Вместе на волне цифровизации

Москва • Экспоцентр
26-27 марта 2019

Дизайн современной корпоративной LAN сети, часть 3

Станислав Рыпалов

Системный инженер, Cisco

Дмитрий Демин

Системный инженер, Cisco

CISCO

Общие рекомендации по дизайну(?)

Как объединить несколько фабрик

CISCO

Fabric Sites & Domains

Подключение нескольких фабрик

Поддерживается с версии SDA 1.2.5

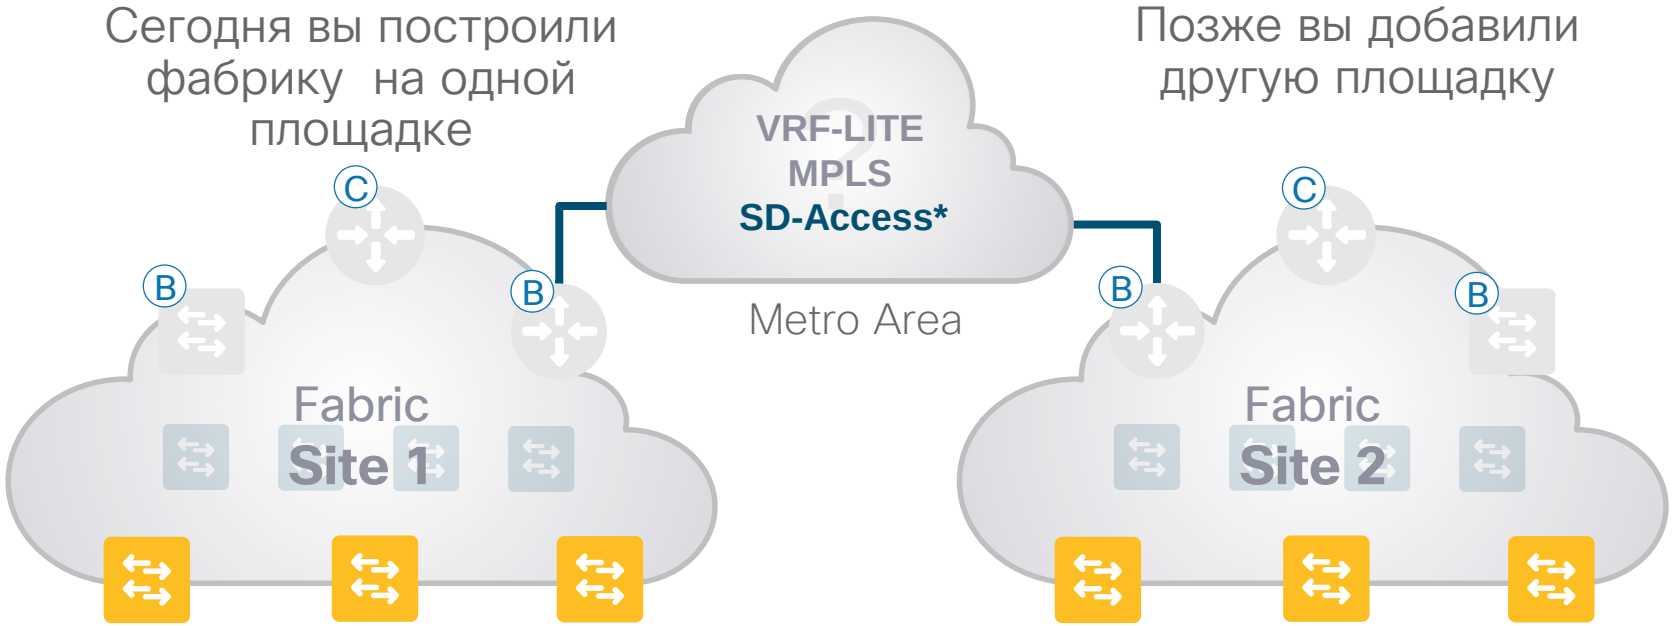

Как их соединить?

SD-Access для распределенного кампуса

Автоматизация настройки связности между площадками

Локальные сервисы на площадке
ISE PSN, WLC и пр.

End-to-End Policy & Сегментация

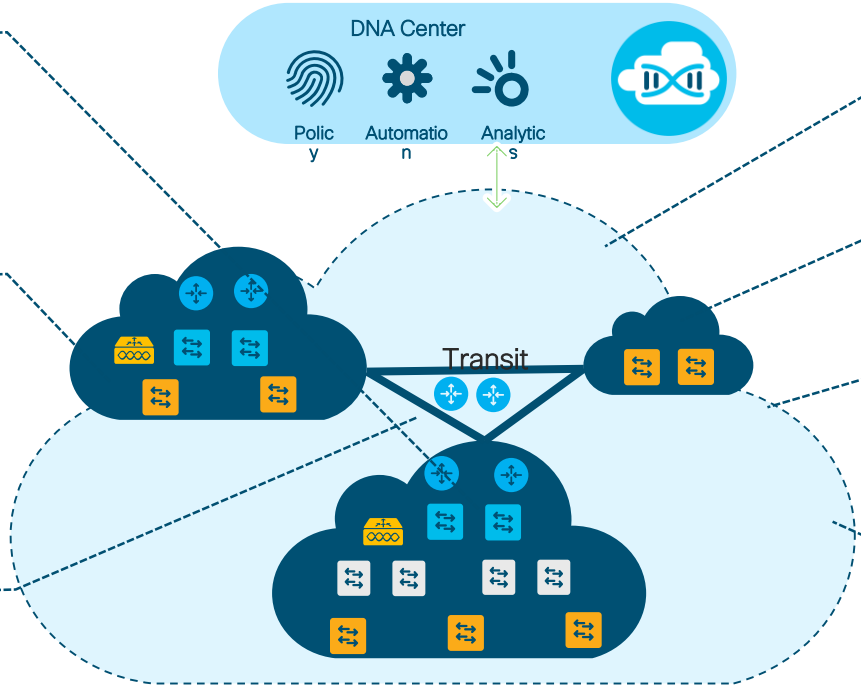

Масштабирование
100 площадок

Гибкость
50 - 100,000 Users/Site

Отказоустойчивость
Локальные Control Planes, Borders

Производительность
Local Breakout

SD-Access для распределенного кампуса

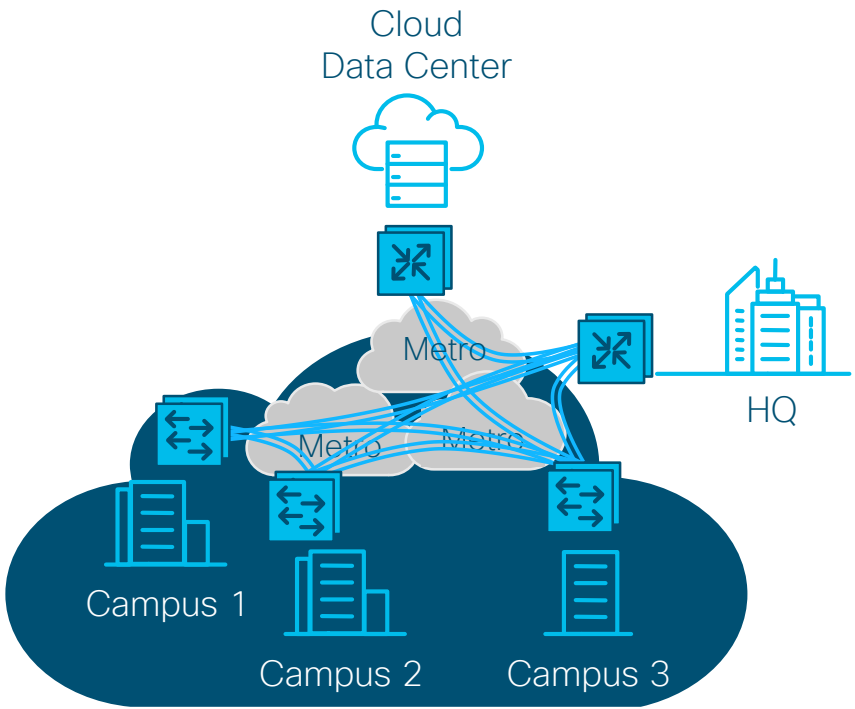

- Сегментация End-to-end
- Централизация Automation & Assurance
- Возможности по интеграции с SD-WAN
  (Viptela SD-WAN в будущих релизах)

SD-Access для распределенного кампуса

Матрица поддержки на устройствах Border

Поддерживается с версии SDA 1.2.5

Узел SDA Border	SD-Access Distributed campus (SD-Access Transit)	SD-Access Distributed Campus (IP Transit)
C9K	Да	Да
ASR1K/ISR4K	Да	Да
C6K	Нет	Да
N7K	Нет	Да

SD-Access для распределенного кампуса

Fabric Site и Domain

Fabric Site это независимое пространство фабрики со своим набором сетевых устройств: Control Plane, Border, Edge, WLC, ISE PSN

Возможно обеспечить разный уровень отказоустойчивости для Site'а с использованием локальных ресурсов: DHCP, AAA, DNS, Internet и пр..

Fabric Site может соотноситься с **одним географическим местом**, **несколькими** или **набором мест**

- Single Location ➔ Branch, Campus or Metro Campus
- Multiple Locations ➔ Metro Campus + Multiple Branches
- Subset of a Location ➔ Building or Area within a Campus

SD-Access для распределенного кампуса
Fabric Site и Domain

Fabric Domain может состоять из одного или нескольких Fabric Site + Transit

Несколько Fabric Site'ов соединяются друг с другом через **Transit Site**

Есть два типа транзита:

- **SD-Access Transit** – Возможность создания нативной SD-Access (LISP,VXLAN,CTS) фабрики с узлом Control Plane уровня домена для взаимодействия между site'ами
- **IP-Based Transit** – Использование для связности традиционных сетей IP (VRF-LITE, MPLS), которые требуют трансляции VRF и SGTs между площадками

SD-Access для распределенного кампуса

Типы транзитов

❖ SD-Access Transit :

- Трафик между площадками использует инкапсуляцию VXLAN-GPO и LISP control-plane.
- Узлы Border используют LISP lookup на узлах domain Control Plane (на Transit site).
- Узлы Border обеспечивают инкапсуляцию VXLAN для отправки трафика на удаленные площадки.
- Политики End to End обеспечиваются метками SGT.
- Transit site не поддерживает performance-based routing.
- Автоматизация настройка через DNAC

❖ IP-Based Transit :

- Трафик между площадками использует существующие шины control и data для IP Transit site
- Узлы Border передают трафик напрямую во внешний домен (VRF-LITE с BGP).
- Трафик передается через обычную IP сеть через внешний домен на удаленные площадки.
- Политики End to End требуют ручной конфигурации.
- Если в качестве транзита используется SD-WAN, возможна поддержка performance-based routing, шифрования и пр.

SD-Access для распределенного кампуса

SD-Access Transit

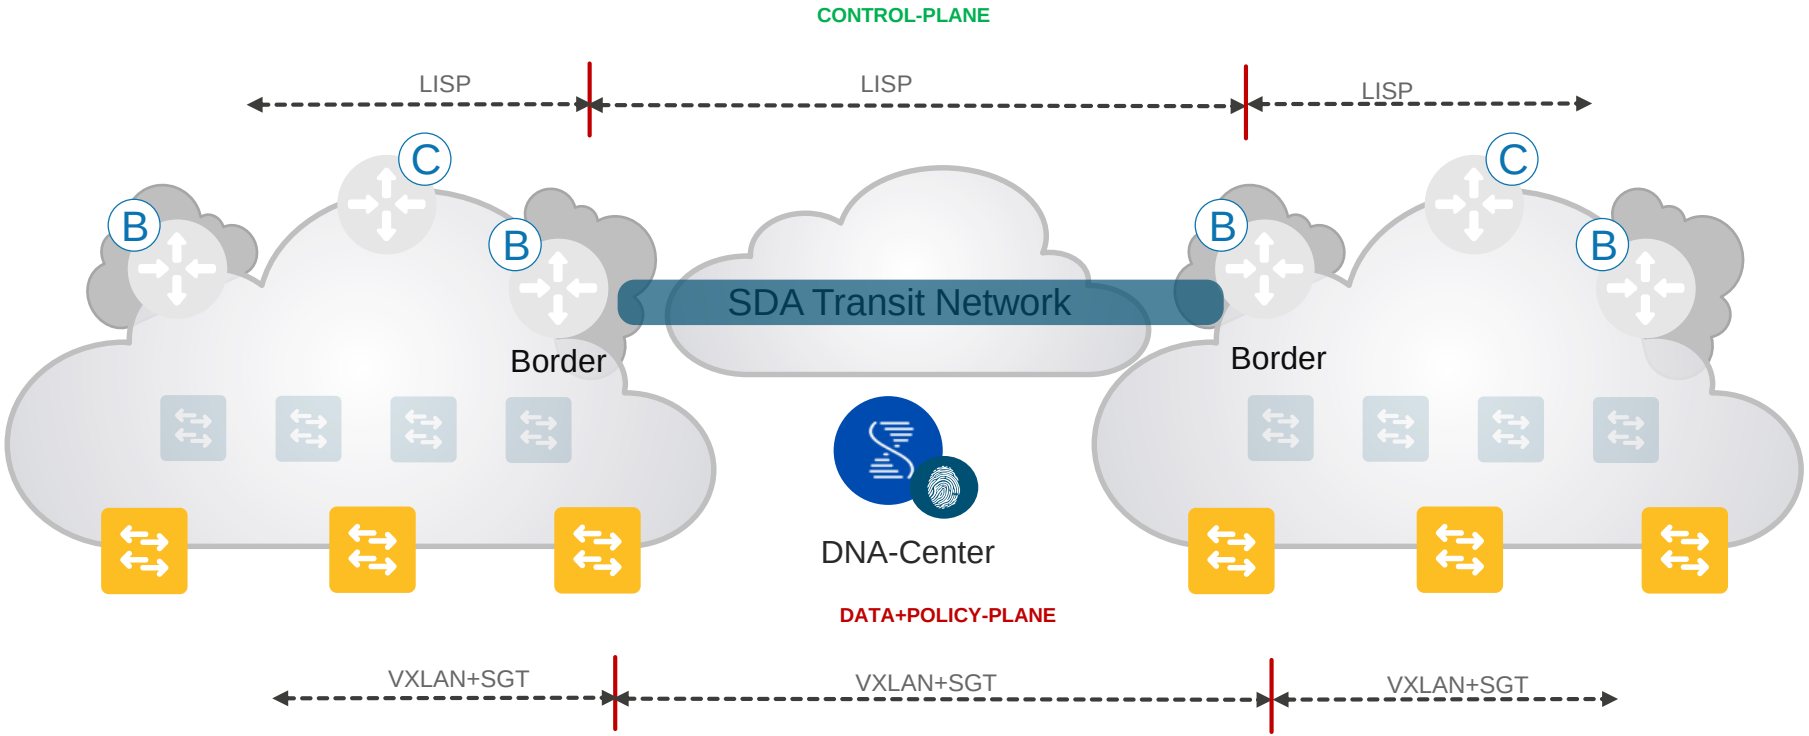

SD-Access для распределенного кампуса

SD-Access Transit

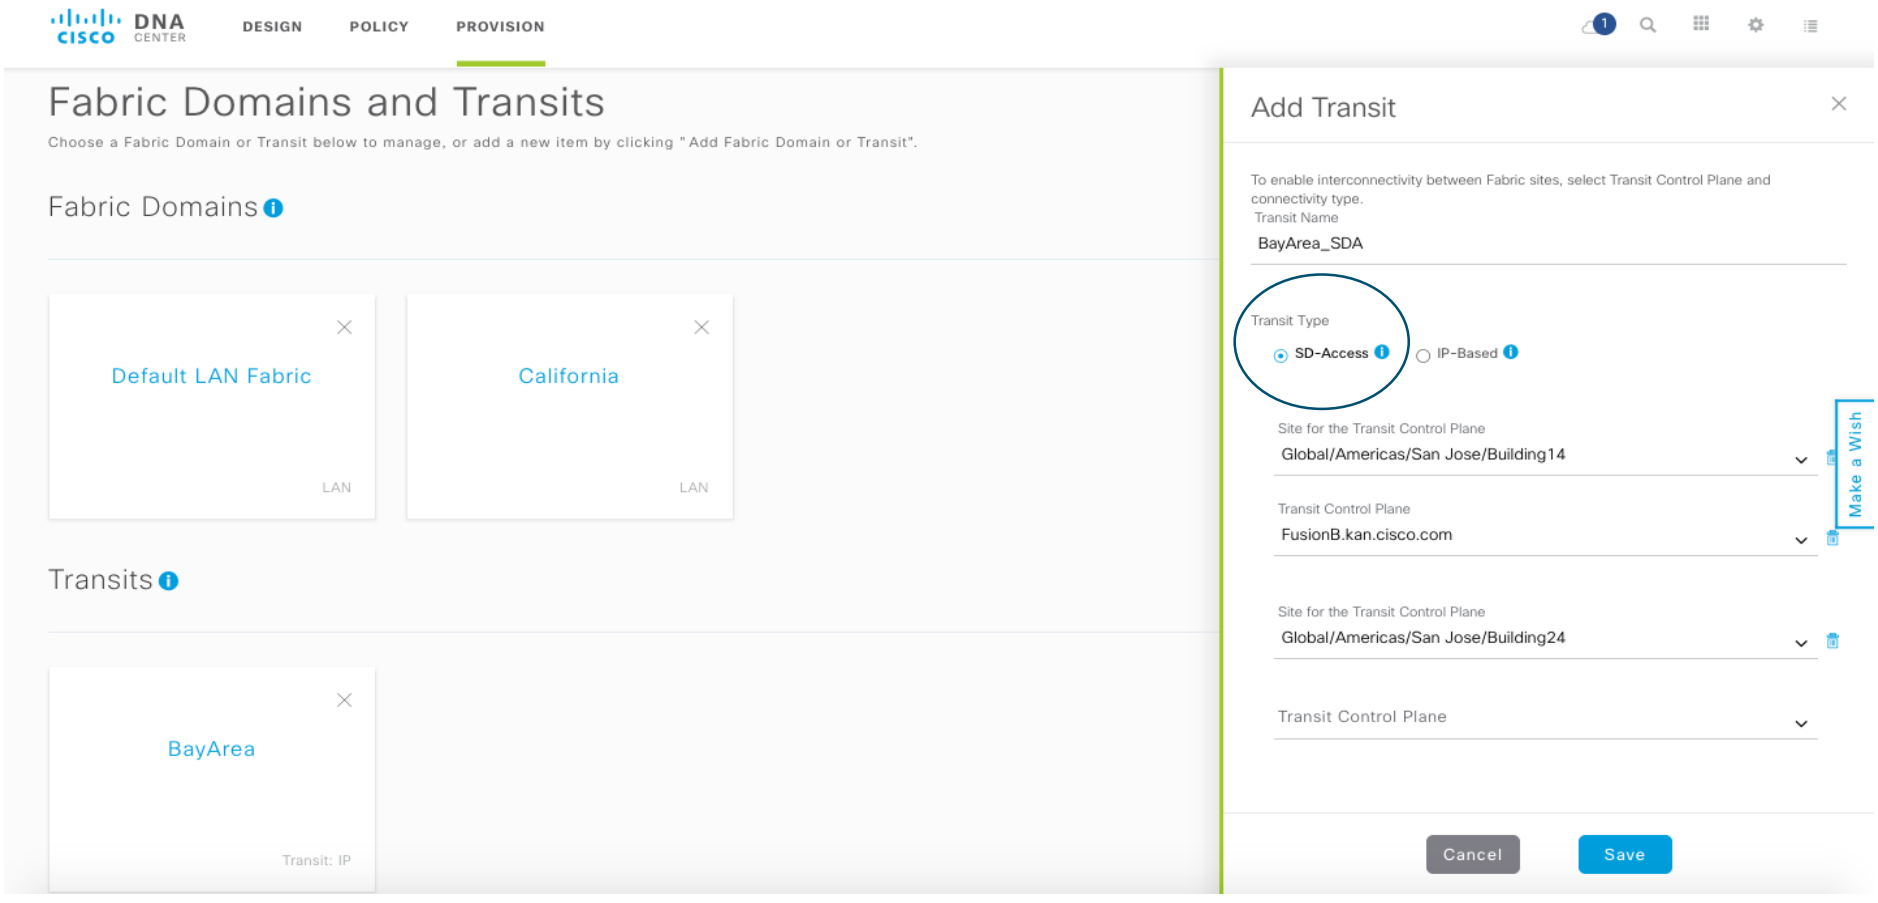

SD-Access для распределенного кампуса

SD-Access Transit

- **Fabric Domain** может состоять из одного или нескольких **Fabric Sites**
- Каждый Site – это уникальная фабрика со своими узлами Control Plane и Border
- Fabric Domain может содержать узел **Transit Control Plane** для взаимодействия между площадками

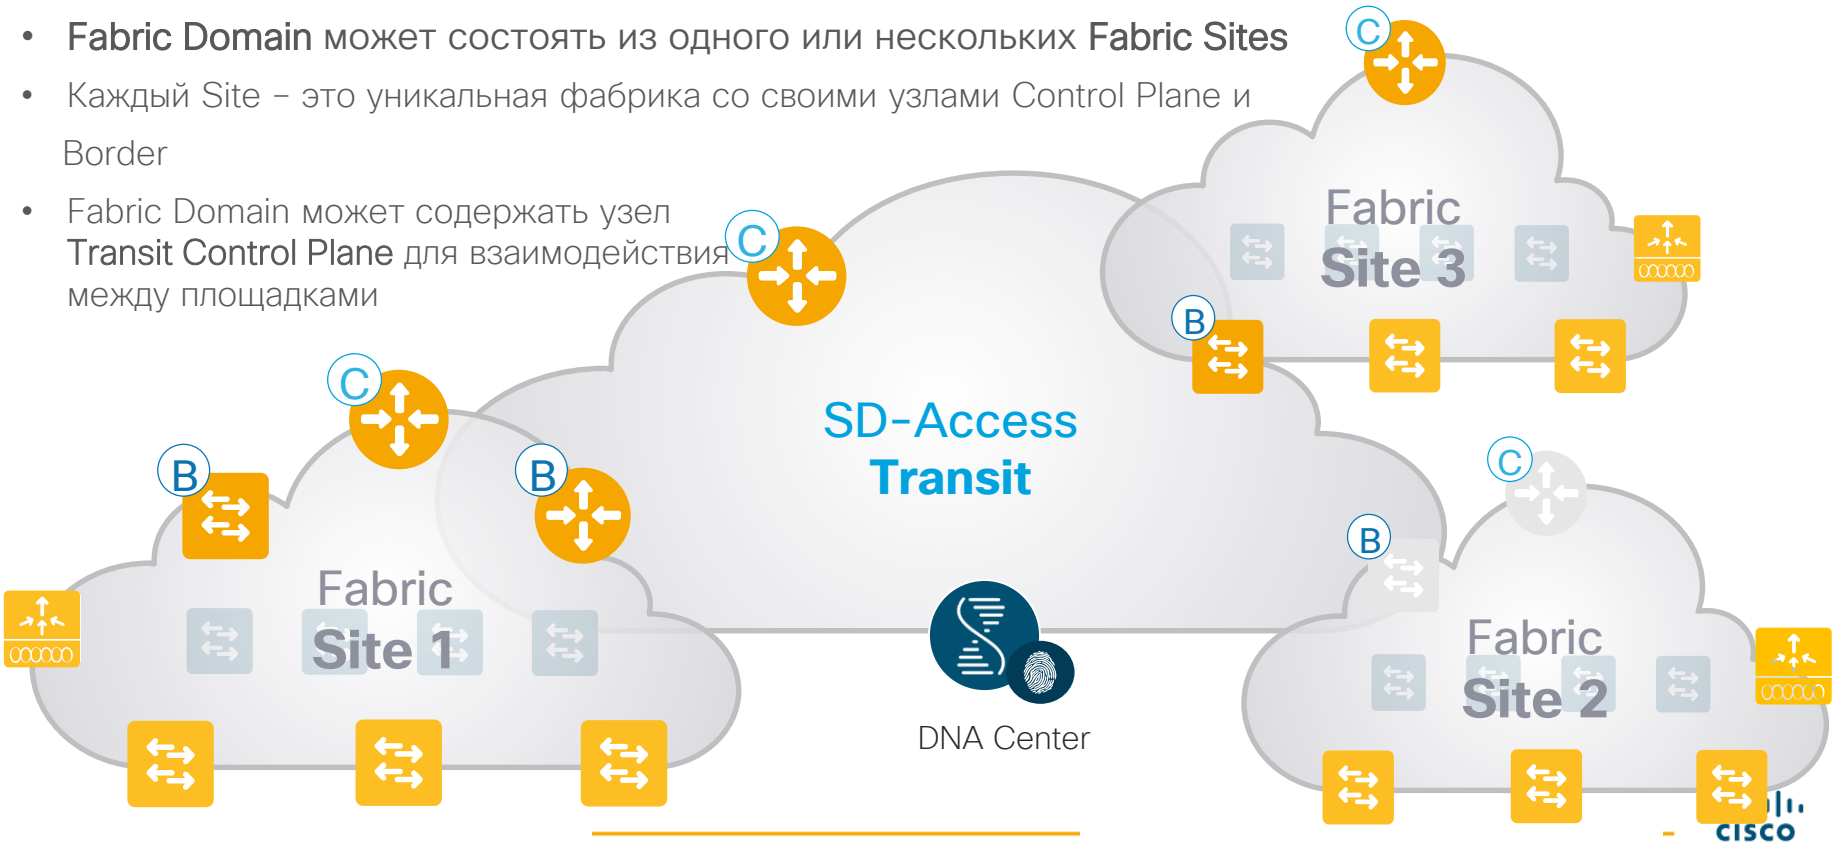

SD-Access для распределенного кампуса

SD-Access Transit

- Узел Border обеспечивает через LISP поиск на узле Transit Control Plane
- Узел Border производит инкапсуляцию VXLAN-GPO в сторону другого узла Border
- Политика End to End обеспечивается передачей меткой SGT
- Настройка SD-Access Transit полностью автоматизирована

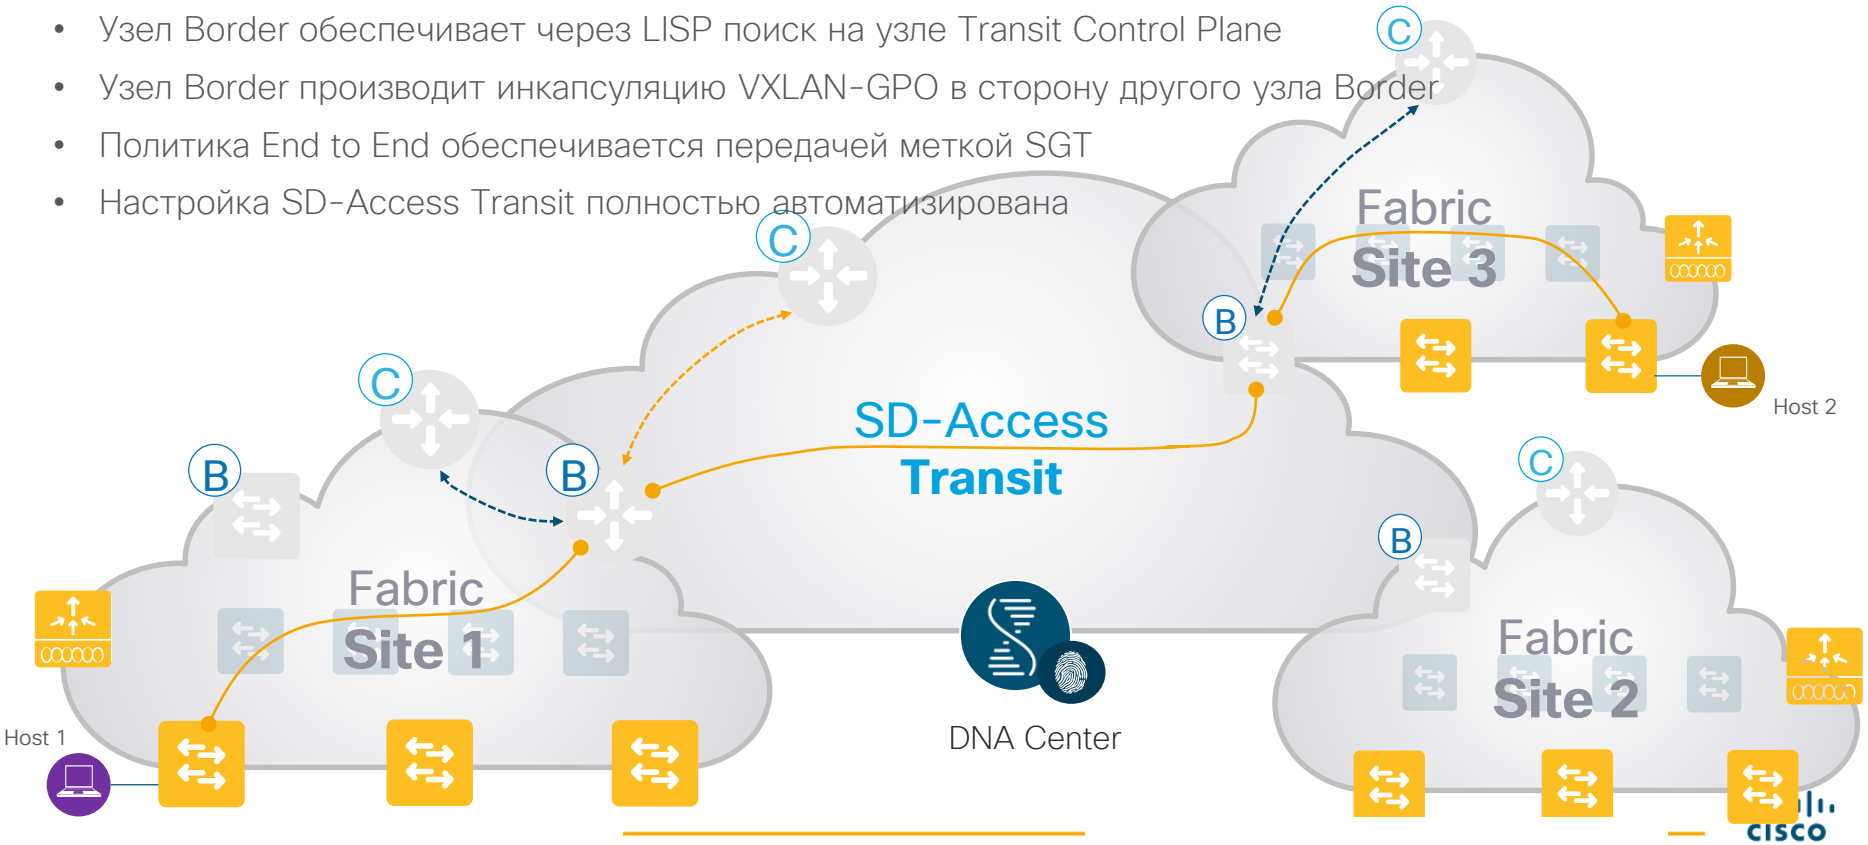

SD-Access для распределенного кампуса

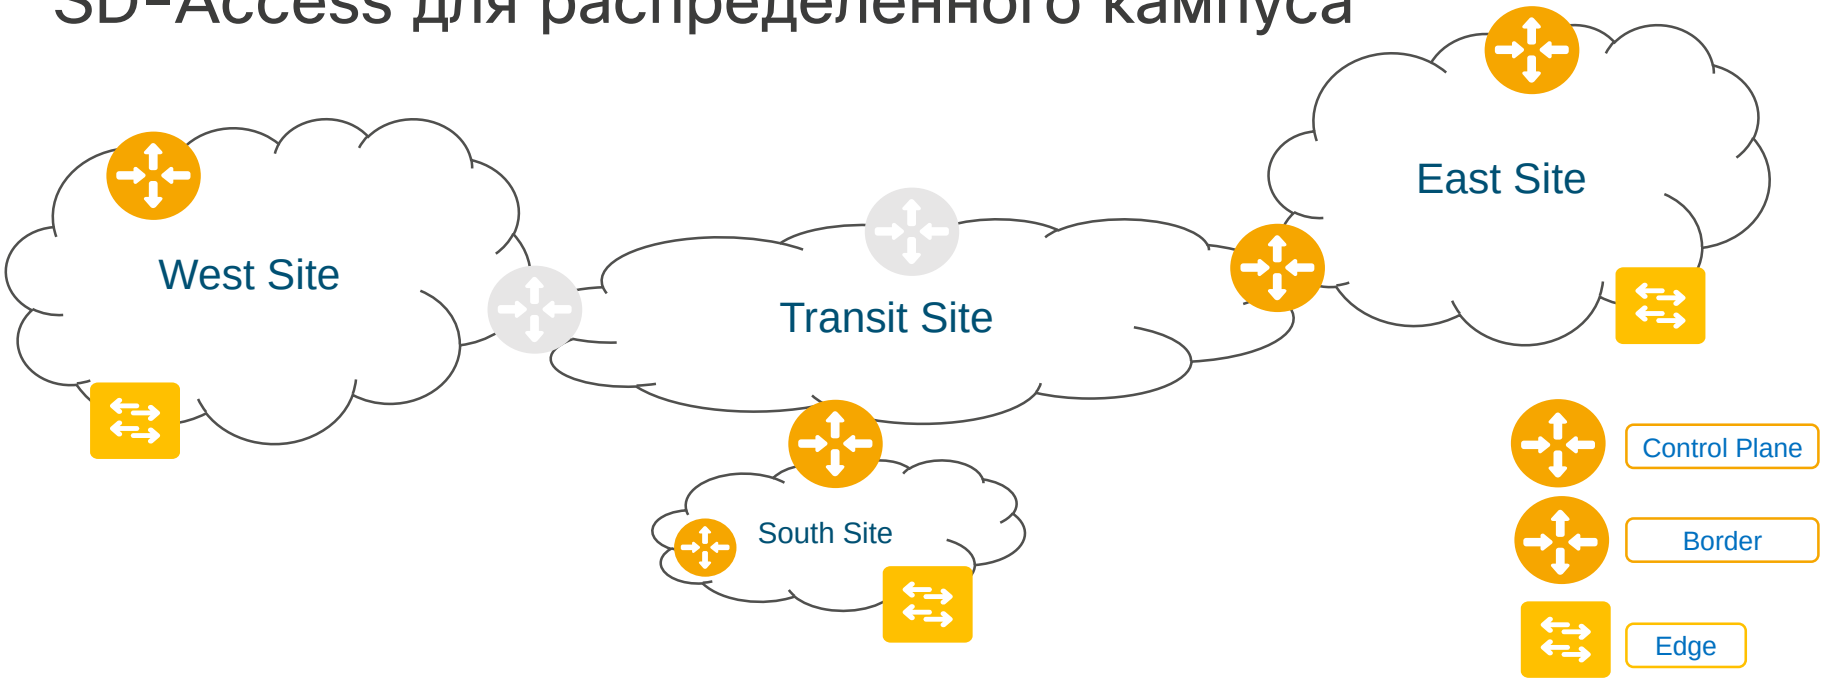

- Каждый site содержит состояние только для своих хостов.
- Трафик за пределы site использует default для транзита.
- Каждый site является автономным доменом сбоя
- Каждый site имеет свой уникальный набор подсетей

