  - Inherent BGP scalability and hierarchy

Route Distinguisher (8 octets)
Ethernet Segment Identifier or ESI (10 octets)
Ethernet Tag ID (4 octets)
MAC Address Length (1 octet)
MAC Address (6 octets)
IP Address Length (1 octet)
IP Address (0 or 4 or 16 octets)
MPLS Label 1 (3 octets)
MPLS Label 2 (0 or 3 octets)

EVPN Planes Of Operation

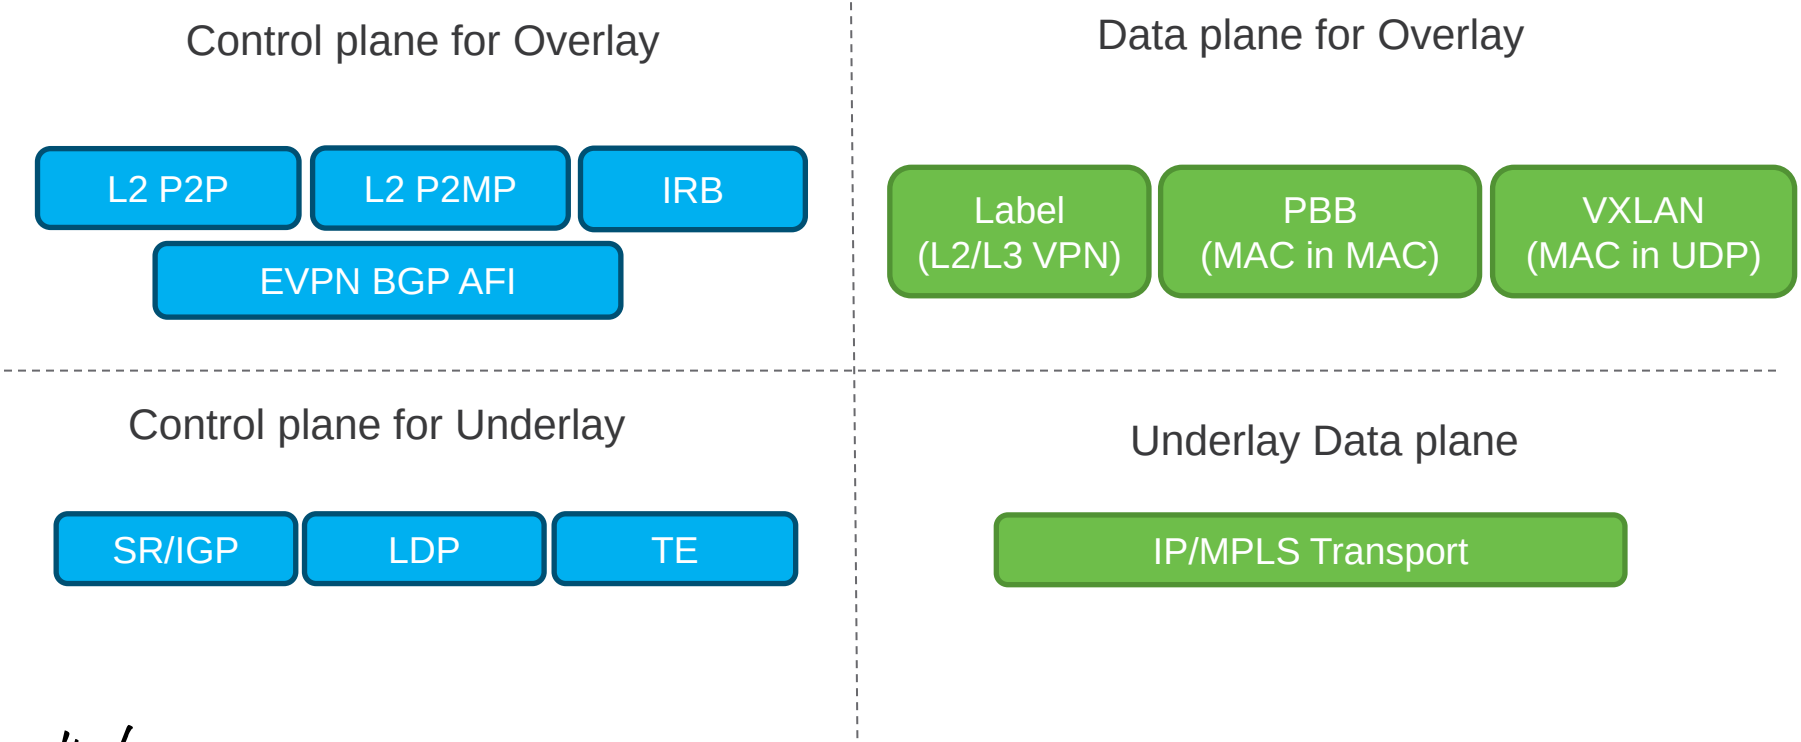

BGP Control Plane for EVPN

```
router bgp 65001
 bgp router-id 6.1.1.1

 address-family l2vpn evpn
 !
 neighbor 6.1.1.10
  remote-as 65001
  update-source Loopback0
  address-family l2vpn evpn
```

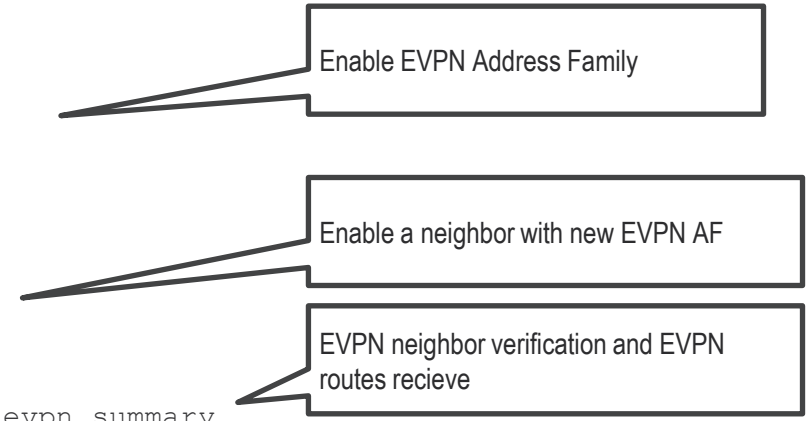


```
RP/0/0/CPU0:R1# sh bgp l2vpn evpn summary
BGP router identifier 6.1.1.1, local AS number 65001

Process       RcvTblVer   bRIB/RIB   LabelVer  ImportVer  SendTblVer  StandbyVer
Speaker               8          8          8          8           8           0

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
6.1.1.10          0 65001    5744    5743        8    0    0    3d23h           1
```

