Design Your Campus Network with EVPN-VXLAN

TANet Conf 2021

Prado Yang

Dec. 2021

Juniper driven by Mist AI

Agenda

Why EVPN-VXLAN in Campus ?

How to build EVPN-VXLAN in Campus ?

Microsegmentation in Campus

Campus Fabric Design using Mist Cloud

Why EVPN VXLAN in Campus?

Problems with Campus Networks Today

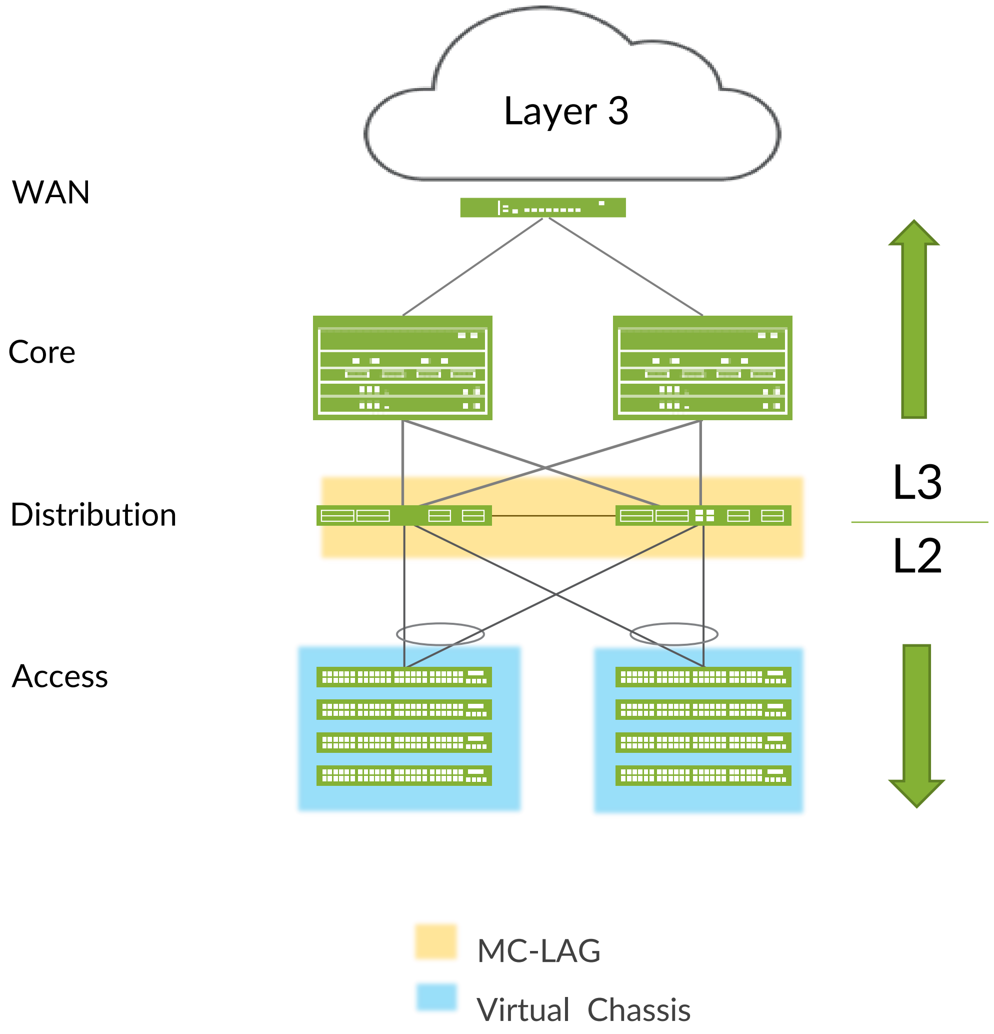

1) Layer 2 user devices connecting to layer 3 Network

2) Proprietary technologies to eliminate STP loops

3) Non-flexible and non-scalable networks

4) Not designed for Mobility and IOT

5) Increasing number of ACLs on every device

EVPN–VXLAN Solves Many Campus Problems

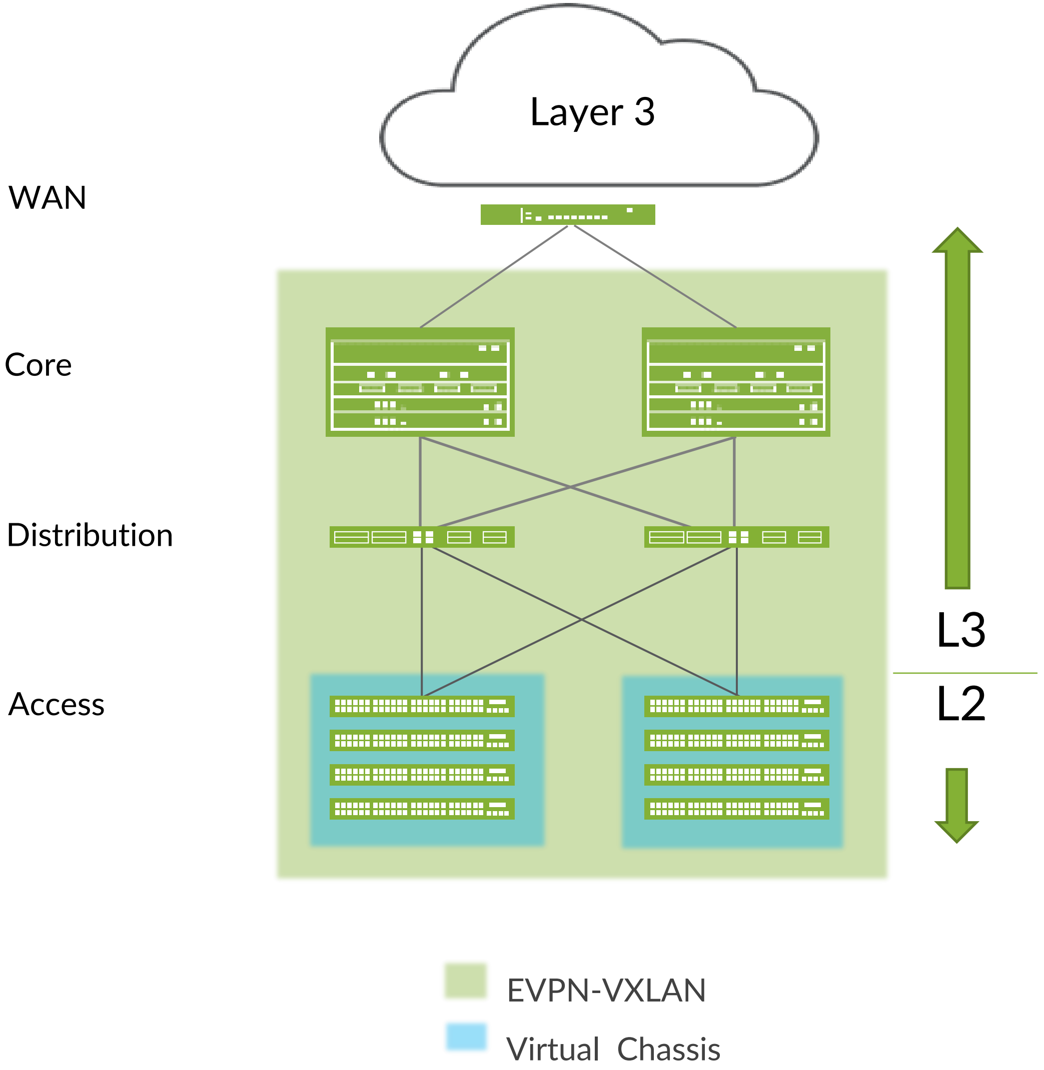

1) Layer 2 user network connecting to layer 3 internet

1) Layer 2 overlay network over layer 3

2) Proprietary technologies to eliminate STP

2) Standards based technology

3) Non-Flexible and non-scalable networks

3) Flexible and scalable

4) Not designed for Mobility and IOT

4) Fast convergence and microsegmentation

5) Increasing number of ACLs on every device

5) Network wide Group based Policies (GBP)

Why VXLAN ?

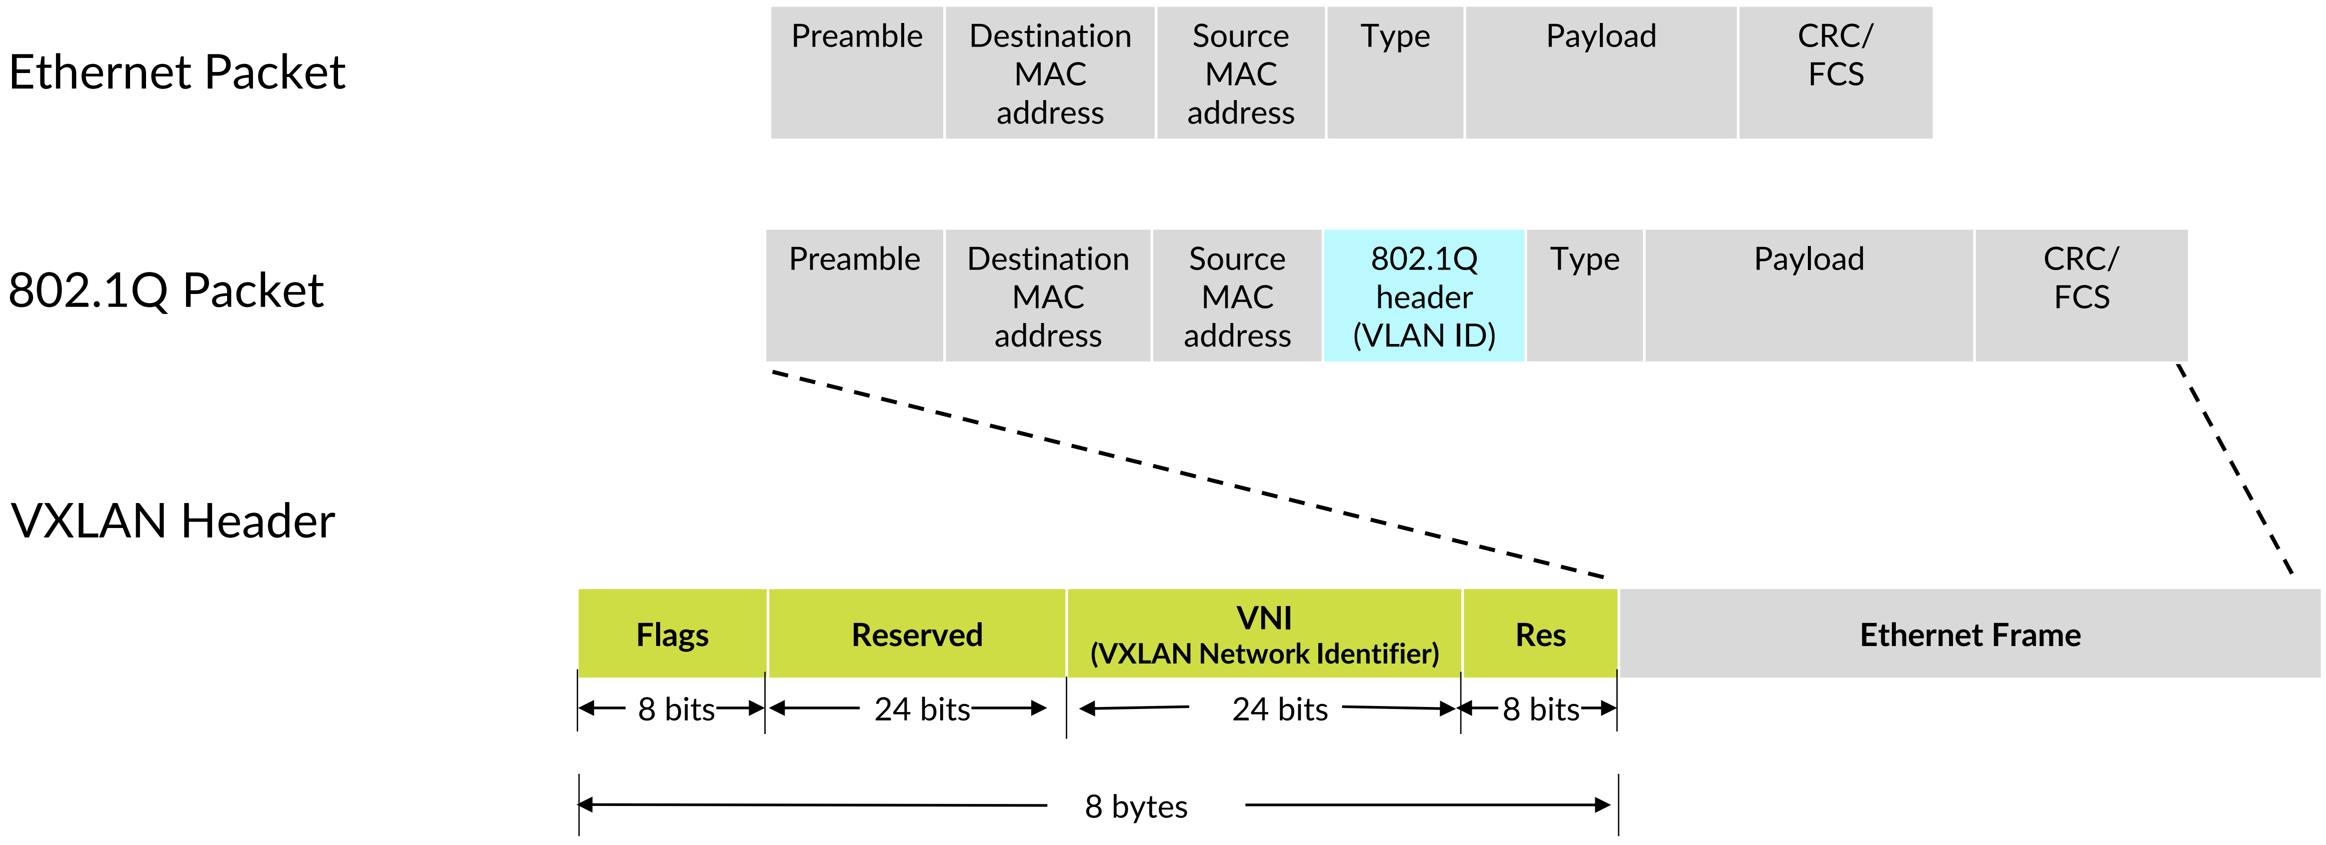

12 bit VLAN id (4K VLANs) versus 24-bit VNI (16 Million addresses possible)

** VNI VXLAN Network Identifier*

How Does VXLAN work?