Масштабирование SD-Access Transit Control Plane

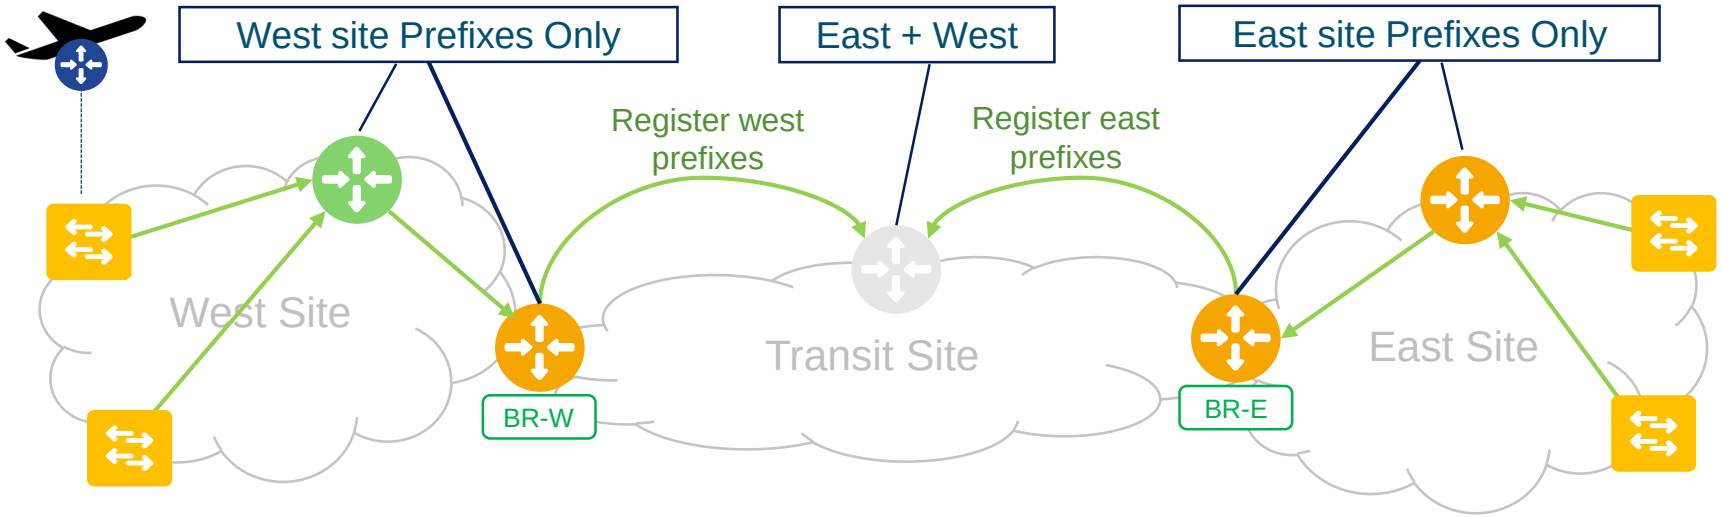

- Устройства Border содержат состояния для локальных префиксов
- Border переходит в состояние передачи трафика строго по требованию
- Control Plane на конечных площадках содержит информацию только о локальных хостах (не удаленных)
- Регистрация суммарных сетей от площадок происходит на Transit Control Plane

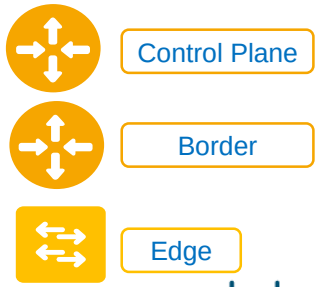

SD-Access для распределенного кампуса

- Какой тип узла выбрать?

Outside world(External)	Подключение к «неизвестным» сетям компании, таким как Интернет – единственная точка выхода из фабрики, подключение к SD-Access Transit
Rest of Company (Internal)	Подключение к «известным» сетям компании, таким как DC, WAN и пр.
Anywhere(Internal +External)	Подключение к Интернет и «известным» сетям компании DC, WAN и пр., подключение к SD-Access Transit

SD-Access для распределенного кампуса

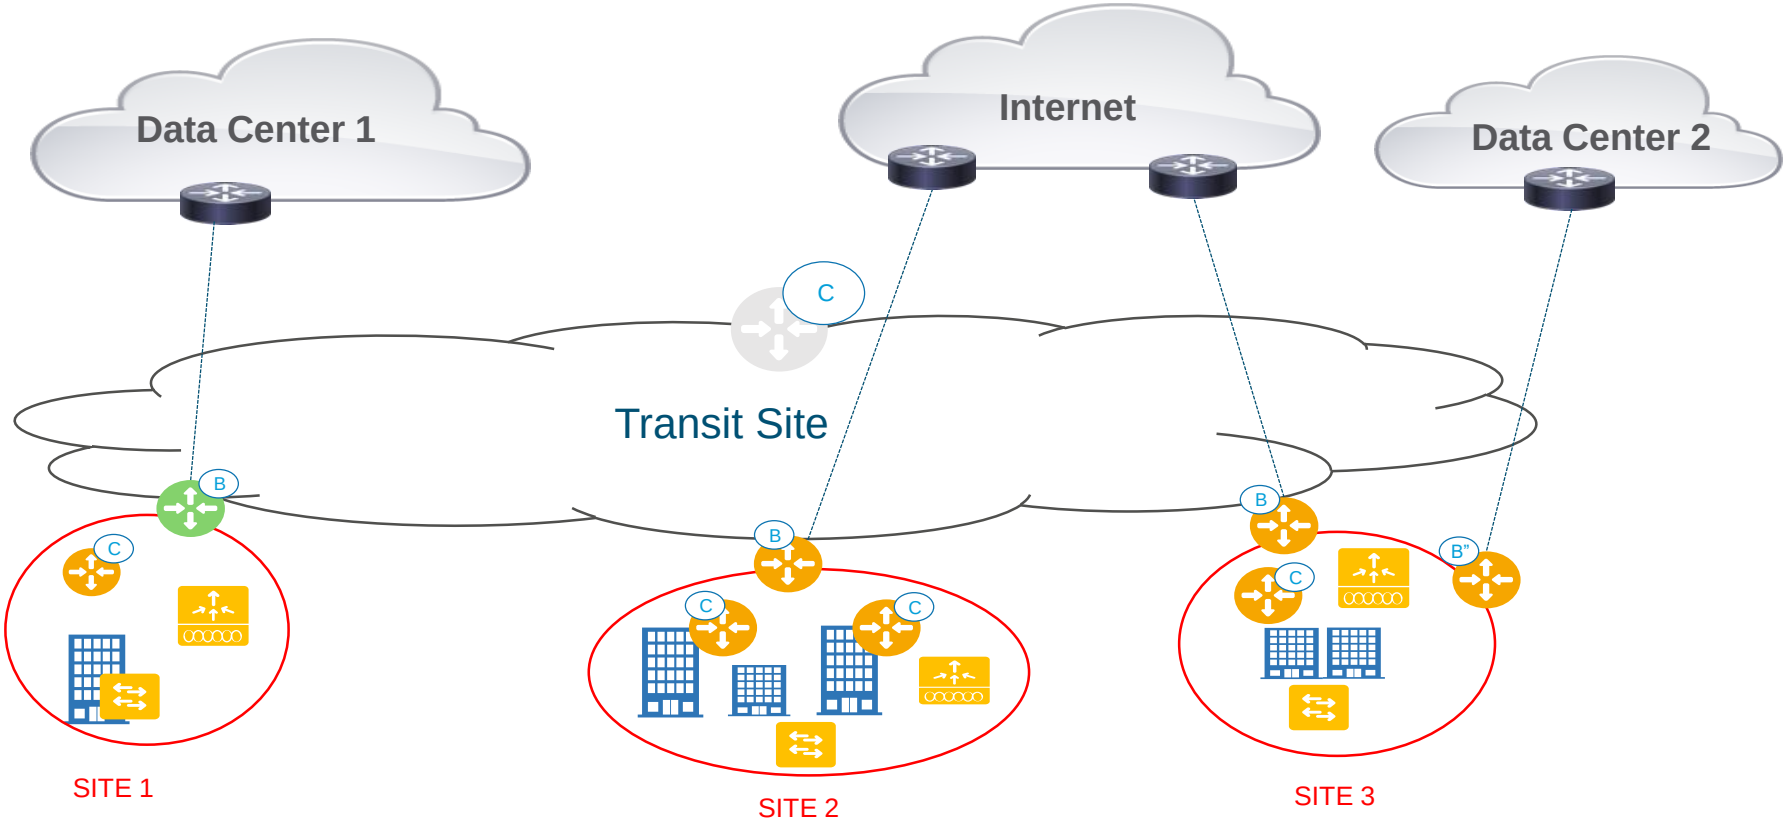

SD-Access для распределенного кампуса

DC/Internet/SD-Access Transit для Site 1

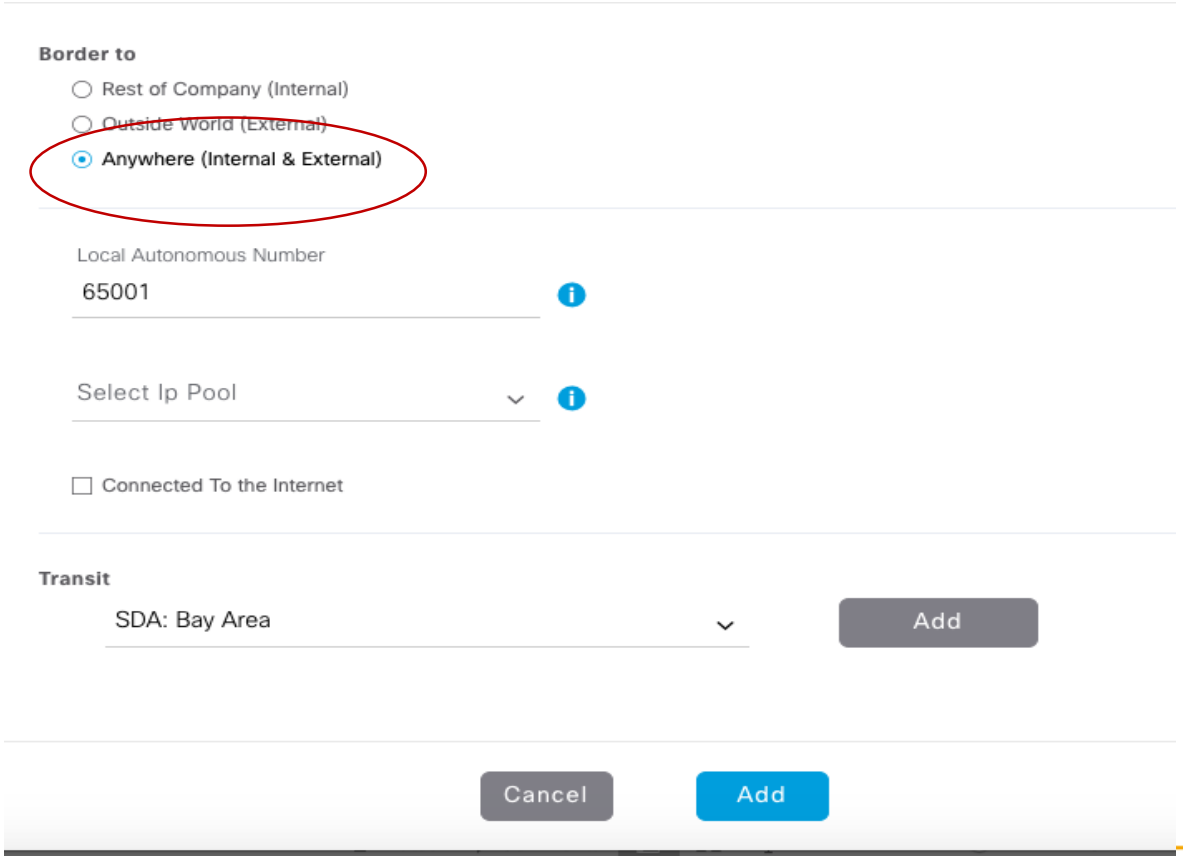

Выбираем Internal + External border

Это необходимо для правильного импорта маршрутов из DC1 и регистрации на transit control plane node.

На слайде приведены настройки на site 1 border для того чтобы site 2 и 3 border могли взаимодействовать с DC1.

SD-Access для распределенного кампуса

Internet/SD-Access Transit для site 2 и 3

Выбираем External border

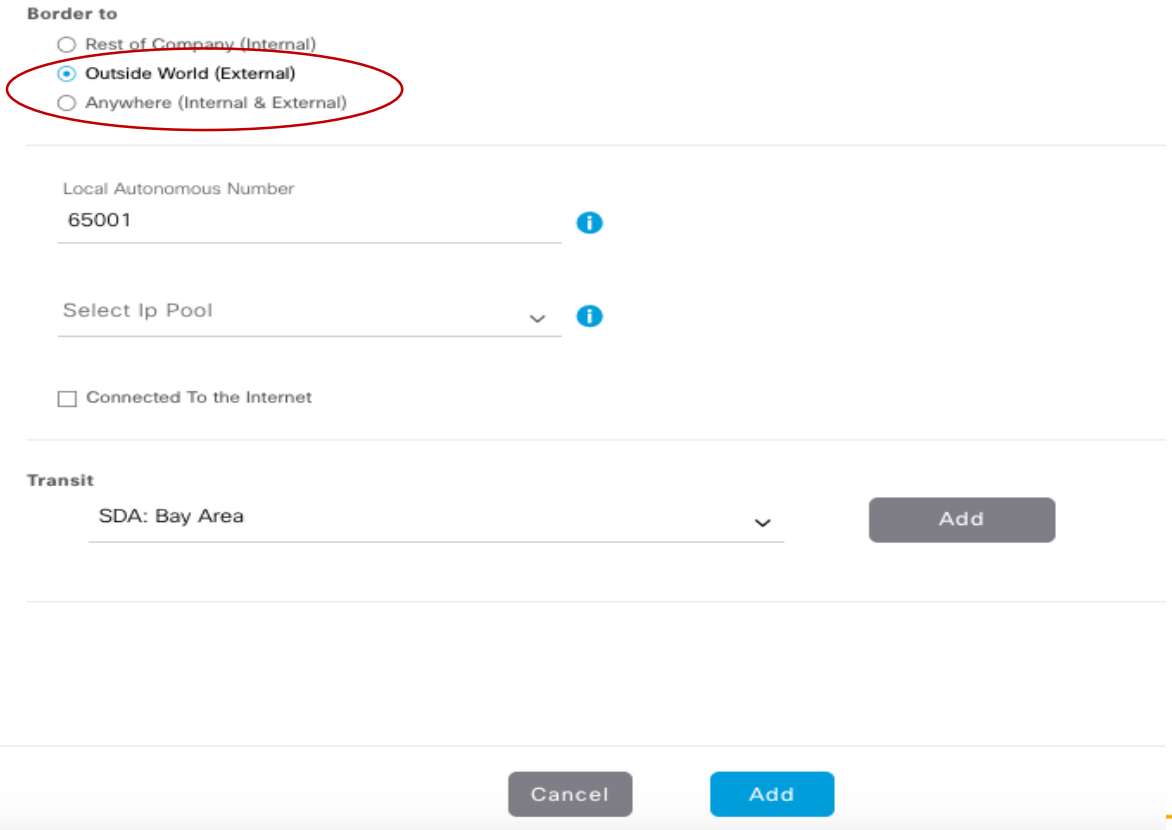

Это необходимо для подключения border устройств к SD-Access transit и Интернет.

На слайде приведены настройки на site 2 и 3 border для подключения к SD-Access Transit.

SD-Access для распределенного кампуса

Доступ в Интернет

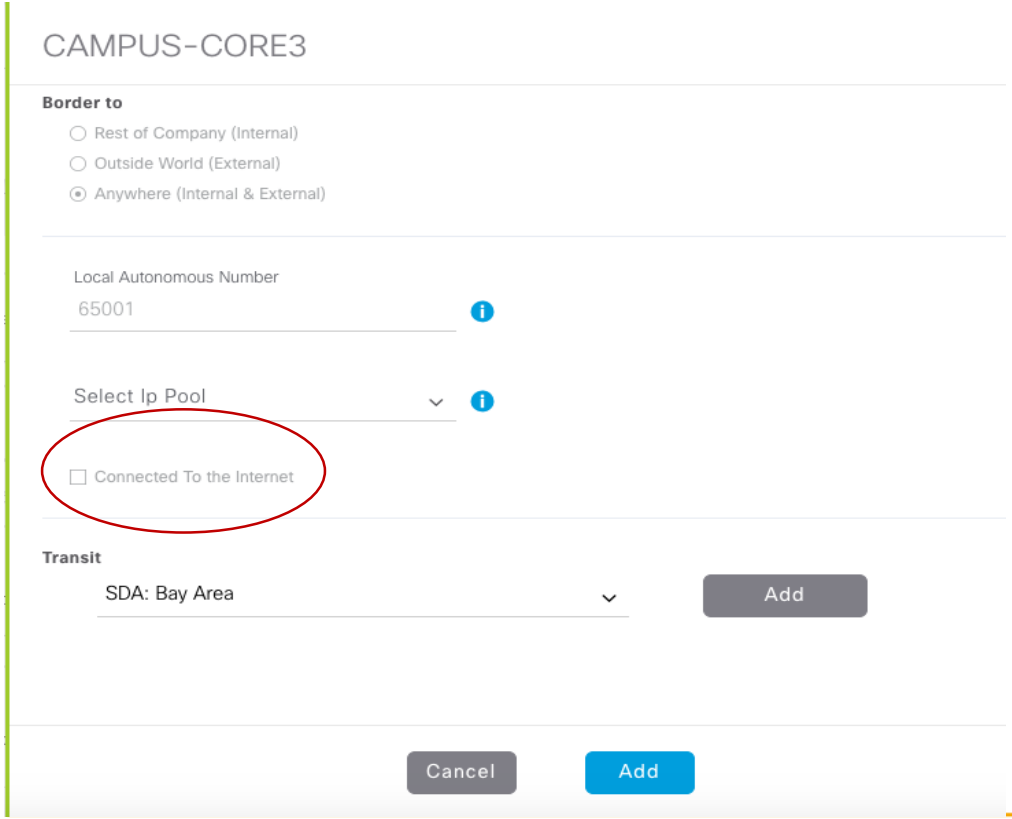

«Галочка» - Connected to Internet

Эта опция обязательно должна быть включена на всех border устройствах, используемых для подключения к internet когда SD-Access используется как транзит.

Необходимо убедиться, что опция включена на border/site устройствах не имеющих прямого доступа к интернет.

«Галочка» должна быть включена на site 2 и 3 border чтобы site 1 подключился к Интернет через SD-Access transit.

SD-Access для распределенного кампуса

Data Center 2 на Site 3

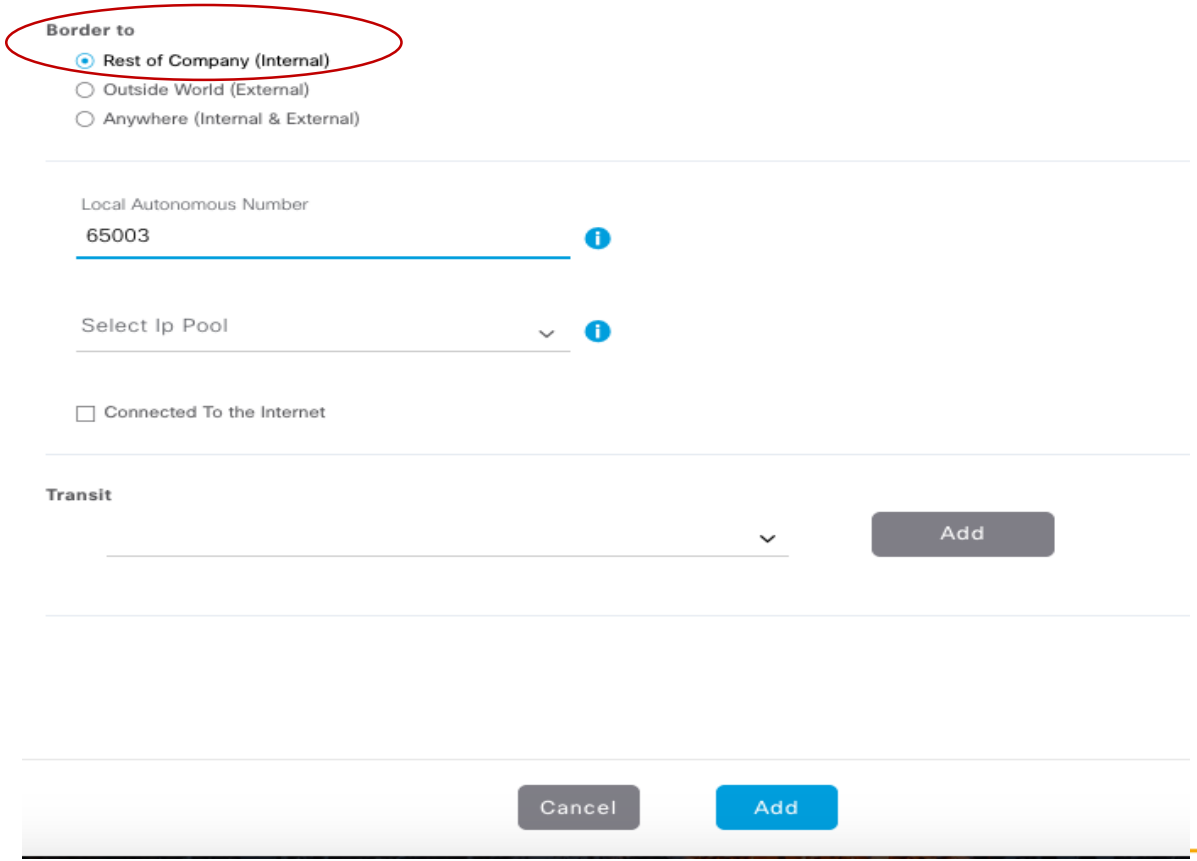

Выбираем Internal border

Это необходимо для правильного импорта маршрутов из DC1 и регистрации на узле control plane.

На слайде приведена настройка на site 3 border (B”) для взаимодействия с DC2.

SD-Access @ DNA Center

Настройка распределенного кампуса в SD-Access

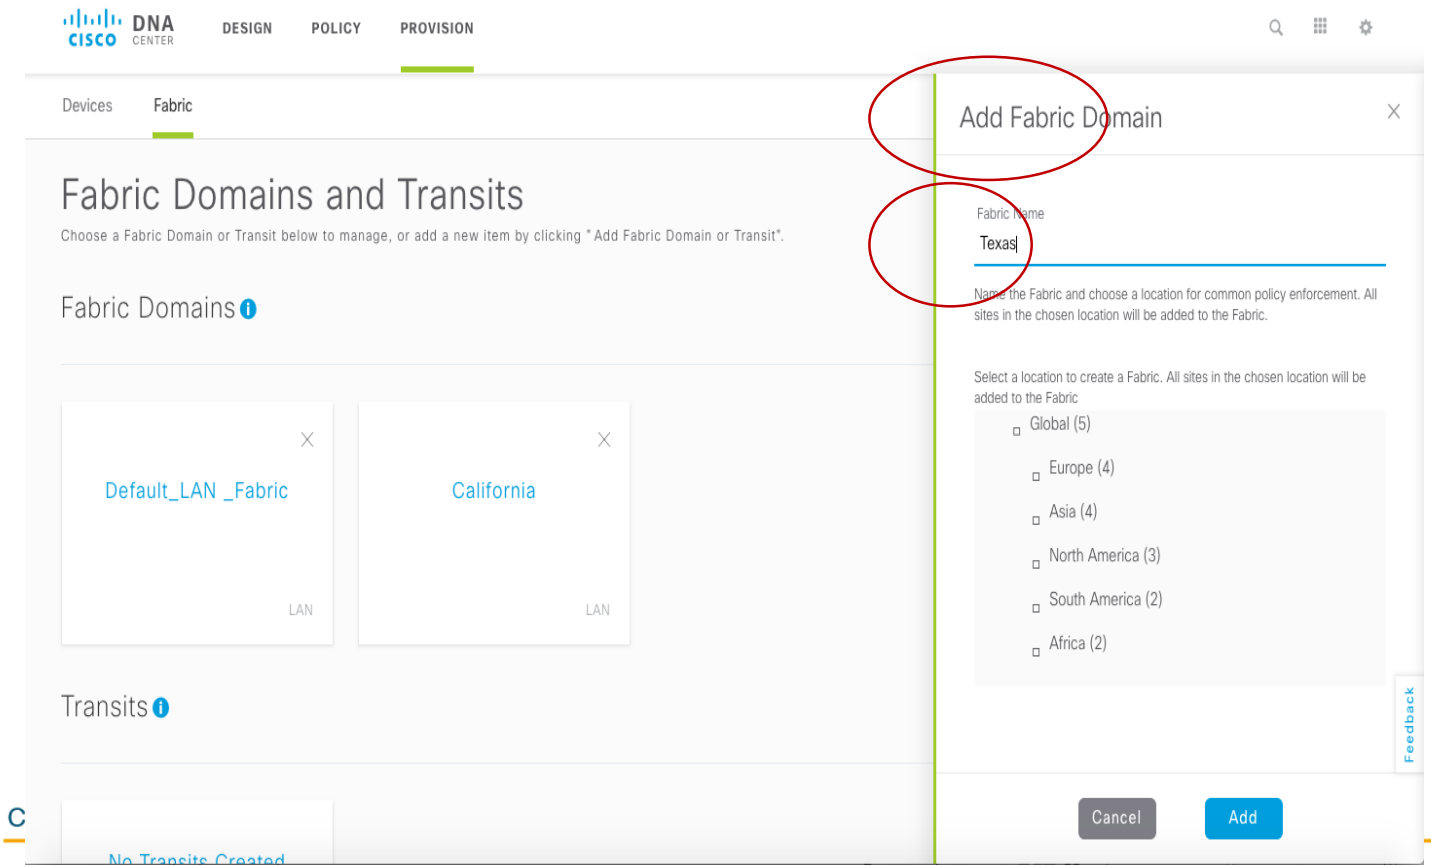

Создаем Fabric Domain

Получим новый fabric domain к которому сможем добавить новые fabric site

SD-Access @ DNA Center

Настройка распределенного кампуса в SD-Access

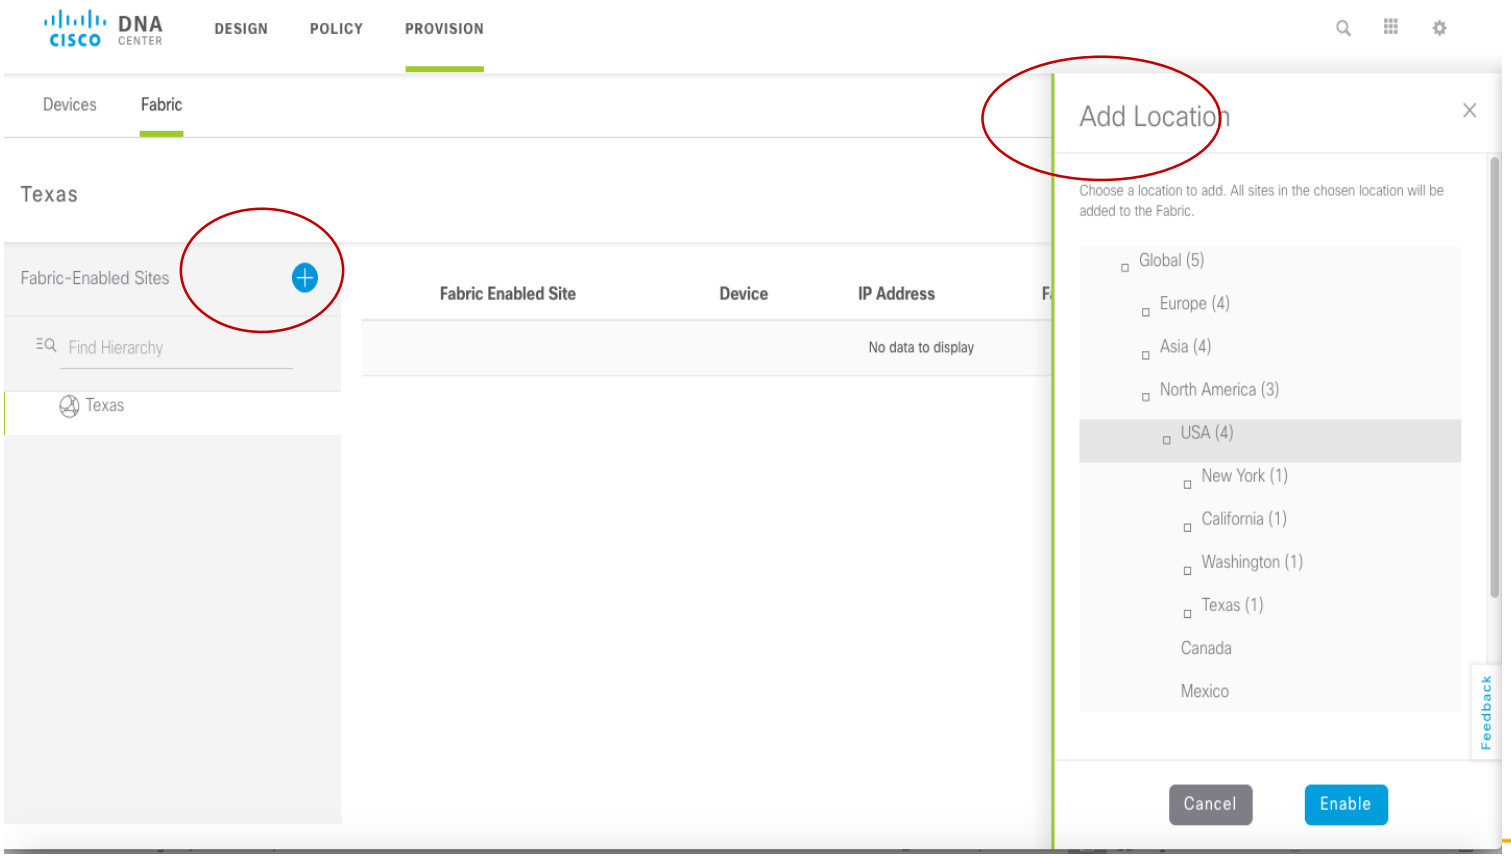

Добавляем Fabric Site к Fabric Domain,

Получаем новый site в рамках fabric domain.