IOS-XR Service Configuration Toolset

Anatomy of a Metro Ethernet Service

EFP Defintiion

- Ethernet Flow Points
- Intelligent Layer 2 Sub-interfaces

Frame Calssification

- Frame Manipulation
- VLAN Tag Rewrite

Frame Forwarding

- Bridge Domain, Bridged Virtual Interface
- Local Connect, Ethernet over MPLS (EoMPLS), Virtual Private LAN Service (VPLS), Ethernet VPN (EVPN)

Ethernet Flow Point

- EVC Infrastructure introduces the concept of an EFP

  interface r/s/module/port.<sub-intf no.> ***l2transport***

  <match criteria commands> **(VLAN tags, MAC, Ether type)**

  <rewrite commands> **(VLAN tags pop/push/translation)**

  <feature commands> **(QoS, ACL etc)**

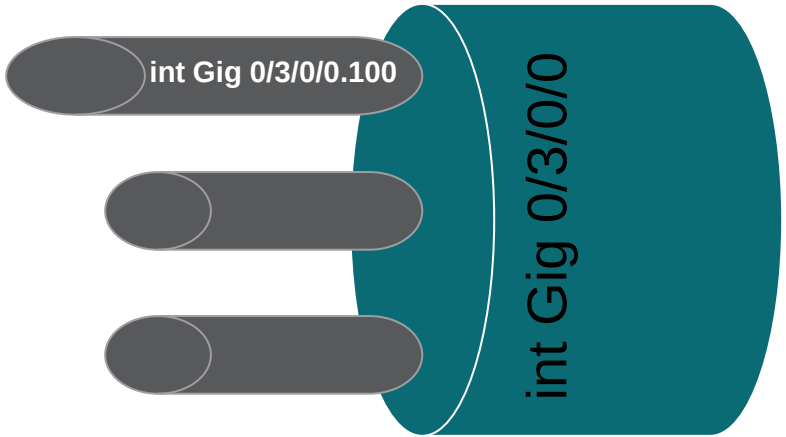

EFP Flexible Tag Classification

The Longest Match Rule and the default option

- Longest match for VLAN tag provides configuration flexibility

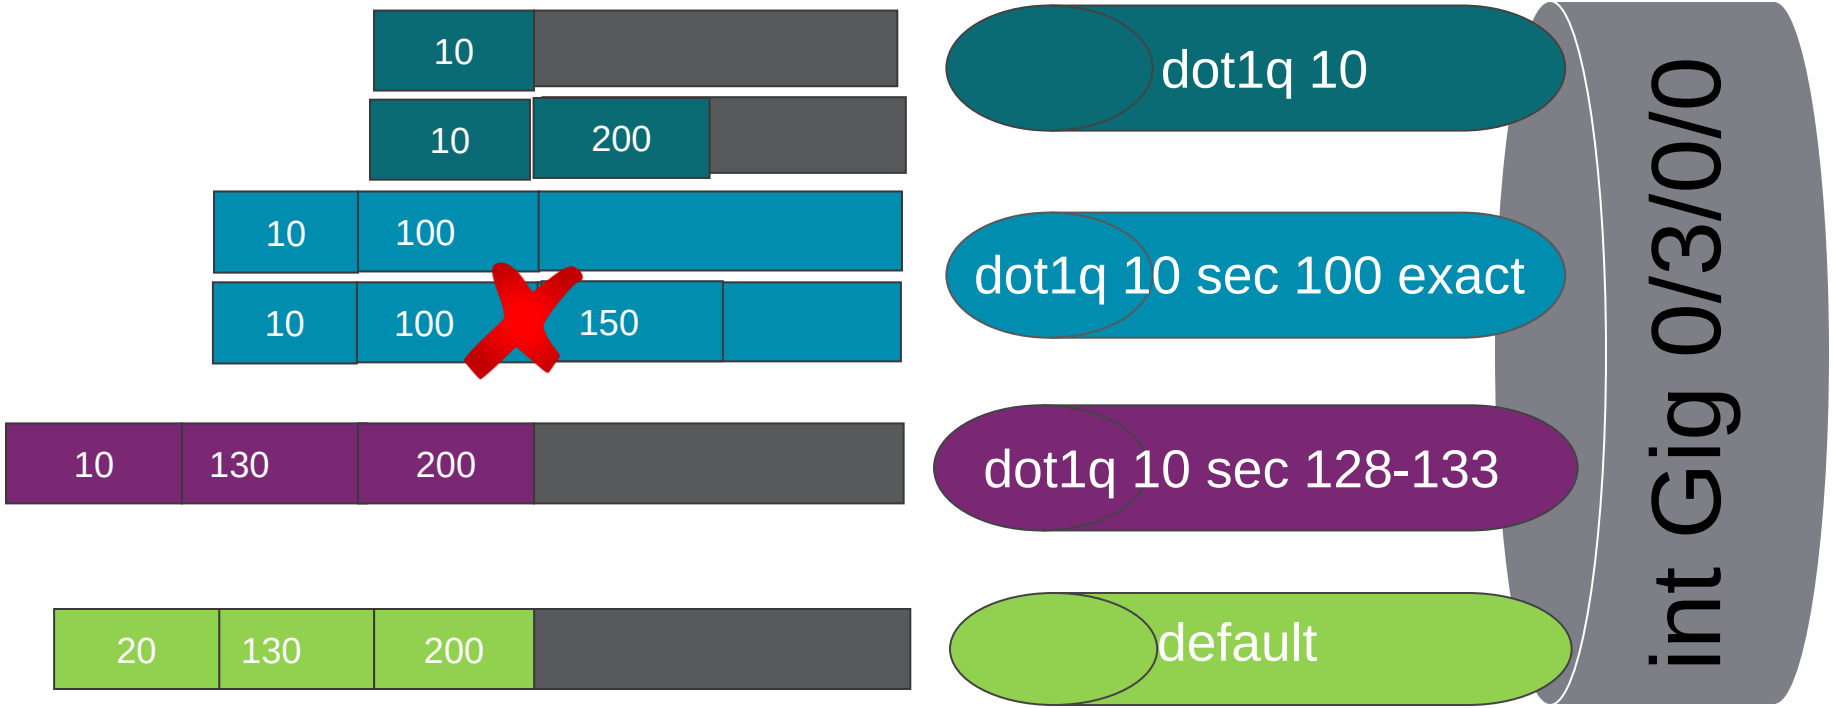

Traffic Forwarding Through an EFP

Traffic Manipulation

- Flexible VLAN Tag Manipulation
  - Push, Pop, Translate
- Any combination up to 2 VLANS