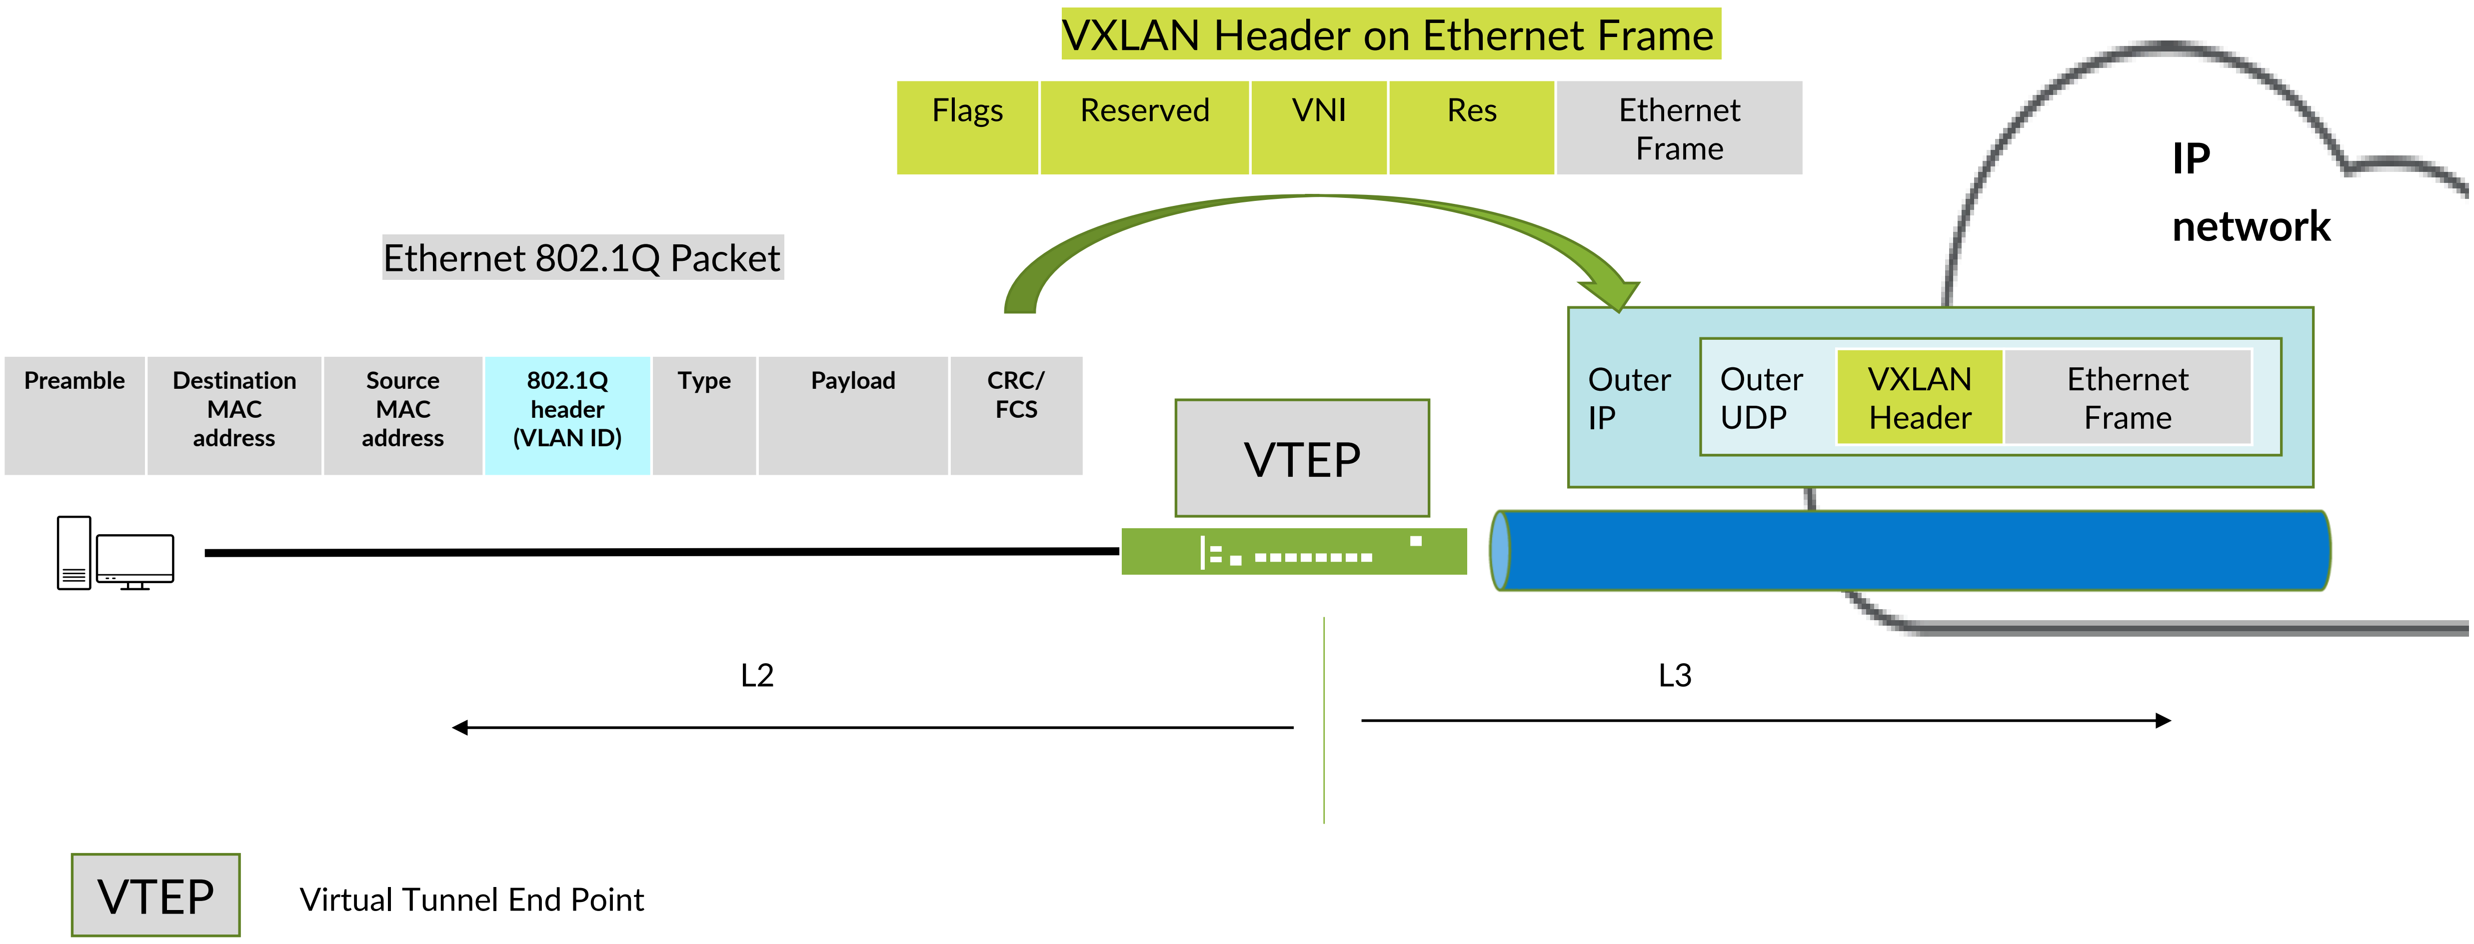

What is EVPN?

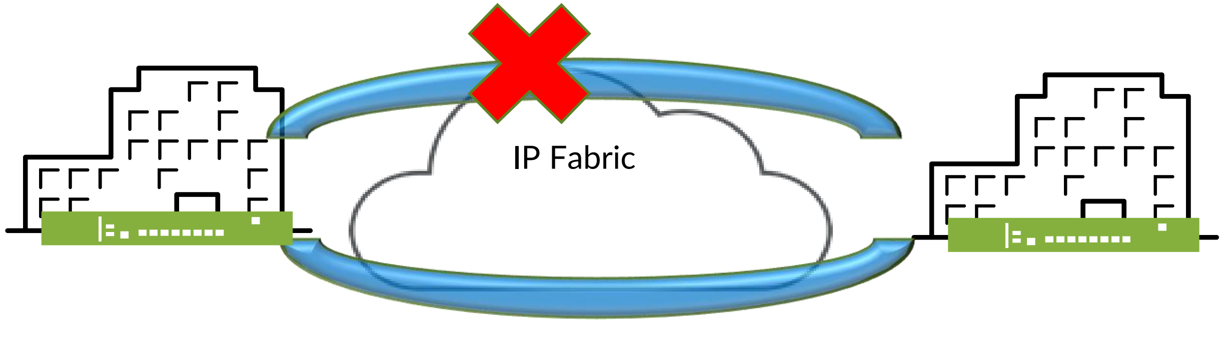

VPLS

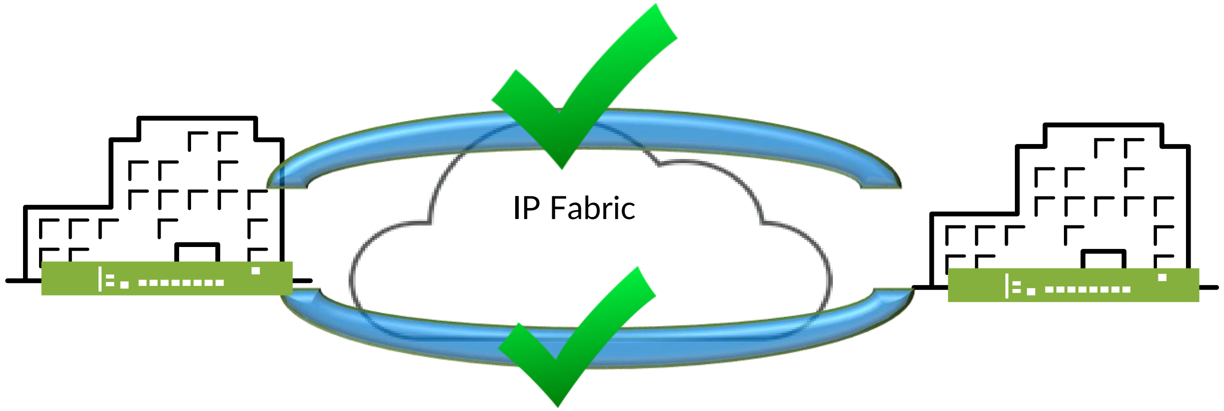

EVPN

Problem Statement

Multi path layer 2 VPN service

Limitation with VPLS

- No support for all active forwarding
- no Multipoint-to-Multipoint LSP
- Required Operators to configure a lot of parameters on top of access configuration

EVPN is BGP extension to transport layer 2 & layer 3 IP information

EVPN Benefits (RFC 7209)

- All active forwarding
- Multipoint to Multipoint LSP
- Minimize flooding of multi-destination frames

EVPN Benefits

Forwarding table

VNI	MAC address
10	a1.b1.c1.d1.e1.f1
11	a1.b1.c1.d1.e1.f2

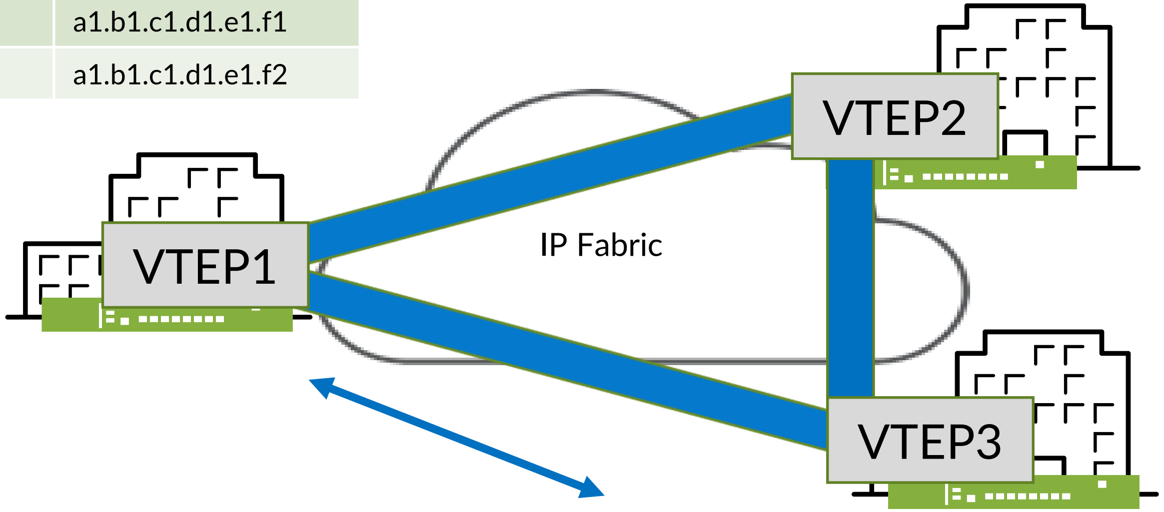

VNI	MAC address
10	a1.b1.c1.d1.e2.f1
11	a1.b1.c1.d1.e2.f2

- All active multi homing
- Multi protocol BGP (MP BGP) as overlay
  - Control plane learning
  - MAC IP bindings distributed over control plane
- MAC & IP Integrated routing and bridging
  - VPLS/VPWS a layer 2 technology
    - A separate L3 gateway needed
  - A pure Layer 3 service creates intra-subnet issues.
  - EVPN optimum for inter-subnet and intra-subnet as the packets have both MAC and IP information
- Reduces provisioning pain
  - Policy driven control on route advertisements
  - Consistent policy-based forwarding

EVPN-VXLAN Benefit in Campus #1: Flexibility

VLAN	VNI	MAC address
1	10	1A.2B.3C.4D.5E.6F
2	20	1A.2B.3C.4D.5E.6F

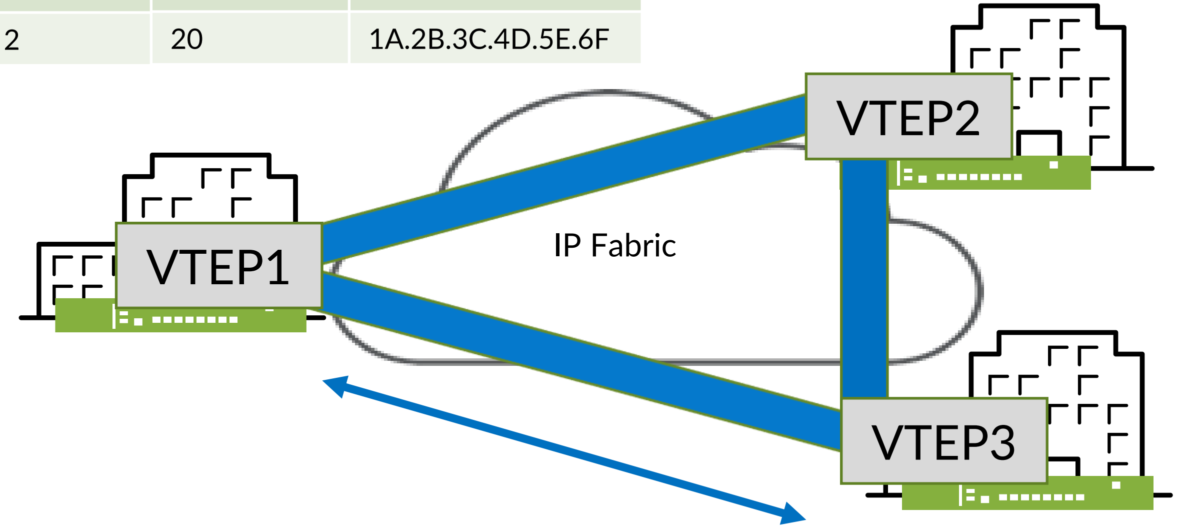

VLAN	VNI	MAC address
1	10	1A.2B.3C.4D.5E.6F
2	20	1A.2B.3C.4D.5E.6F

- Need for consistent VLANs across locations
- Current solutions inadequate
  - GRE tunnels: No redundancy
  - VPLS: No active active paths
- EVPN-VXLAN
  - No extra config needed
  - Active-active paths

EVPN-VXLAN Benefit in Campus #2: Scalability

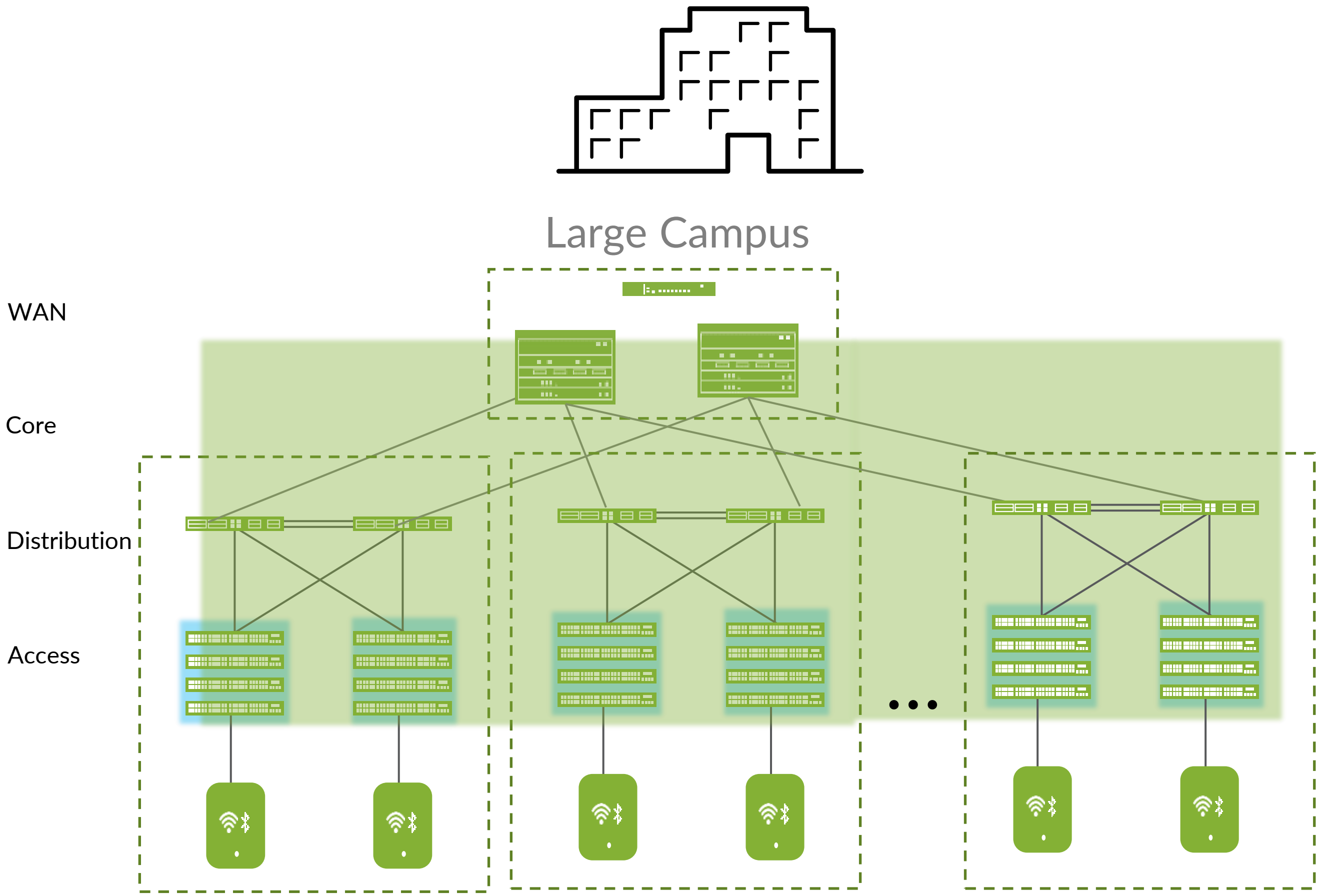

- Large namespace in overlay (16M)
- Enhanced support for both layer 2 and layer 3
- Forwarding decision made by scalable control plane (BGP)
- Integrated routing/bridging for optimized forwarding in overlay
- Fine grained policy control for better network utilization