Используем иерархическую модель – например, страна и локации в рамках страны New York , California , Texas etc. ,будут добавлены к fabric site

SD-Access @ DNA Center

Настройка распределенного кампуса в SD-Access

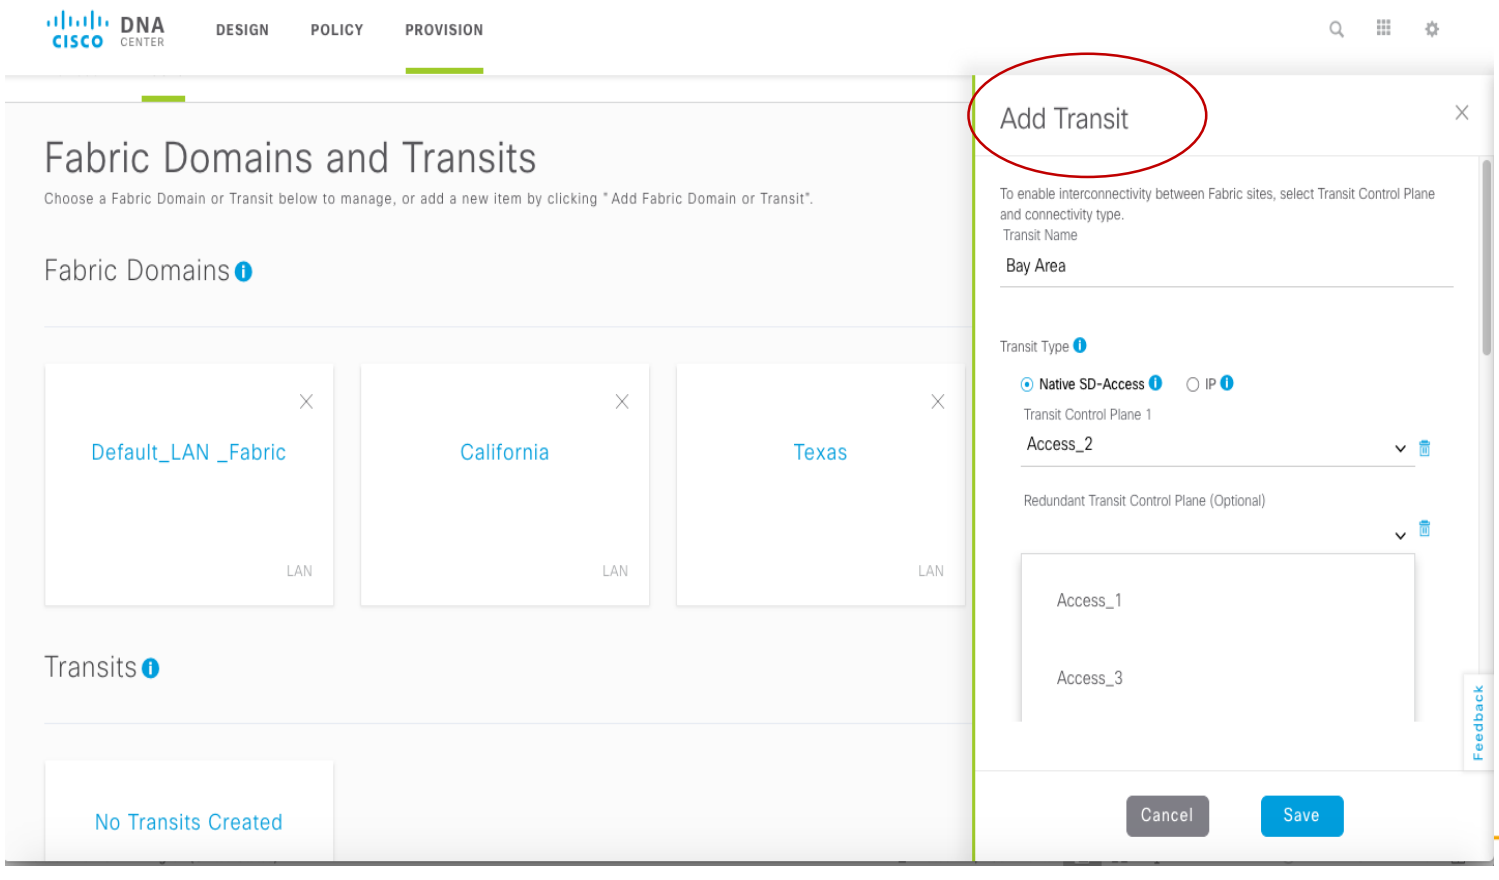

Создаем Transit

Добавляем узлы Transit CP к Transit

SD-Access @ DNA Center

Настройка распределенного кампуса в SD-Access

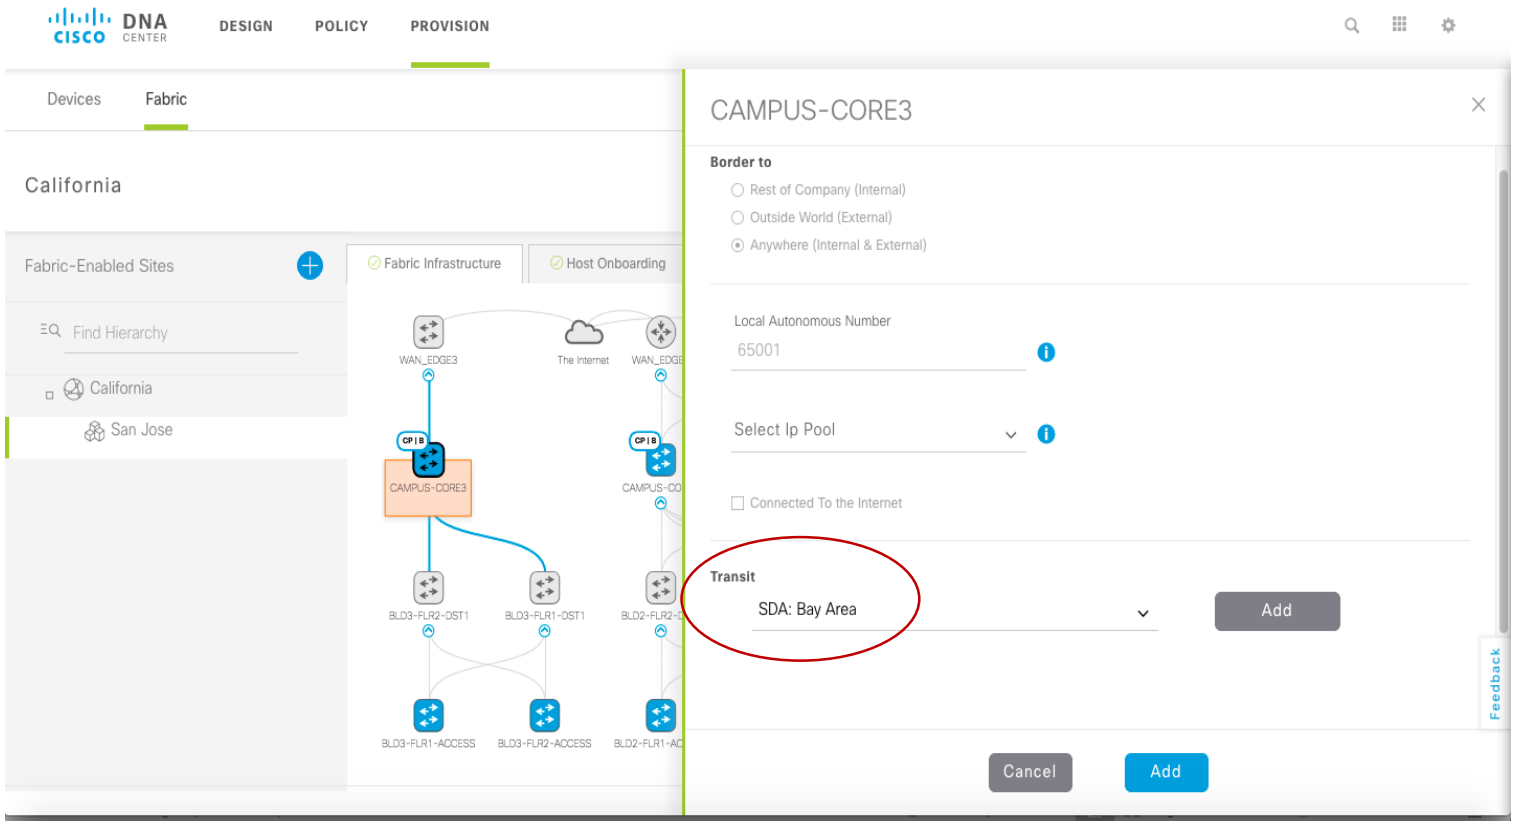

Добавляем выбранные узлы Border к Transit.

SD-Access для распределенного кампуса – Политики

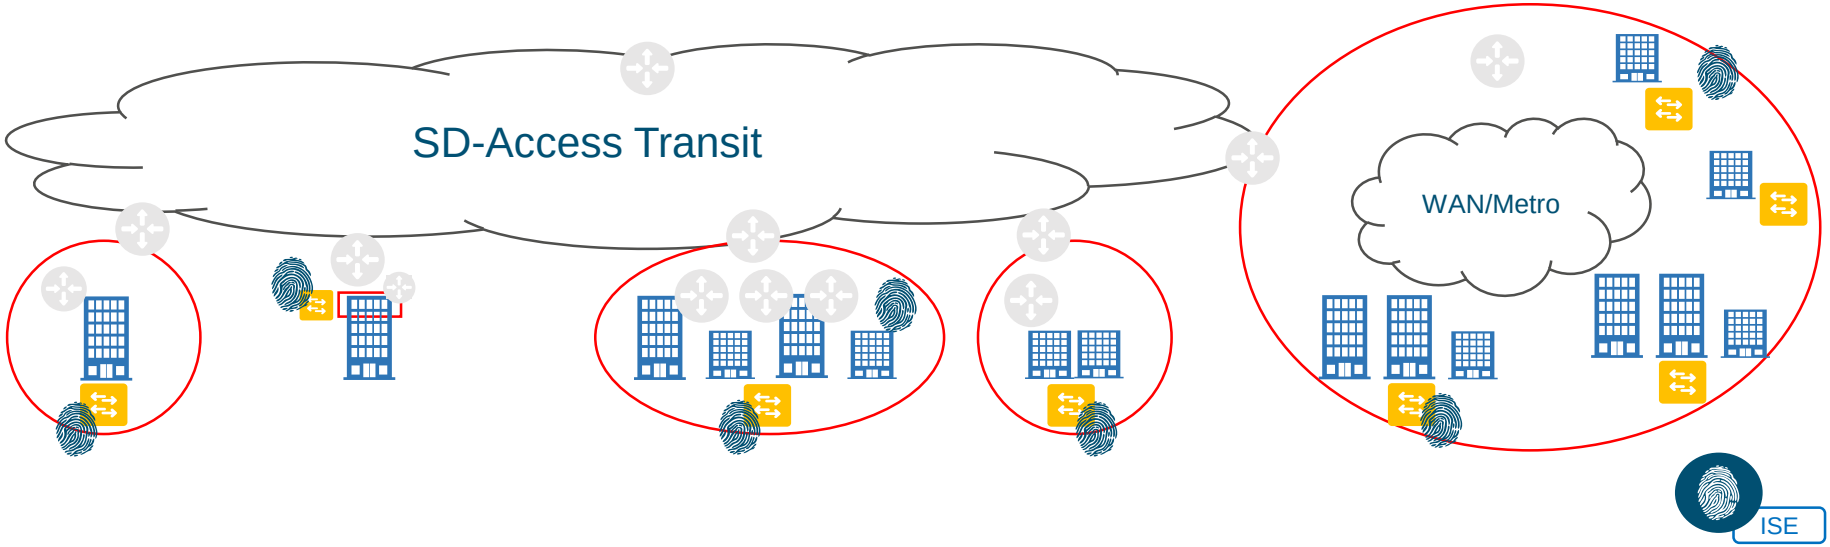

- Группы и VN’ы сохраняются между площадками
- Политики End to end применяются на базе меток SGT

SD-Access для распределенного кампуса – Политики

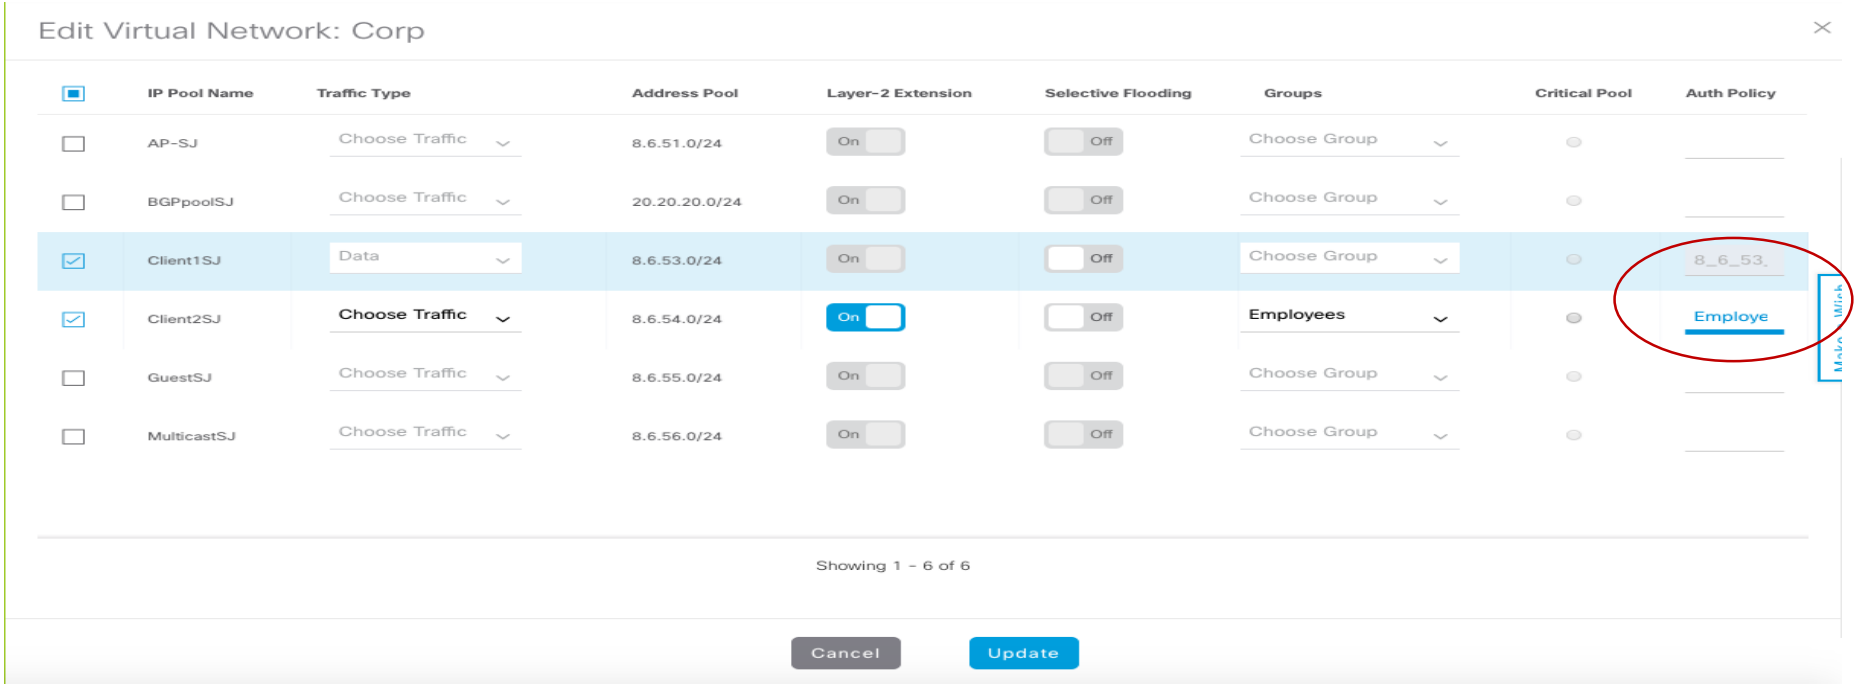

- Разные подсети между площадками могут иметь одинаковые Auth Policy/Vlan
- Общие политики поддерживаются для общего набора пользователей

SD-Access для распределенного кампуса – беспроводная сеть

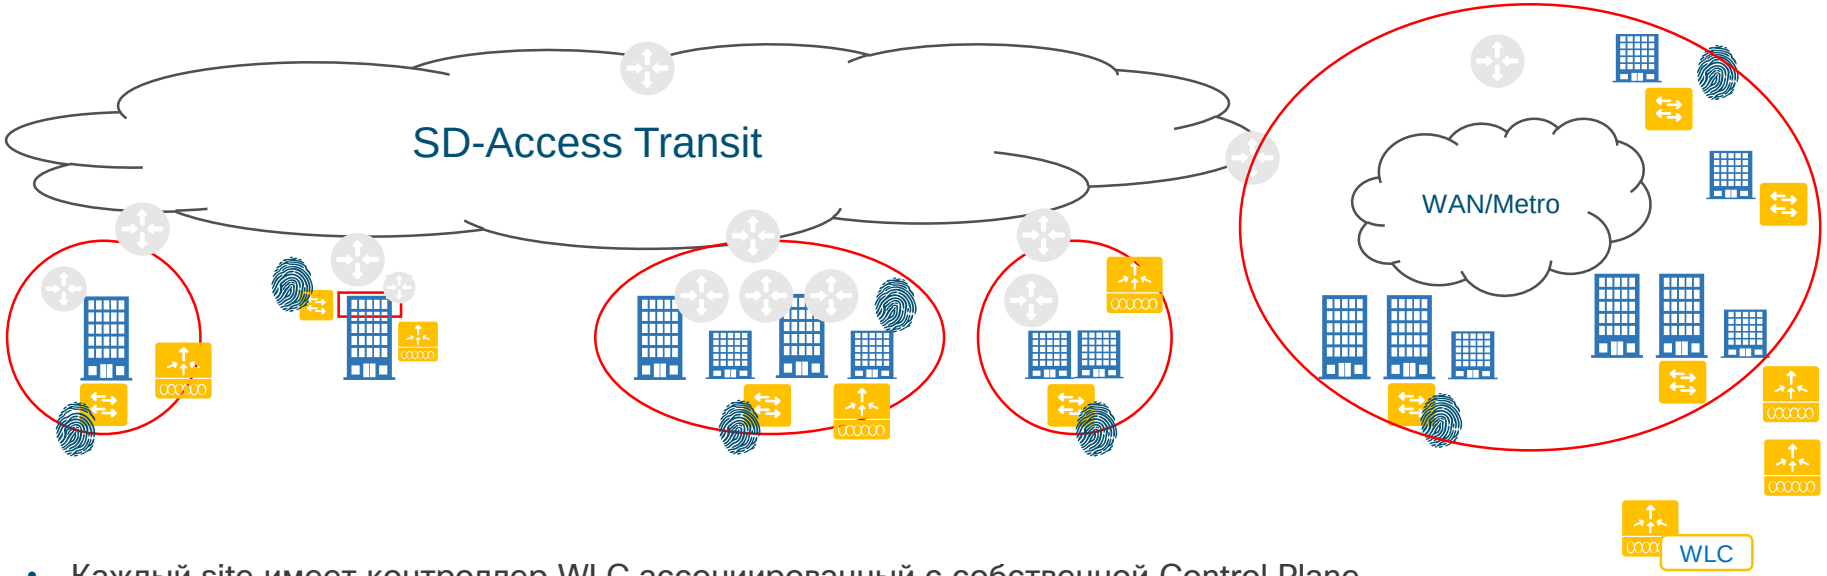

- Каждый site имеет контроллер WLC ассоциированный с собственной Control Plane
- Небольшие локации в рамках одного сайта могут иметь общий WLC

ISE распределенное внедрение – вариант 1

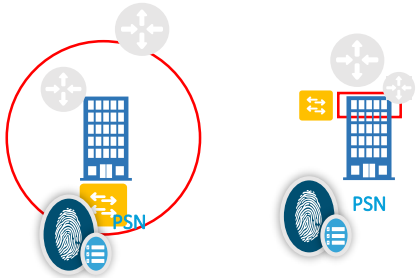

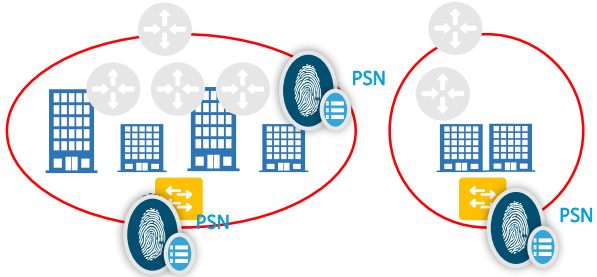

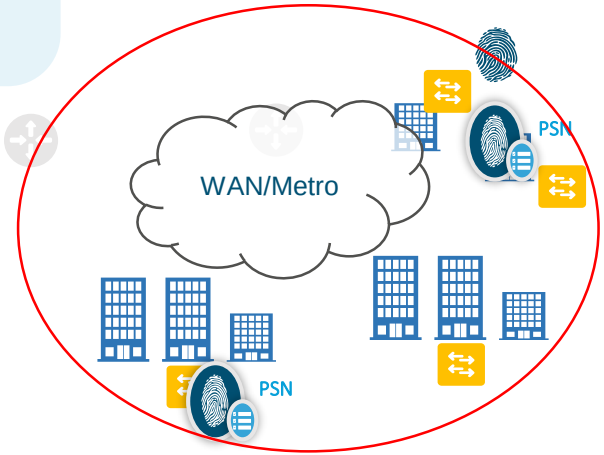

- Узлы PSN на каждой площадке
- Максимум 2 PSN на площадку
- Узлы PAN в центральном ЦОД

ISE распределенное внедрение – вариант 2

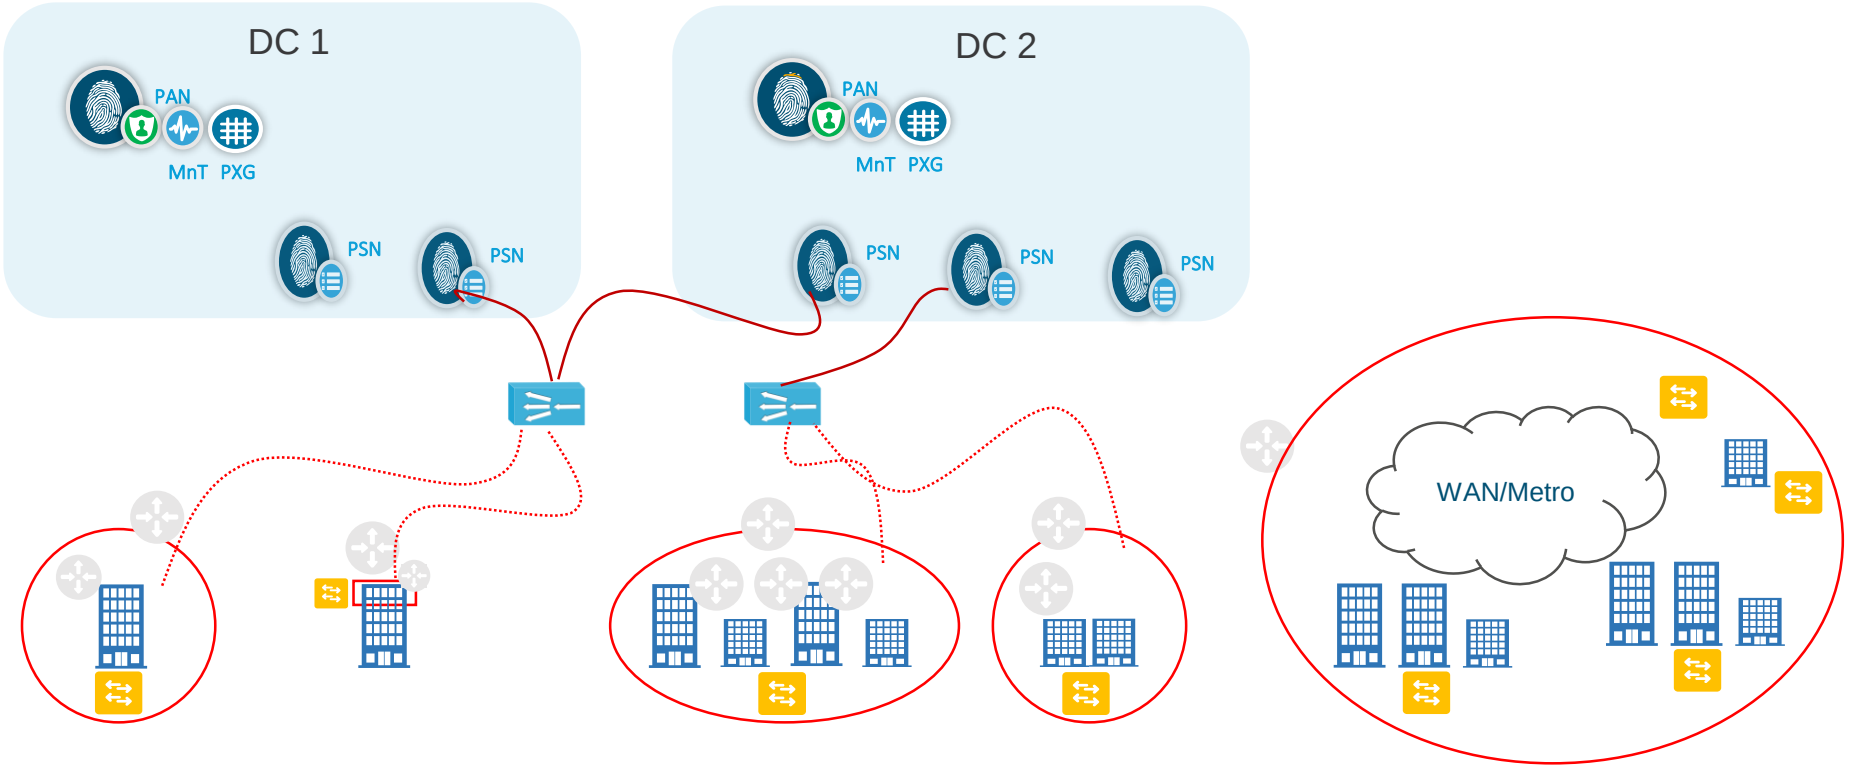

- Узлы PSN за балансировщиками нагрузки
- В интерфейсе DNAC указывается IP на балансировщике нагрузки (VIP)

Рекомендации по развертыванию и миграции

Cisco Validated Designs

…provide a framework for design and deployment guidance based on common use cases.

"Classic" and Modular CVDs, White Papers, and Tools

cisco.com/go/cvd/campus

Connectivity Services

Where do I place Cisco DNA Center?

Local DC or Services Block

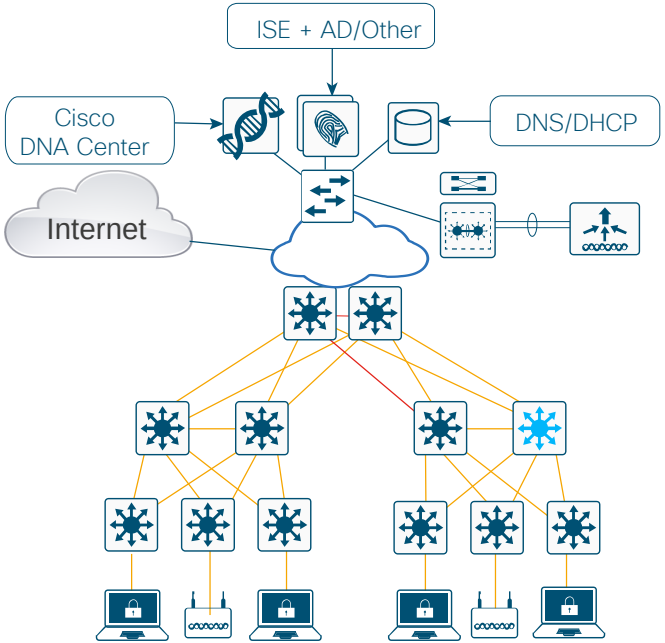

Remote DC

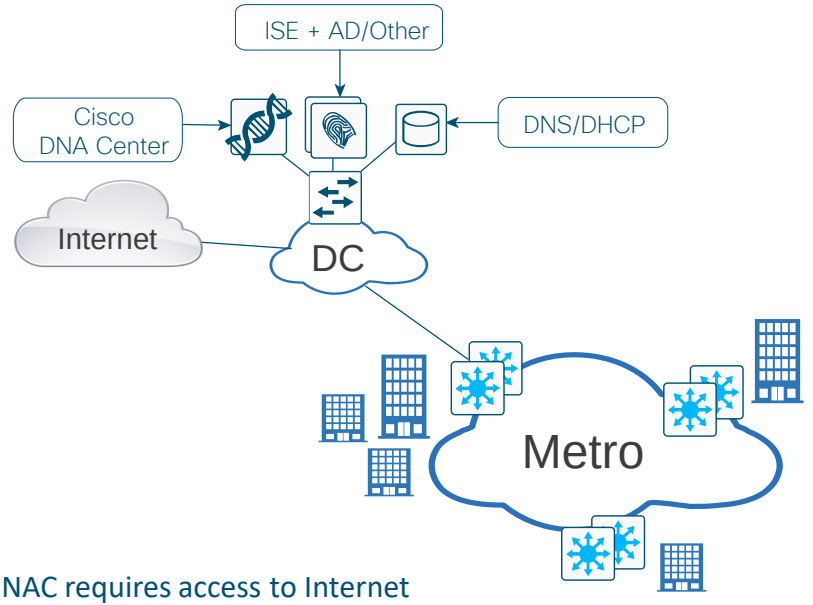

NOTE: Cisco DNAC requires access to Internet

Network Infrastructure – Underlay

Cisco SD-Access underlay options

Manual Underlay

- Routed Network
- System MTU: 9100
- Loopback 0 with /32 subnet
- Resiliency – BFD, ECMP, NSF
- Multicast – SSM, sparse-mode
- CLI, SNMP credentials
- Discover & Manage network device
- Upgrade Software version

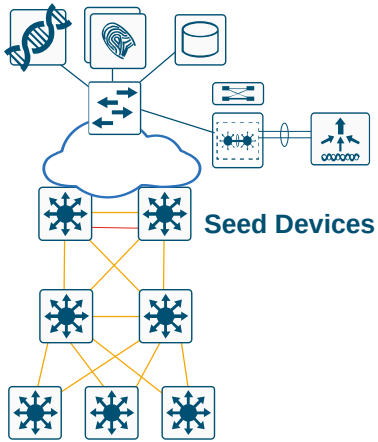

Automated Underlay

- Discover Seed Device
- IP Address Pool
- LAN Automation
  - ✓ Discover the network device
  - ✓ Onboard the network device
  - ✓ Underlay Configuration
  - ✓ Upgrade software
  - ✓ Manage Device in Cisco DNA-Center

Cisco DNA Center Appliance Connections

Physical Interface & Connections

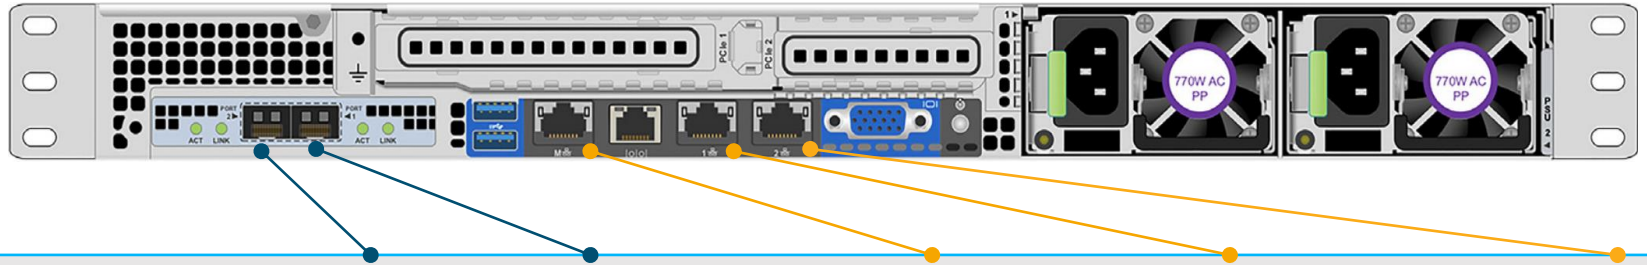

DN1-HW-APL	PORT 2 MLOM SFP+ 10Gb	PORT 1 MLOM SFP+ 10Gb	M Integrated RJ45	1 Integrated RJ45	2 Integrated RJ45
Wizard Name	enp10s0	enp9s0	—	enp1s0f0	enp1s0f1
Use	Intra-cluster communications	Enterprise network infrastructure	CIMC out-of-band server appliance management	Optional dedicated management network for web access	Optional isolated enterprise network
Example Cluster VIP Address	192.168.0.1	10.4.249.250	—	100.119.104.200	—
Example Interface Address	192.168.0.2 255.255.255.248	10.4.249.241 255.255.255.0	10.204.49.25 255.255.255.0	100.119.104.241 255.255.255.0	Unused in CVD

- NTP, DNS is required
- Service Subnet & Cluster Service Subnet is required (/21 subnet)