  - 1 to 1, 1 to 2, 2 to 1 or 2 to 2
- Uses “rewrite” keyword
- Symmetric Application

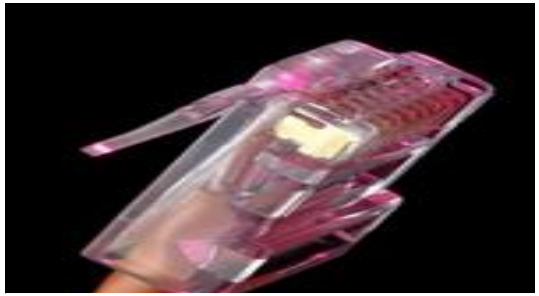

Point to Point Forwarding

- MEF defined E-LINE services
- Allows 2 sites to be connected via EFPs
- Two Primary Mechanisms
  - Local Connect
  - EoMPLS

Multipoint Forwarding

- E-LAN, E-TREE services
- Allows 2+ sites to connected via EFPs
- MAC based Forwarding
- Bridge-Domain, BVI
- VPLS, H-VPLS
- EVPN, PBB-EVPN

Configured under “l2vpn” CLI

Bringing Everything Together

IOS-XR Flexible Ethernet SW Infrastructure

EFP (**E**thernet **F**low **P**oint) or sub-interface

Flexible VLAN tag classification

Flexible VLAN tag rewrite

Flexible Ethertype (.1Q, QinQ, .1ad)

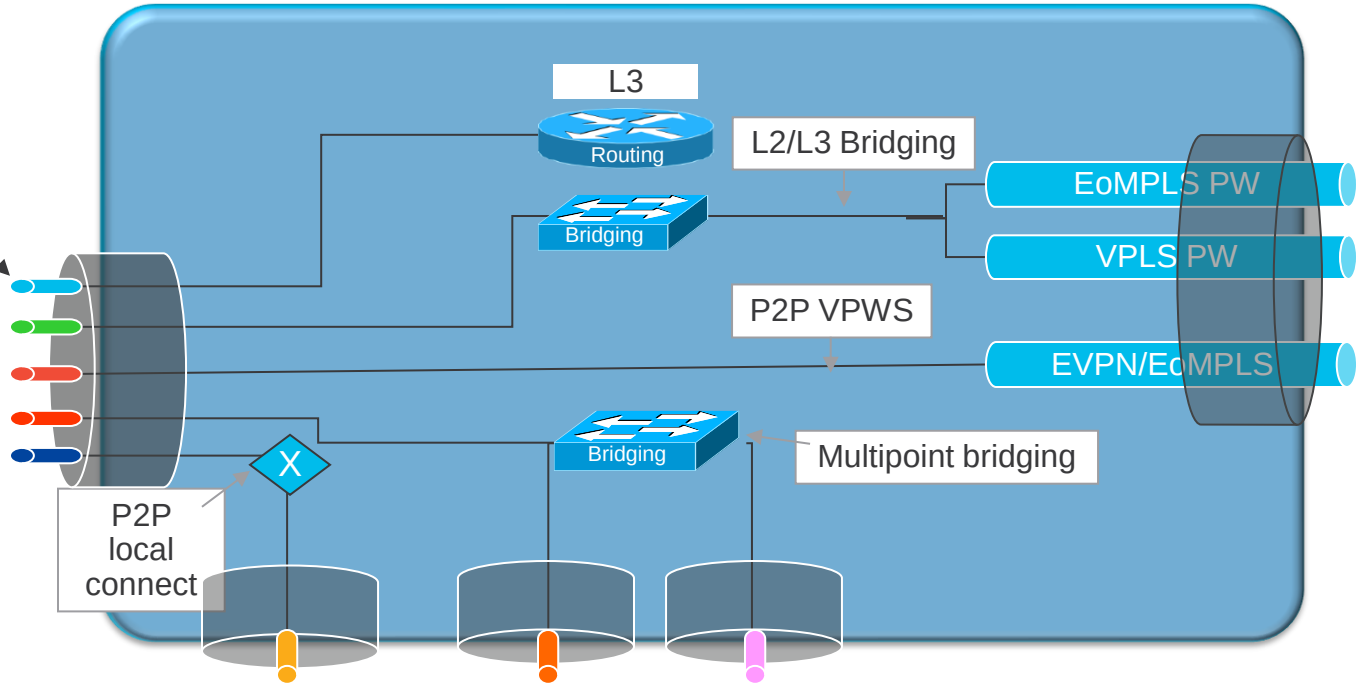

Flexible service mapping and multiplexing

L2 and L3, P2P and MP services concurrently on the same port

Metro Fabric Services Configuration

Metro Fabric Services Configuration

Service	Technology
Point to Point	• EoMPLS • EVPN VPW (Single Homed and Multi Homed)
Multi Point (L2/L3 services)	• VPLS • MP Switching • EVPN • PBB-EVPN