EVPN-VXLAN Benefit in Campus #3: Micro Segmentation

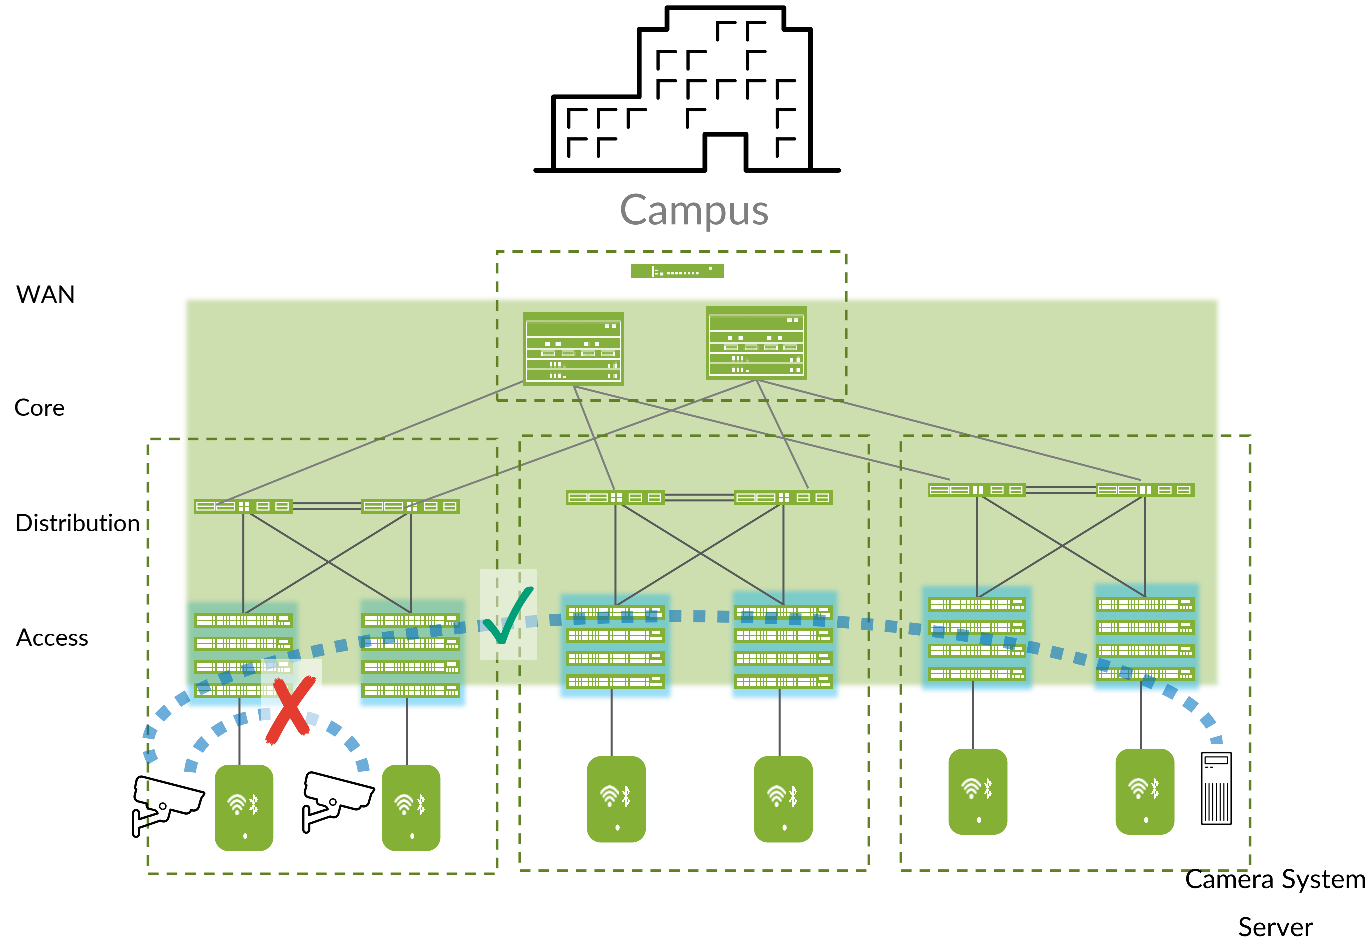

1. Replaces device specific ACLs
2. Network wide Group based policy
3. Micro segmentation
4. Macro Segmentation

EVPN-VXLAN Benefit in Campus #4: Standards Based

	Standards based
Stacking	✗
VPC	✗
Fabric Path	✗
MC-LAG	✗
Instant access	✗
EVPN-VXLAN	✓

- Lack of standards-based Technology in campus
- Previous technologies
  - Stacking
  - VPC
  - Fabric path
- Standards based
  - EVPN: RFC 7209
  - VXLAN: RFC 7348

How to build EVPN-VXLAN in Campus ?