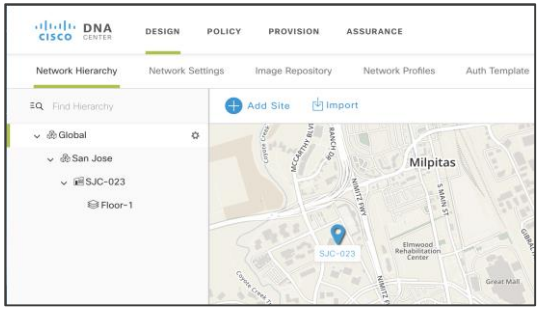

Network Hierarchy

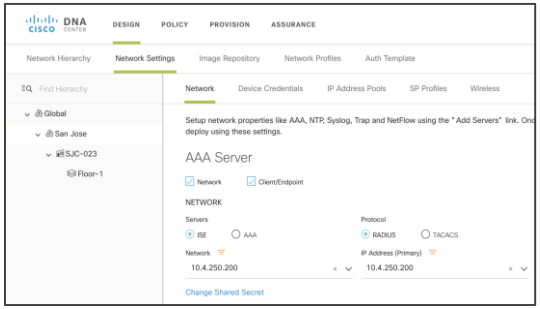

Network Settings – Network Settings

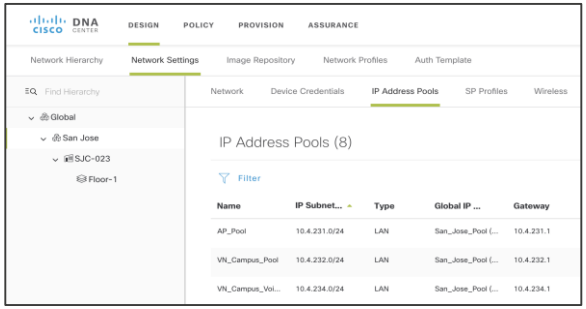

Network Settings - Address Pools

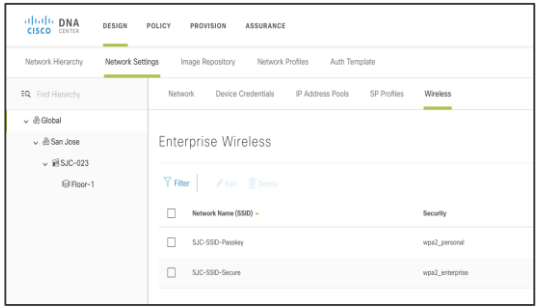

Network Settings - Wireless

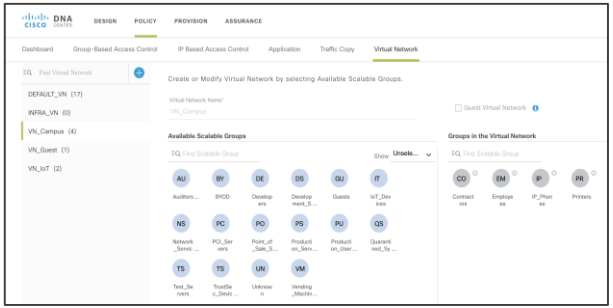

Policy – Macro Segmentation

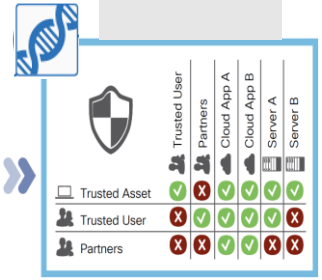

Policy – Micro Segmentation

Cisco SD-Access

before Onboarding Endpoints – Fusion Configuration

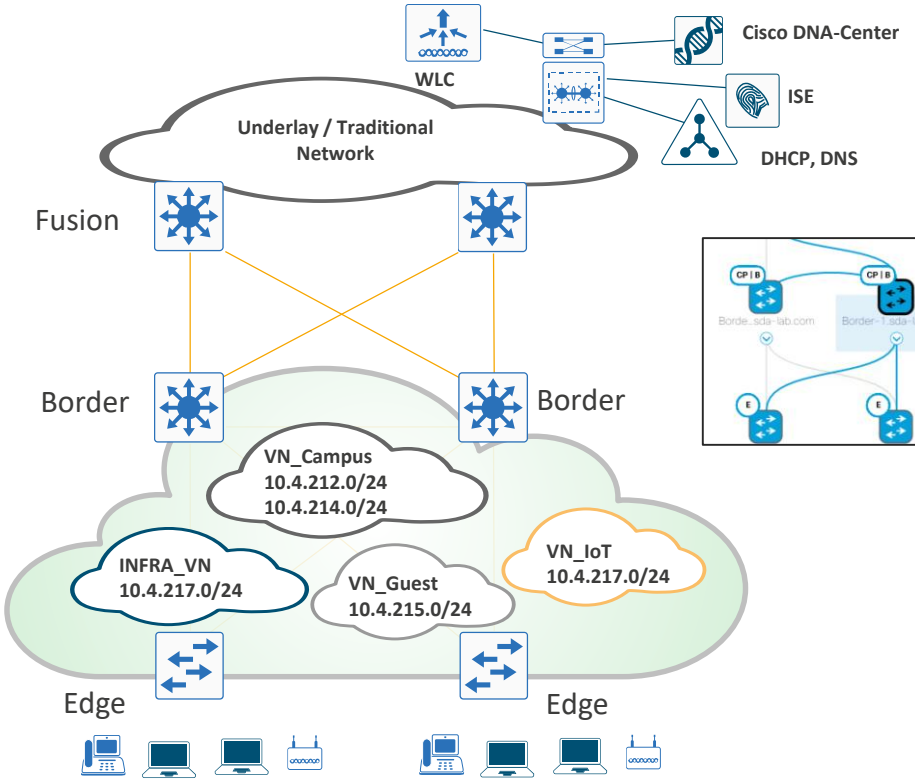

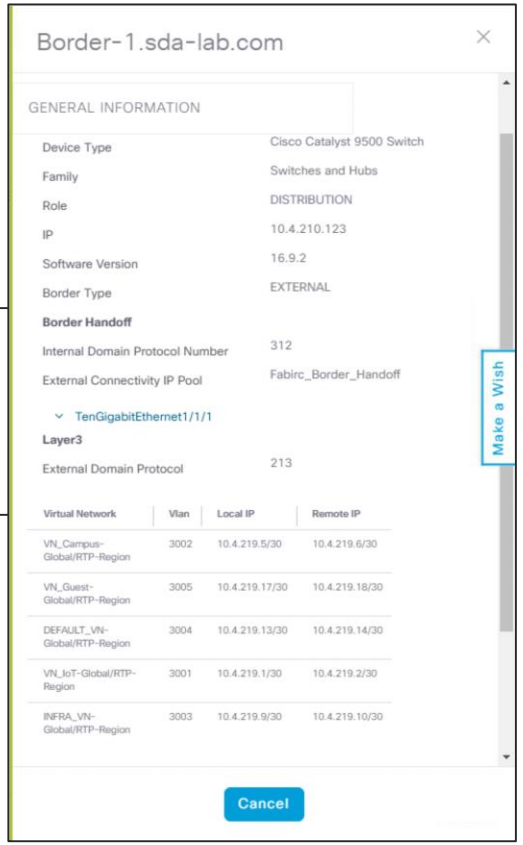

Fusion Configuration

Connecting Fabric to Traditional Infrastructure

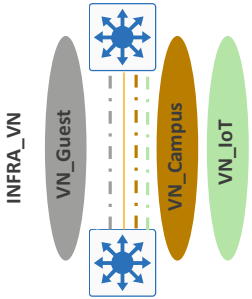

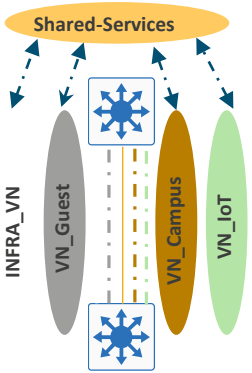

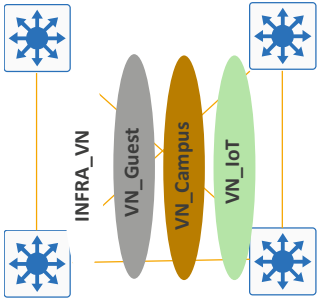

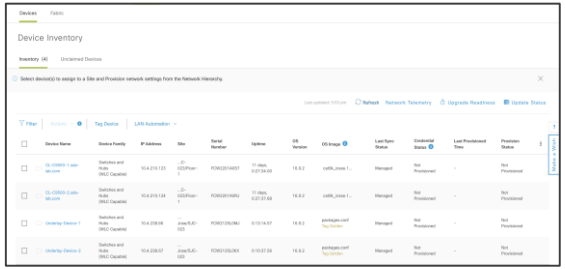

Provision

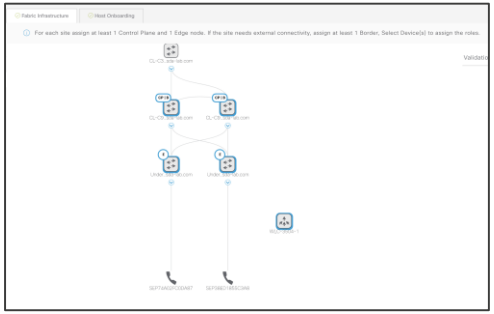

Fabric Provision – Transit Site
Fabric Site

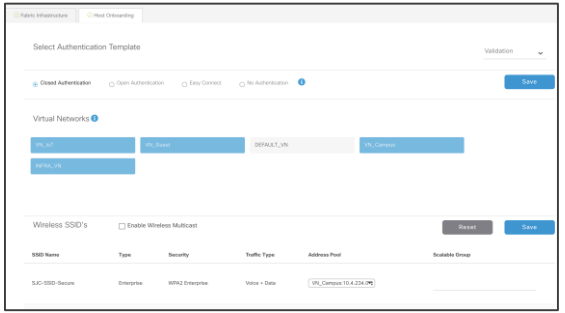

Fabric Provision Host Onboarding
Default Authentication Template
Address Pool Assignment

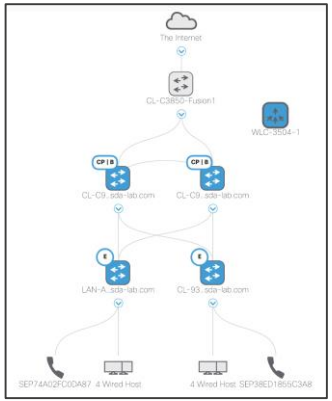

Fabric Network

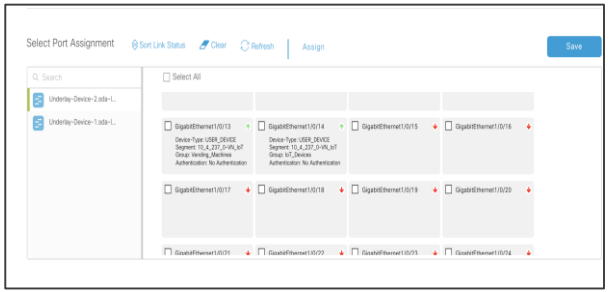

Fabric Provision Host On-boarding
Port Assignment

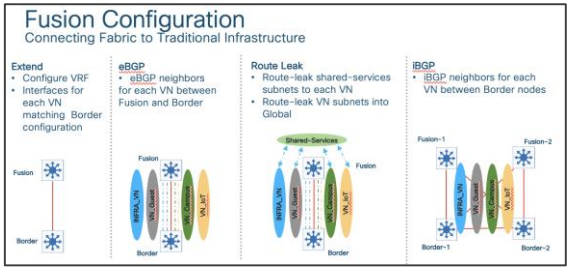

Fusion Configuration

На что обратить внимание

Existing Network MTU

- **VXLAN** adds **50 bytes** to the Original Ethernet Frame in the Overlay
- Avoid Fragmentation by adjusting the network MTU
- Ensure Jumbo Frame support on switches in the underlay network

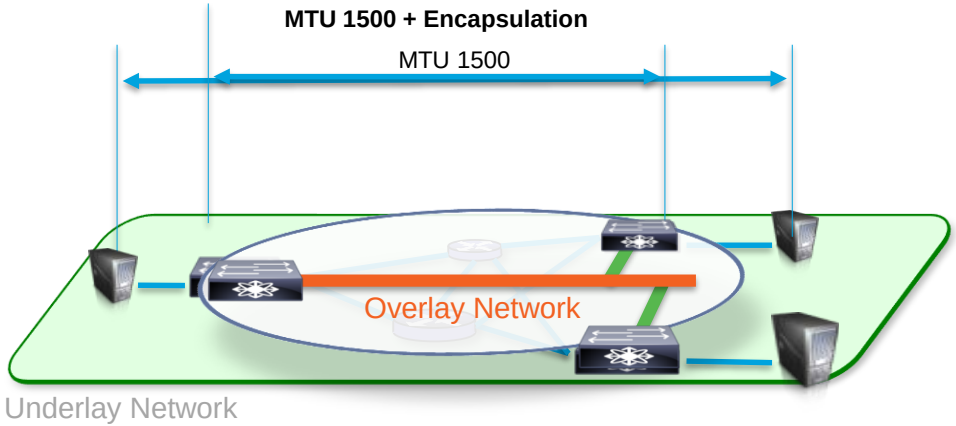

TCP MSS

- TCP MSS adjust is supported in 16.9.1s and later
- Available only on Catalyst 3K and 9K only and **works only on TCP based applications**
- Applied to the overlay SVI on Fabric Edges via Template Editor
- PMTUD is being explored as a solution for UDP traffic.

As of now, Jumbo MTU is mandatory on all switches.

Re-configuration of Access Layer

- Layer-2 **Switched Access** today

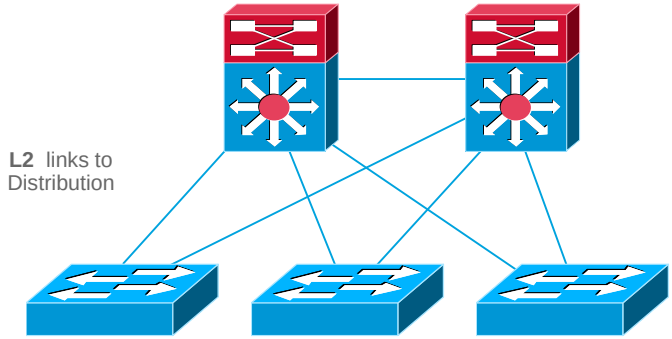

- **Routed Access** tomorrow

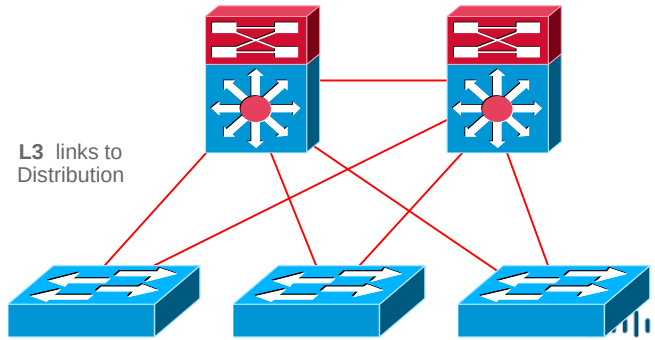

Physical Network Topology

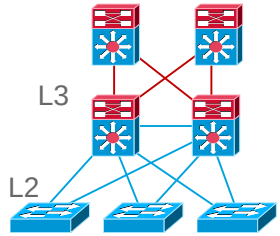

3-Tier Hierarchical

Cisco SD-Access fabric runs over most topologies:

- Traditional 3-tier hierarchical network
- Collapsed core/aggregation
- Routed access
- U-topology

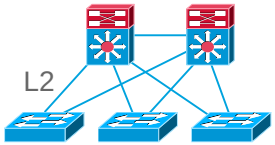

Collapsed Core

- **Ideal to start with routed access** – allows fabric to extend to very edge of campus network with minimum impact.
- Ensure that all switches have IP reachability to infrastructure elements

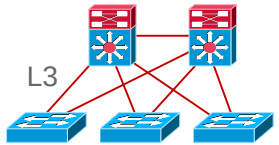

Routed Access

follow campus CVDs with routed access:
www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Campus/routed-ex.html

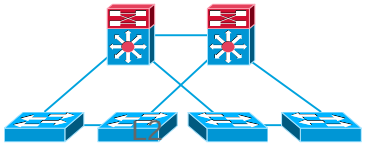
U-Topology

IP Addressing for Underlay and Overlay

Know your IP addressing and IP scale requirements

- **IPv4 only** (today)
- Fabric uses Loopback 0 as Source-Interface for Encapsulation
- Best to use single Aggregate for all Underlay Links and Loopbacks

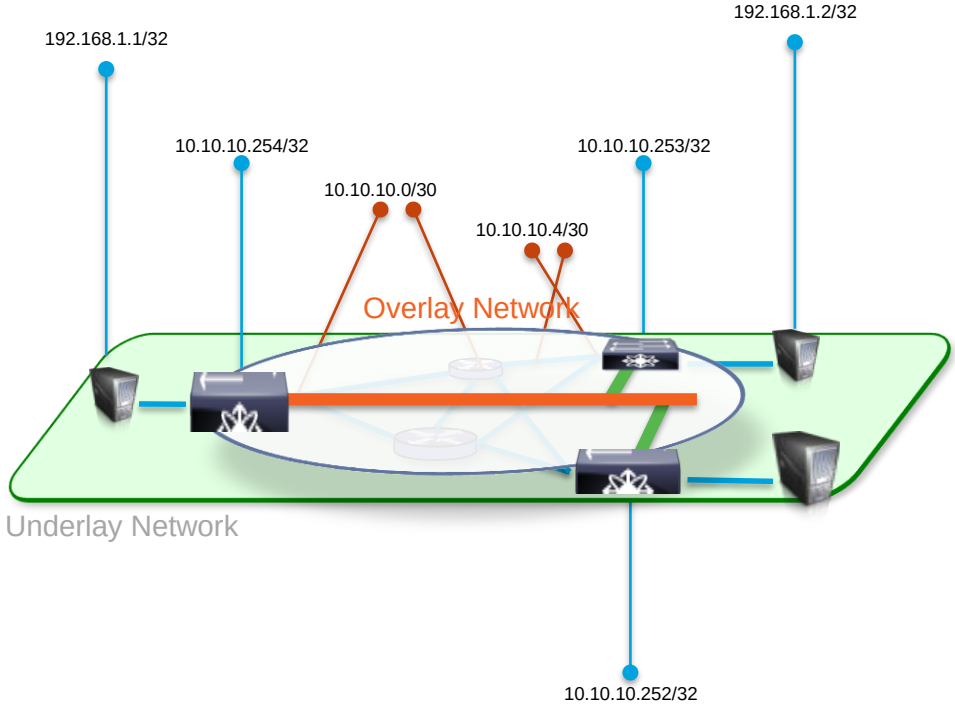

Location of Shared Services Infrastructure

Are the Shared Services in VRF or Global Routing Table?

- Cisco SD-Access fabric can leverages traditional infrastructure services
- IP reachability from underlay/overlay to DNS, DHCP, etc. required

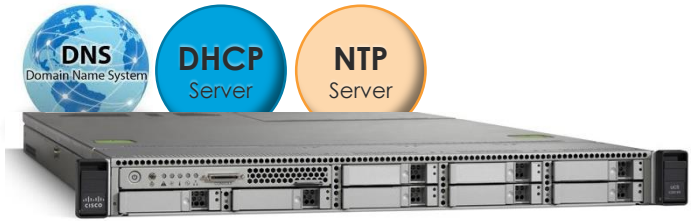

Location of Shared Services Infrastructure

- Larger deployments have infrastructure services hosted in Data Centre
- Hybrid model also possible (mix of distribution/core/Data Centre)
- The Shared Services **should be outside** the fabric for Cisco SD-Access

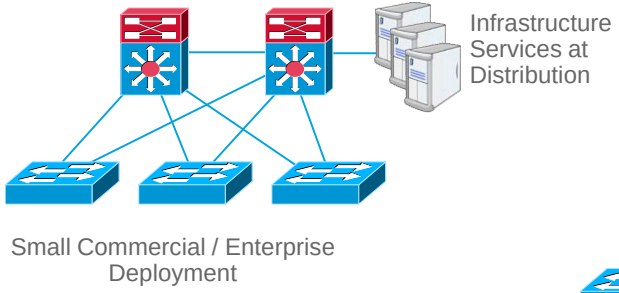

Small Commercial / Enterprise Deployment

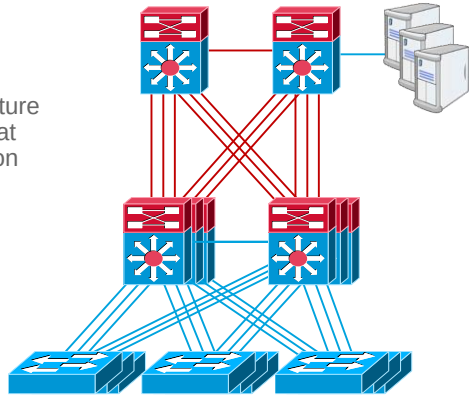

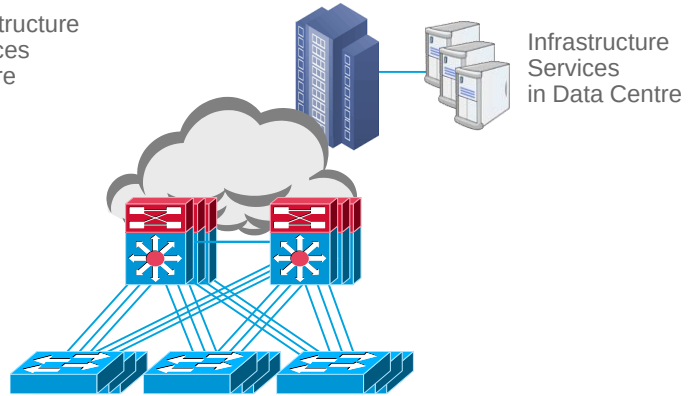

Large Enterprise Deployment

Features enabled today

Where are policies applied today?

- For example, features like QoS, NetFlow, Policy-based Routing, IP ACLs?
- **Need to move** the policy enforcement point(s) down at the **Access layer** or **outside the fabric**

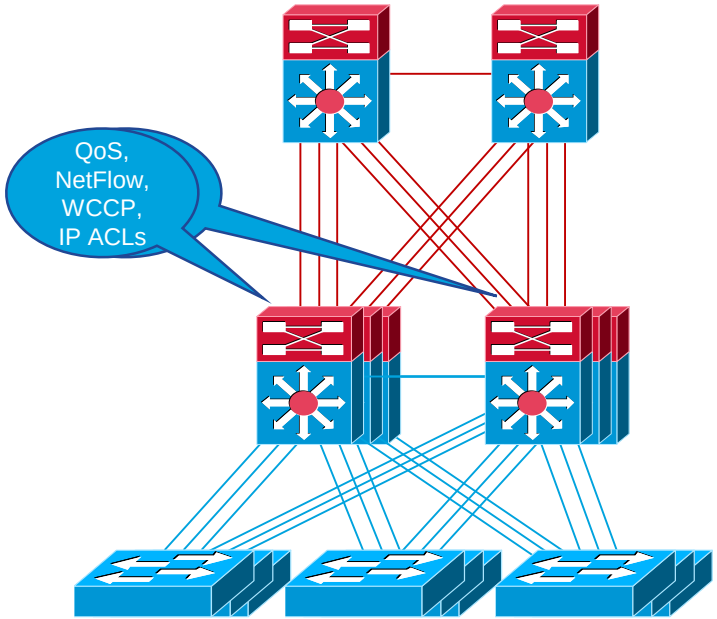

Move the features to different points in the fabric network

- **Move some** Policy enforcement point(s) **down to the Access Layer.** For example, IP ACLs, QoS, NetFlow can be applied at the Access layer
- **Move some** Policy enforcement point(s) **outside the SD-Access fabric.** For example, PBR, WCCP can be applied external to the fabric.

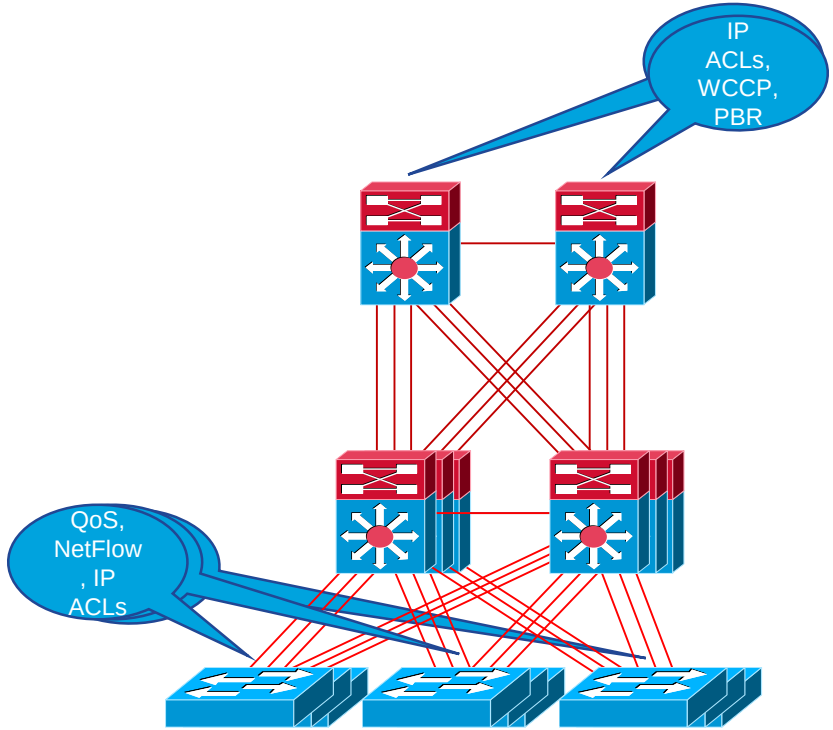

Two Basic Types of Deployments

- Campus Networks/(Large Sites)
- Branch Networks/(Small Sites)

Typical Campus Networks

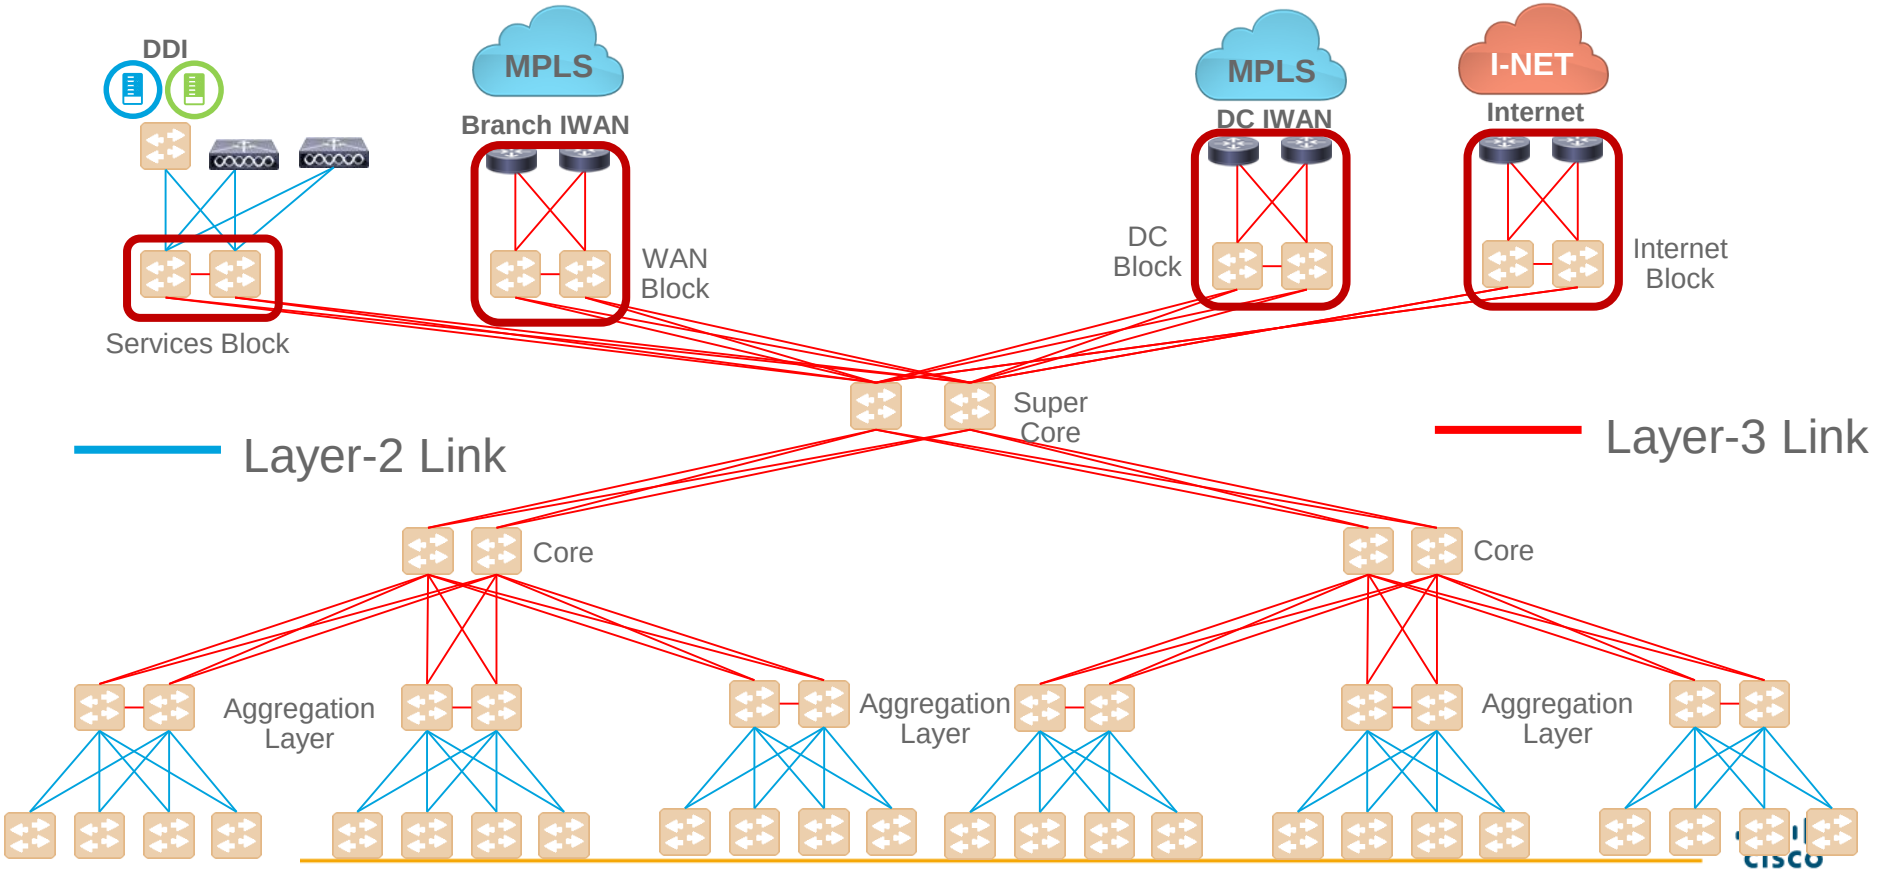

Typical Branch Networks

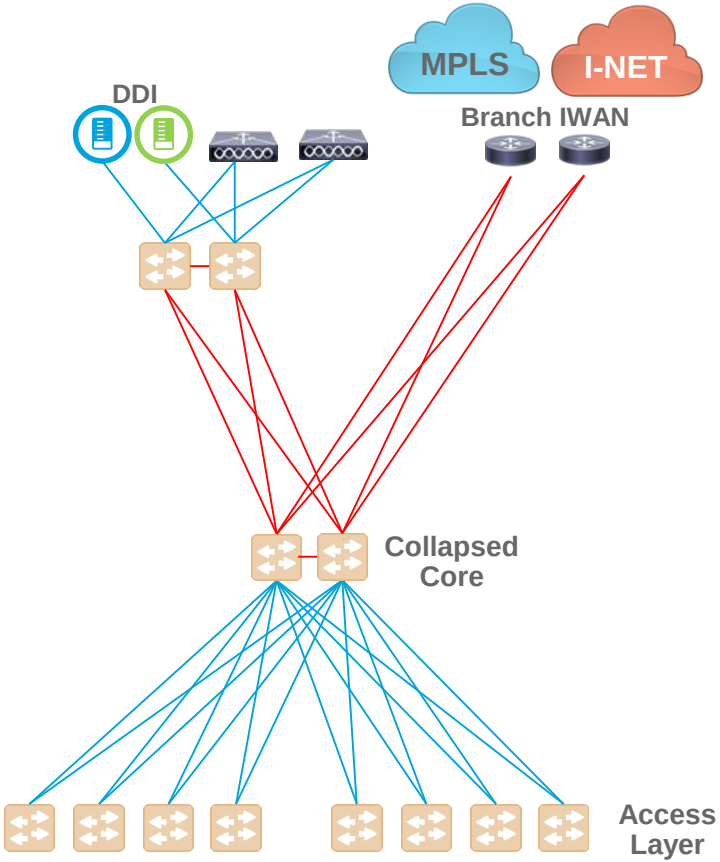

Two Basic Approaches to Migration

- **Parallel Deployment** (all at once)
- **Incremental Deployment** (one at a time)