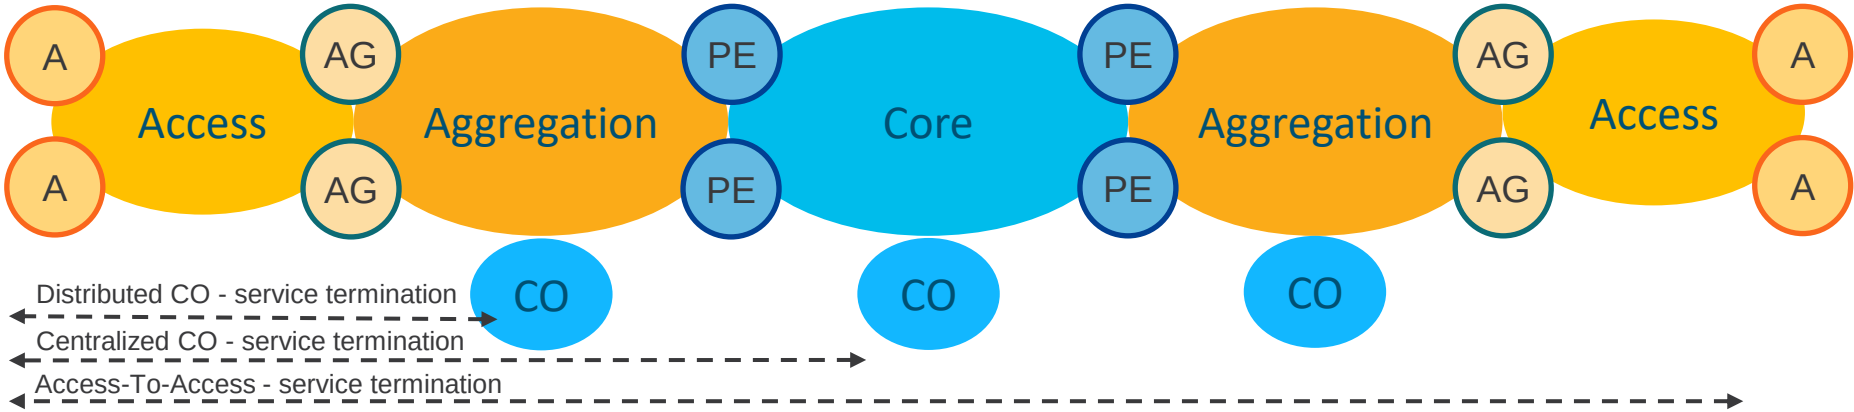

Point to Point Services Configuration
EoMPLS and EVPN-VPWS

```
interface gig 0/0/0/4.1 l2transport
  encapsulation dot1q 100
  rewrite ingress tag pop 1 symmetric
l2vpn
  xconnect group CISCO
     p2p EoMPLS_VPWS
       interface Gig 0/0/0/4.1
       neighbor 1.1.1.1 pw-id 100
```

```
interface gig 0/0/0/4.1 l2transport
  encapsulation dot1q 100
  rewrite ingress tag pop 1 symmetric
l2vpn
  xconnect group CISCO
     p2p EVPN_VPWS
       interface Gig 0/0/0/4.1
       neighbor evpn evi 100 target 10 source 10
```

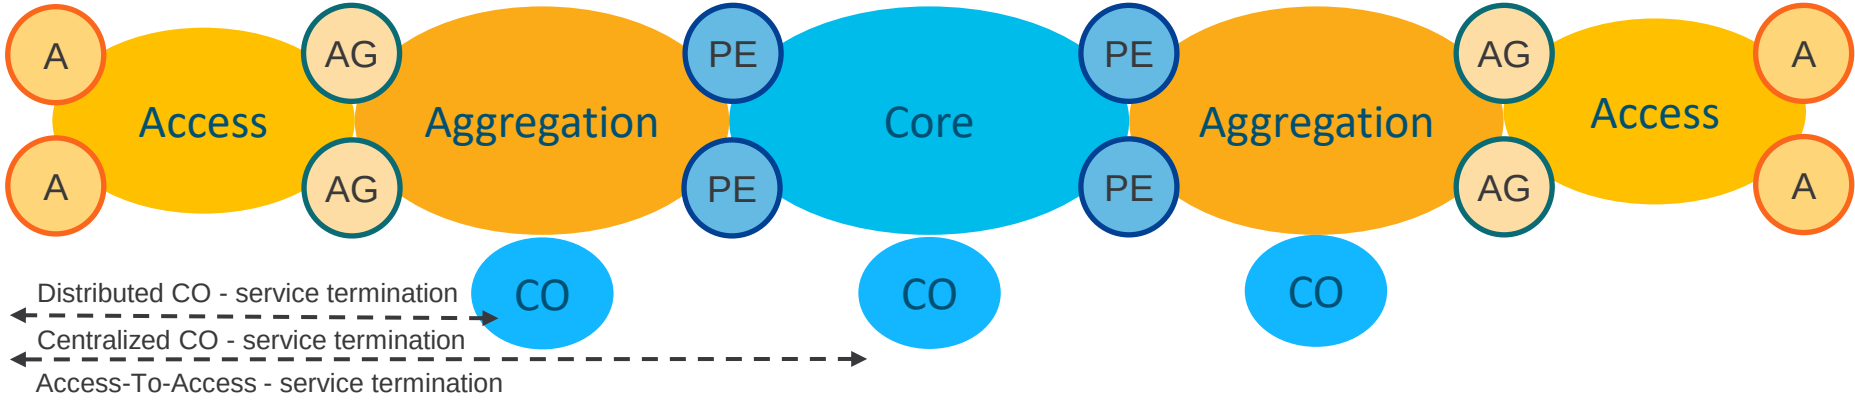

Multipoint L2/L3 Services Configuration Example With VPLS


```
l2vpn
  bridge group CISCO
    bridge-domain BD1
      interface TenGigE0/0/0/0.1
      interface TenGigE0/0/0/1.1
      routed interface bvi 20 ← BVI
      neighbor 11.1.1.1 pw-id 20
      vfi CISCO-VFI
        neighbor 1.1.1.1 pw-id 21
        neighbor 22.2.2.2 pw-id 21

Interface bvi 20 ← BVI
  ipv4 address 1.1.1.1 255.255.255.0
```