Campus Architectures

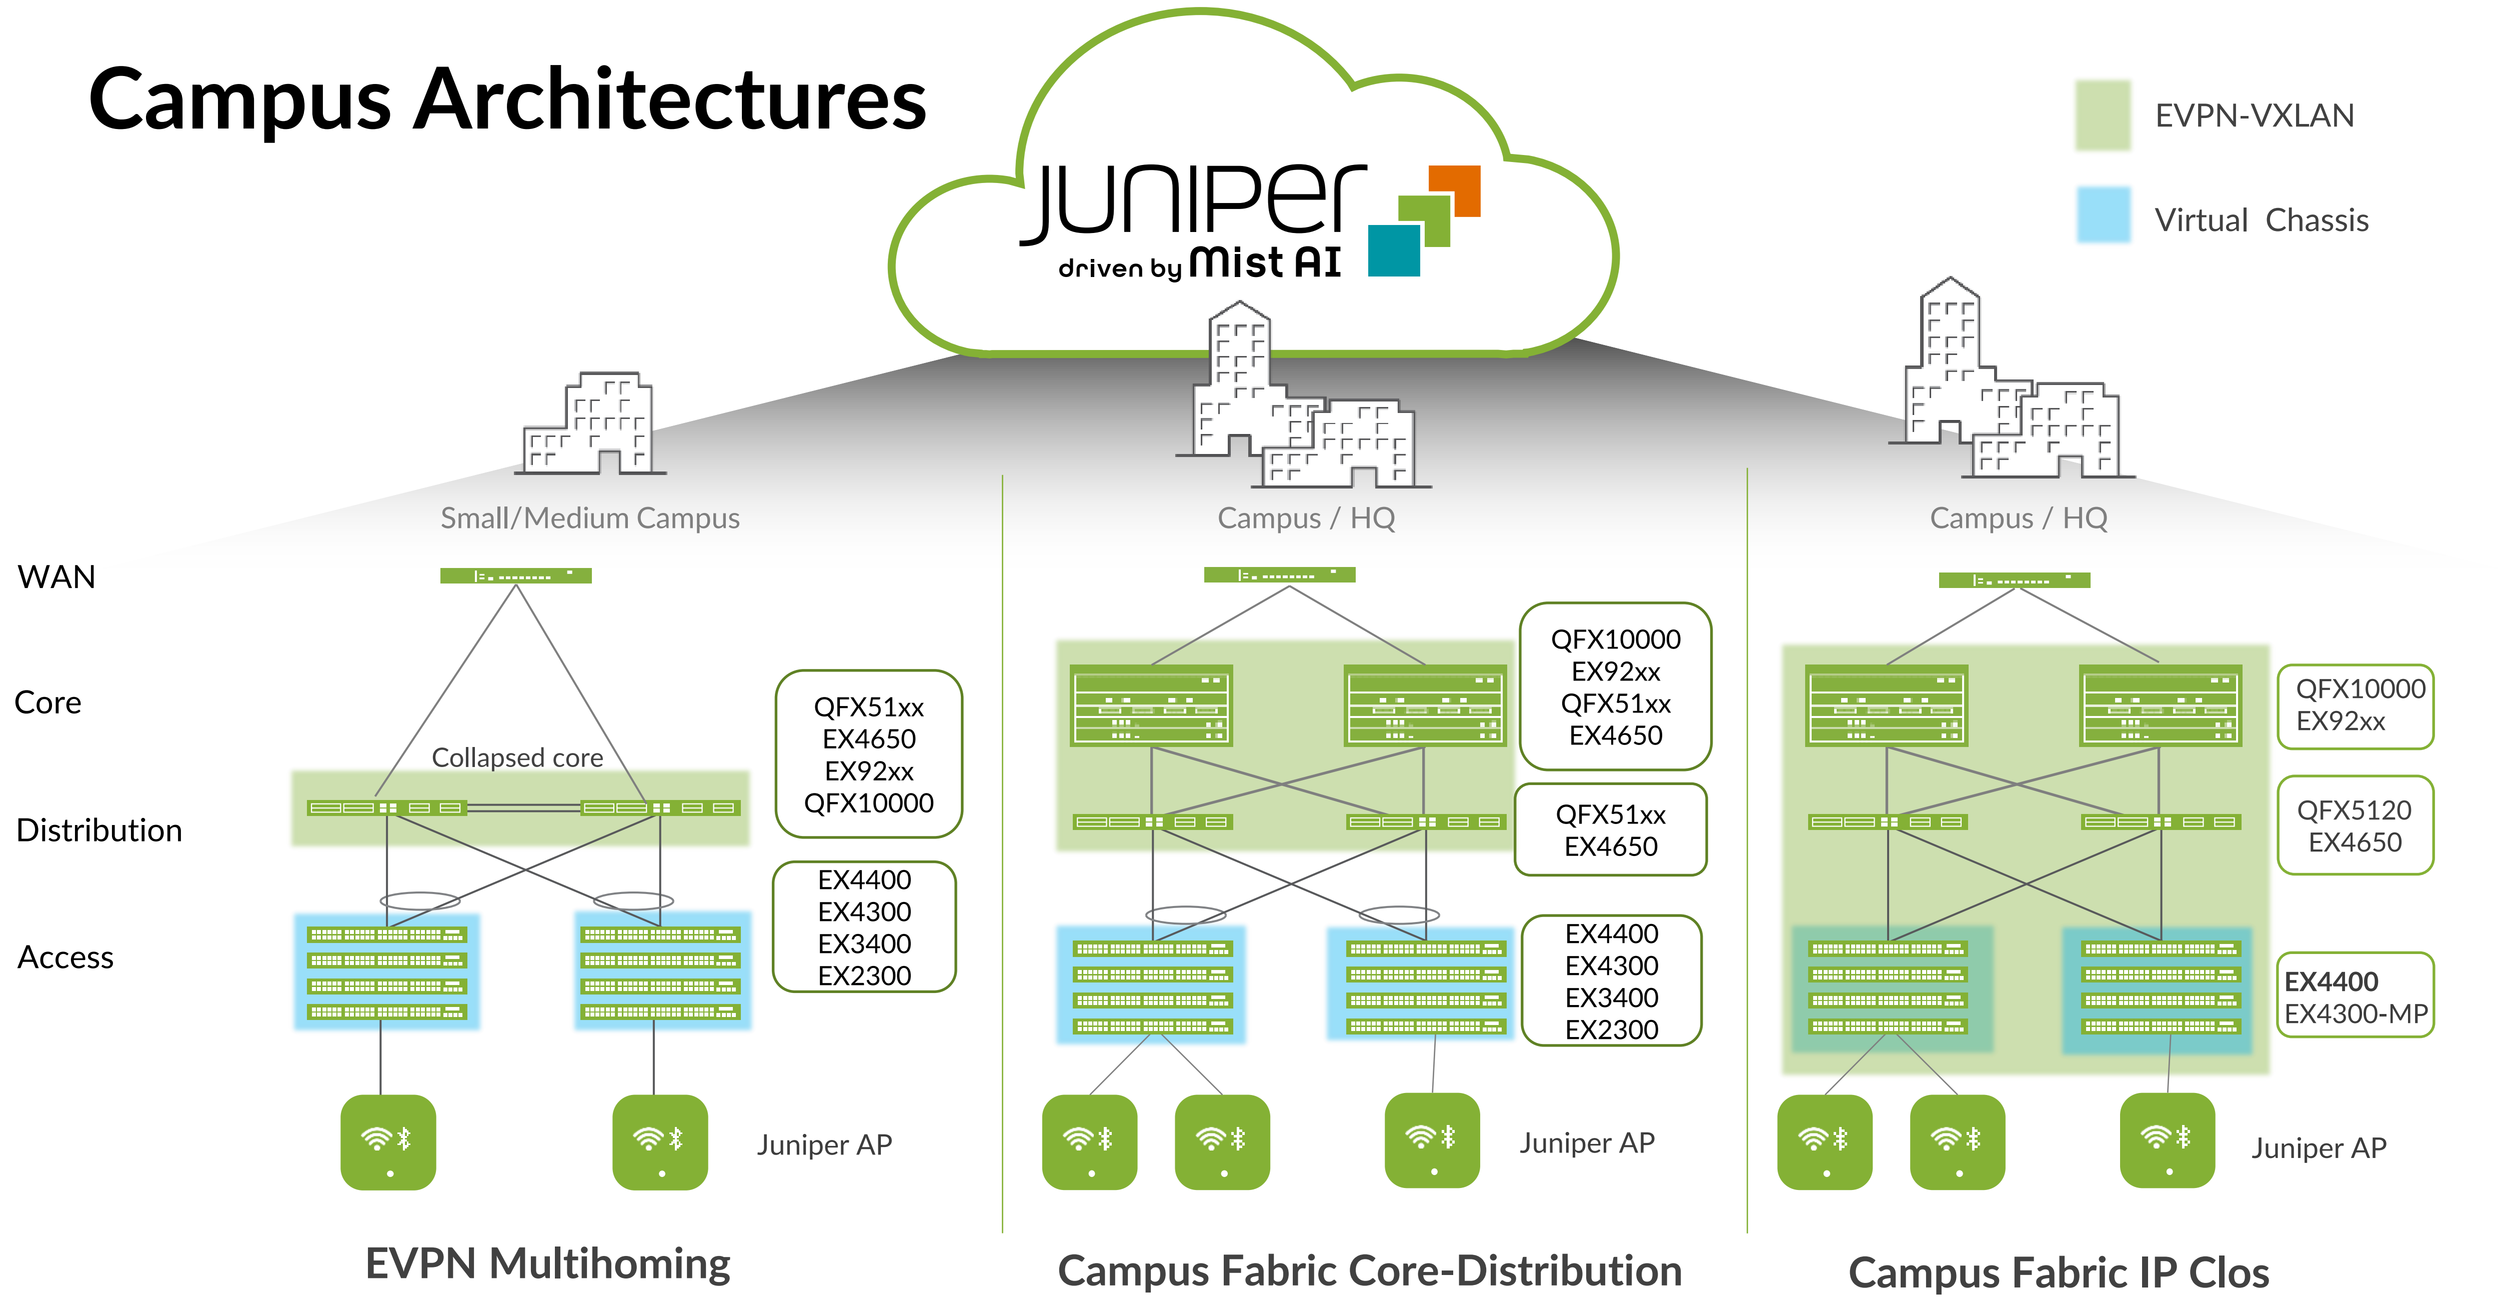

Campus Fabric Core-distribution Using CRB

Centrally-Routed Bridging

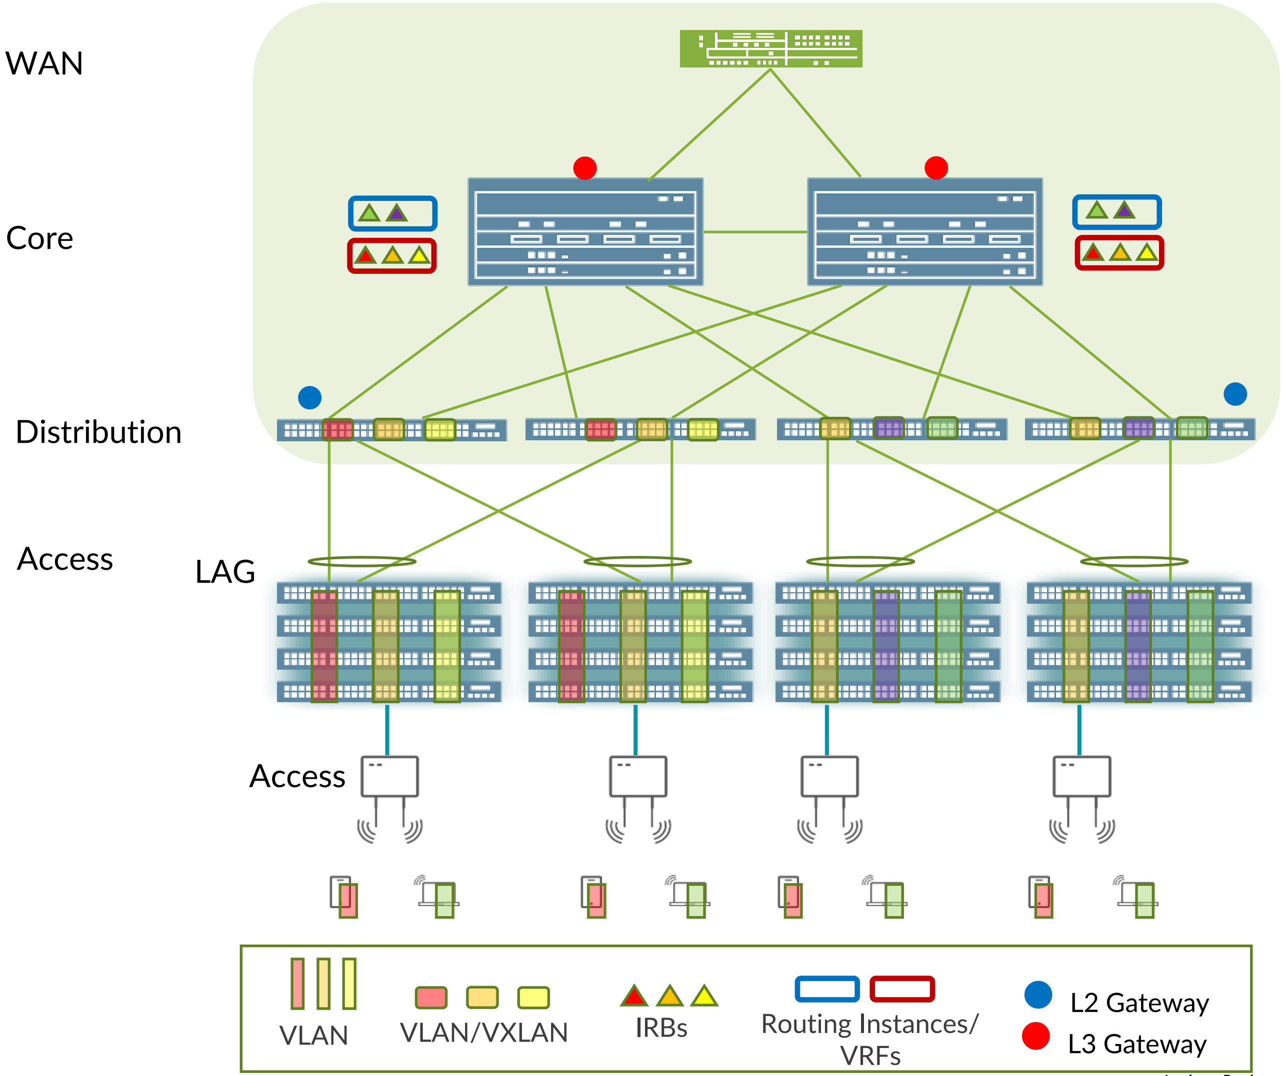

- L3 VXLAN gateway on core switches, L2 VXLAN gateway on core and distribution switches
- IRBs at the core provide L3 routing services
  - Simpler configurations as IRBs are only defined at the core
- Traffic is placed in the appropriate VLAN/VXLAN in the distribution layer
  - Enables location agnostic endpoint connectivity
  - Same default gateway address for a given L2 domain anywhere in the campus or across campuses

[NCE: Configuring an EVPN-VXLAN Campus Fabric with CRB]

Campus Fabric Core-distribution Using ERB

Edge-Routed Bridging

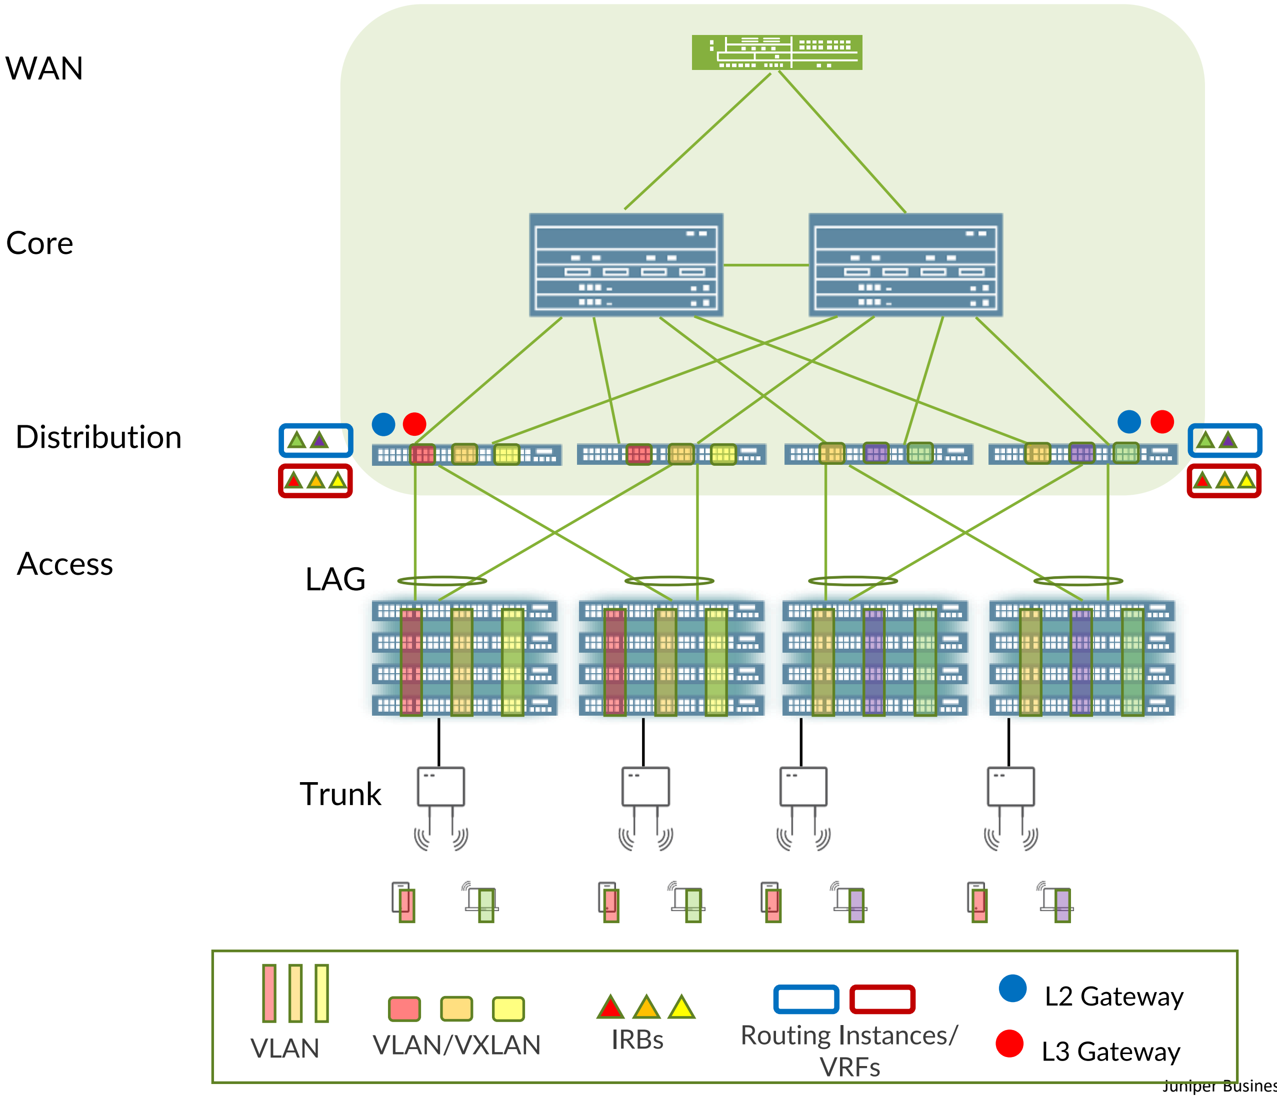

- L2/L3 VXLAN gateways are configured on distribution devices
  - IRB interfaces for VLANs/VXLANs are defined at distribution to provide L3 routing services
- Core layer provides IP underlay routing only
- Traffic is placed in the appropriate VLAN/VXLAN at the distribution layer
  - Enables location agnostic endpoint connectivity
  - Same default gateway address for a given L2 domain anywhere in the campus or across campuses

[NCE: Configuring an EVPN-VXLAN Campus Fabric with ERB]

Campus Fabric IP Clos

End-to-end EVPN-VXLAN

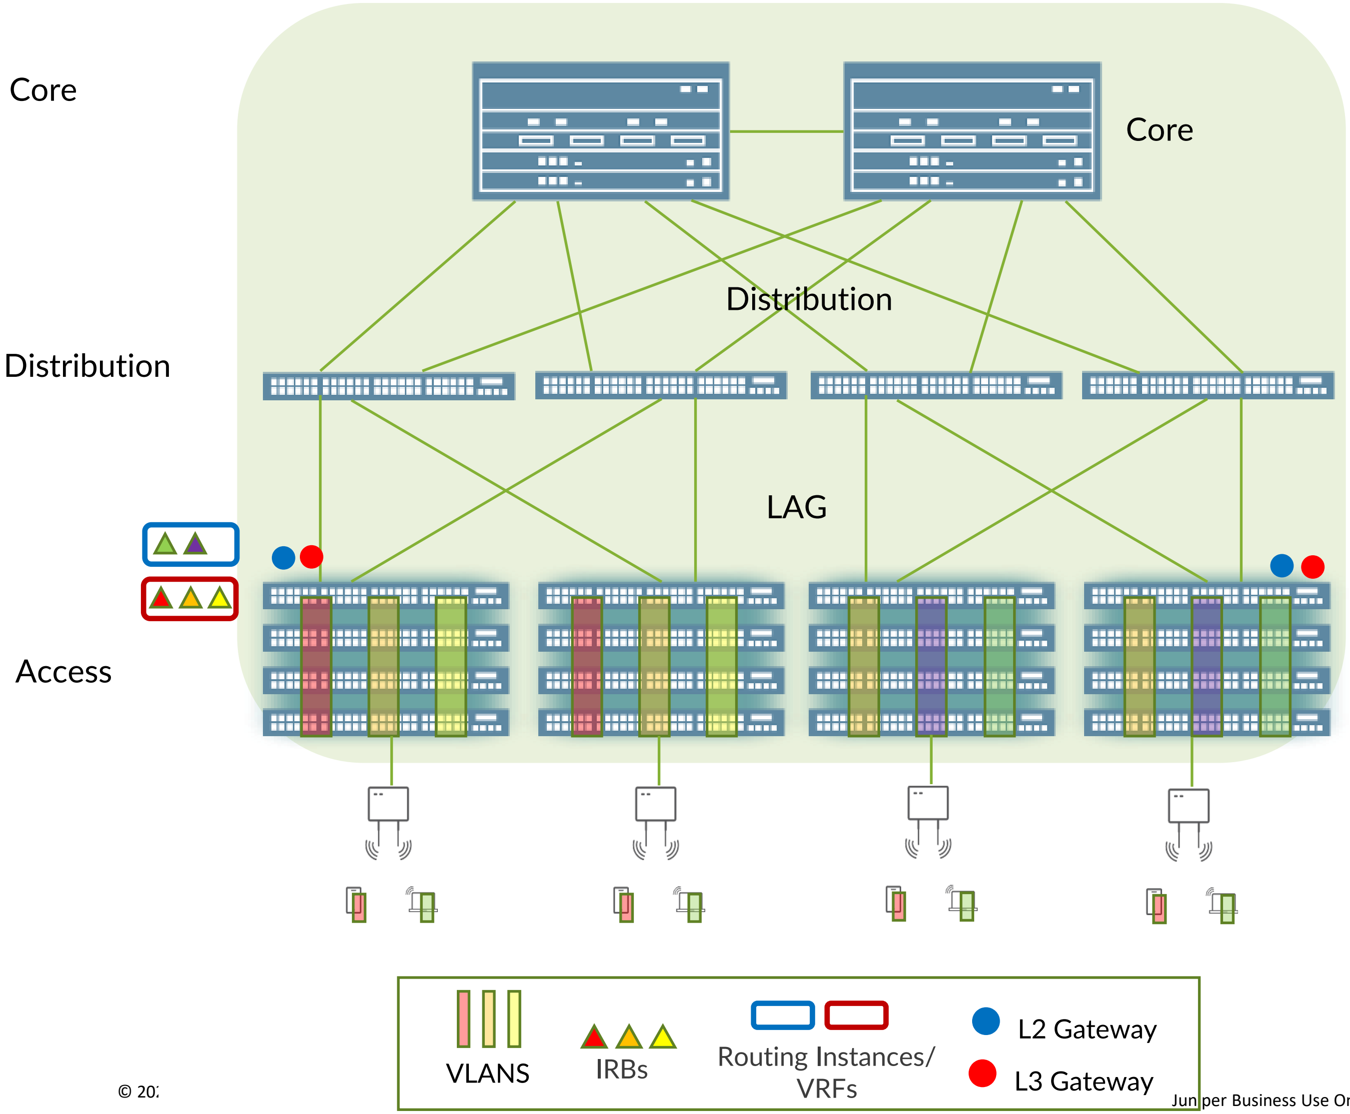

- VXLAN L2 gateway extended to the access layer with the new EX4400
- L2/L3 VXLAN gateway at the access layer
  - Access switches can be part of a Virtual Chassis
- Traffic is placed in the appropriate VLAN/VXLAN at the distribution layer
  - Enables location agnostic endpoint connectivity
  - Same default gateway address for a given L2 domain anywhere in the campus or across campuses