Migration Approaches: Parallel vs Incremental

Parallel	Incremental
Best for Branch (small scale) deployments	Best for Campus (any size)
Requires cable runs to create a new parallel network	Requires a couple of cables from new access and distribution switches
Power and outlets for the parallel network	Incremental power and outlet requirement
Legacy hardware in existing network	Legacy hardware in existing network
Upgrade most of the network infrastructure	Upgrade most of the network infrastructure
Clean slate (leaving behind any complexity in the old design)	Will need to carry forward the constraints of the old design in the underlay
Test users in a complete new network	Test of functionality is partial
Easy Rollback of migrated users	Easy Rollback of migrated users

IMPLEMENTATION

RESOURCES

RESOURCES

IMPLEMENTATION

Parallel Install not feasible for Campus Networks

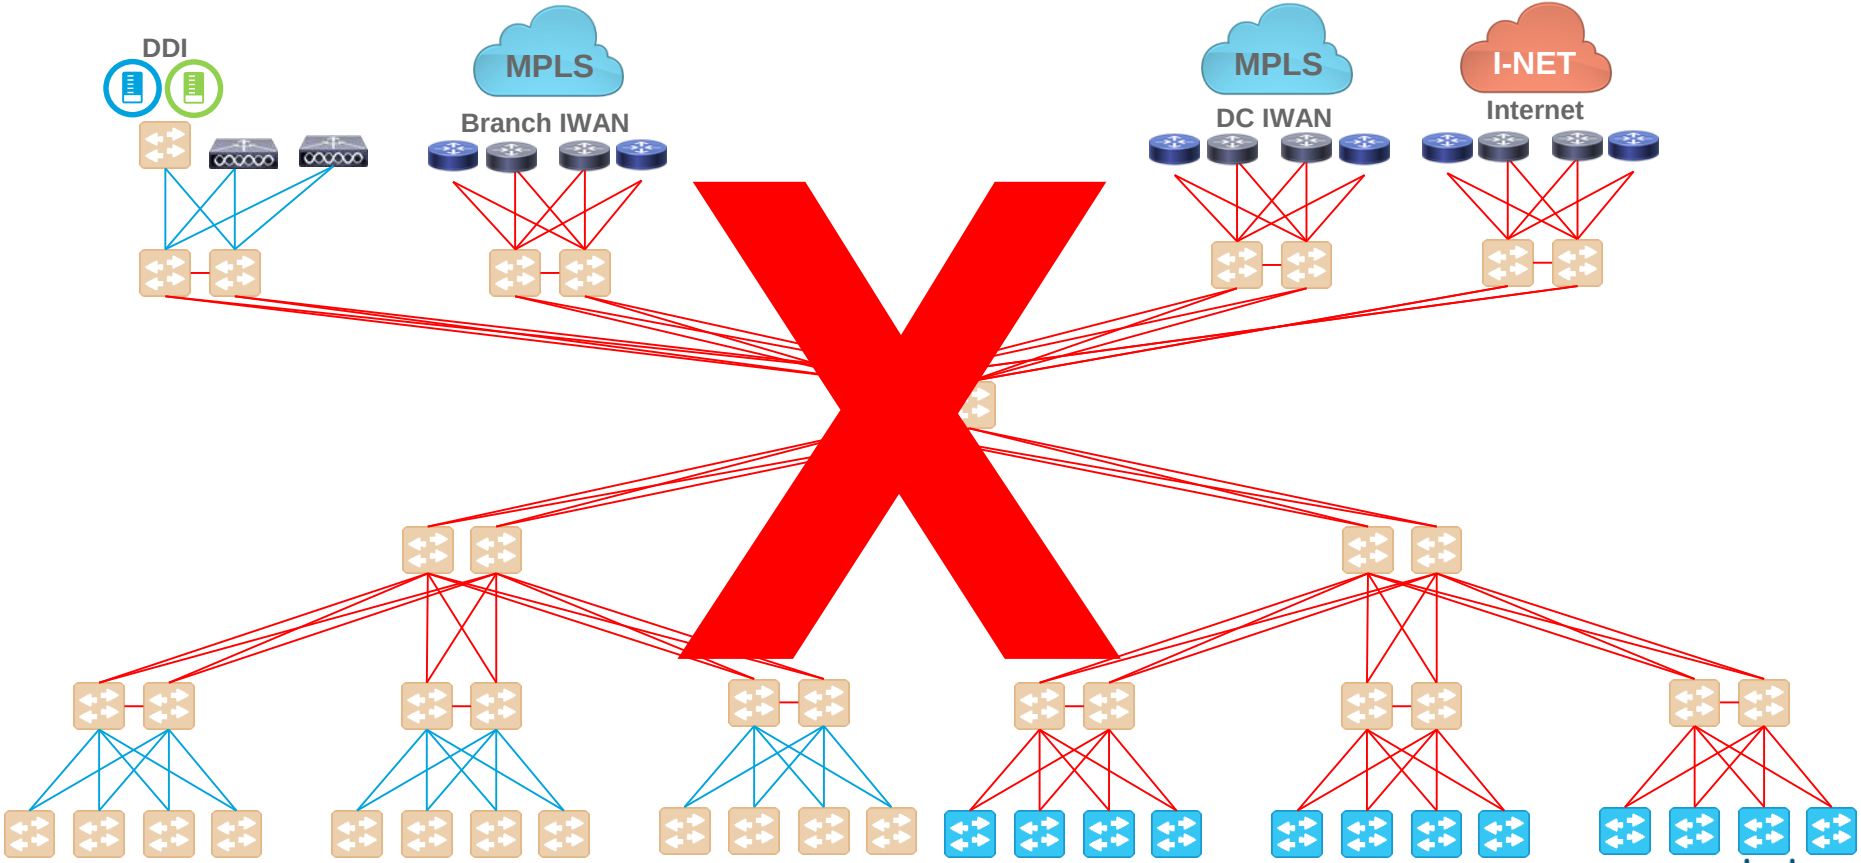

Parallel Install for Branch Networks

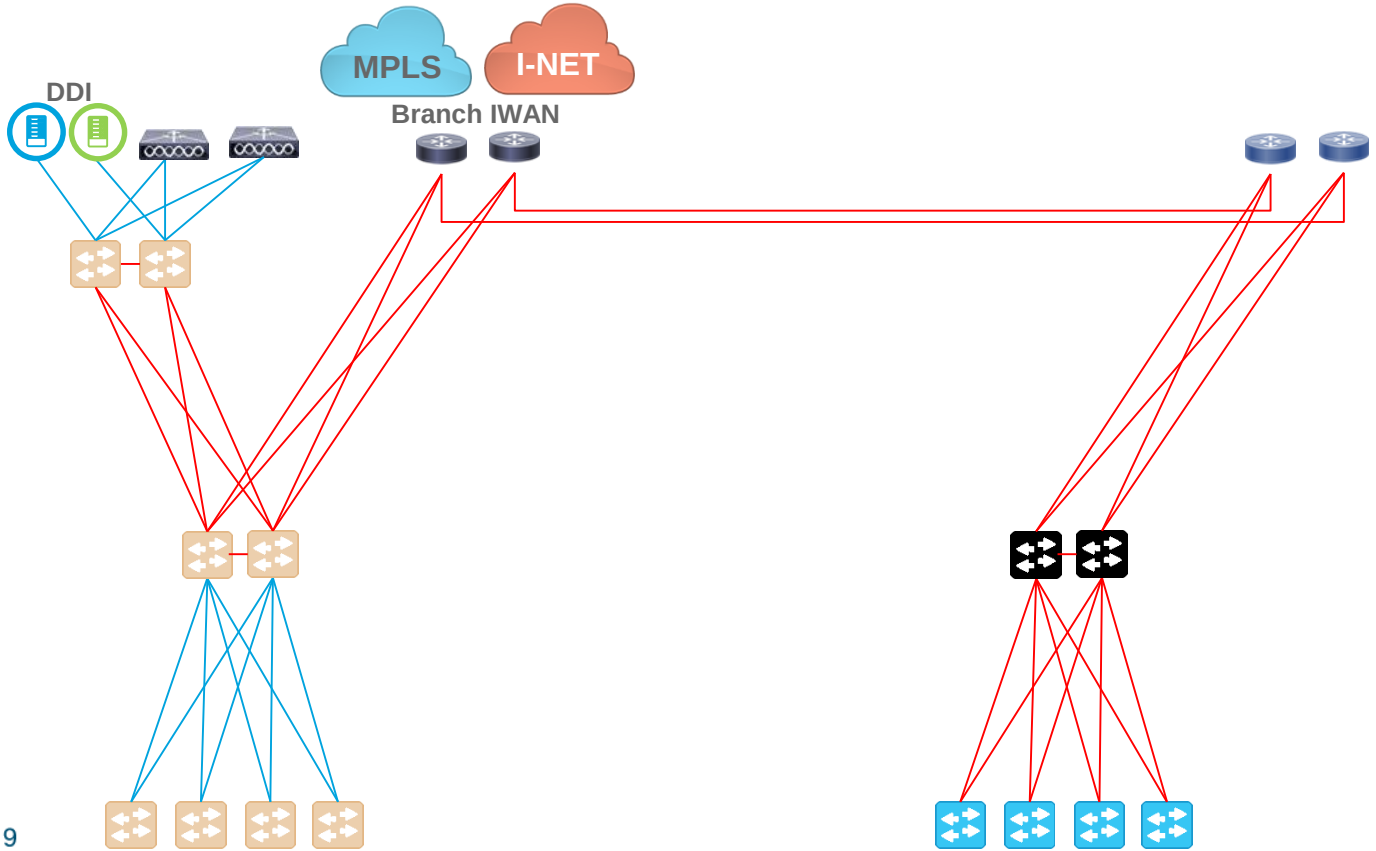

Migration with New subnets

Incremental Migration – High Level concept

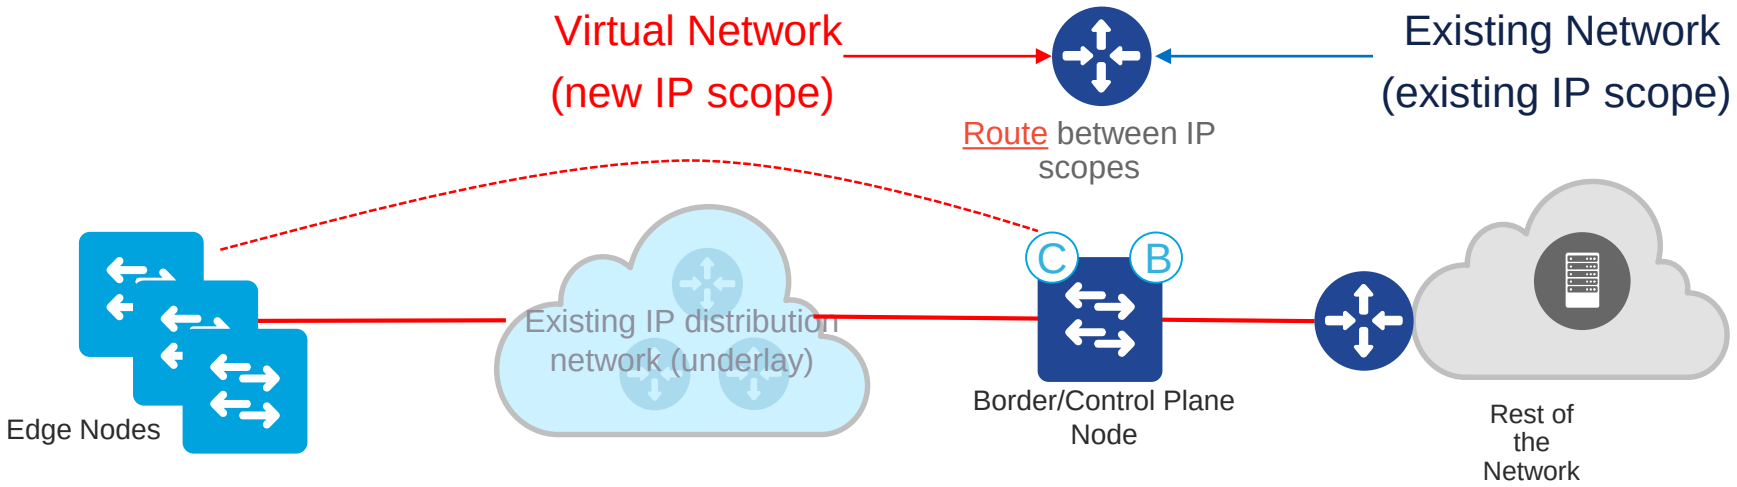

- Deploy a **Border/Control Plane node** and an **Edge node**
- A virtual network with new address is formed over the existing network
- **Incrementally** add Fabric Edge nodes
- The virtual network connects to the existing/external network via the border

Considerations for using new subnets to transition

- Immediately realize the advantages of bigger subnets, but lesser subnets that are optimized for Cisco SD-Access
- Design for the present and the future
- Add DHCP scope and size
- Update existing firewall rules for that one big subnet
- Not a big issue for endpoints with IP stacks that work well with DHCP

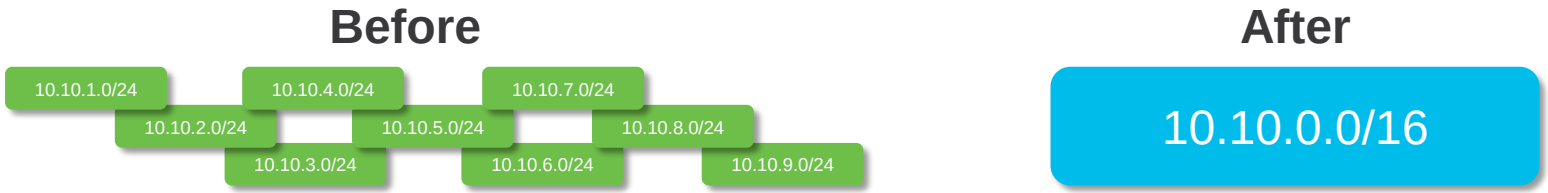

Reference Network Topology to begin Migration

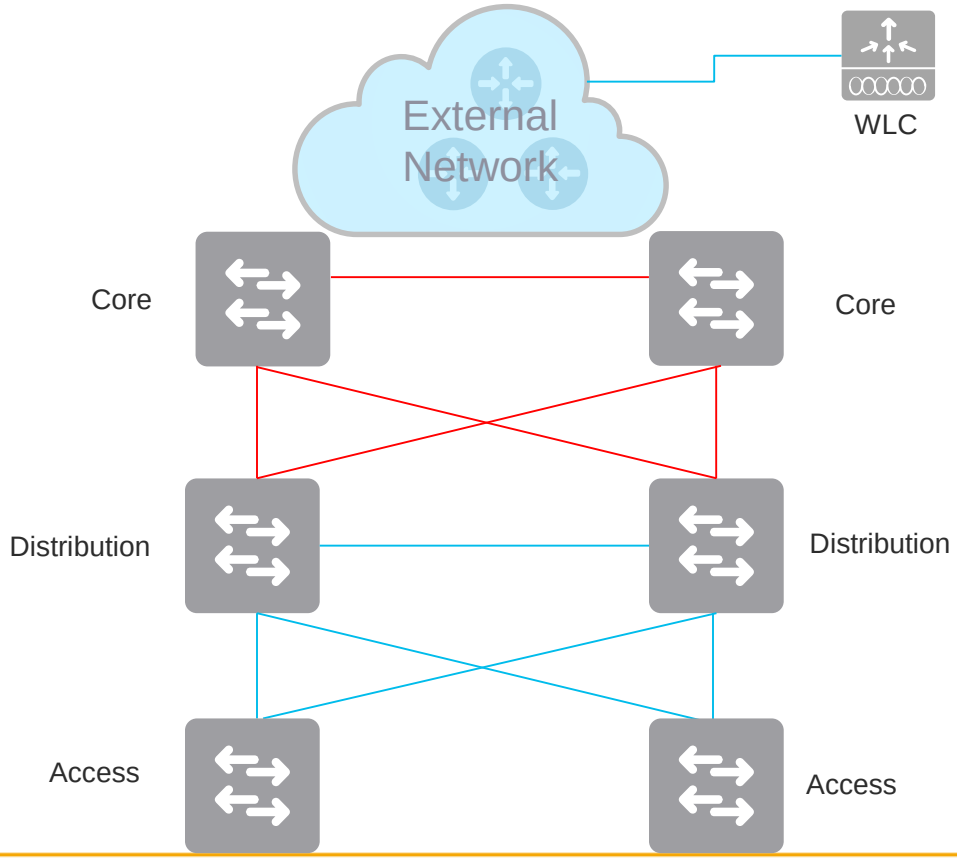

Connecting Default Fabric Border

- Option 1: Reconfigure Existing Core

You can reuse an existing Core switch if it supports Fabric functionality

NOTE: This may require software upgrade, and adding new fabric overlay configurations

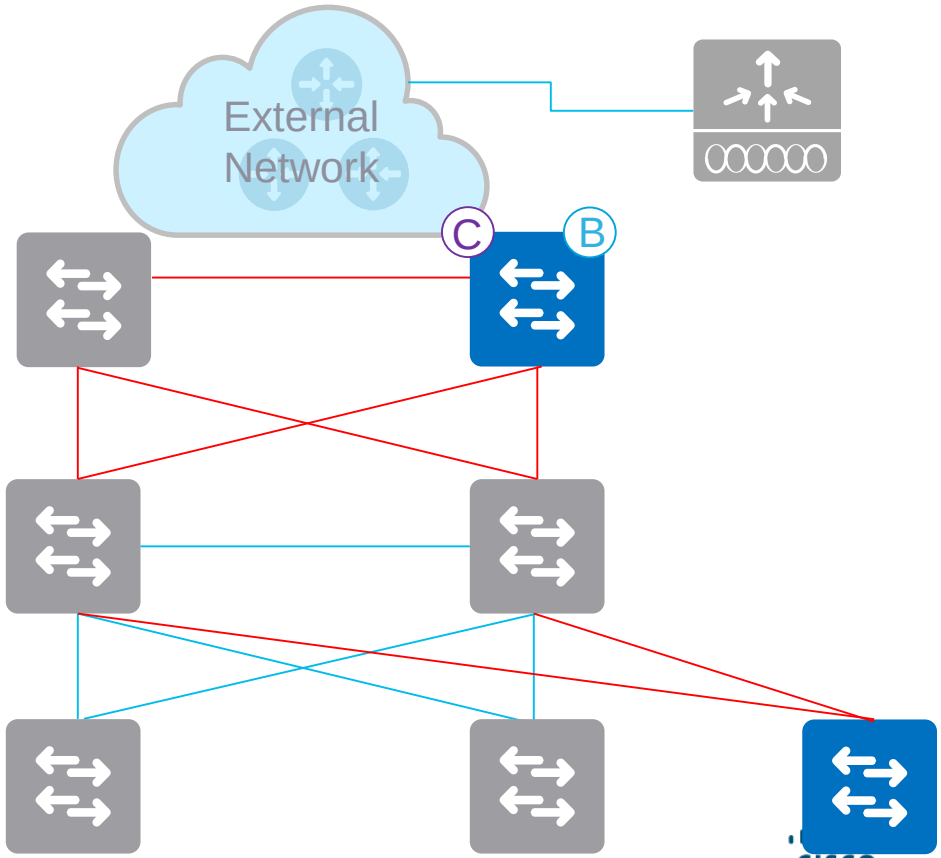

Connecting Default Border

- Option 2: Connect new switch to the existing core

If the existing core does not support Fabric functionality,

Connect a new switch to the existing core layer that will be a B/CP

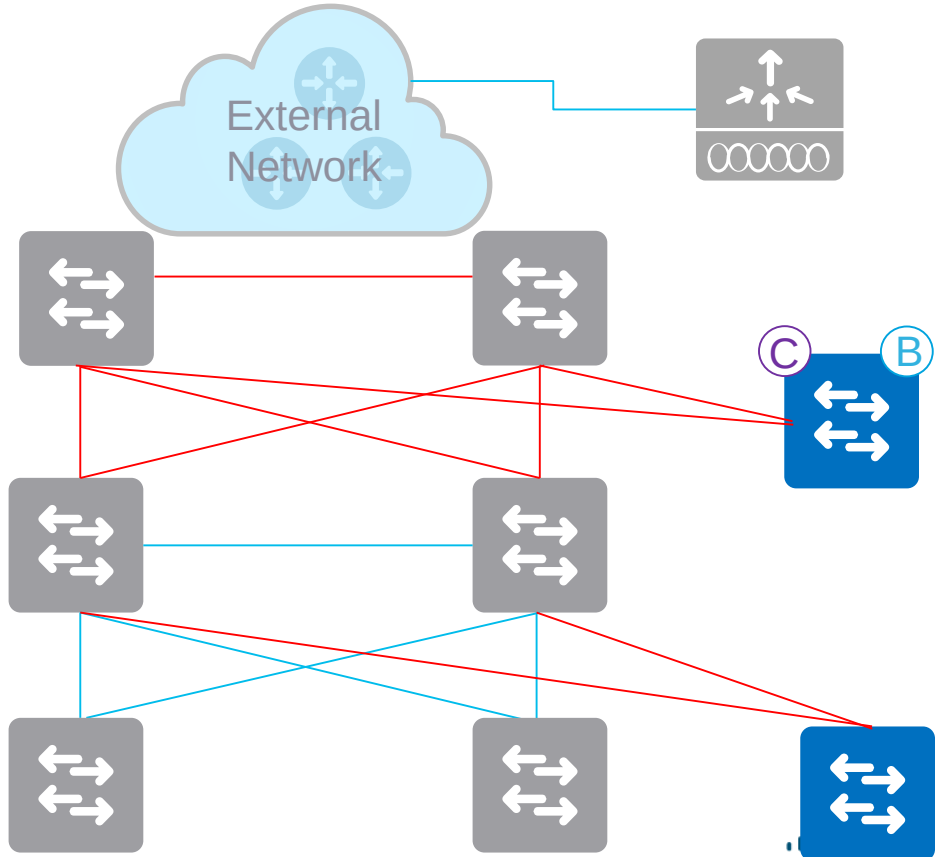

Prepping the Switch

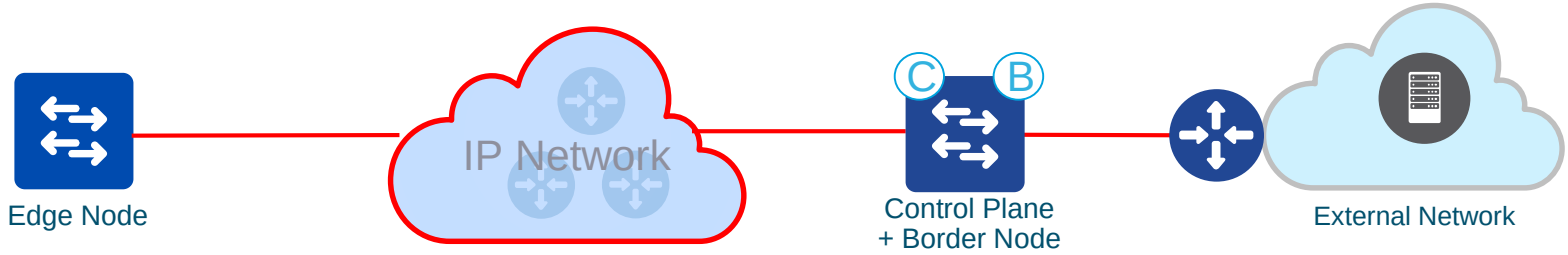

Set following on the Fabric nodes and other nodes in the underlay

- Set MTU to 9100 on the switch and the existing network.
- Configure ‘ip routing’
- Set ‘username’ and ’password’ for device access
- Configure VTY and console lines for device access
- Configure NTP
- Configure SNMP, syslog
- Configure Loopback0 (/32) for RLOC, and underlay IP addresses

Getting Started Steps – ISIS as an IGP

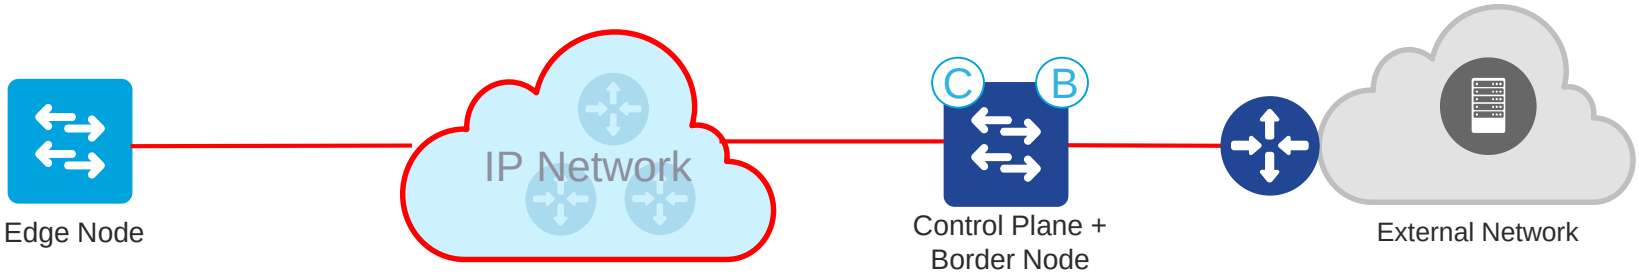


```
router isis
 passive-interface Loopback0
 net 49.0001.XXXX.XXXX.XXXX.00
 is-type level-2-only
 ispf level-2
 log-adjacency-changes
 metric-style wide level-2
 no hello padding
 authentication mode md5 level-2
 authentication key-chain ON
```

```
interface GigabitEthernet x/x
 ip router isis
 isis network point-to-point
 isis metric <metric> level-2
 isis circuit-type level-2-only
 isis authentication mode md5 level-2
 isis authentication key-chain ON
 carrier-delay ms 0
 dampening
```

Getting Started Steps – OSPF as an IGP

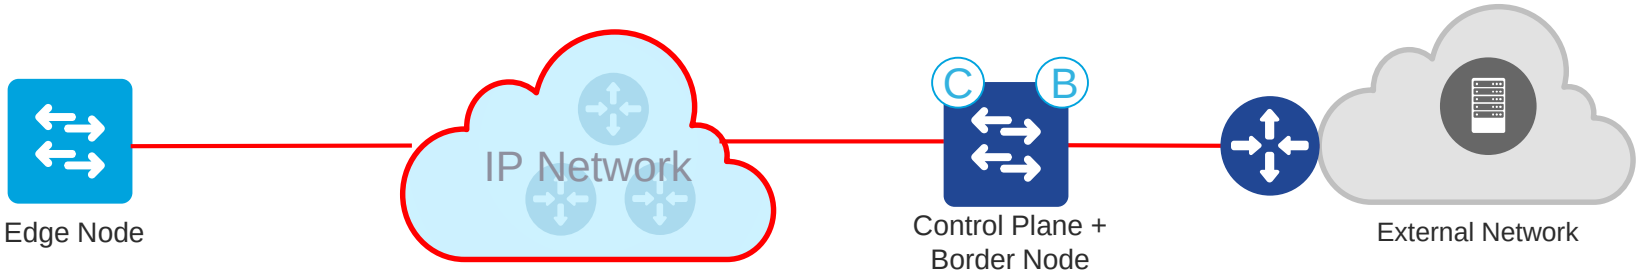


```
interface GigabitEthernet1/1/1
 no switchport
 ip address 192.168.22.58 255.255.255.252
!
interface GigabitEthernet1/1/2
 no switchport
 ip address 192.168.22.38 255.255.255.252
!
interface Loopback0
 ip address 192.168.21.9 255.255.255.255
 ip ospf network point-to-point
```

```
router ospf 1
 router-id 192.168.21.9
 passive-interface default
 no passive-interface GigabitEthernet1/1/1
 no passive-interface GigabitEthernet1/1/2
 network 192.168.21.9 0.0.0.0 area 0
 network 192.168.22.38 0.0.0.0 area 0
 network 192.168.22.58 0.0.0.0 area 0
```

Graphical Migration

- Using Cisco DNA Center

Logging in to Cisco DNA Center

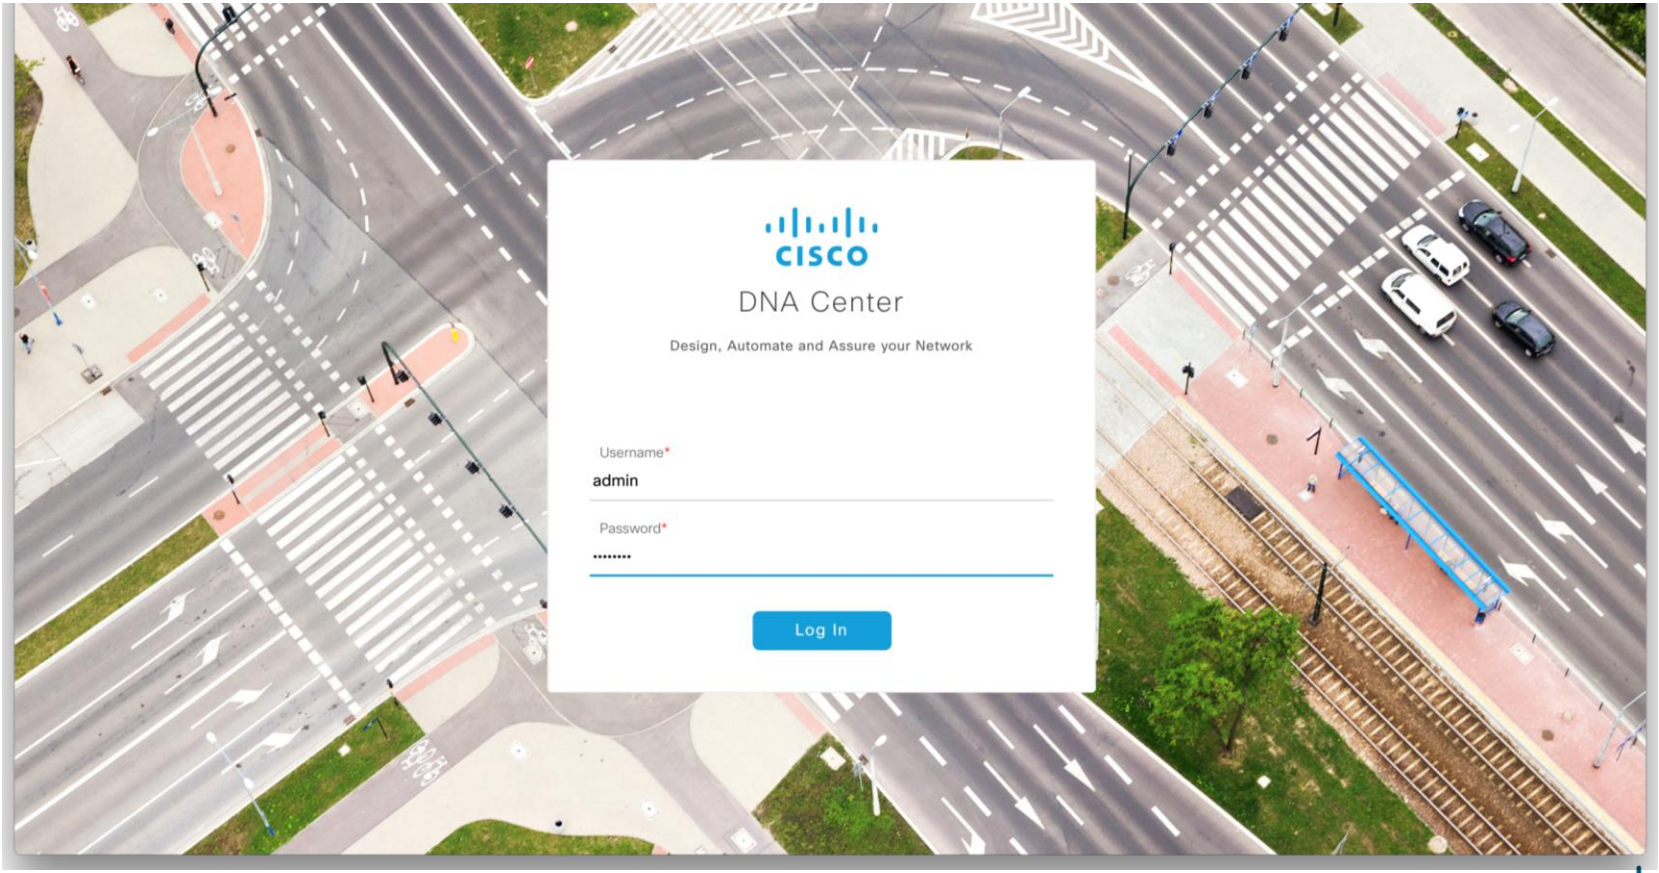

Steps in DNAC GUI

Existing Network Topology

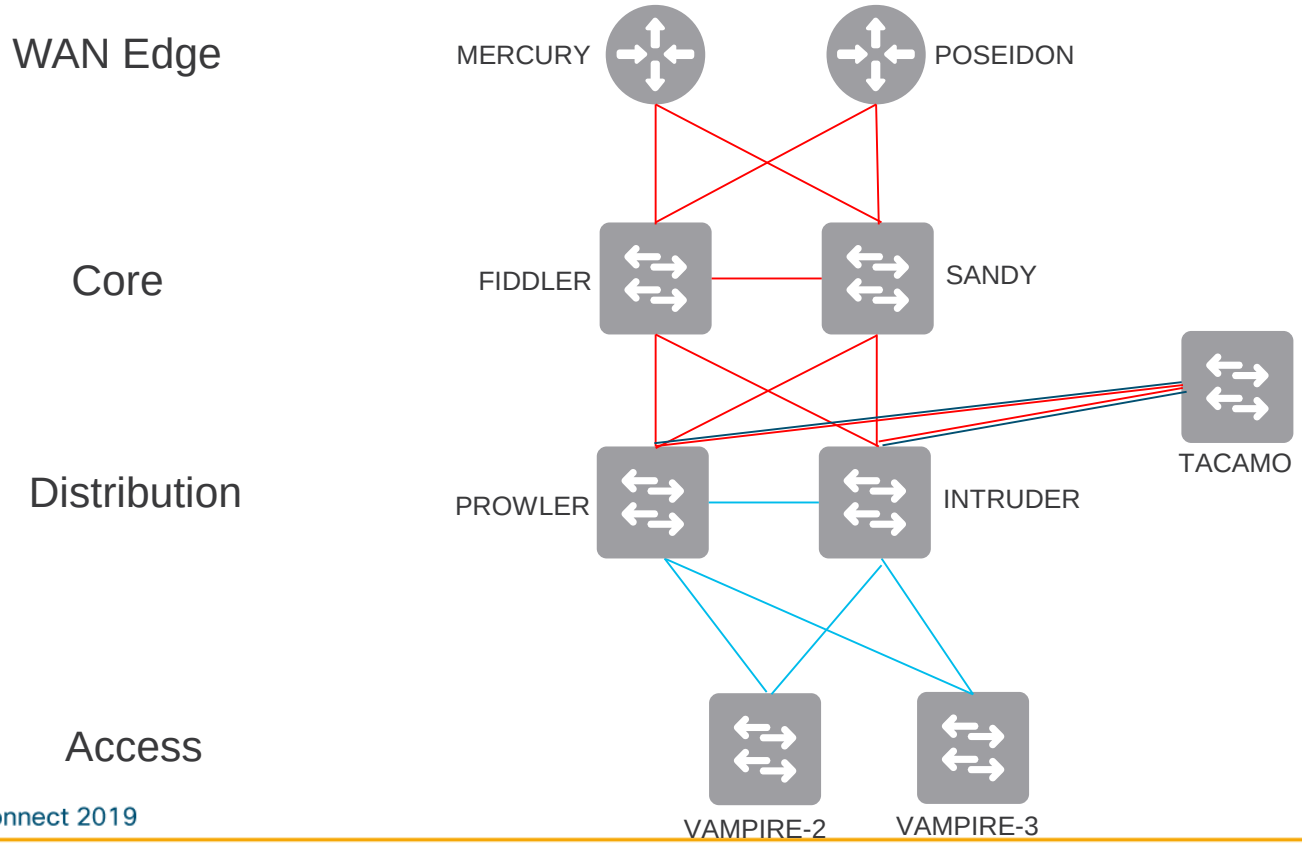

Minimum number of IP Pools

- Border Pool – for leveraging Border automation
- Client Pool – for onboarding clients (Wired)
- AP Pool – If APs are in the network, then create an IP Pool for APs
- Wireless endpoints – If you want different addresses for wireless endpoints