Multipoint Bridging between EFPs, if required

Routed Interface integration for L3 termination

VPLS configuration with/without Split Horizon for multipoint services through MPLS/SR netwrok

Layer 3 Interface configuration

Multipoint L2/L3 Services Configuration Example With VPLS

Split Horizon Rules

- Traffic forwardig across pseudowires may cause unwanted flooding and loops.
- Split horizon is utilized to:
  - Disable traffic forwarding on “Full Mesh” VPLS PW’s
  - Enable traffic forwarding on spoke PWs for H-VPLS to enable scale
- Split-Horizon is **disabled** for neighbors outside the vfi
- Split-Horizon is **enabled** for neighbors within the vfi

```
l2vpn
  bridge group CISCO
    bridge-domain BD1
      ... ...
      ... ...
      neighbor 11.1.1.1 pw-id 20
      vfi CISCO-VFI
        neighbor 1.1.1.1 pw-id 21
        neighbor 22.2.2.2 pw-id 21
```

Multipoint Services Configuration Example with EVPN IRB


```
evpn
 evi 100
  bgp
   route-target import 65001:100
   route-target export 65001:100
  !
  advertise-mac

interface BVI100
 host-routing
 vrf Tenant
 ipv4 address 30.10.12.1 255.255.255.0
 mac-address 1000.1000.1001

l2vpn
 bridge group bg100
  bridge-domain bd100
   interface Bundle-Ether100.1
   routed interface BVI100
     evi 100
```

Ethernet Virtual Instance
Layer2-vrf or MAC-VRF id.
Unique for each tenant.
Globally significant RT's

Bridged Virtual Interface for
L2/L3 Termination
Makes EVPN IRB Possible

EVPN Service Configuration

Dual Homed CE Configration Example with EVPN

- Ethernet Segment (ES) is a set of links that connect one tenant site to one of more PEs.
- Should be unique (10 Octets) for each segment (a segment can a pair of links from a dual-homed Host

```
evpn
 interface Bundle-Ether100
  ethernet-segment
   identifier type 0 11.11.11.11.11.11.11.11.11
   bgp route-target 1111.1111.1111
```