Charles Clos – 1953

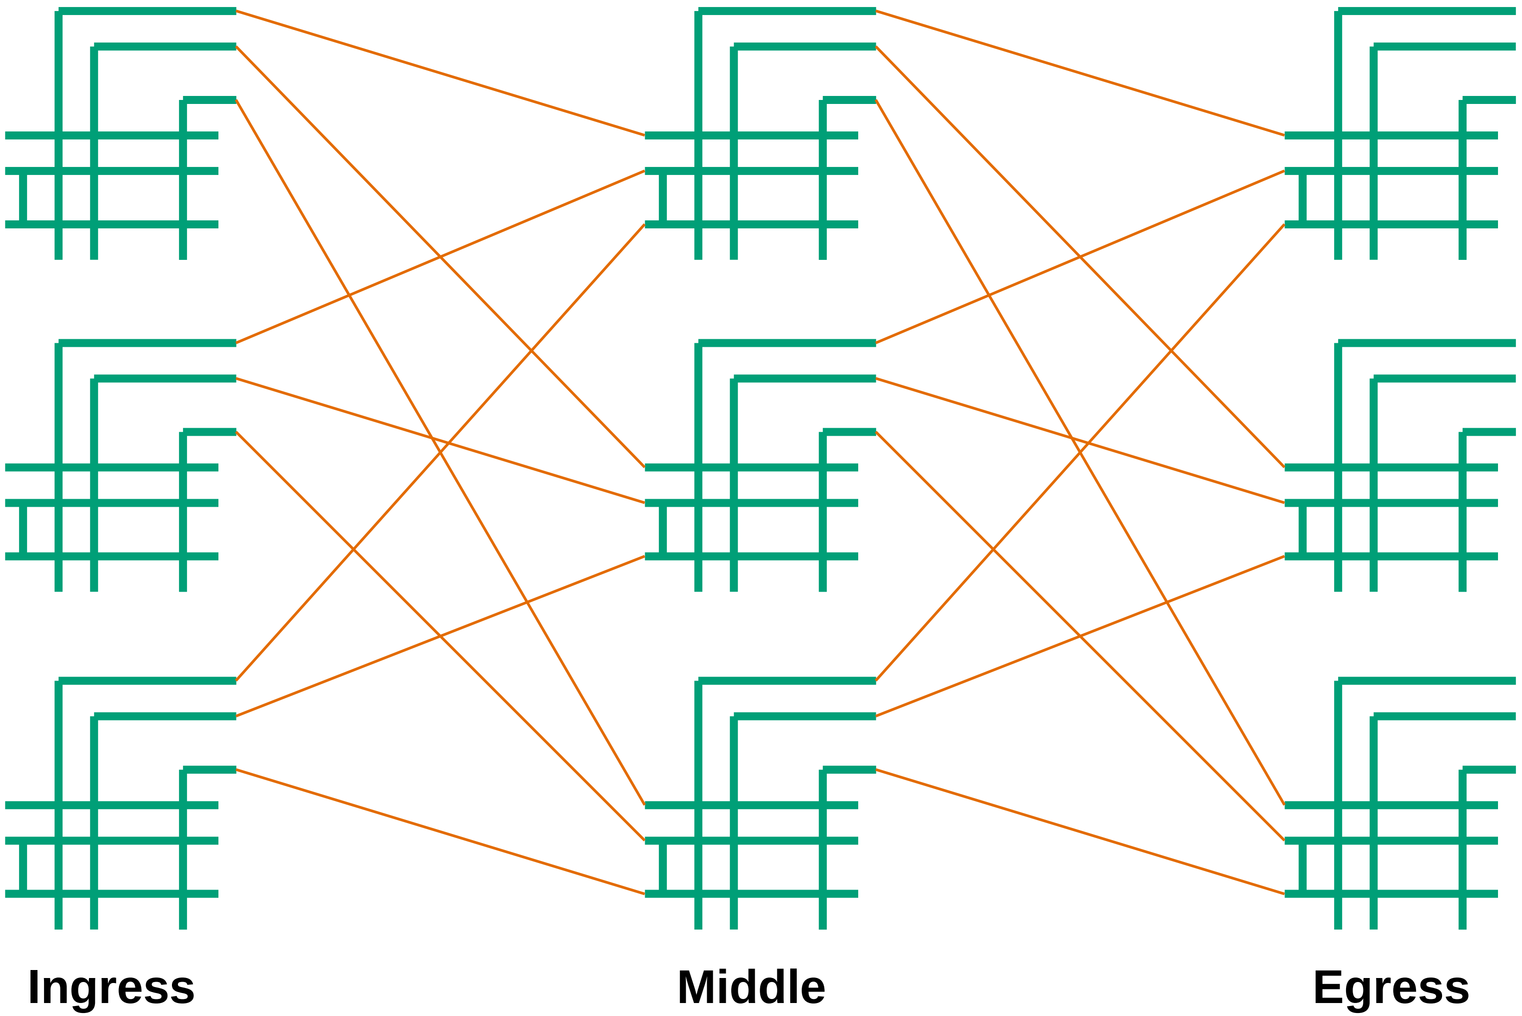

Spine and Leaf

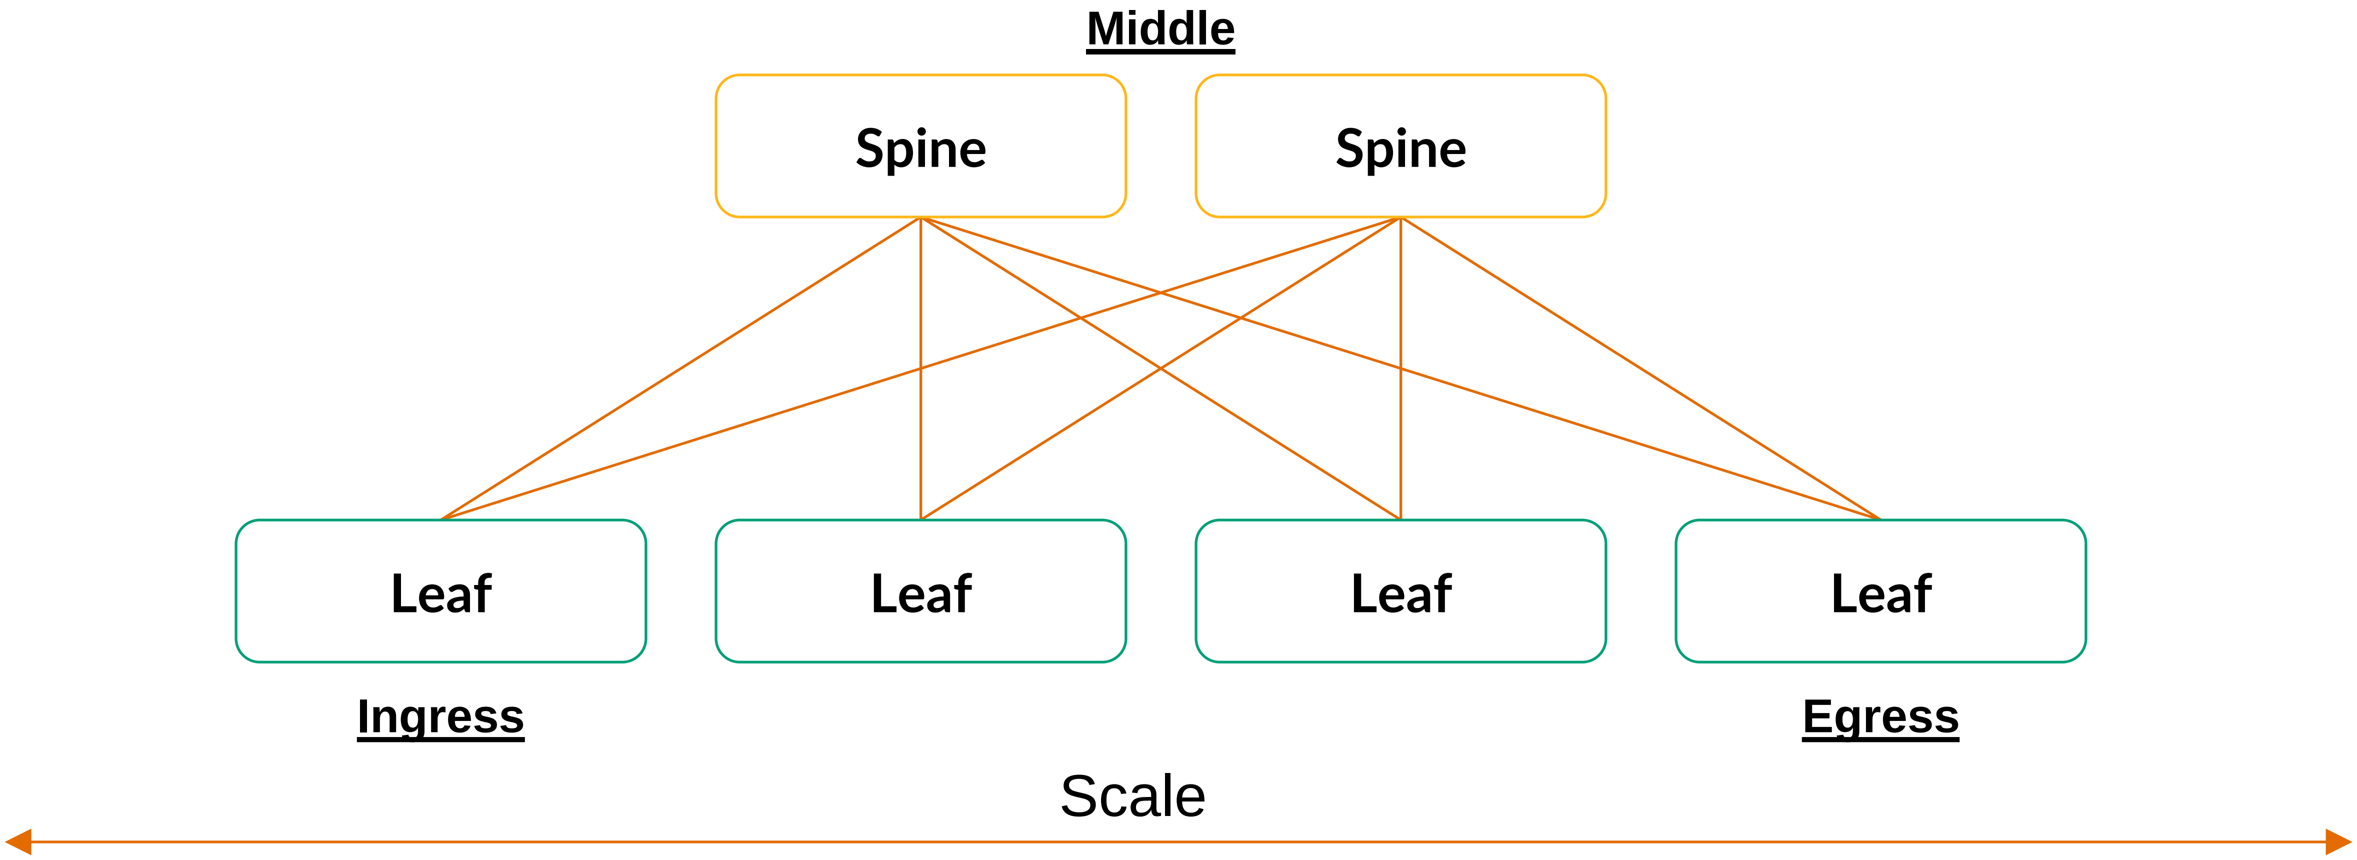

IP Clos Network Requirements

Requirement	OSPF	IS-IS	BGP
Advertise prefixes	Yes	Yes	Yes
Scale	Limited	Limited	Yes
Traffic Engineering	Limited	Limited	Yes
Traffic Tagging	Limited	Limited	Yes
Multi-Vendor Stability	Yes	Yes	Even more so

5 Steps to Building EVPN-VXLAN in Campus

EVPN-VXLAN Campus Fabrics

Underlay **L2 Gateway** **LAG to Fabric**

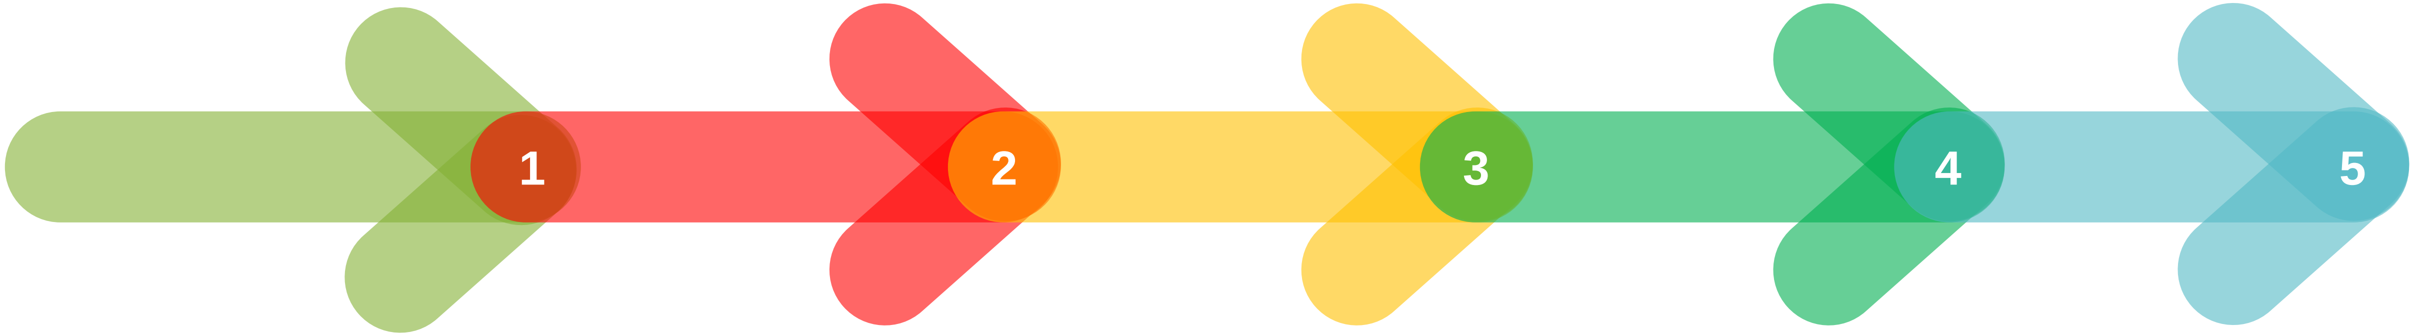

Overlay **L3 VXLAN Gateway**

EVPN-VXLAN Campus Fabrics

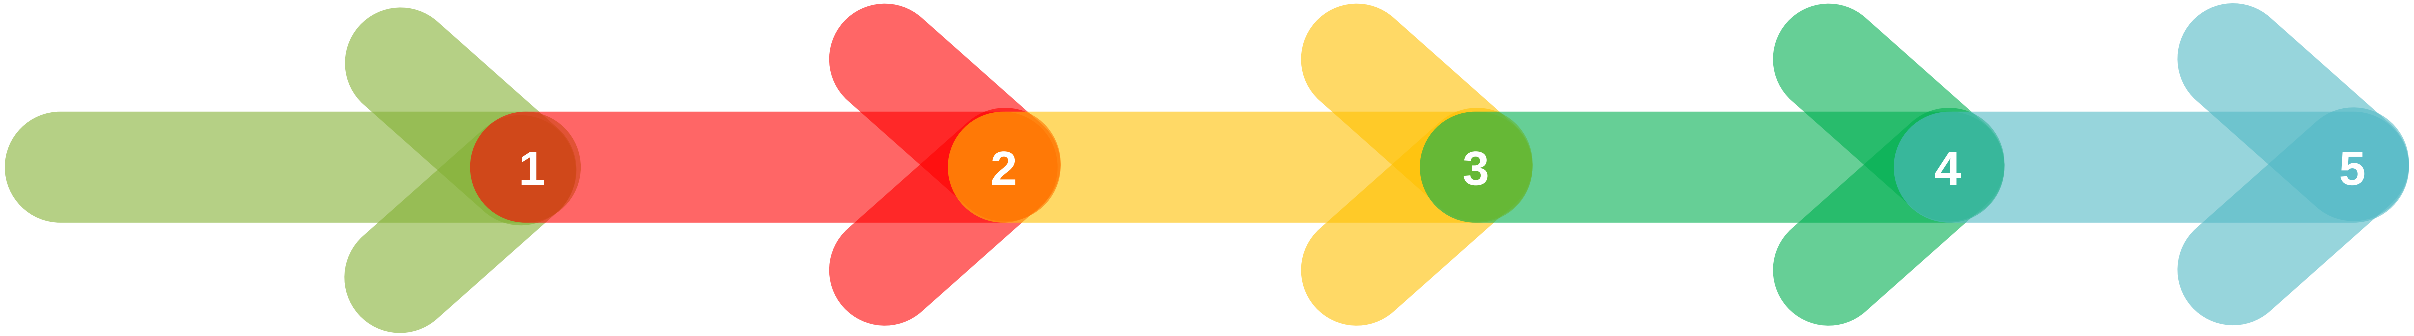

1. Underlay

- OSPF
- eBGP*
- ISIS

* preferred

2. Overlay

- iBGP peering
- Route Distinguisher (RD)
- EVPN v4 address

3. L2 Gateway

- VLAN to VXLAN mapping

4. L3 Gateway

- IRB definition
- VRF definition
- VXLAN to VXLAN routing
- VXLAN to VLAN routing