Steps in DNAC GUI

Provision only the potential Fabric Nodes

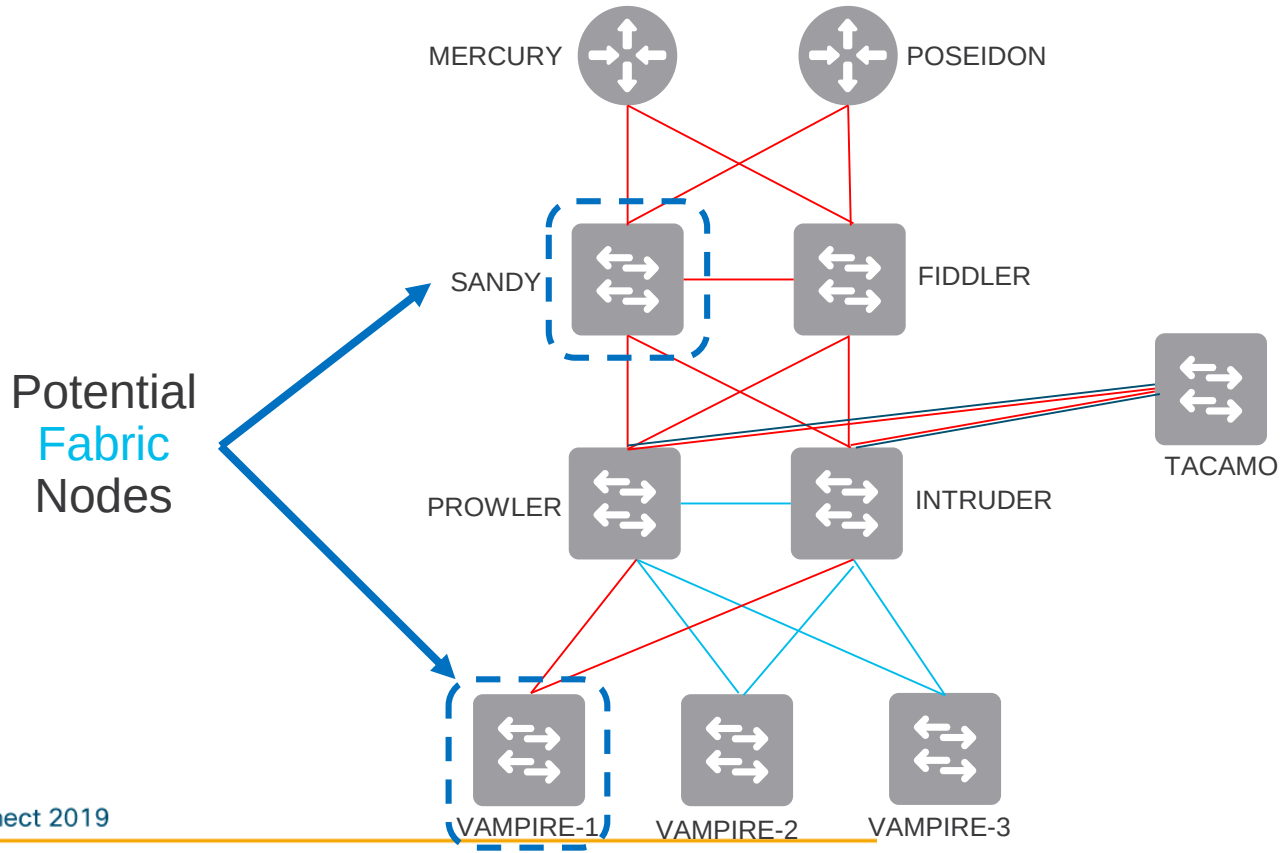

Steps in DNAC GUI

East-West Communication: Fabric Border is Exchange Point with Fusion Router

Un-encapsulated packet

VXLAN encapsulated packet

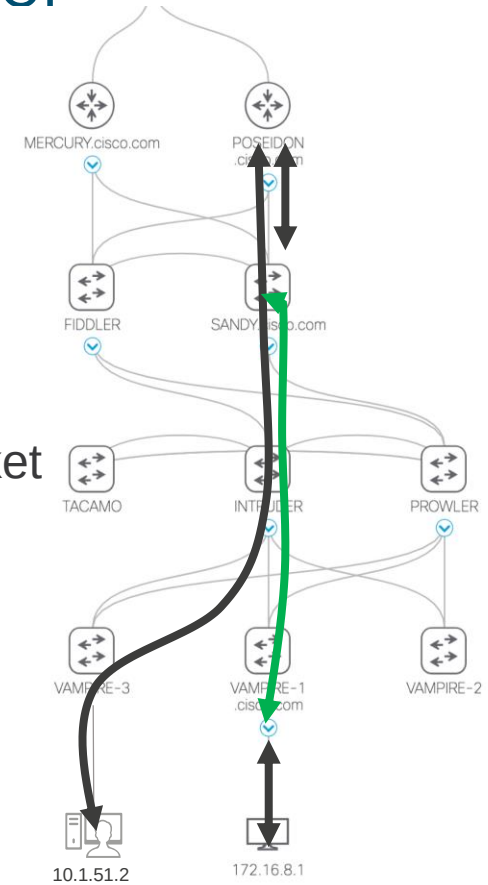

Current State of the Network

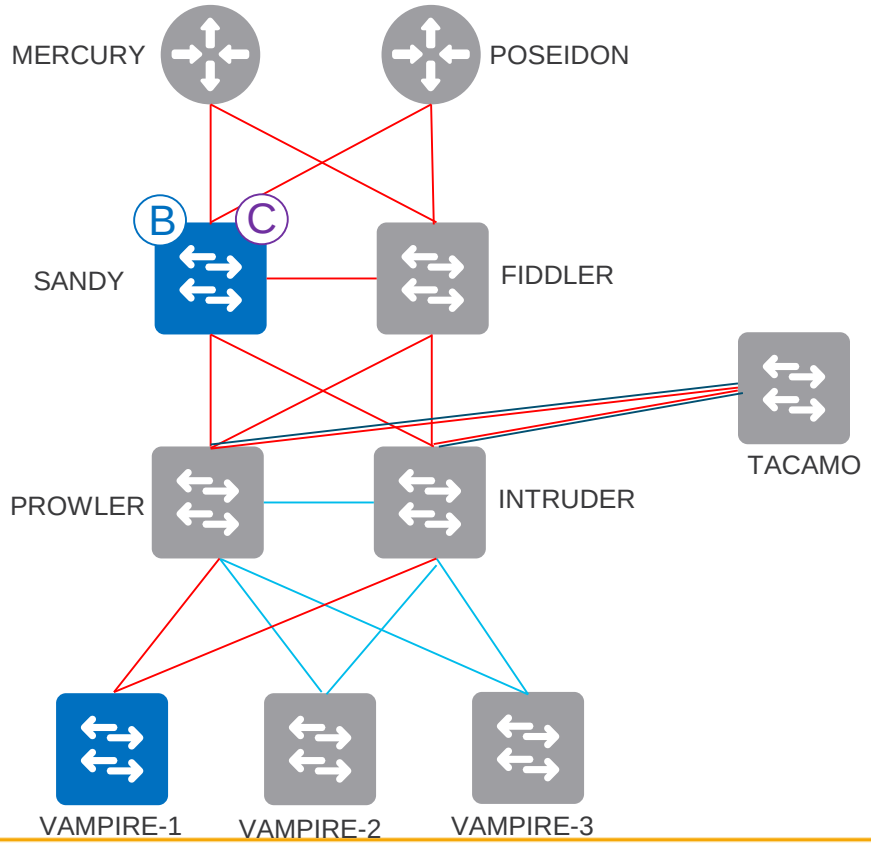

Re-configure Links: L2 to L3 Routed Links

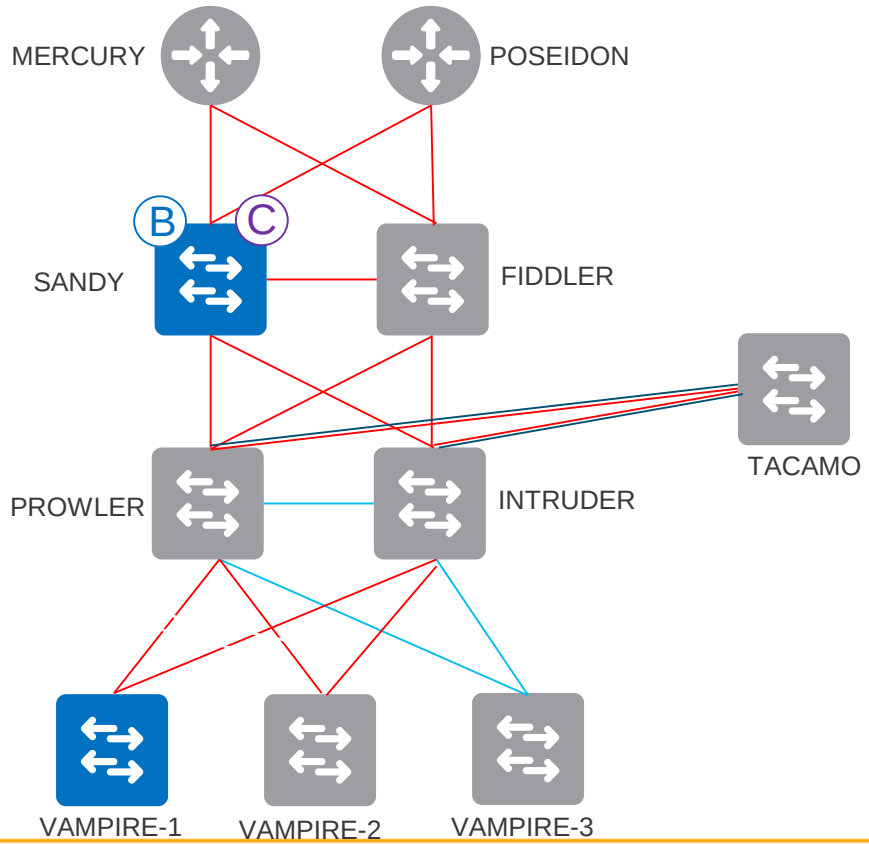

Configure Fabric Edge on Access Switch

Configure Fabric Edge

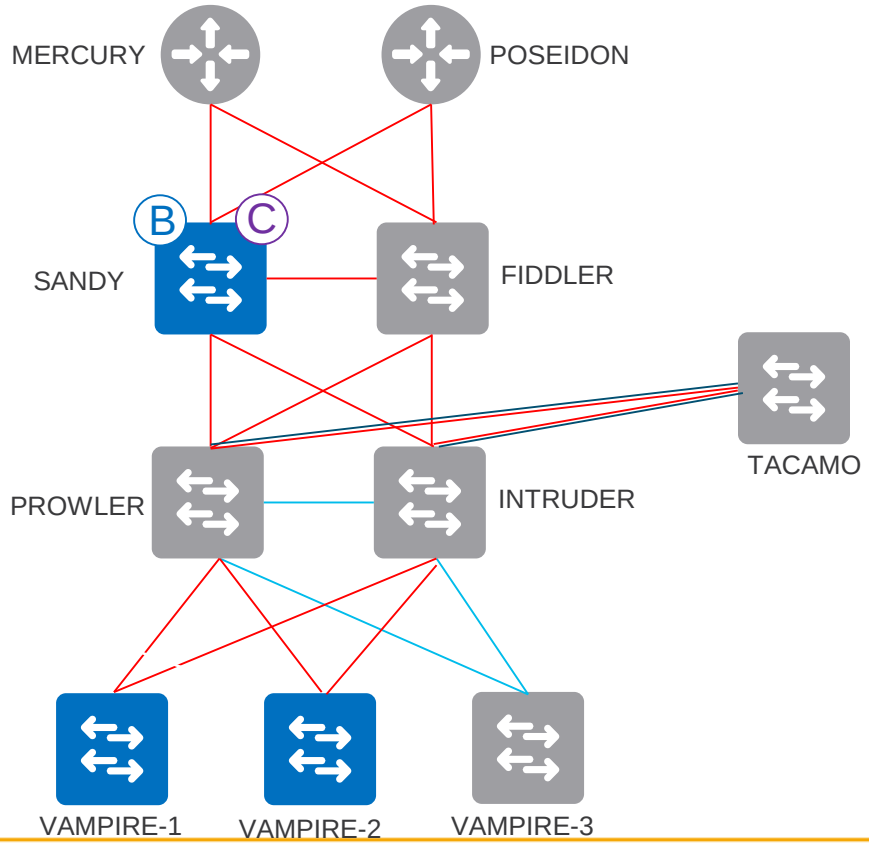

Redundant Fabric Border/Control Plane node

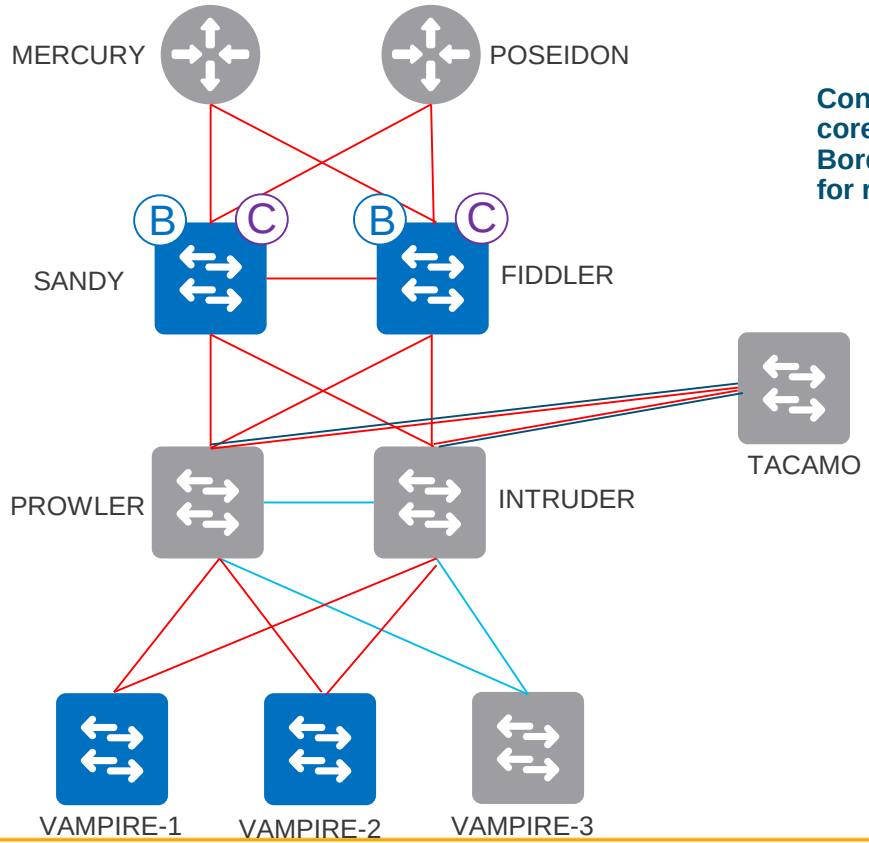

Configure second core as Fabric Border/Control Plane for redundancy

Reconfigure Links: L2 to L3 Routed links

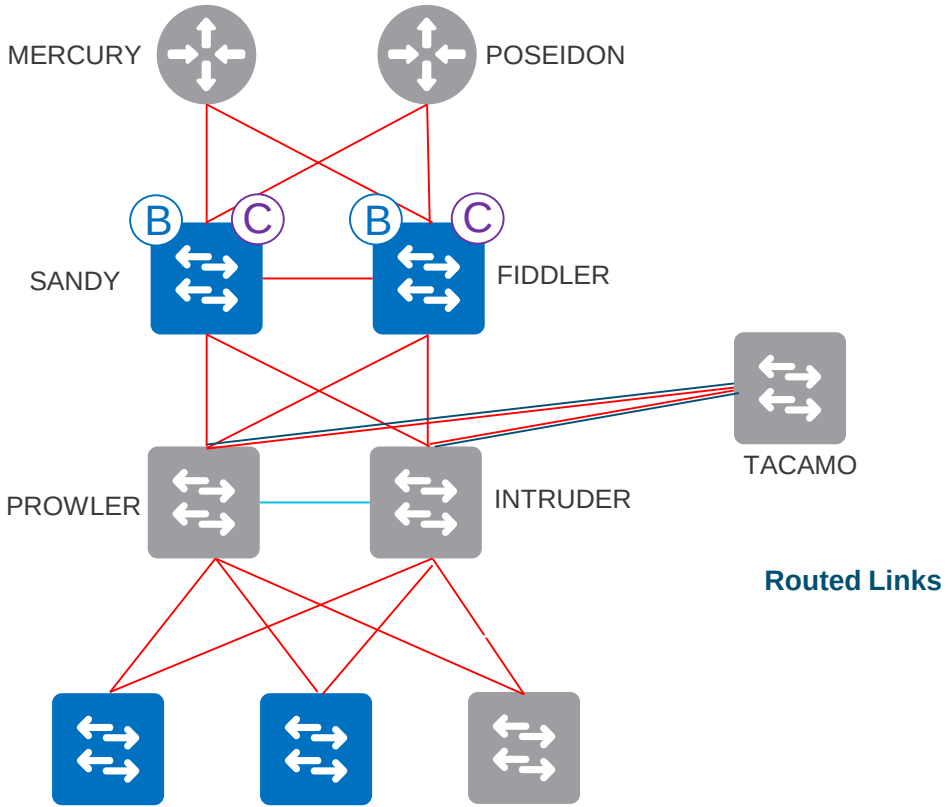

Configure Fabric Edge on Access

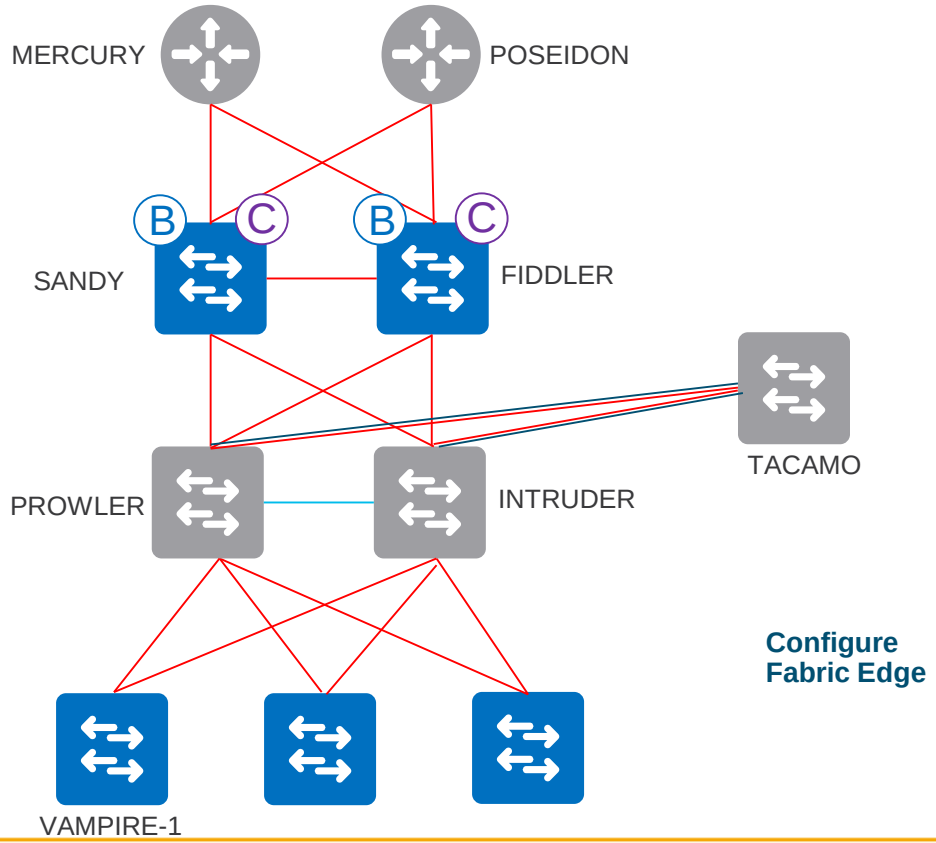

Configure Fabric Edge

Distribute Control Plane node for Scale

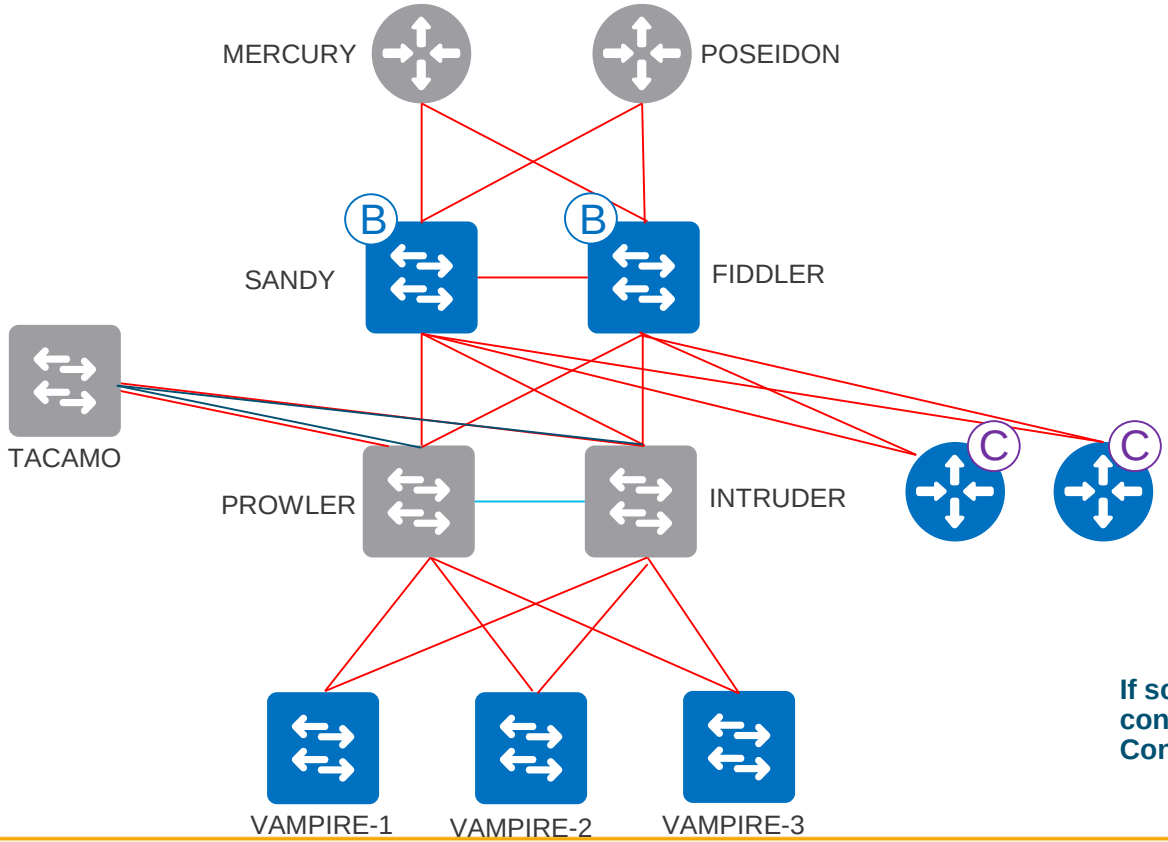

If scale demands configure dedicated Control Plane nodes

Migration with existing subnets

Incremental Migration – High Level concept

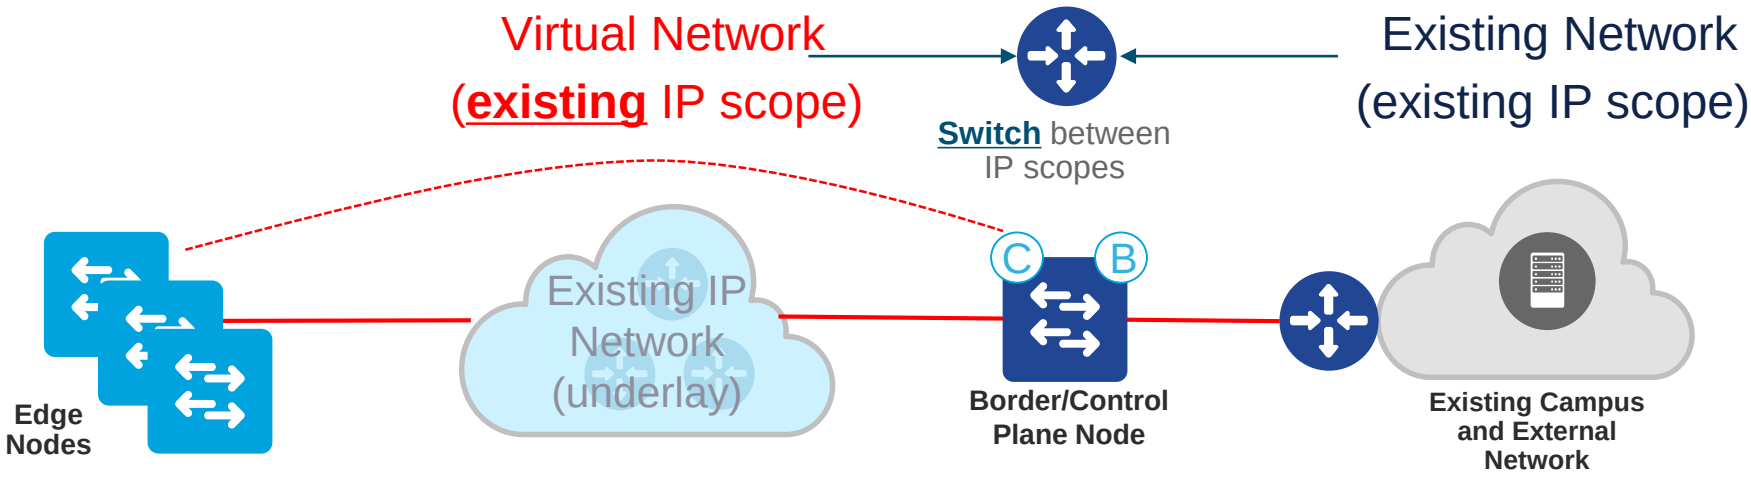

- Deploy a Border node and incrementally add Edge Nodes
- A virtual network is formed over the existing (underlay) network
- The virtual network(s) uses same subnet address as existing network
- The virtual network connects to the external network through the border

Migrating to Cisco SD-Access retaining existing subnets

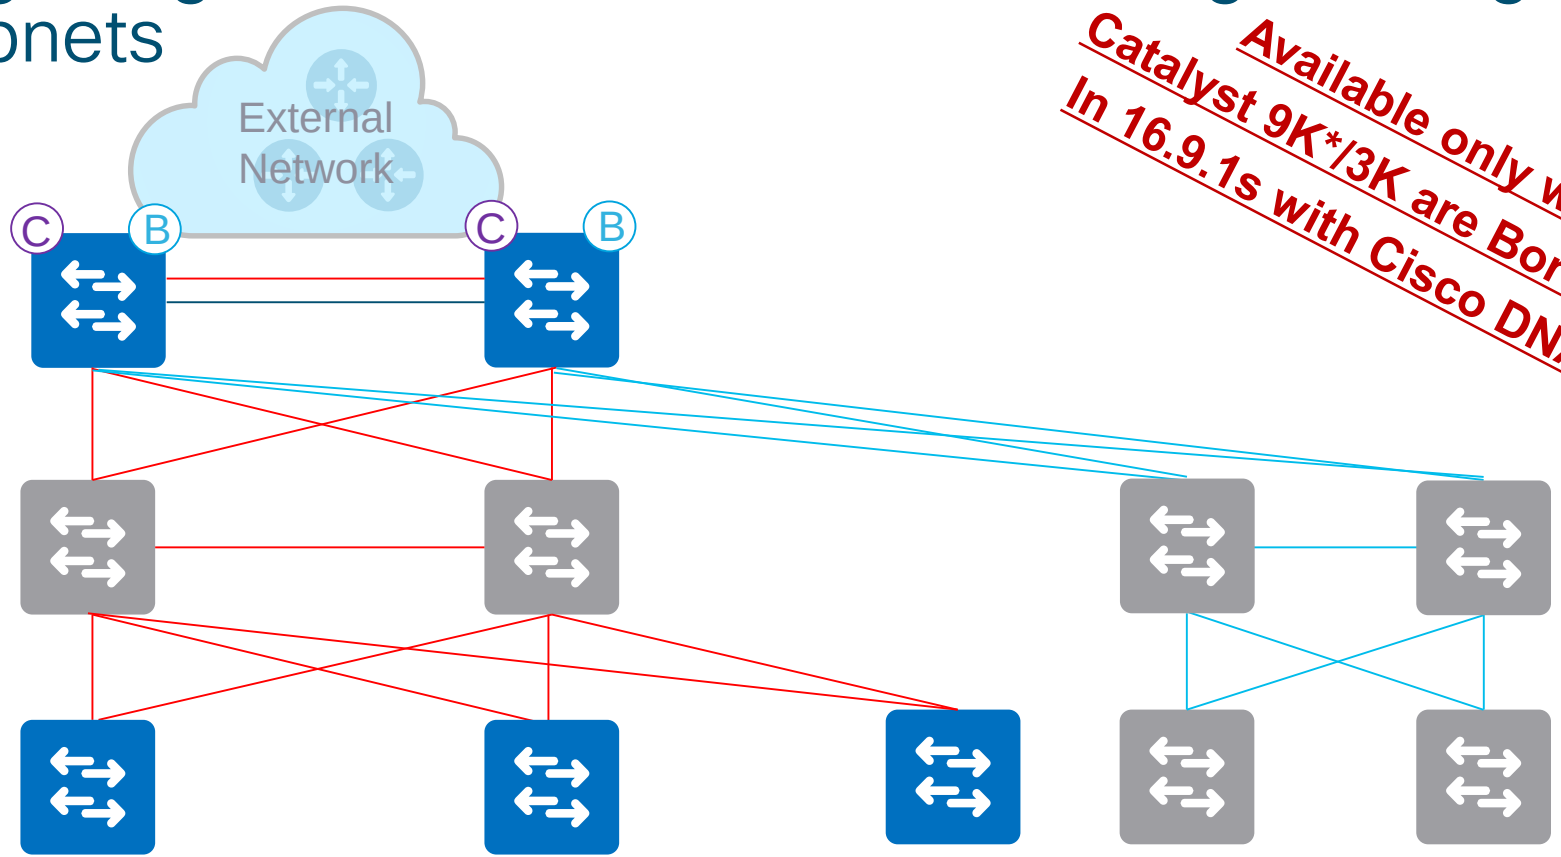

Available only where Catalyst 9K*/3K are Border Nodes In 16.9.1s with Cisco DNAC 1.2.5