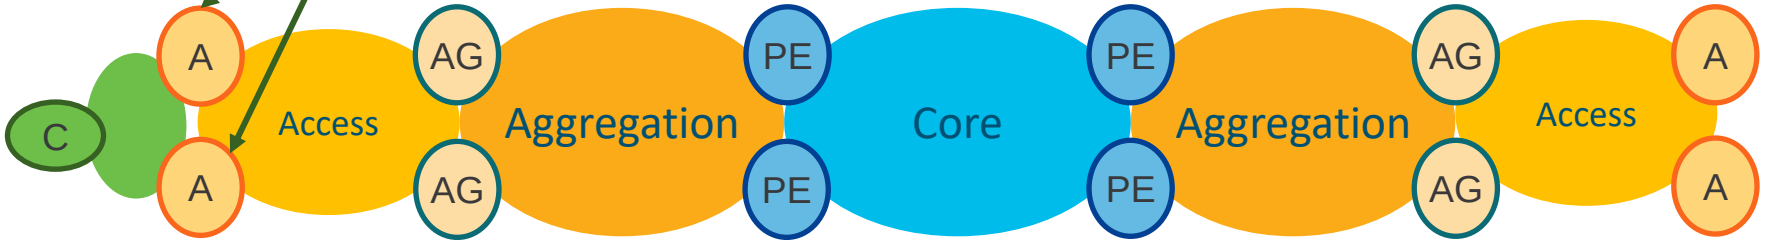

Services Orchestration and Management

Deployment model: Existing vs NSO

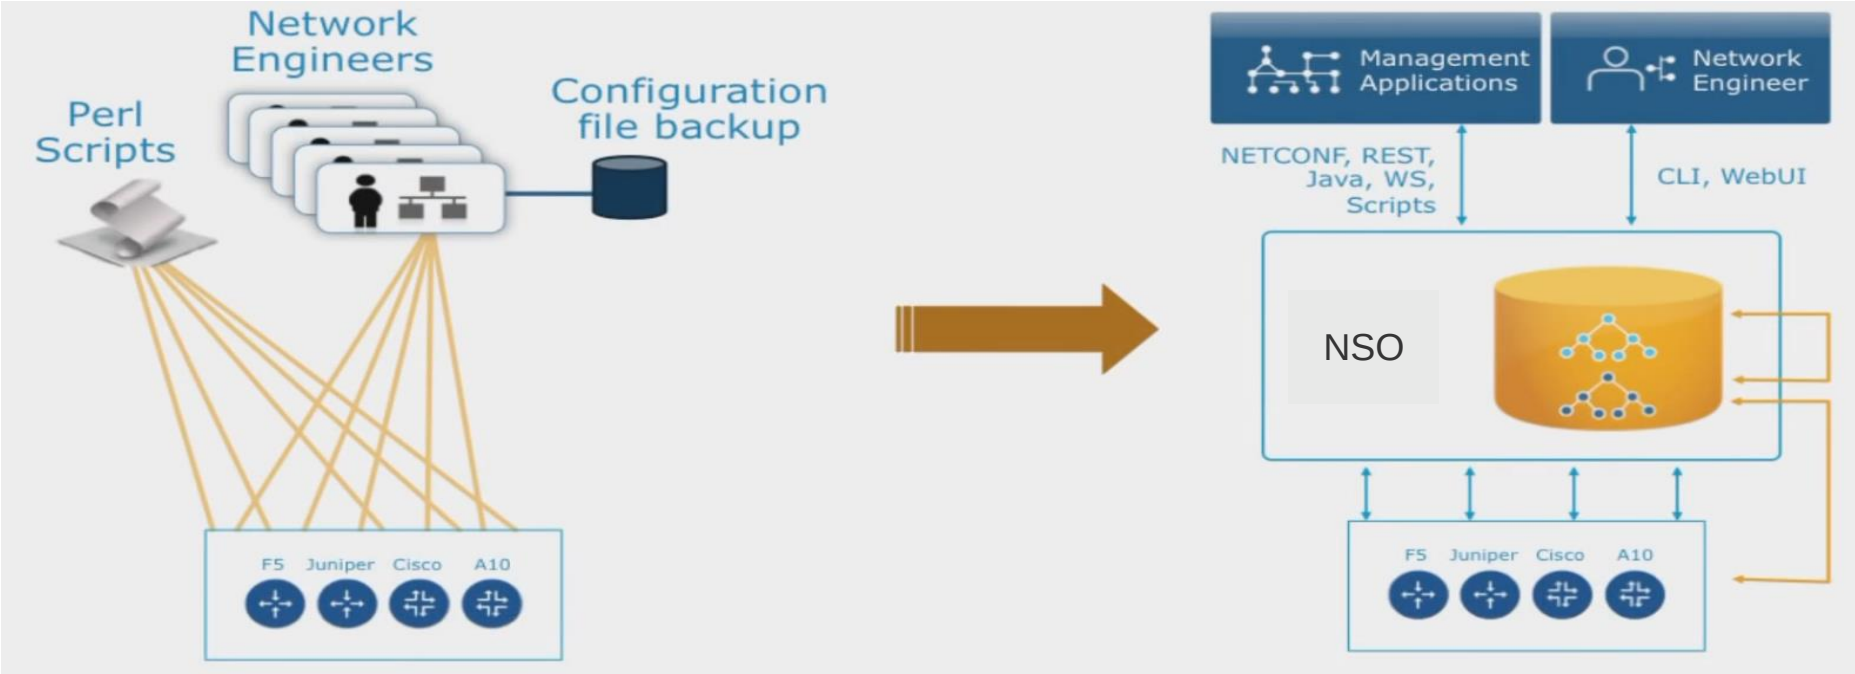

NSO Main Features

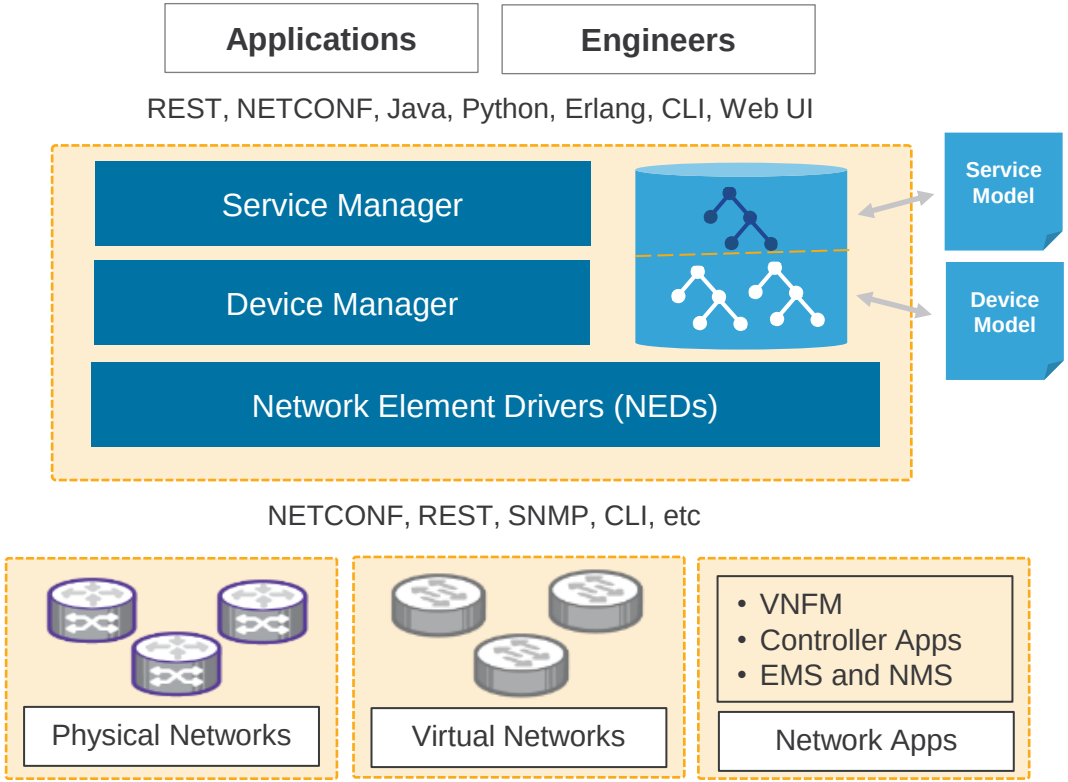

- Logically centralized network services
- Data models for data structures
- Structured representations of:
  - Service instances
  - Network configuration and state
- Mapping service operations to network configuration changes
- Transactional integrity
- Multiprotocol and multivendor support

NSO Network Element Driver (NED)

- Management support for devices - major bottleneck
- NSO uses Network Element Drivers (NED) to Communicate to any management interface
- Built in NED support for:
  - Cisco IOS
  - Cisco IOS XR
- NED packages available for many other multivendor products