5. LAG to Fabric

- LAG from access to fabric
- ESI-LAG from fabric to access switches

1. Simple IP Fabric Underlay

Underlay | Overlay | L2 Gateway | L3 VXLAN Gateway | LAG to Fabric

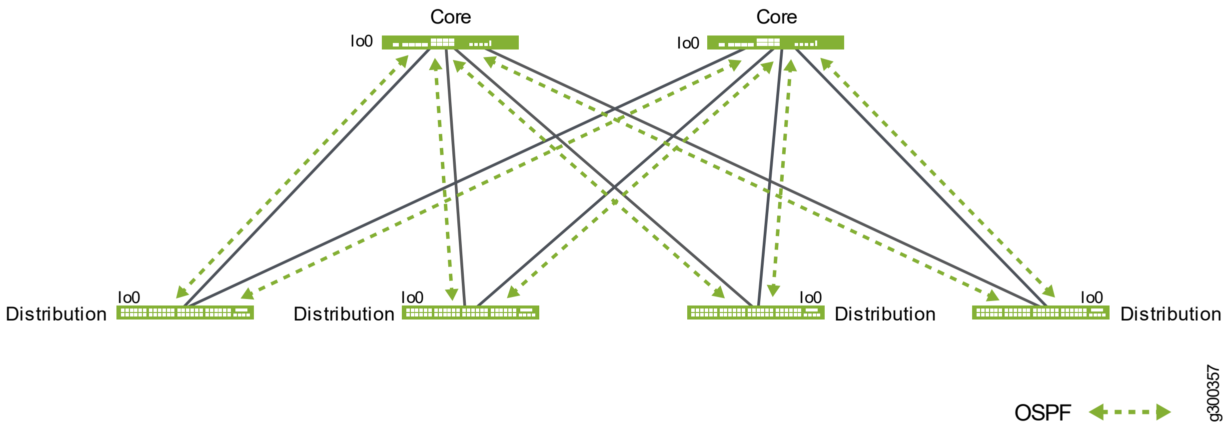

- Simple Layer 3 fabric at the core and distribution layer
- No Spanning Tree or proprietary L2 multi-chassis technologies
- Topology agnostic
  - IP-Clos topology recommended
    - Consistent scale out architecture
    - Predictable performance and scaling
- Use OSPF or eBGP to enable loopback reachability between all boxes

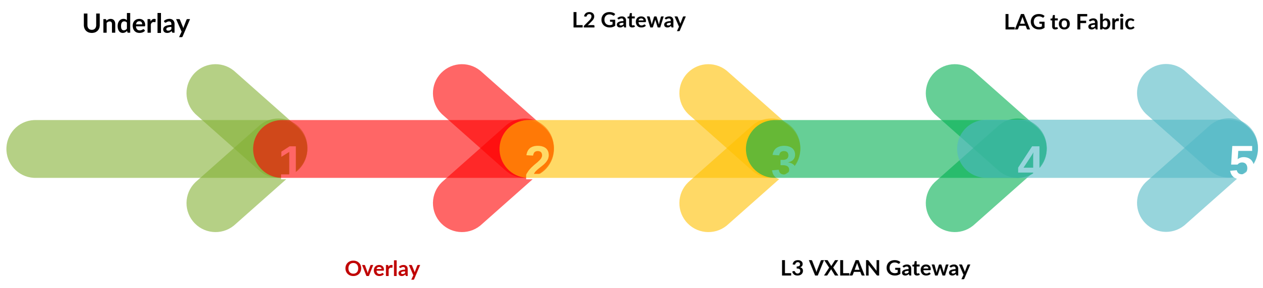

2. Overlay Control Plane

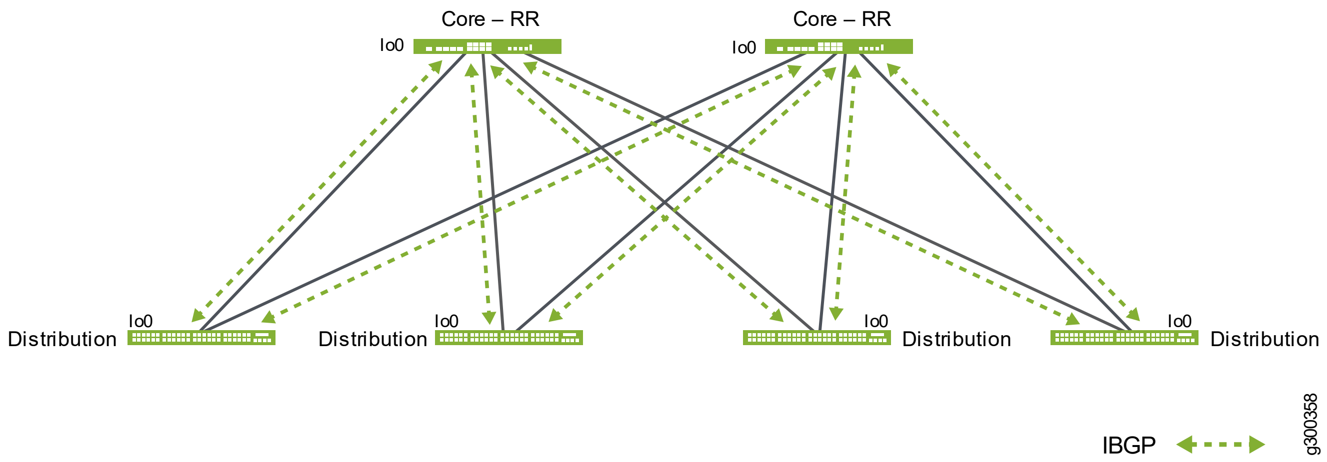

- MP-BGP EVPN control plane
- iBGP between the loopbacks
  - Core to core
  - Distribution to core
- Core switches act as Route Reflectors
  - Eliminates need for full mesh BGP
  - Consistent BGP configuration on all distribution switches

3. L2 Gateway Config

Core Isolation	Core isolation function working in conjunction with LACP, automatically brings down all Layer 2 Ethernet Switch Identifier (ESI) link aggregation group (LAG) interfaces on the switch. This should be disabled.
VXLAN Encapsulation	Configures a VXLAN encapsulation type. VNI list establishes which VXLAN virtual network identifiers (VNI) can be propagated over the L3 overlay.
Mapping VLAN to VNI	Map VLAN to a unique VNI

4. L3 VXLAN Gateway

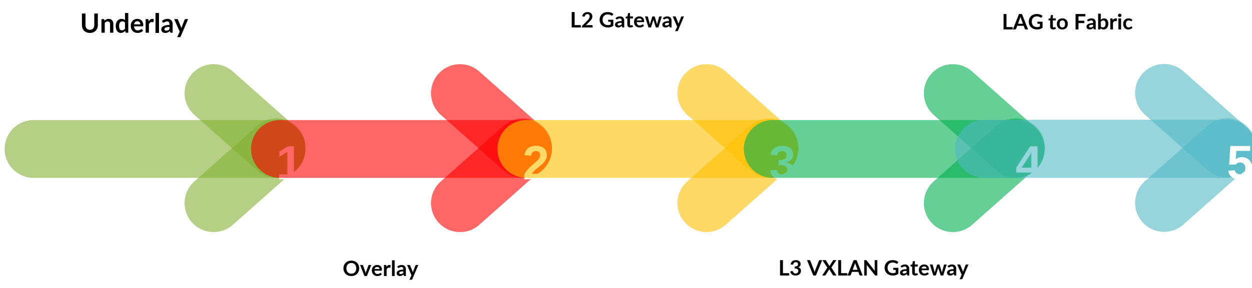

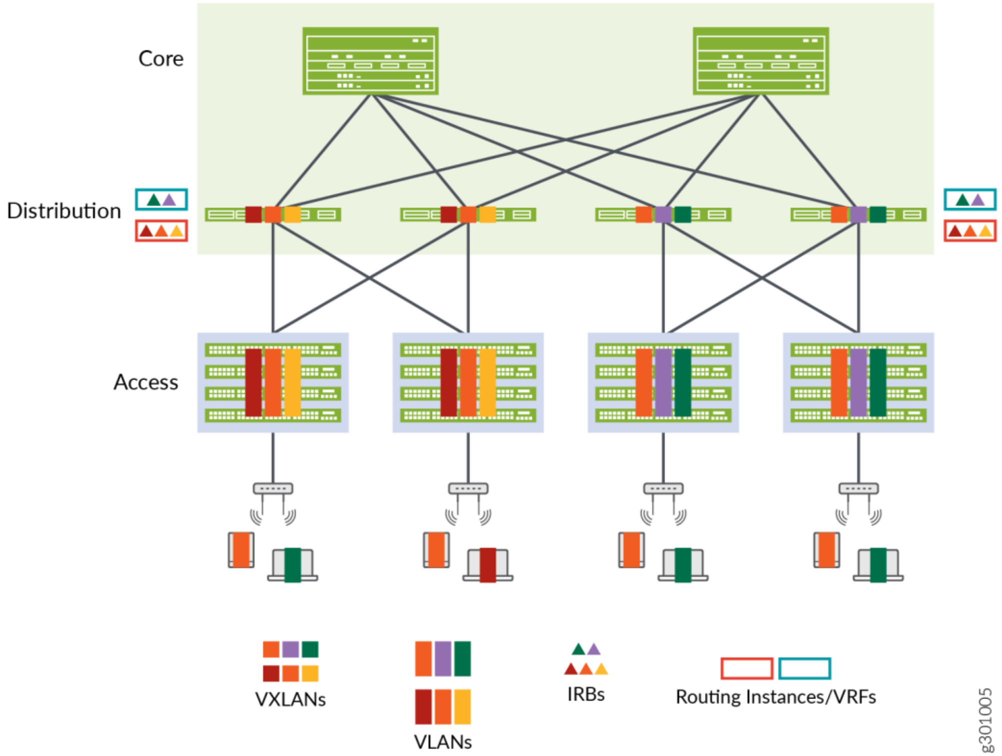

- IRBs can be placed in the same VRF
  - All subnets in a single routing table instance and have reachability to each other
- IRBs can be placed in different VRFs
  - Subnets part of the same VRF will have a single routing table instance and will have reachability to each other
  - Subnets part of different VRFs will have separate routing table instances and can communicated with each other only if routes are explicitly leaked between the VRFs
  - Inter-VRF traffic can also be forced to be routed through a stateful firewall for advanced security between VRFs

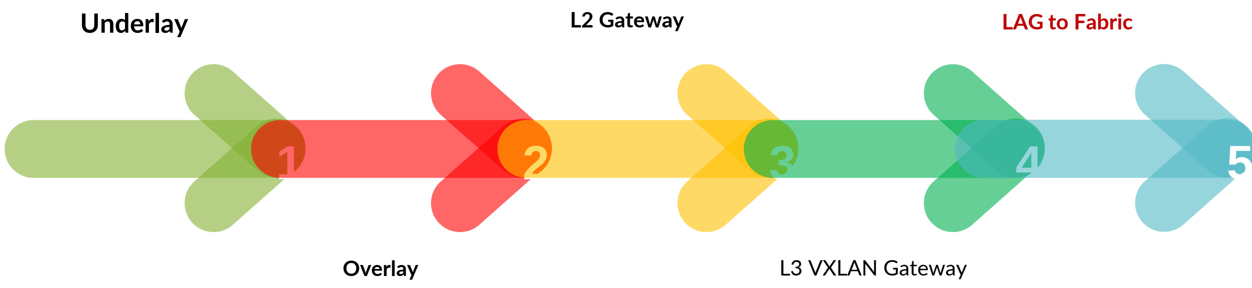

5. LAGs to EVPN Fabric

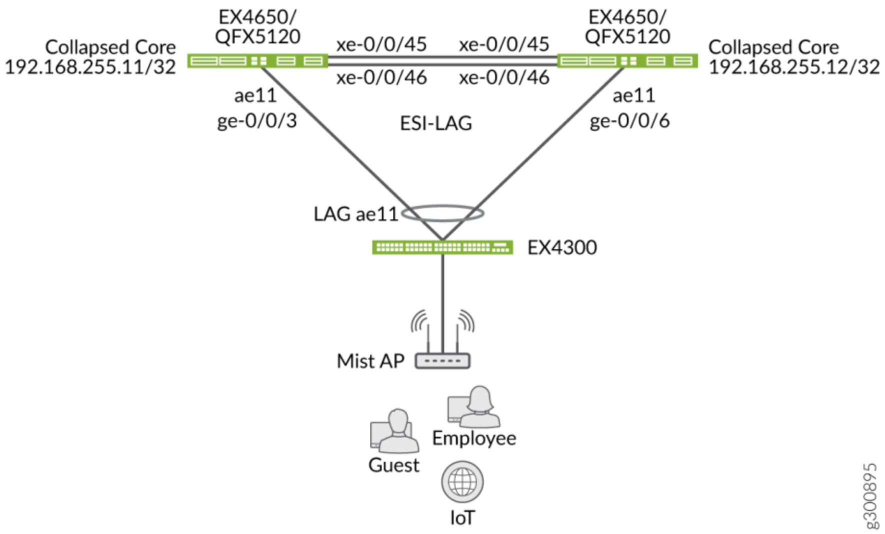

- EVPN supports N-way "scale-out" Ethernet multihoming
- No ICL link required
- Flexible overlay supports layer 2 and layer 3 services
- Active-Active Multihoming
- Multi-homed Access switches are identified in the overlay by unique Ethernet Segment ID (ESI)
- Any access layer switch

EVPN Multihoming (ESI-LAG)