10.1.1.0/24 VLAN 1021 L2 VNI = 8188

10.1.1.0/24 VLAN10

Migrating to Cisco SD-Access retaining existing subnets

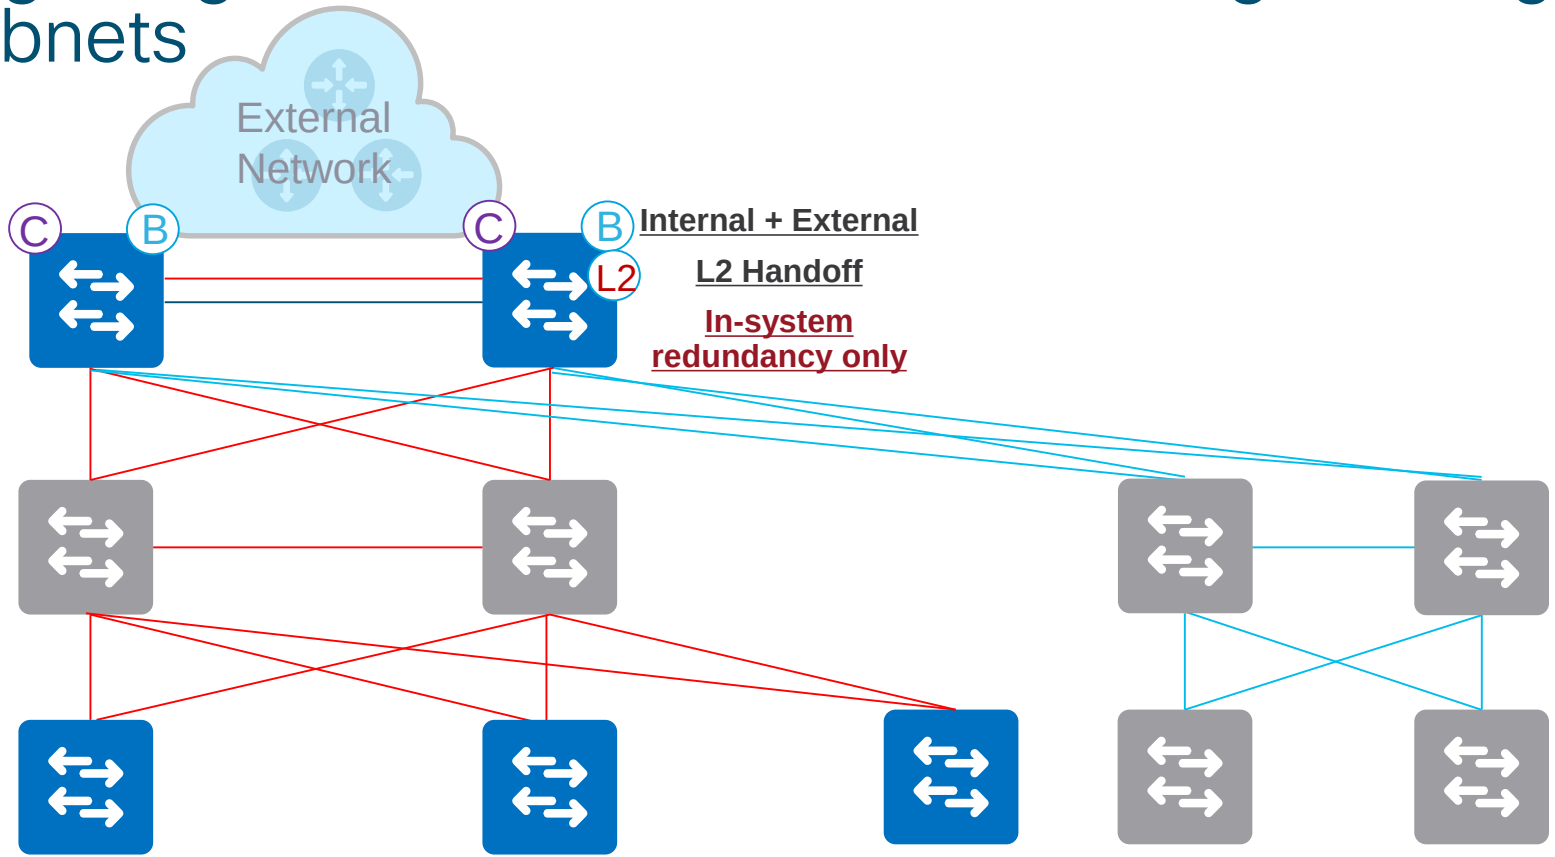

10.1.1.0/24 VLAN 1021 L2 VNI = 8188

Migrating to SD-Access retaining existing subnets

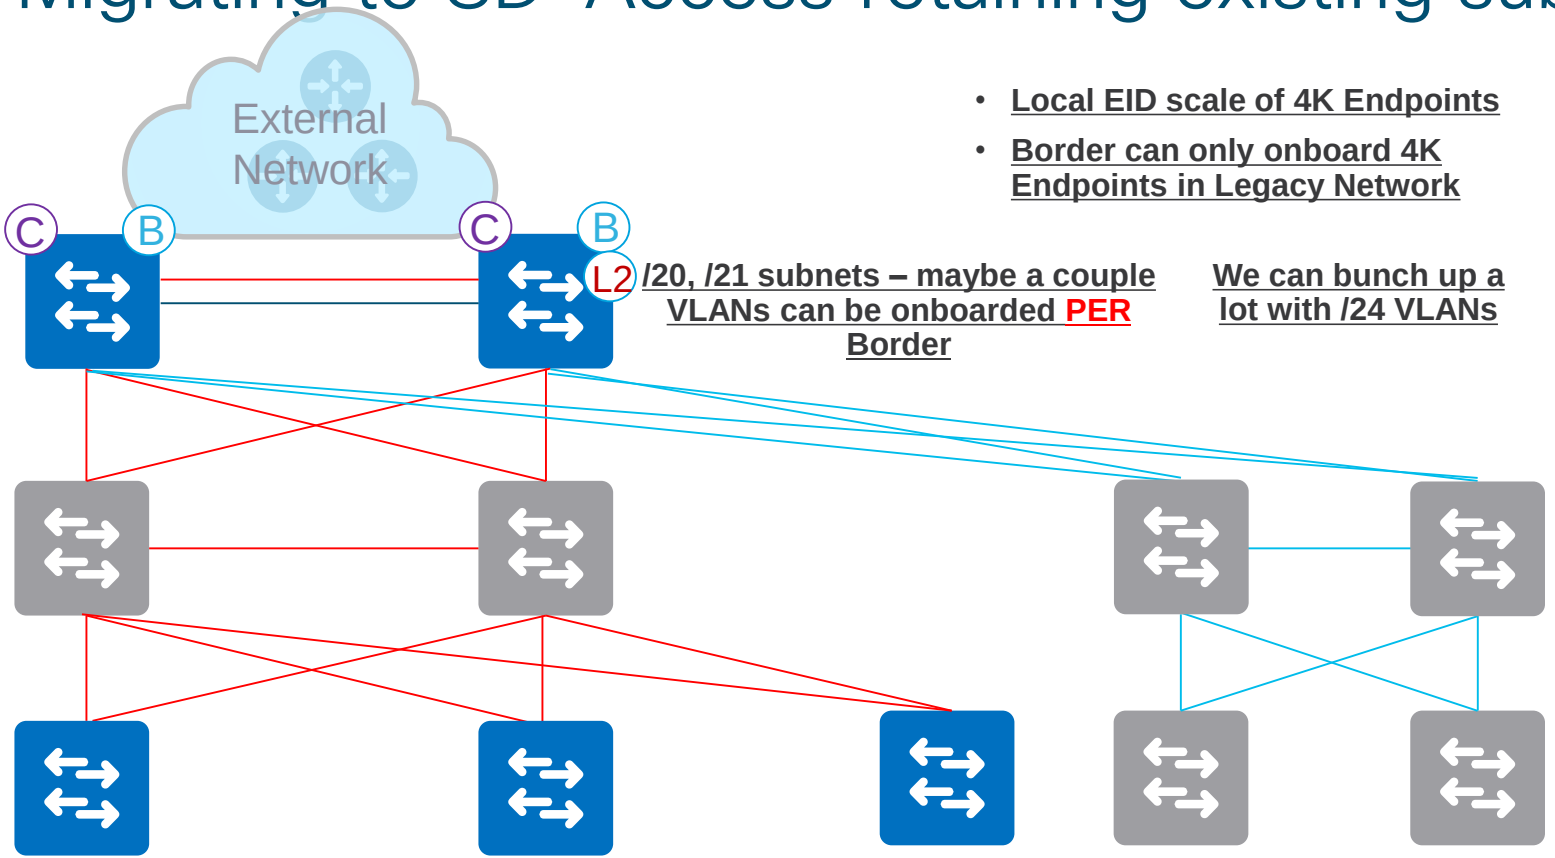

- **Local EID scale of 4K Endpoints**
- **Border can only onboard 4K Endpoints in Legacy Network**

/20, /21 subnets – maybe a couple VLANs can be onboarded PER Border

We can bunch up a lot with /24 VLANs

10.1.1.0/24

VLAN10

10.1.1.0/24 VLAN 1021 L2 VNI = 8188

Separate L2 Border recommendation

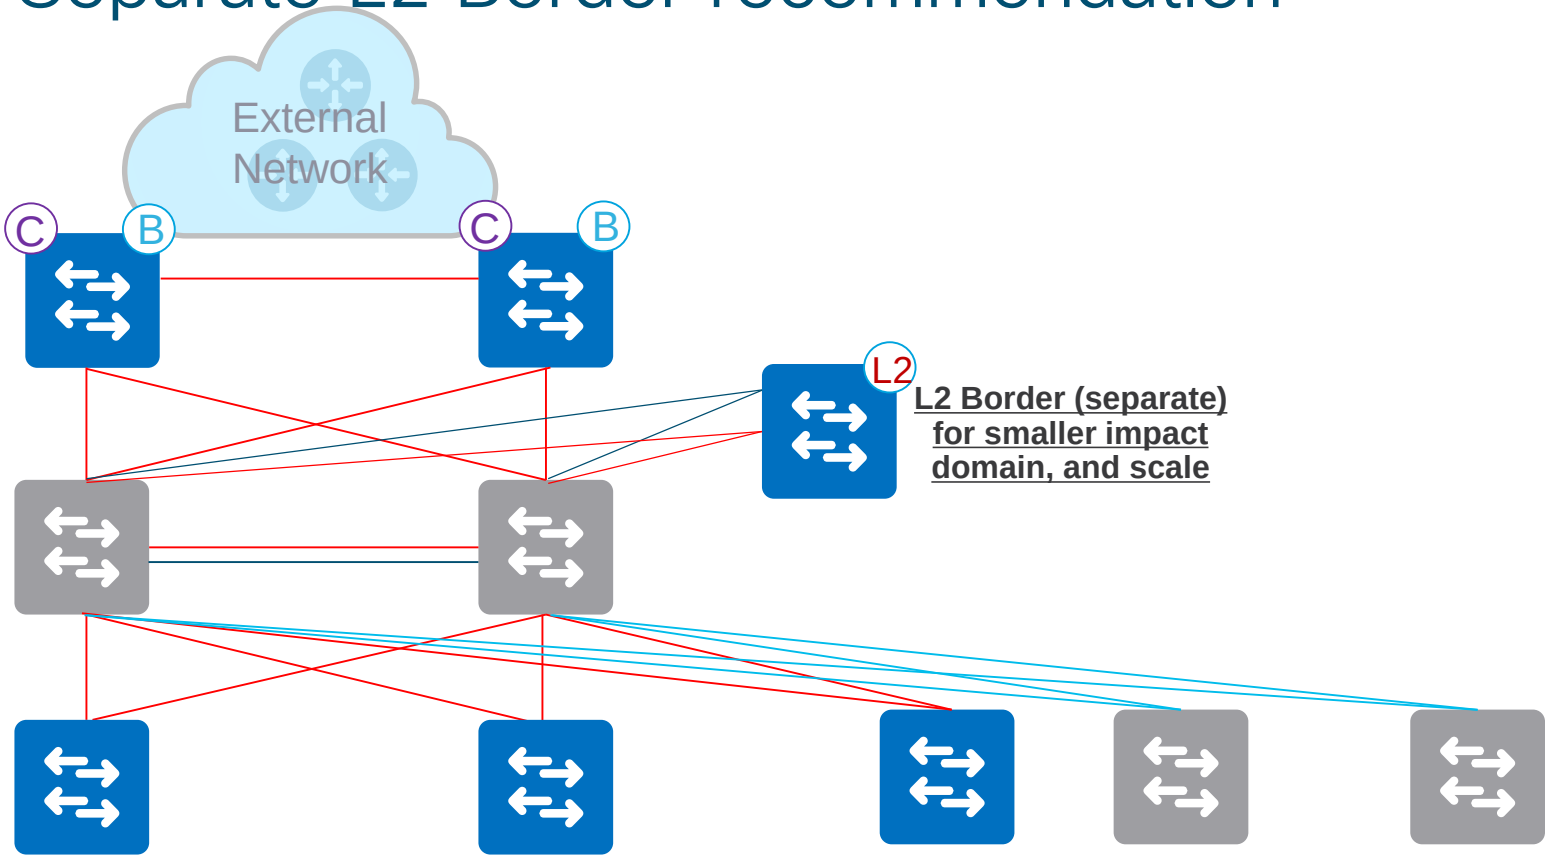

10.1.1.0/24 VLAN 1021 L2 VNI = 8188

10.1.1.0/24 VLAN10

Separate L2 Border recommendation

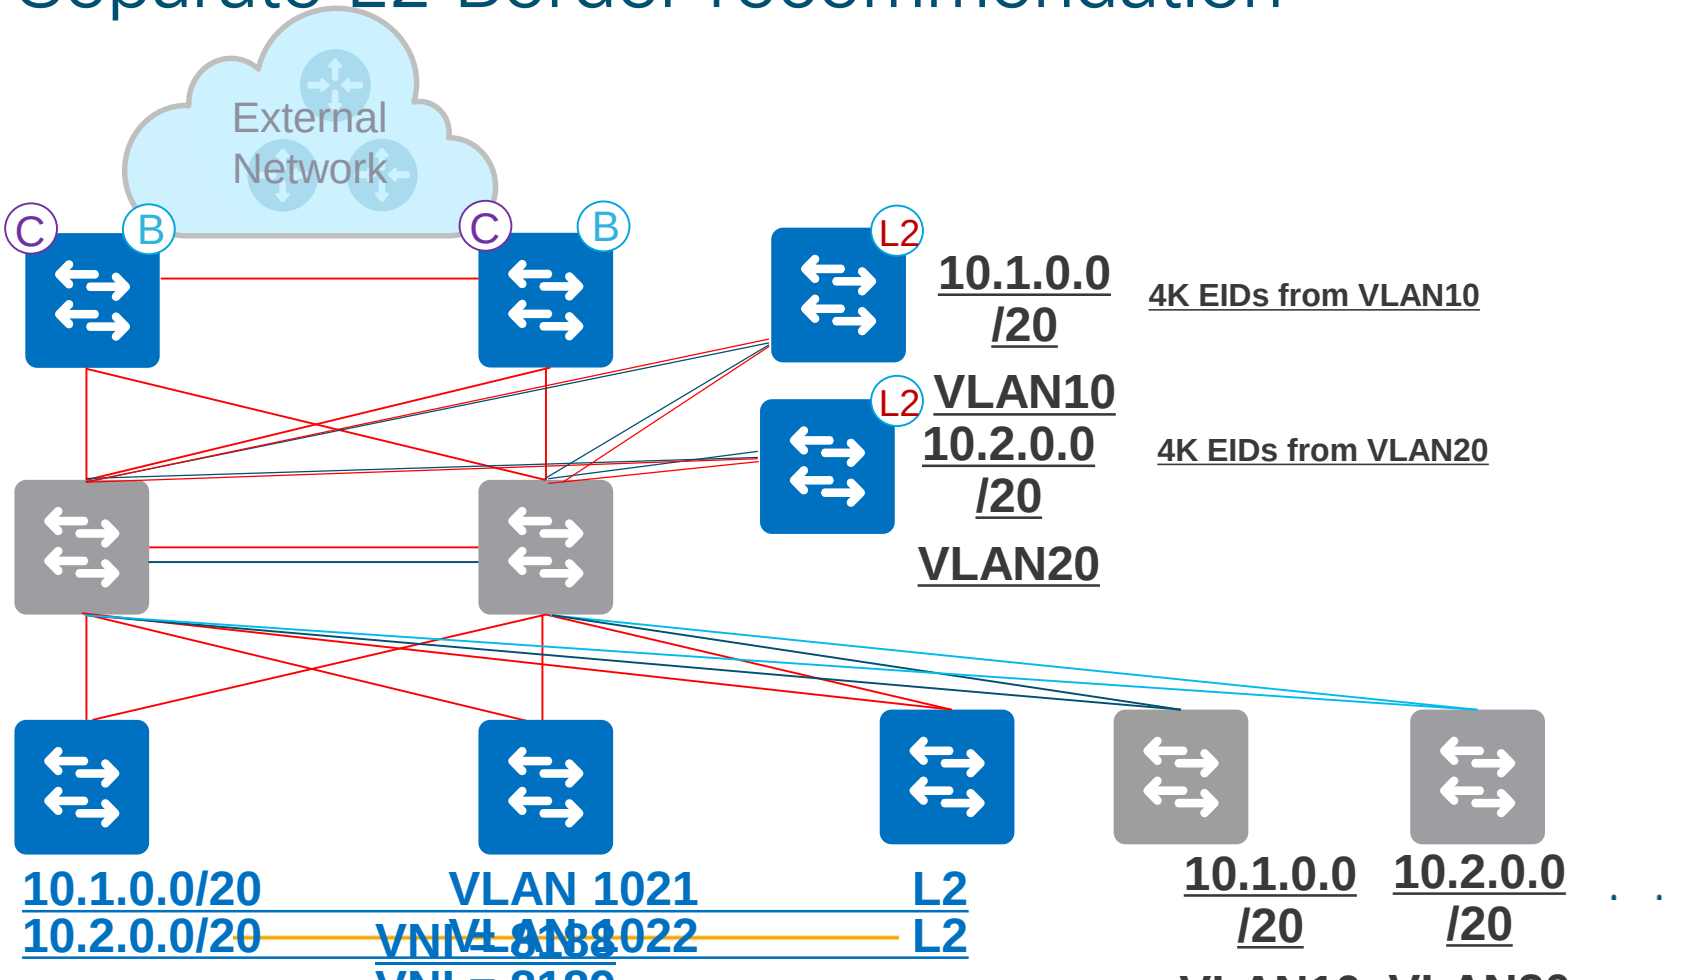

Flash-cut Existing Access to Fabric Edge

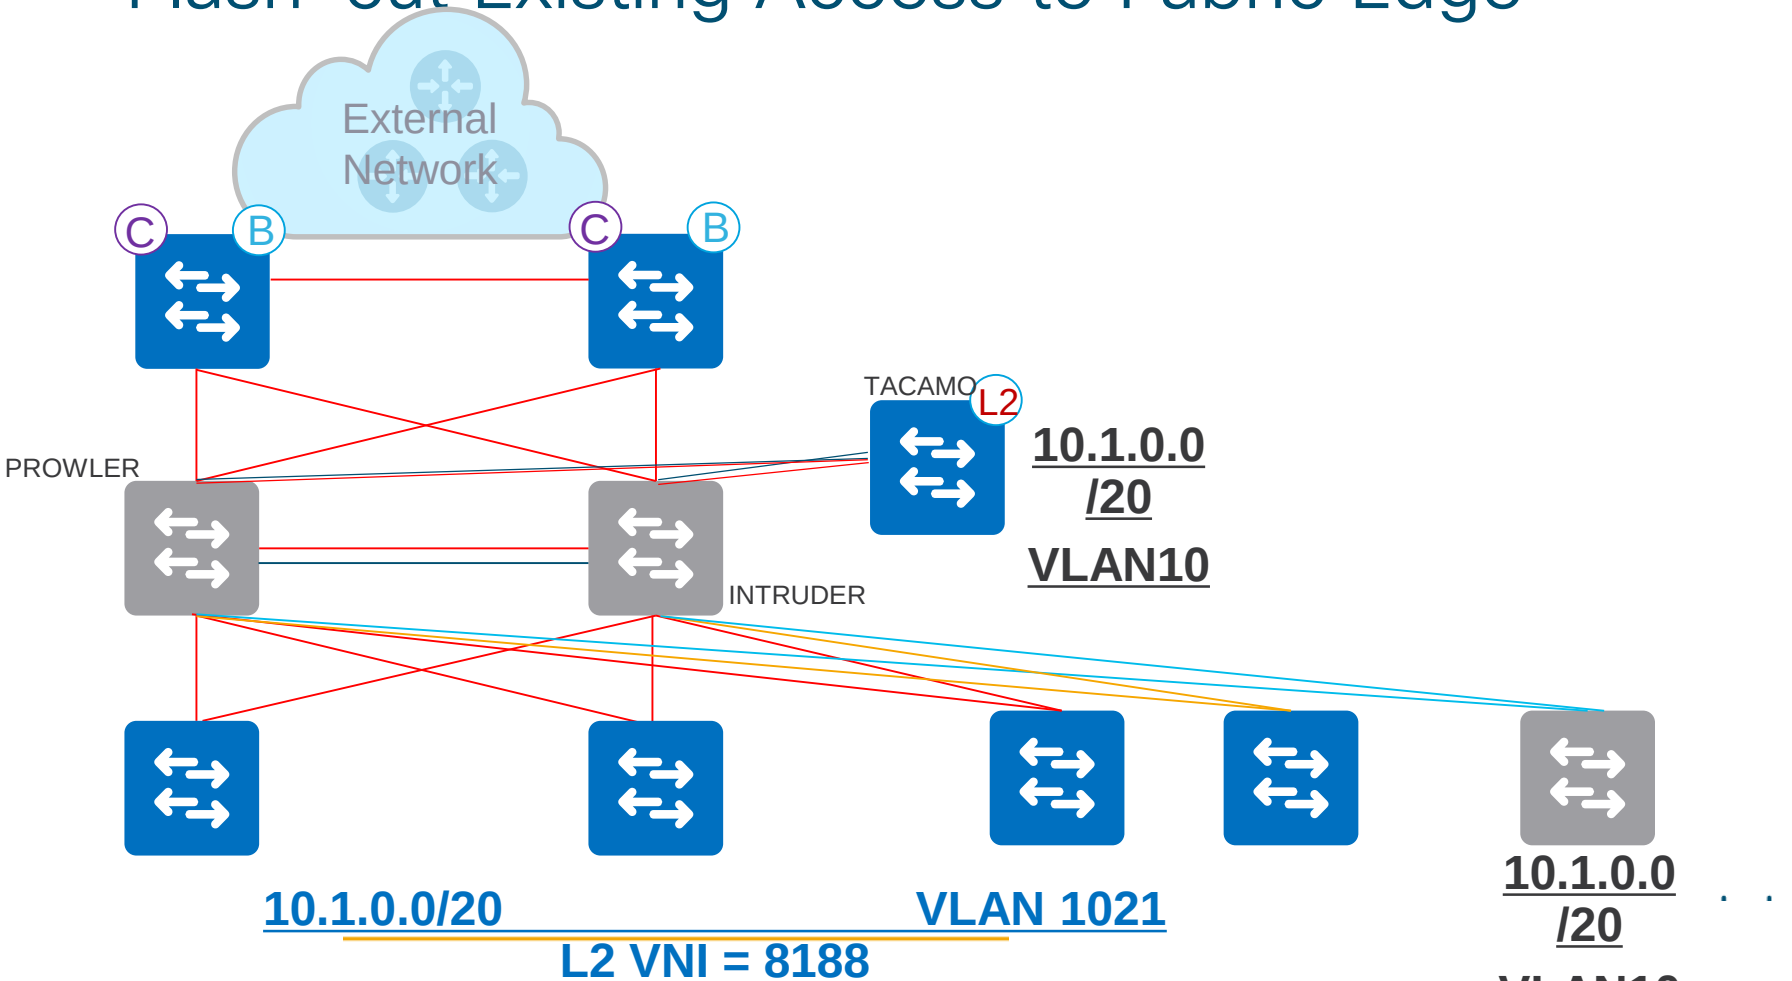

Steps in DNAC GUI

This has to be simultaneous from here on.....

SINGLE PHASE

FROM NOW ON HERE

Steps in DNAC GUI

East-West: Hosts in same subnet, inside and outside fabric

Bridged packet

Un-encapsulated packet

VXLAN encapsulated packet

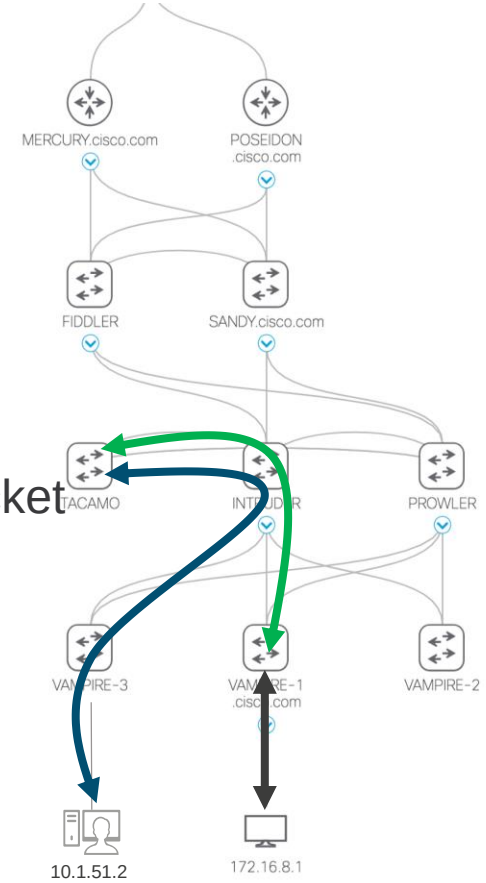

Migration with routed access

Migrating Routed Access to Cisco SD-Access

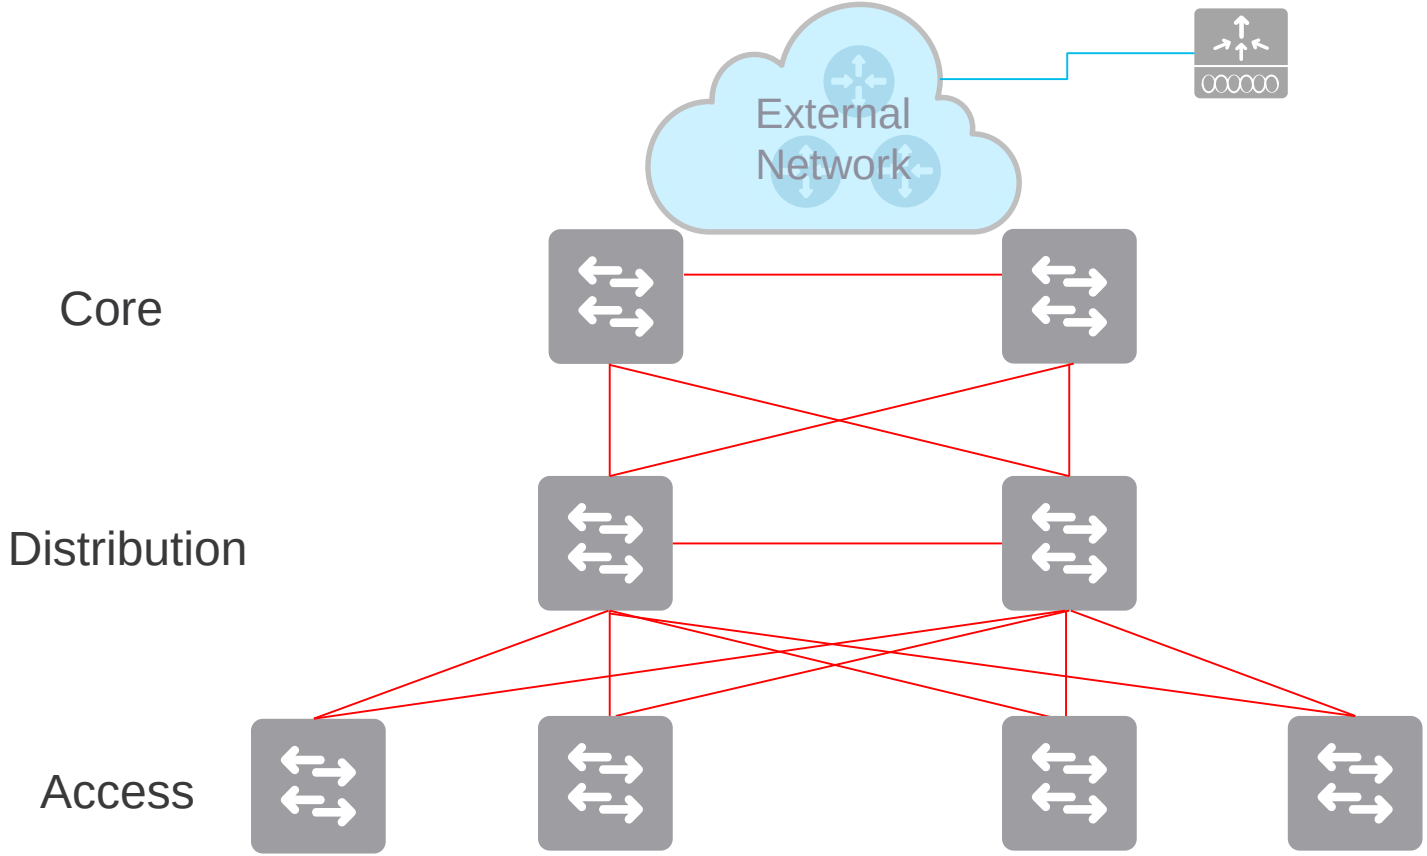

Routed Access Design Considerations

- Can re-use the existing subnets to migrate into Cisco SD-Access
- No changes to existing DHCP scope and subnet size
- No changes to existing firewall or other policies that are based on IP-ACL
- Old network design is retained for familiarity
- Cannot realize the advantages of bigger subnets, but lesser subnets that are optimized for Cisco SD-Access

Routed Access Migration to Cisco SD-Access

- Shutdown existing SVI (Vlan10 in this case)
- Provision existing subnet from Cisco DNA-Center (10.1.1.0/24 in this case)
- Cisco DNA-Center will provision Vlan1021 with 10.1.1.0/24
- Move hosts to fabric-enabled IP Pool
- Verify connectivity