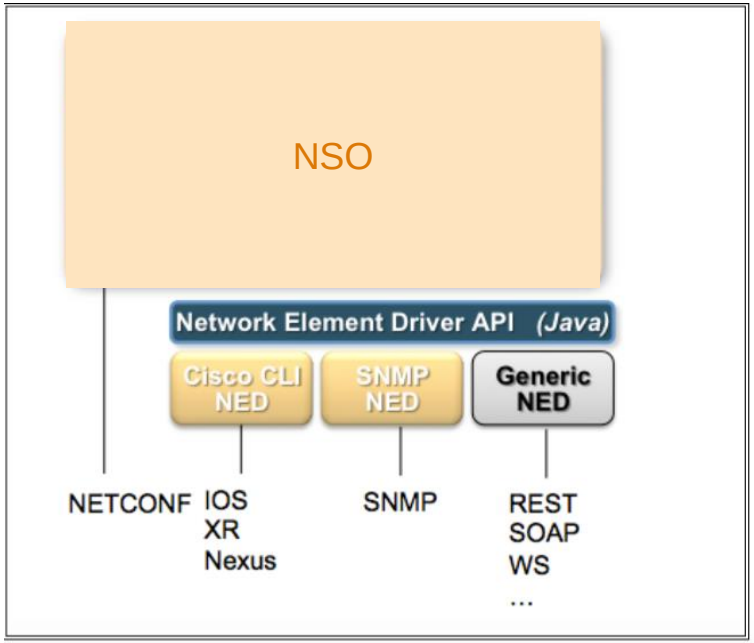

Services Automation Framework

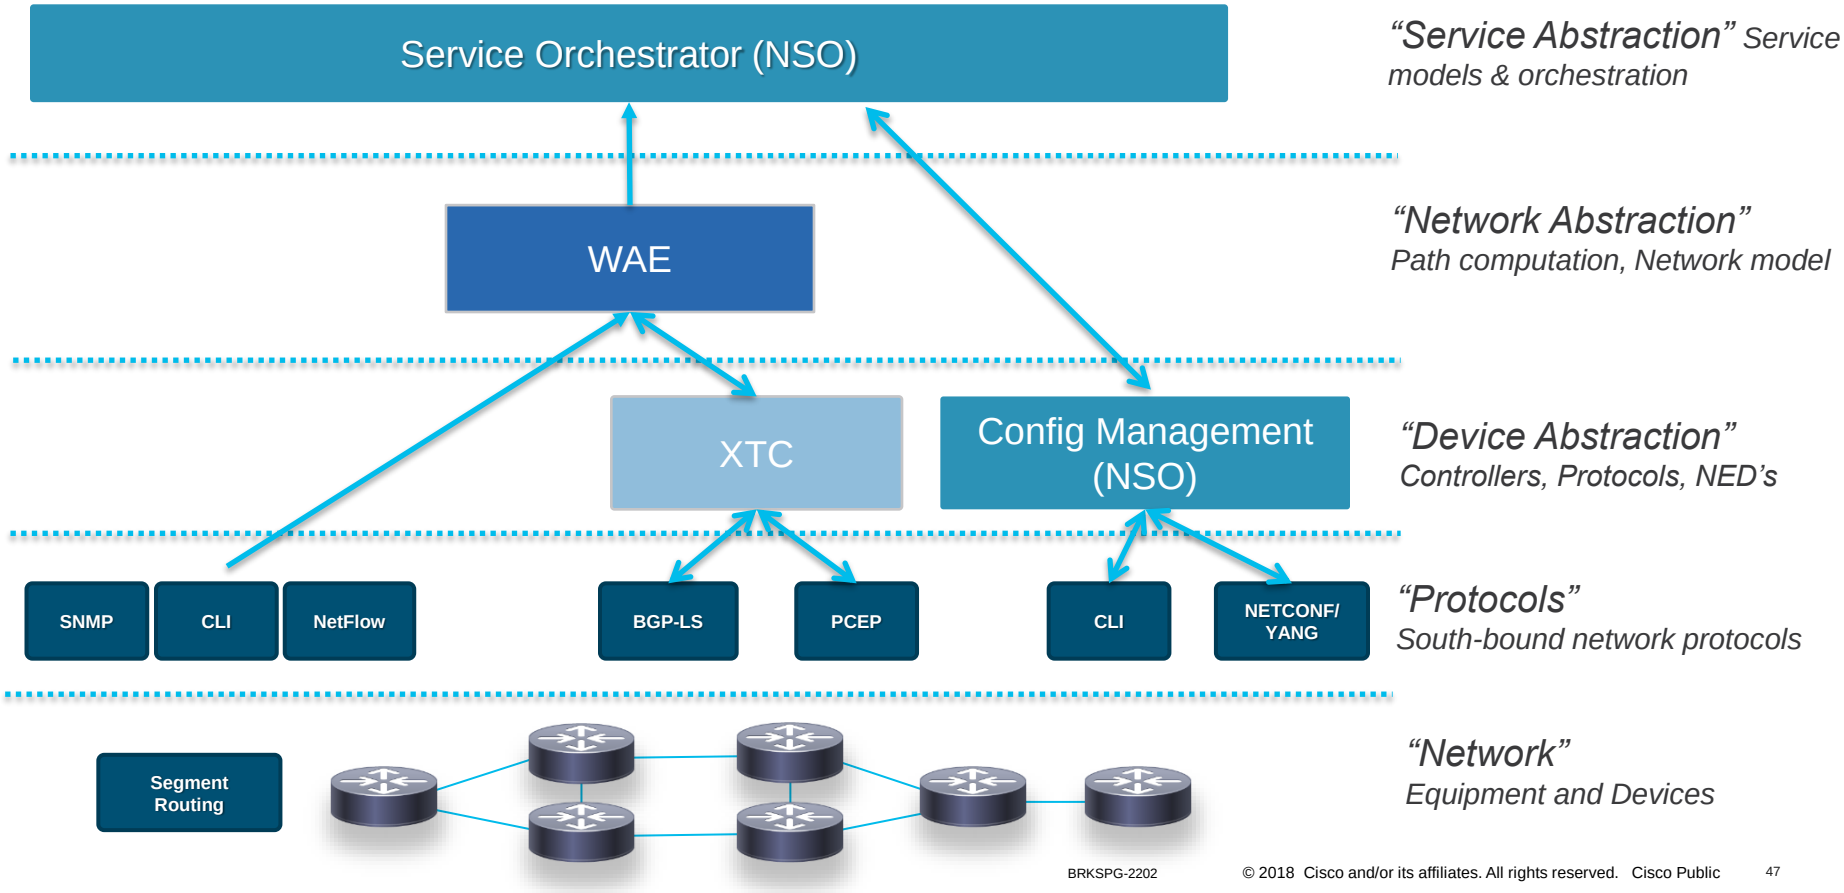

In Conclusion …

In Conclusion …

- Cisco Metro Fabric Provides a validated reference design for Metro Services Deployment
- Underlay Configuration using ISIS/OSPF for SR and BGP for EVPN
- Point to Point and Multipoint services
- Flexible, versatile Configuration toolset
- Various flavors of EVPN – EVPN VPWS, EVPN IRB – and EoMPLS/VPLS for Metro Services configure
- Services Orchestration and Management framework makes automation easier

Cisco Spark

Questions?

Use Cisco Spark to communicate with the speaker after the session

How

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

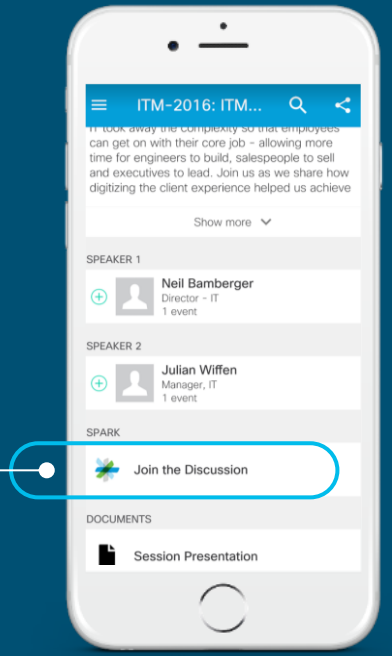