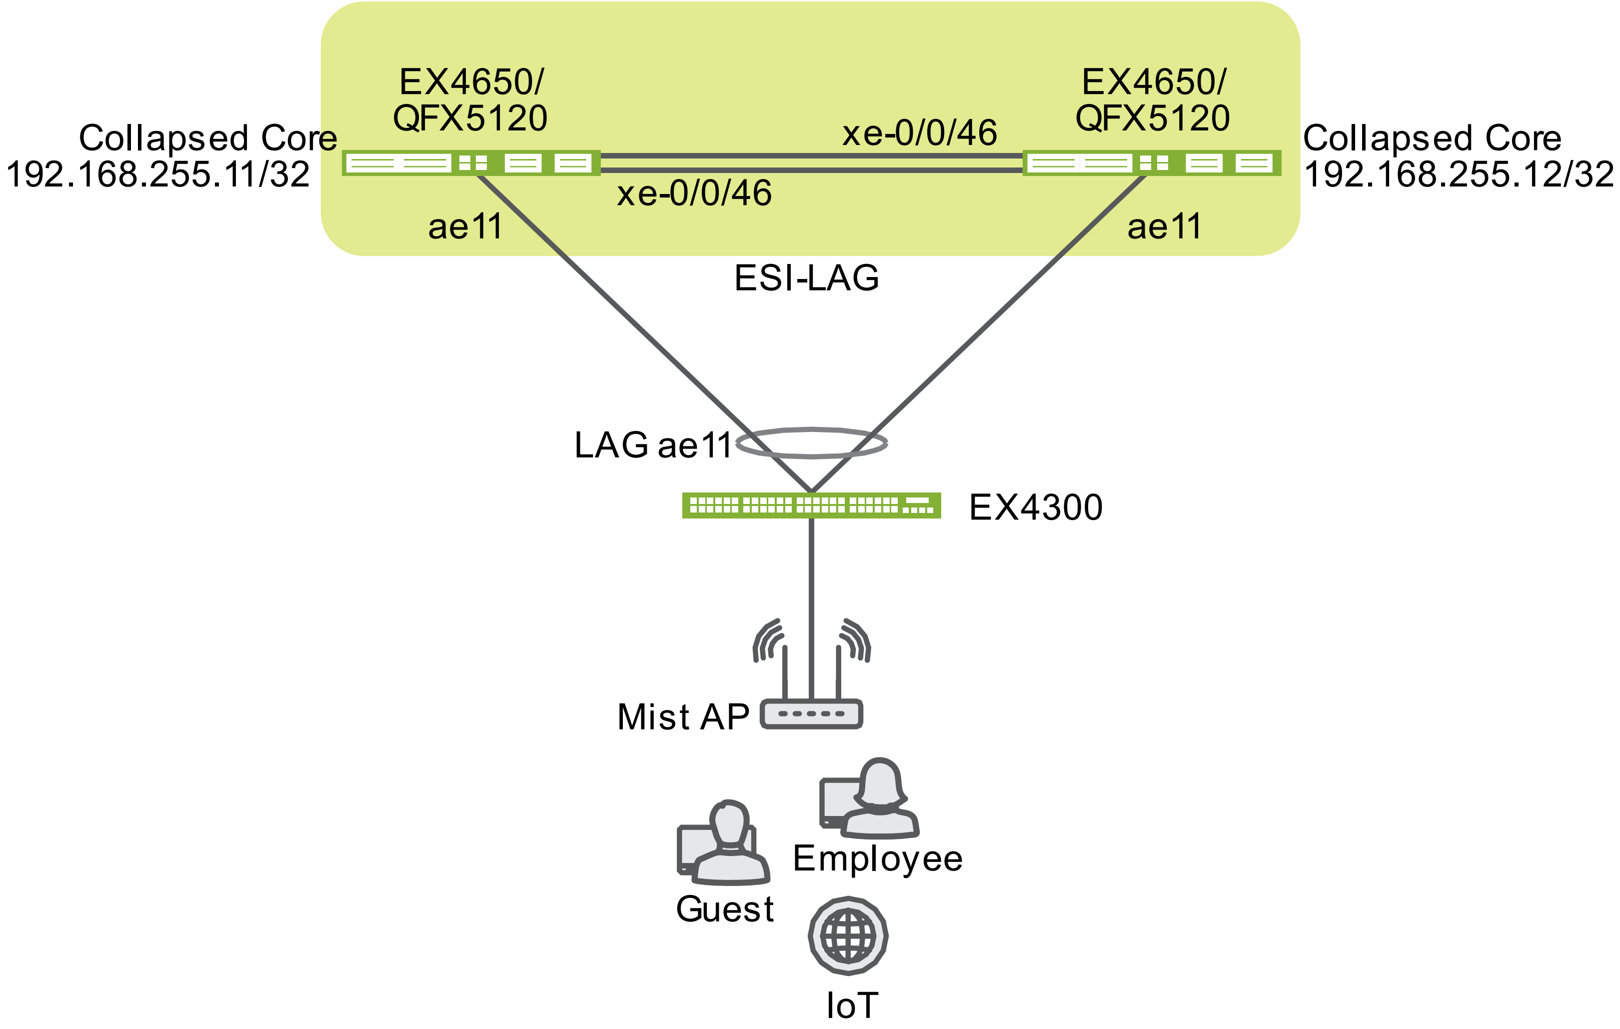

Campus Fabric Core-Distribution Using CRB

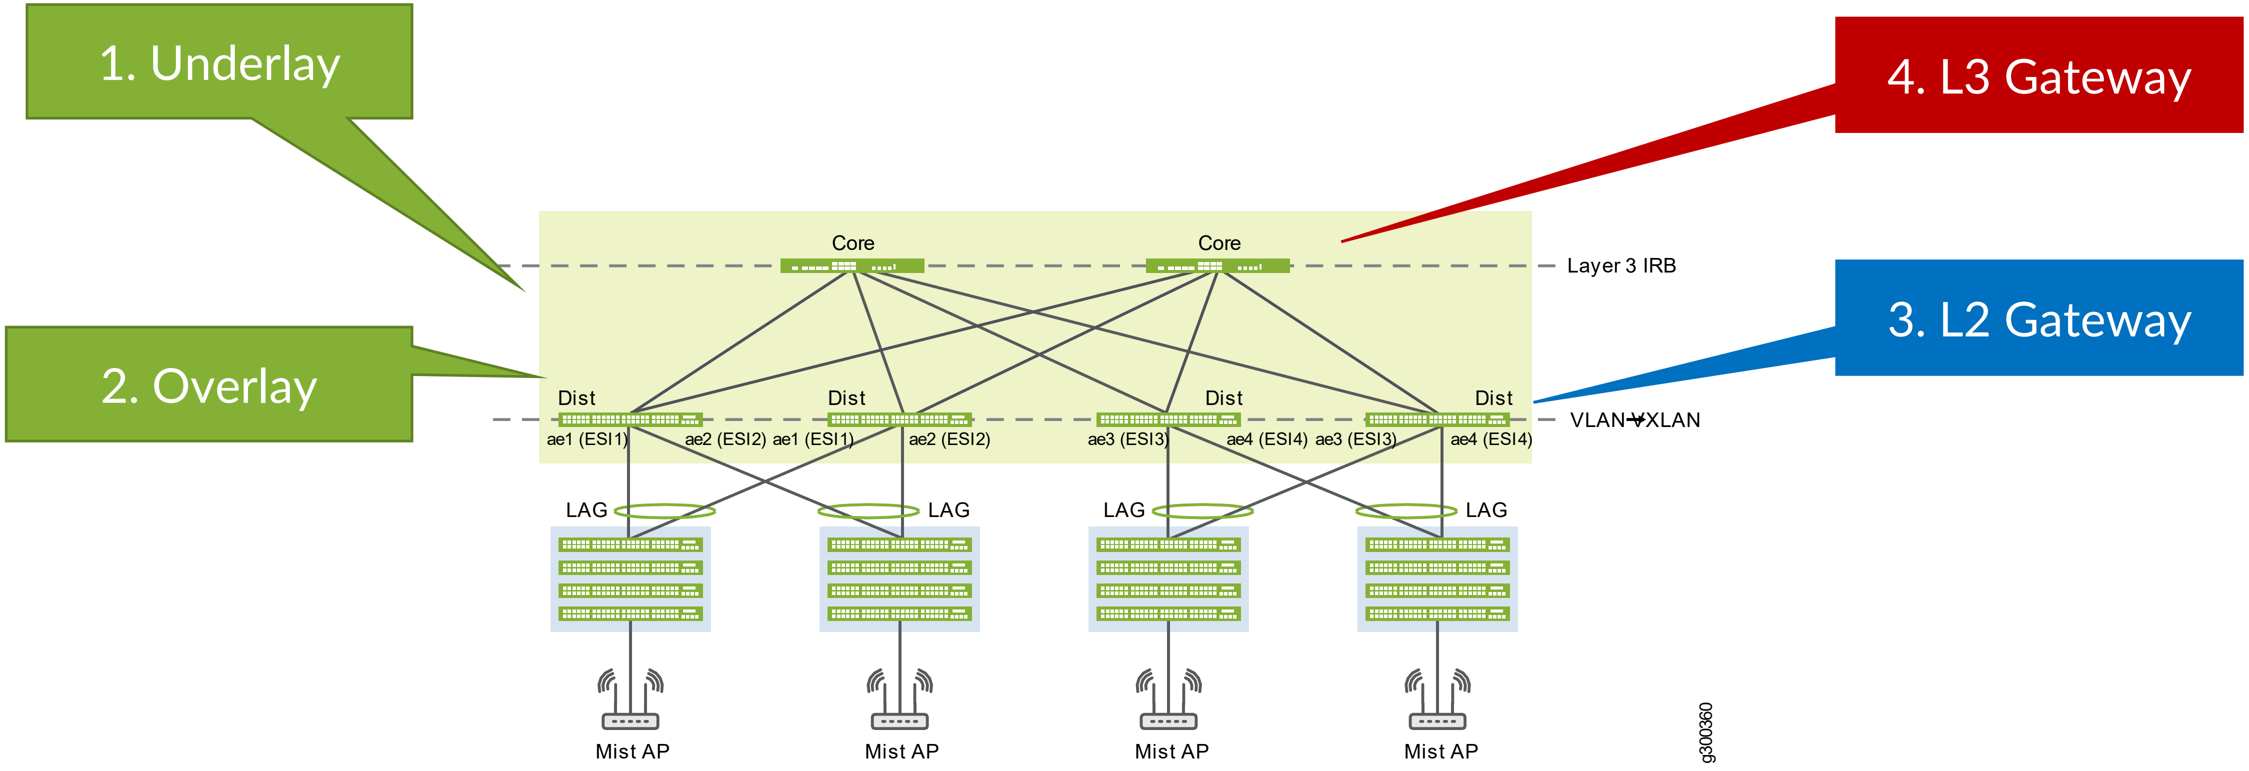

How to Configure an EVPN-VXLAN Fabric for a Campus Network With CRB

Campus Fabric Core-Distribution Using ERB

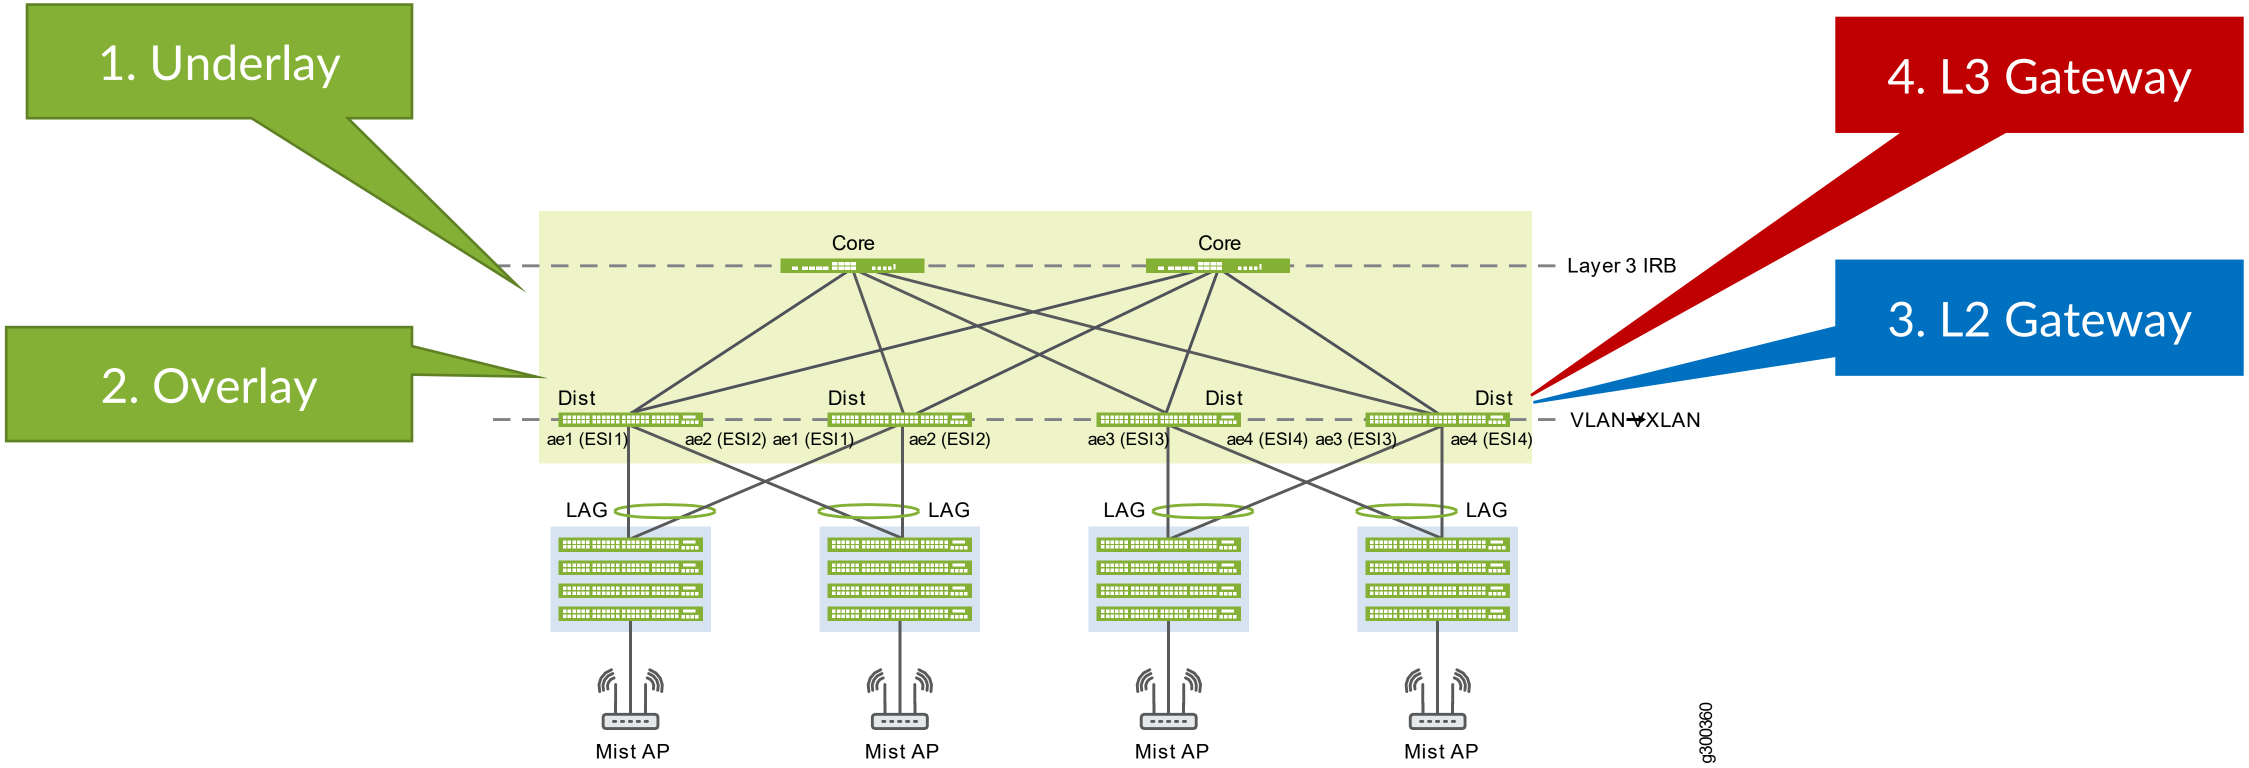

How to Configure an EVPN-VXLAN Fabric for a Campus Network With ERB

Campus Fabric IP Clos Config Steps

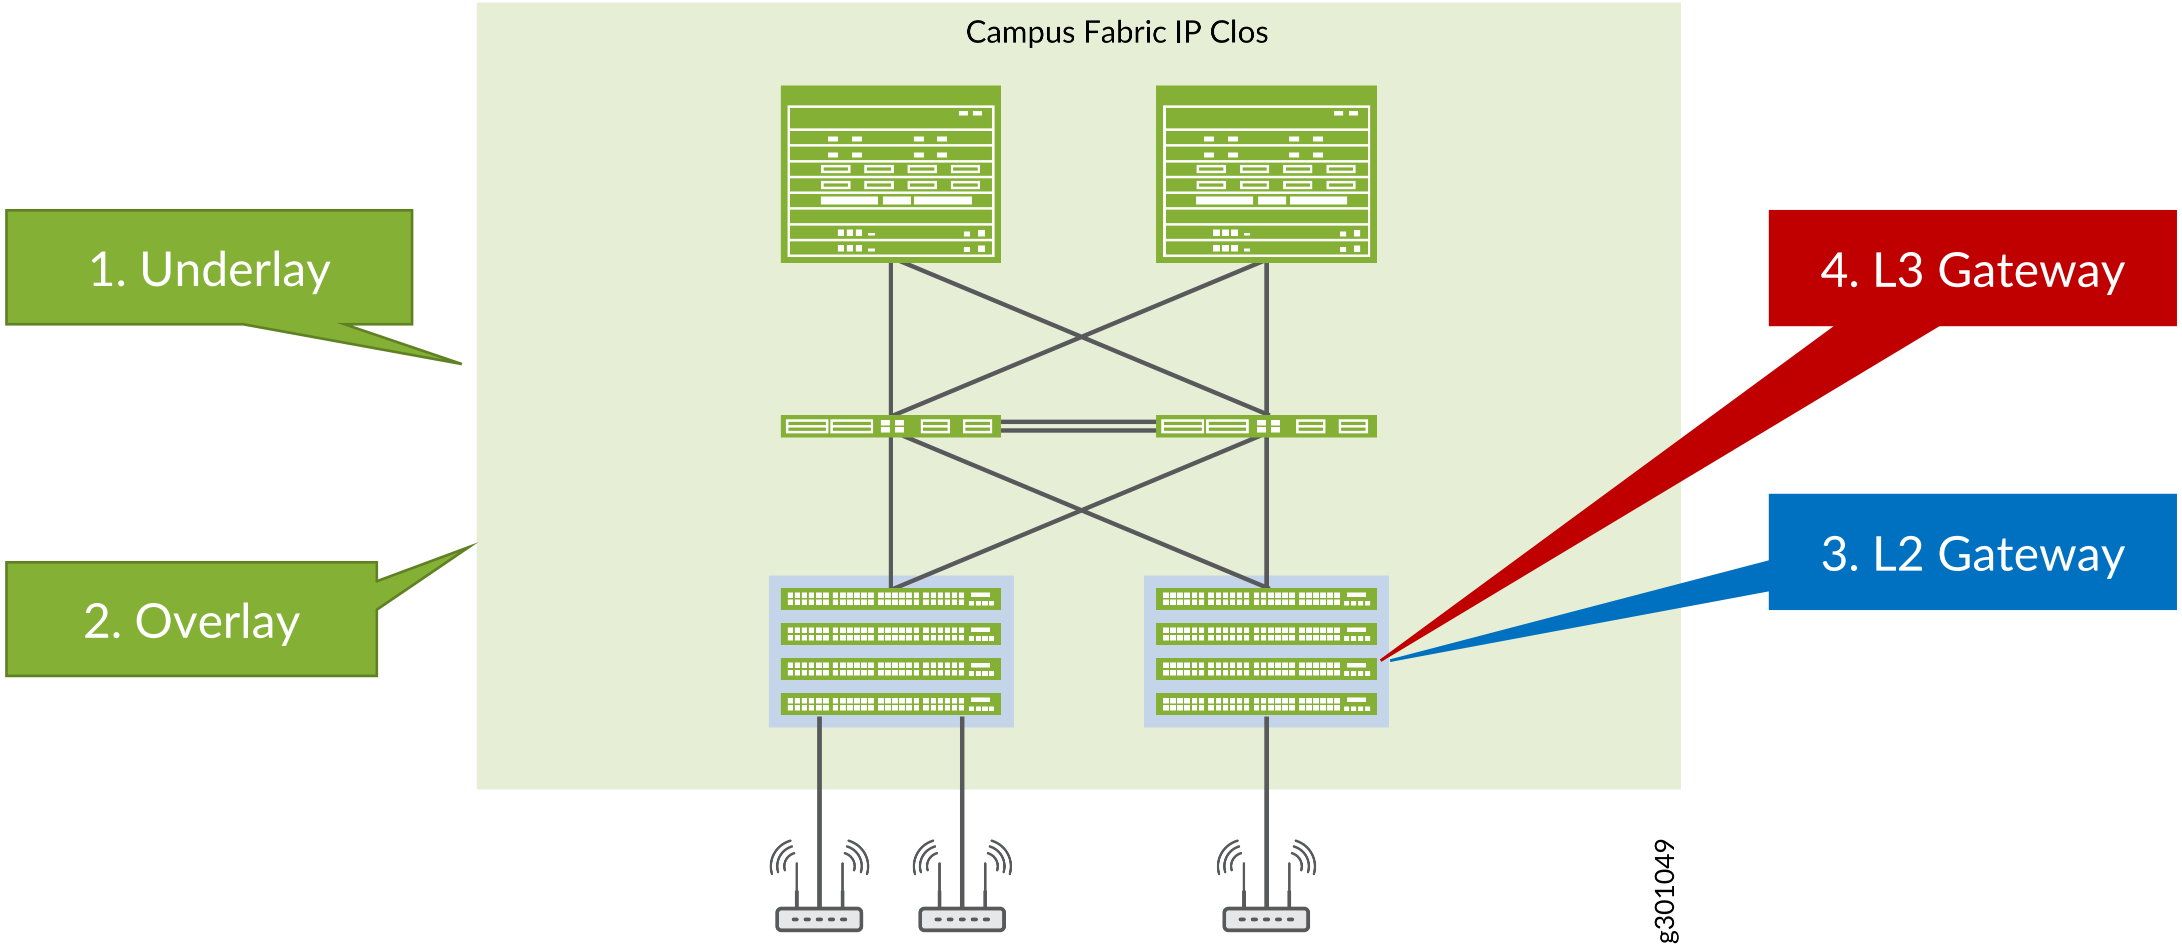

[How to Configure an IP Clos Fabric for a Campus Network]