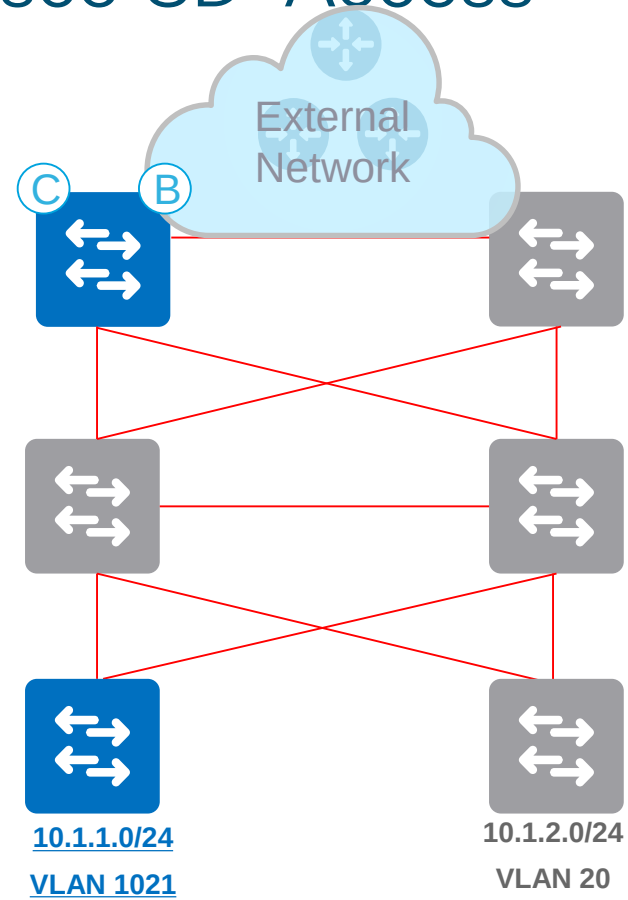

Routed Access Migration to Cisco SD-Access

- Repeat the process for other VLANs on the Fabric Edge
- Repeat the same process on other access switches in converting them to Fabric Edge
- Migration is One-Switch—At-A-Time – **NOT** – One-Vlan-At-A-Time

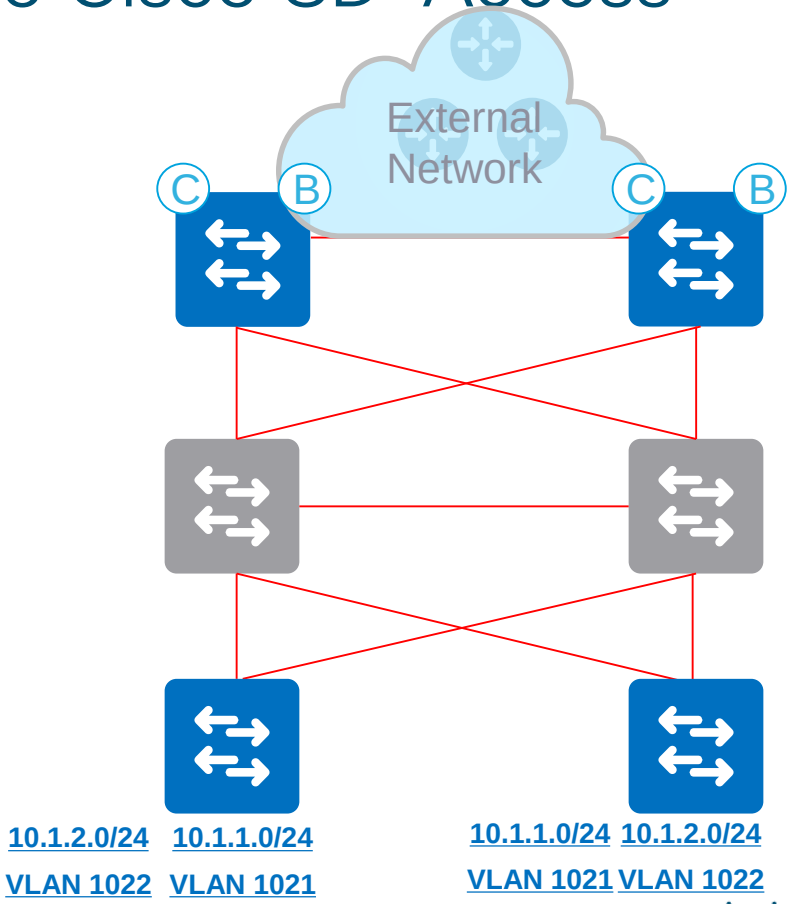

Migrating Wireless

Cisco SD-Access Wireless Adoption

- Greenfield Building

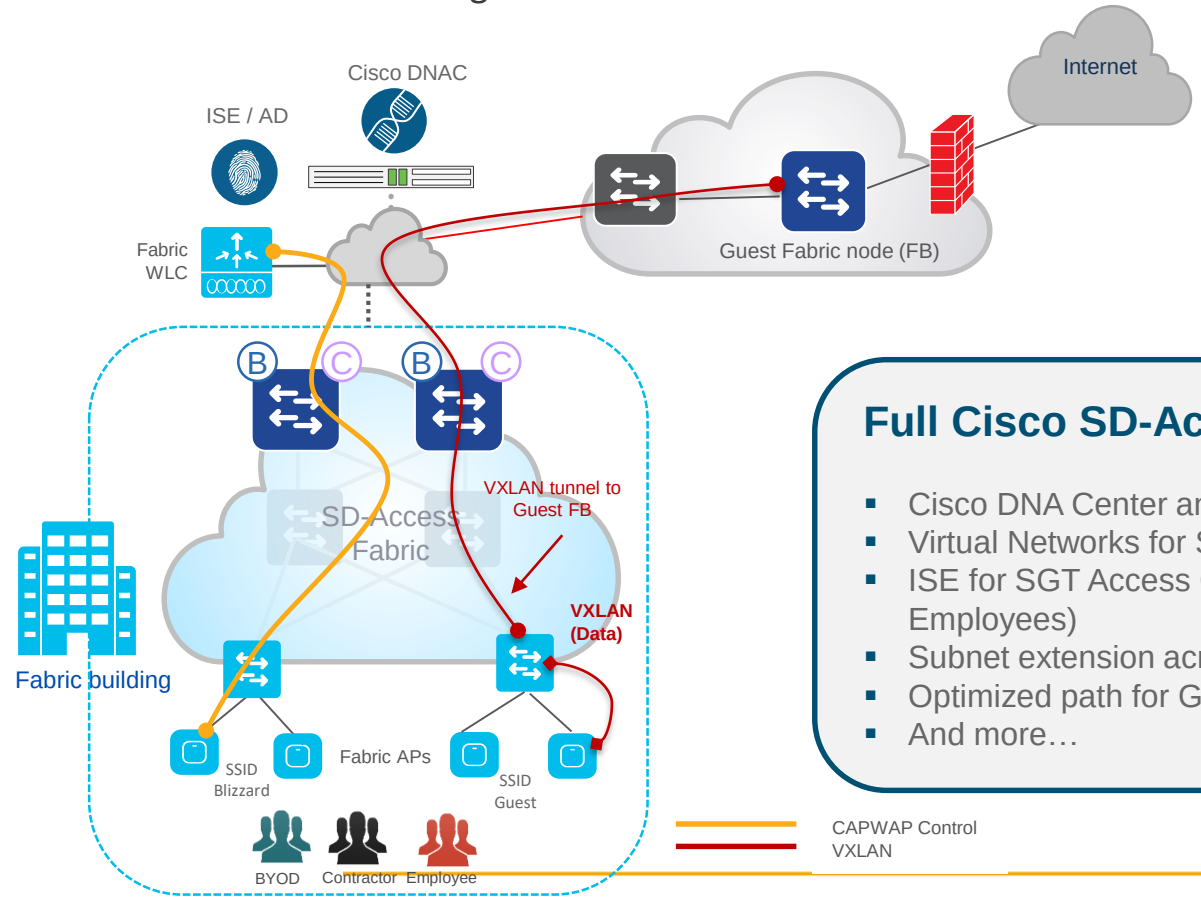

Full Cisco SD-Access Wireless value

- Cisco DNA Center and NDP for Automation & Assurance
- Virtual Networks for Segmentation (ex Employee, IoT, Guest)
- ISE for SGT Access Control within VRF (ex. Contractor, BYOD, Employees)
- Subnet extension across Campus with distributed data plane
- Optimized path for Guest and no Anchor WLC
- And more…

Migrating to Cisco SD-Access Wireless from CUWN

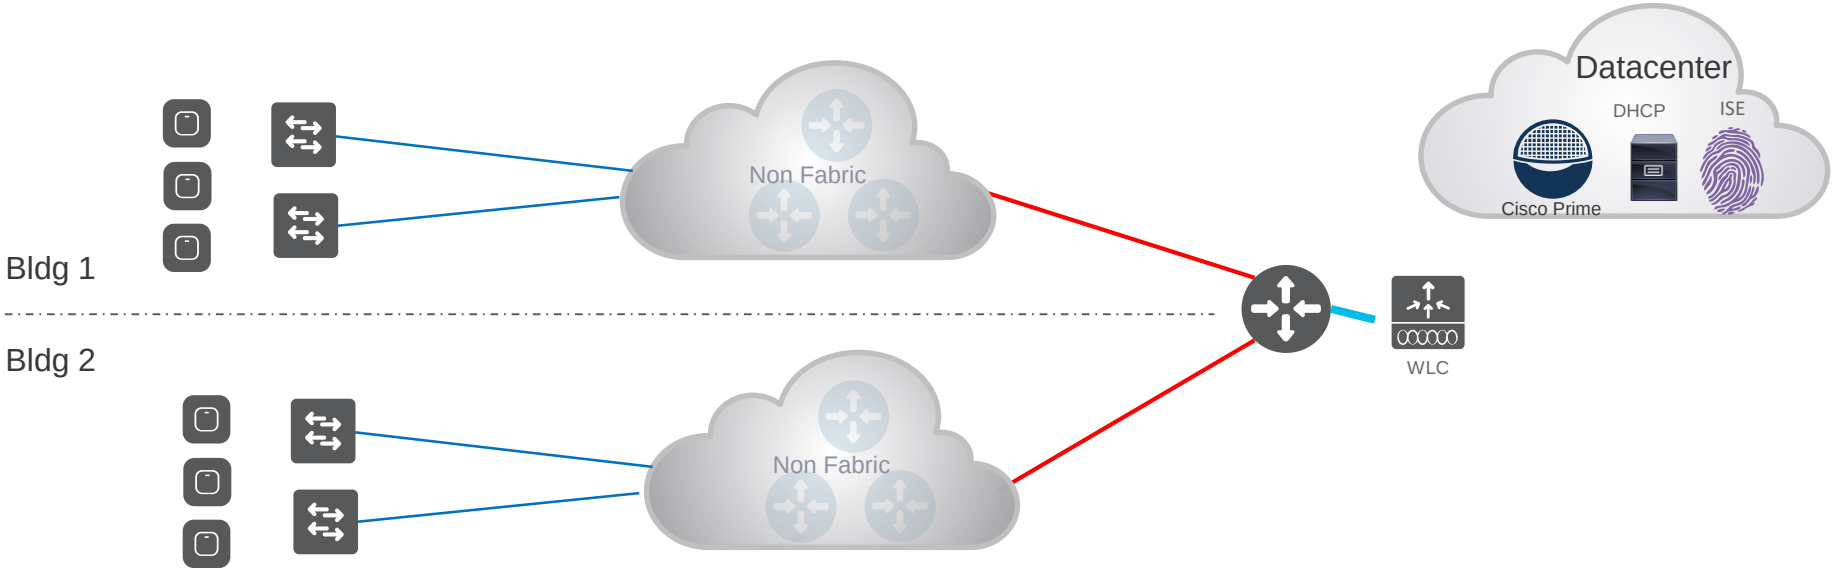

- Customer has a site with AireOS Centralized wireless
- Assumptions:
  - Migration to Fabric happens in a single area (e.g. building) at the time and **migration is in one shot**
  - **No need for seamless roaming** between new SDA area and the existing wireless deployment

Cisco SD-Access Wireless Adoption

- Migration for an existing CUWN deployment

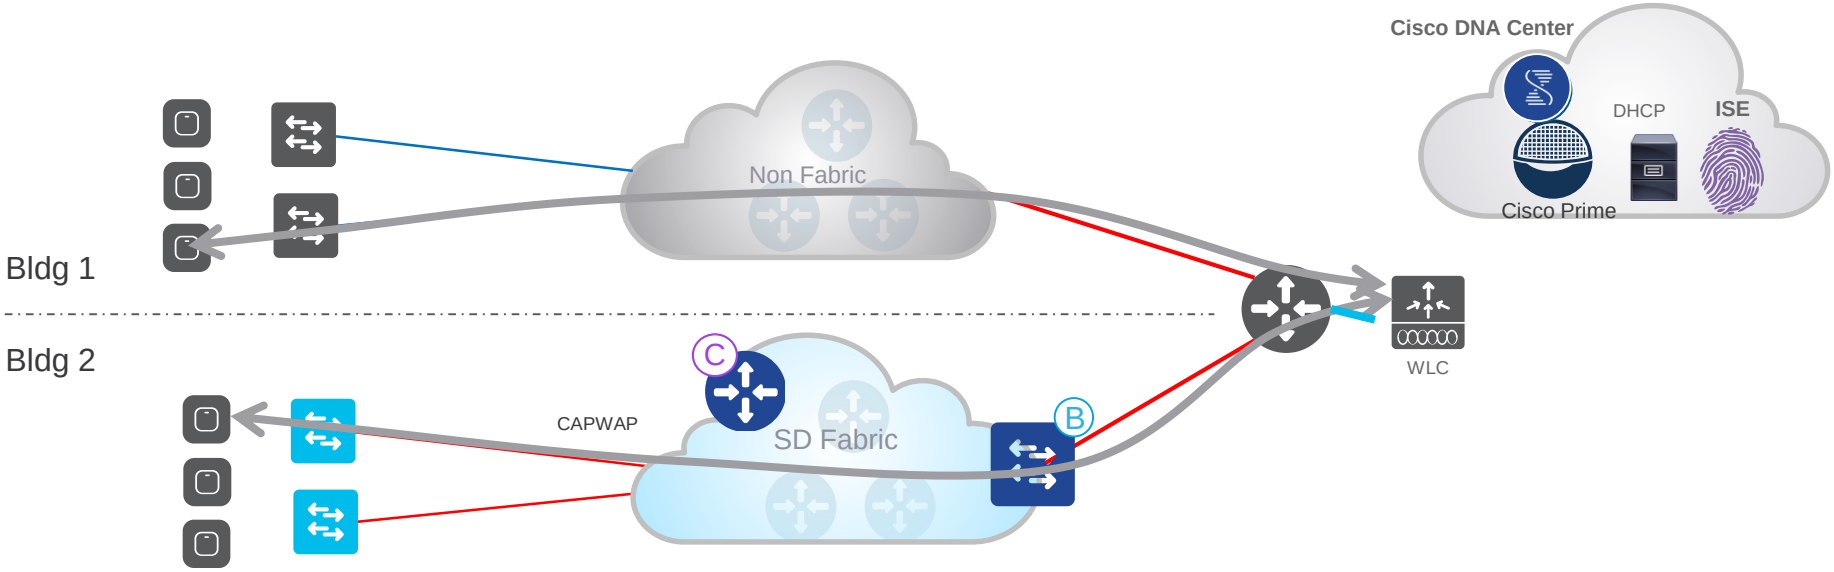

1. Add Cisco DNAC and ISE (if not present already)
2. Migrate wired network to Fabric first
3. Wireless is over the top

CAPWAP

Cisco SD-Access Wireless Adoption

- Migration for an existing CUWN deployment

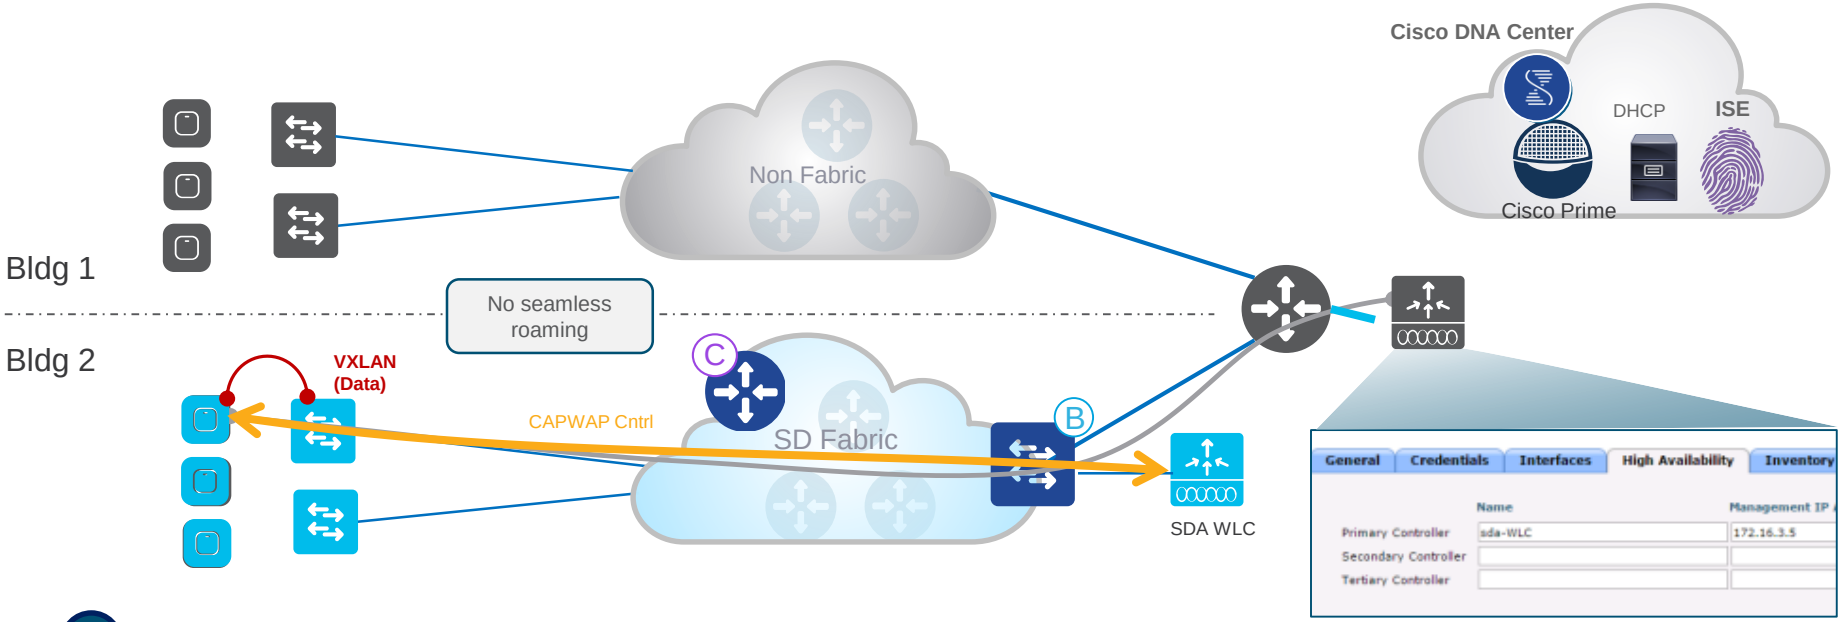

1. Add a dedicated WLC for Cisco SD-Access and configure it with same SSIDs
2. on CUWN WLC, configure the APs in the area to join the new Fabric WLC
3. Traffic now goes through the Fabric

CAPWAP Control
VXLAN

Cisco SD-Access Wireless Adoption

- Migration for an Existing CUWN deployment

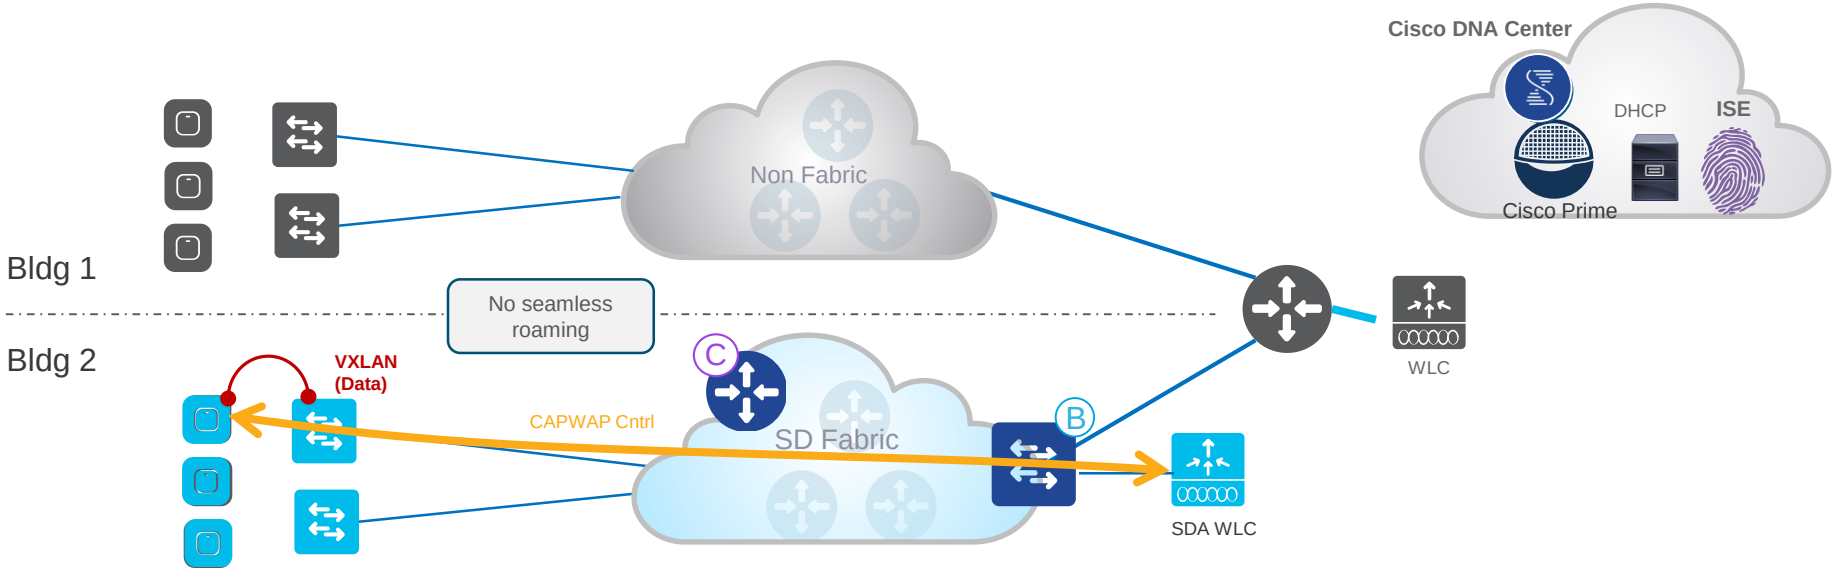

Recommendations

- Prime for CUWN areas, Cisco DNAC for SDA areas
- Dedicated WLC for Cisco SD-Access Wireless
- Same SSIDs on Fabric and non-Fabric
- Same RF Groups for CUWN WLC and SDA WLC
- WLCs in different Mobility Group (no seamless roaming between areas)

С чего начать?

Ways to build a PoC

- Start in a lab – isolated, controlled environment
- How do I connect a lab to the production network if I want to validate use cases
- Enable archive logging on the lab device to see commands being pushed by the Cisco DNA Center.

Connect PoC to Production

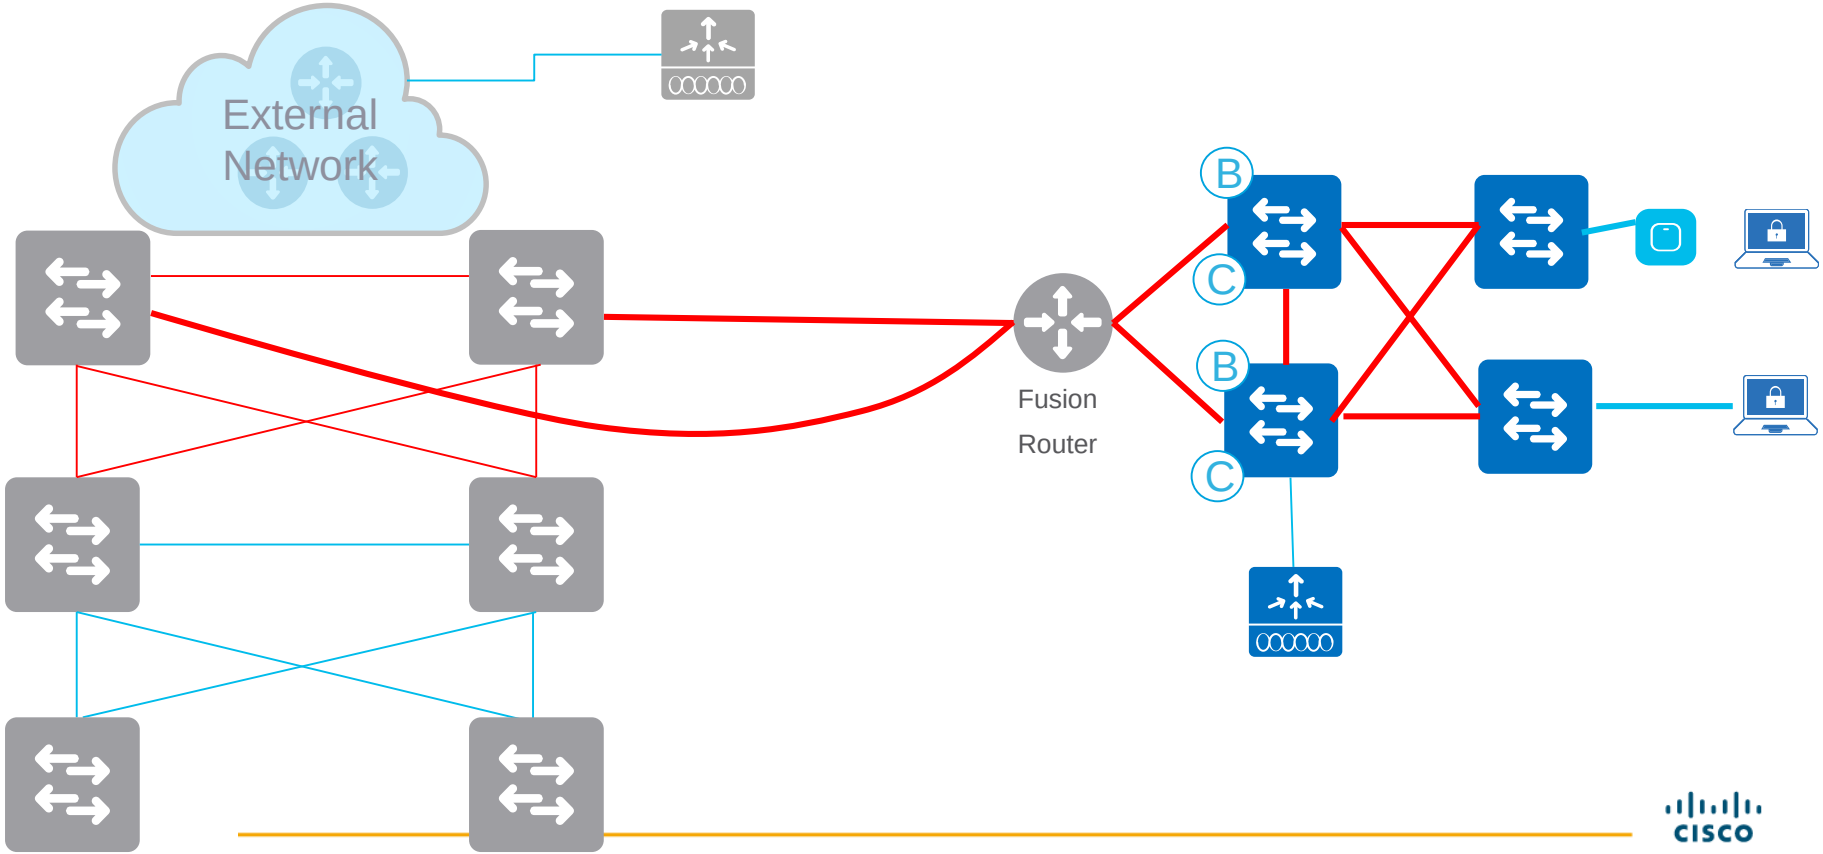

Layer-3 Routing Protocol Normalization to IGP

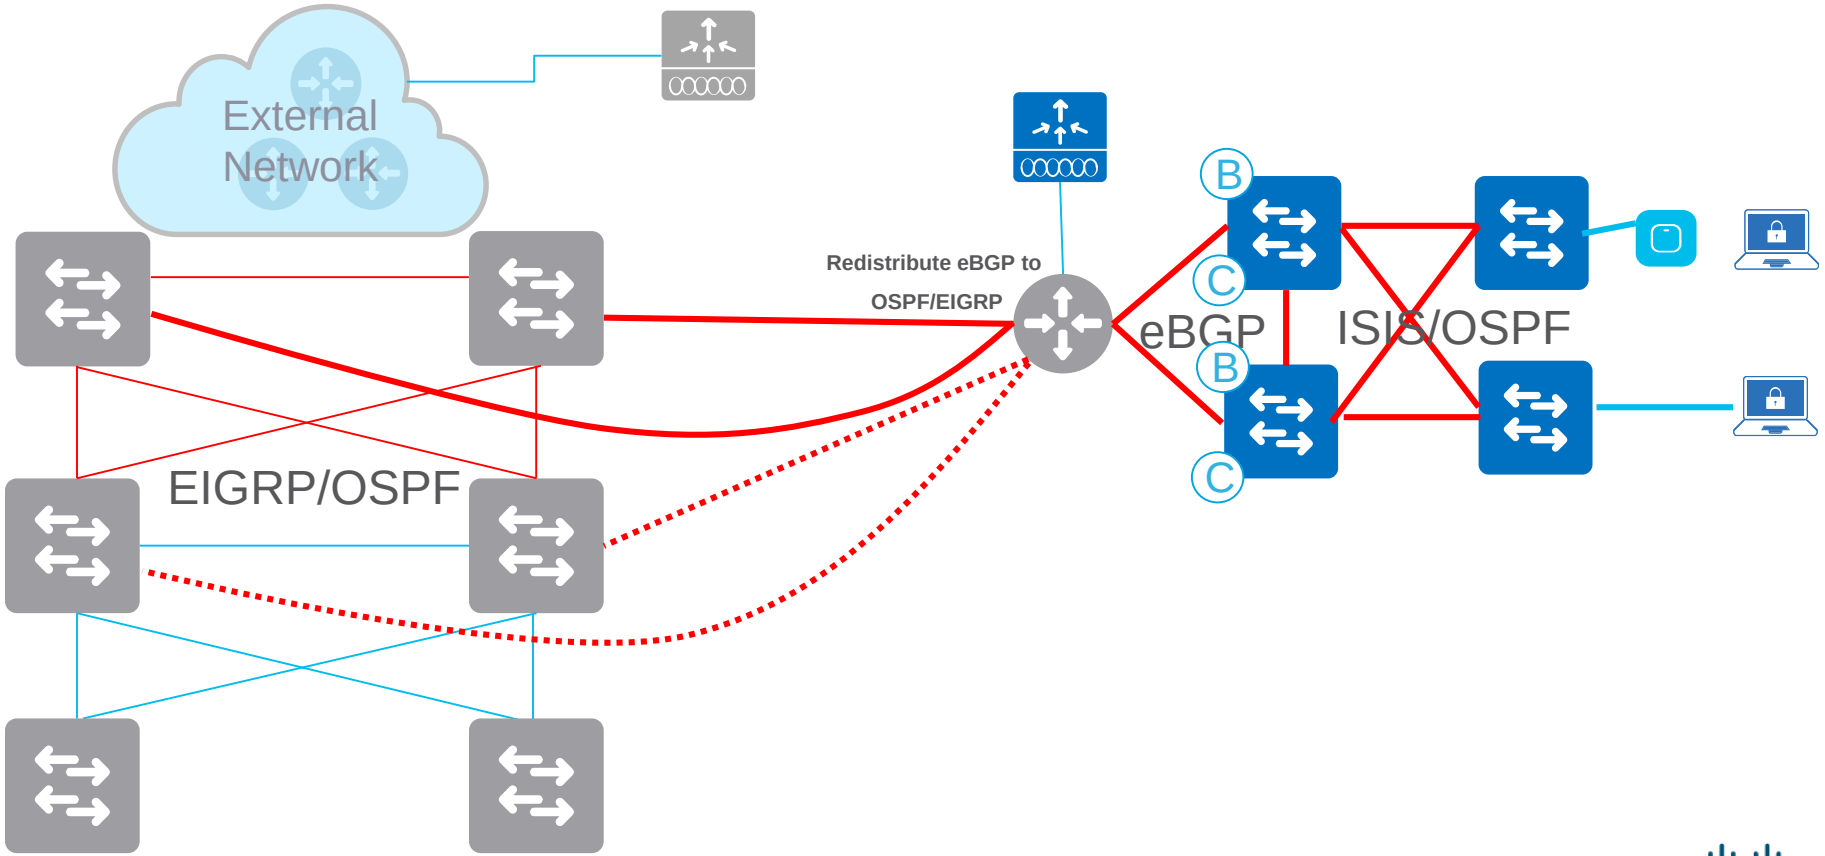

Заключение

Key Takeaways

- **Can migrate existing** network topologies to SD-Access
- **Can migrate existing subnets** into Cisco SD-Access
- Supports migration of
  - **Layer-2** as well as,
  - **Routed Access** designs
- **Automation** support makes it **easy** for migration
- Considerations in migration
  - **PoC** in labs
  - Verify every use-case works as per requirement
  - **Start small**, small/medium Campus/Branch locations