cs.co/ciscolivebot#BRKSPG-2002

Complete Your Online Session Evaluation

- Please complete your Online Session Evaluations after each session
- Complete 4 Session Evaluations & the Overall Conference Evaluation (available from Thursday) to receive your Cisco Live T-shirt
- All surveys can be completed via the Cisco Live Mobile App or the Communication Stations

Don’t forget: Cisco Live sessions will be available for viewing on-demand after the event at www.ciscolive.com/global/on-demand-library/.

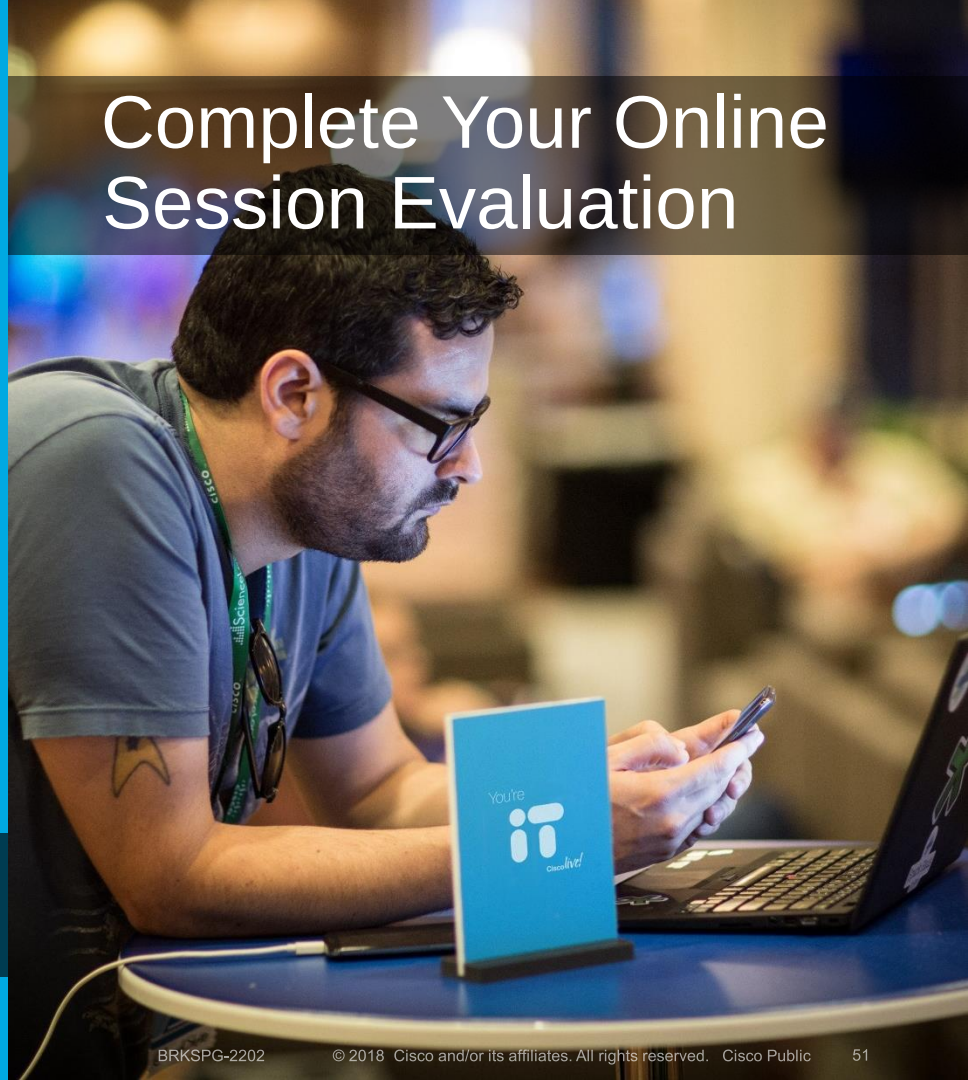

Continue Your Education

- Demos in the Cisco campus
- Walk-in Self-Paced Labs
- Tech Circle
- Meet the Engineer 1:1 meetings
- Related sessions

Thank you


cisco
You're
Cisco live!