Microsegmentation in Campus using EVPN-VXLAN

Standards-based microsegmentation

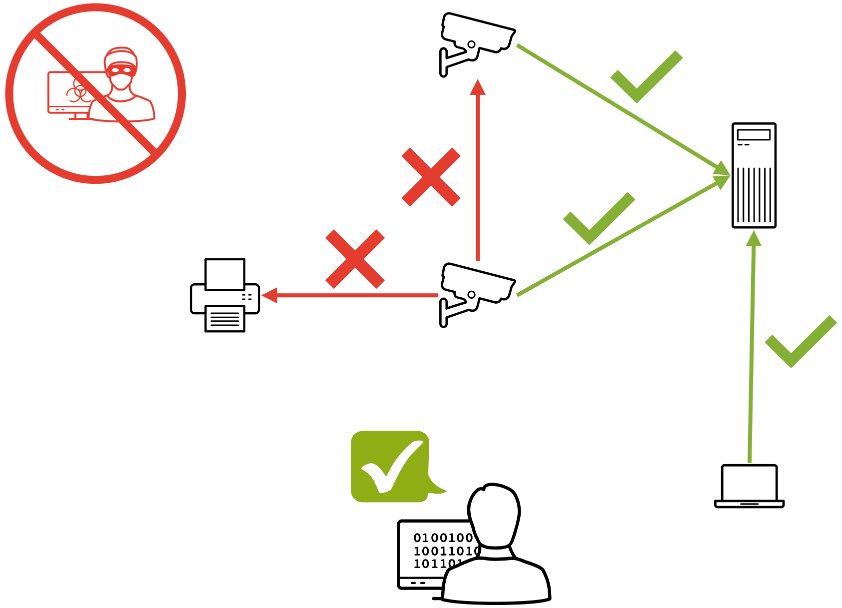

Group Based Policies (GBP)

- Leverage underlying VXLAN infrastructure
- Granular access policy & controls
- Location agnostic end to end security

Outcomes

- Consistent security policies across the network
- Ability to block lateral threats
- Reduce ACLs

GBP: IOT to Server

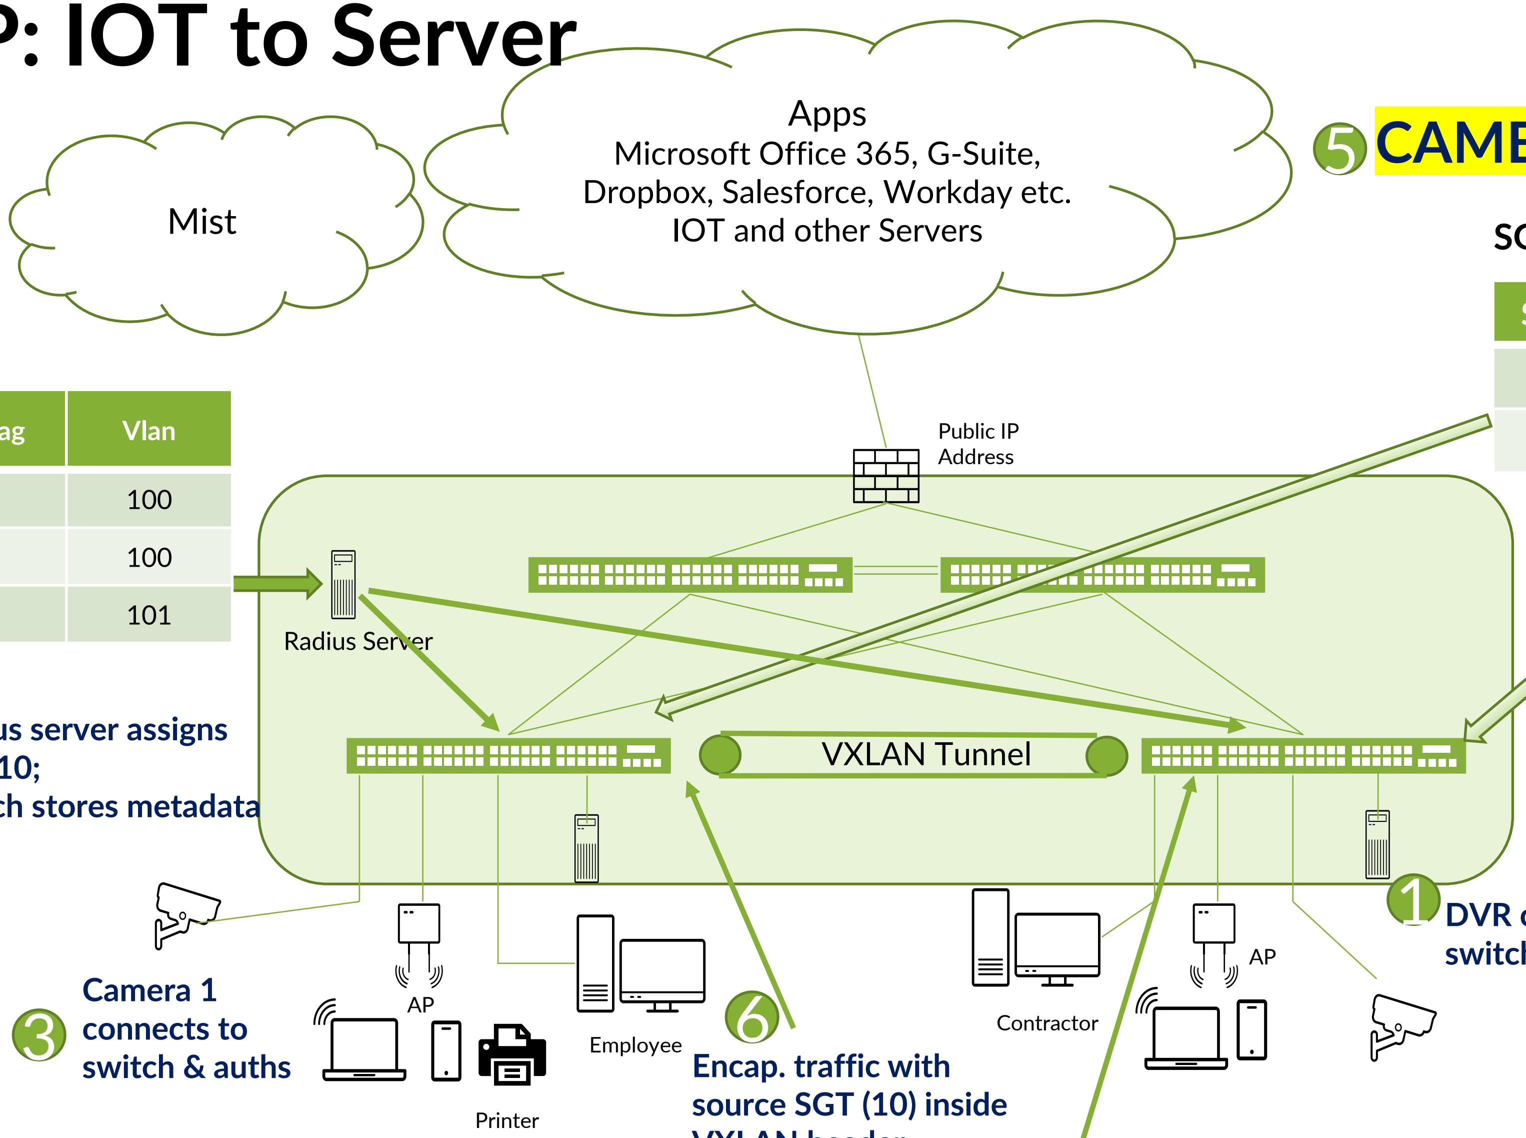

Device	SG Tag	Vlan
Camera 1	10	100
Camera 3	10	100
DVR	11	101

5 **CAMERA 1 tries to access the DVR**

SG ACLs

Src. Group Tag	Dest. Group Tag	Policy
10	11	Allow
10	10	Deny

1 **DVR connects to switch & auths**

2 **Radius server assigns SGT 11; Switch stores metadata**

3 **Camera 1 connects to switch & auths**

4 **Radius server assigns SGT 10; Switch stores metadata**

6 **Encap. traffic with source SGT (10) inside VXLAN header**

7 **Lookup (Source SGT (10), Dest. SGT (11)) security policy ALLOW as per policy**

Source-group tags can be defined based on
- MAC address
- Port
- VLAN
- Port, VLAN
- Subnet/IP Address

GBP Building Blocks

Configure EVPN-VXLAN IP Clos

Policy definition and creation

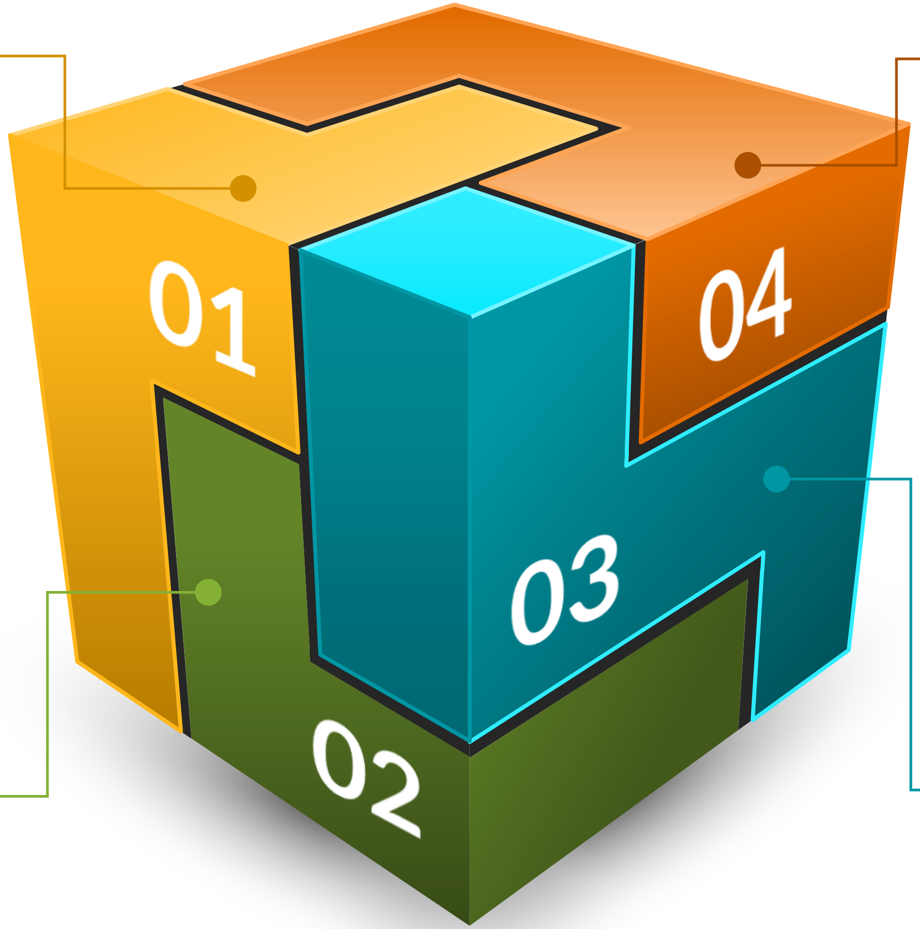

User/device to SGT mapping

RADIUS server configuration

Campus Deployment using Juniper Mist Cloud

Cloud native architecture for campus networks

Describe the WHAT | **Software delivers the HOW** | **Know WHEN and WHY**

Architect

Operator

Intent

Closed-loop Automation and Assurance

Analytics

Day 0 – Design
- Campus fabric
- Distributed enterprise

Day 1 – Deploy
- ZTP
- Templates

Day 2+ – Operate
- SLEs
- Marvis actions
- Conversational interface

Campus Fabric Deployment

1

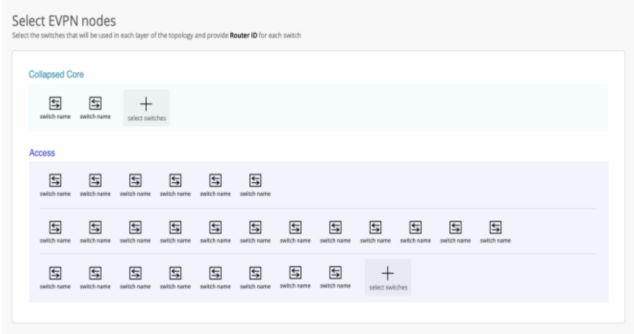

Choose the topology and allocate device roles

- Define the intent for the topology definition (IP-Clos, Multi-homing etc)
- Choose device roles – access, distribution, core

2

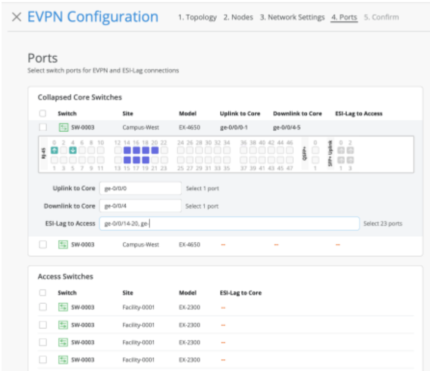

Define Physical Connections

- Provide the physical connectivity between – core/distribution and access devices

3

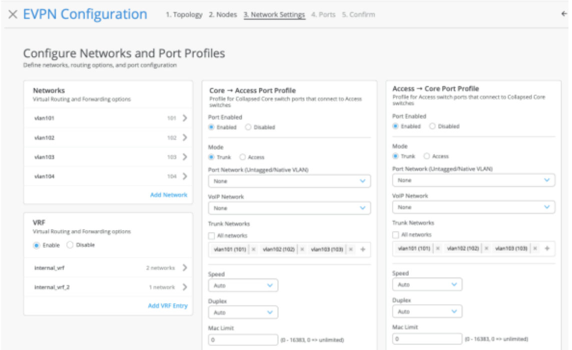

Define Networks of Interest

- Configure the user networks

4

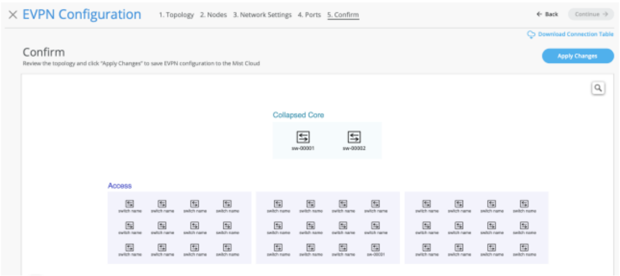

Apply the intent

- Verify, apply and confirm the intent of provisioning the fabric

Thank you