Possibilities

#CiscoLive

Real World Route/Switch to Cisco SD-Access Migration Tools and Strategies

Jerome Dolphin
Technical Marketing Engineer
CCIE#17805 (R&S, SEC), CCDE#2013::3

DGTL-BRKENS-3822

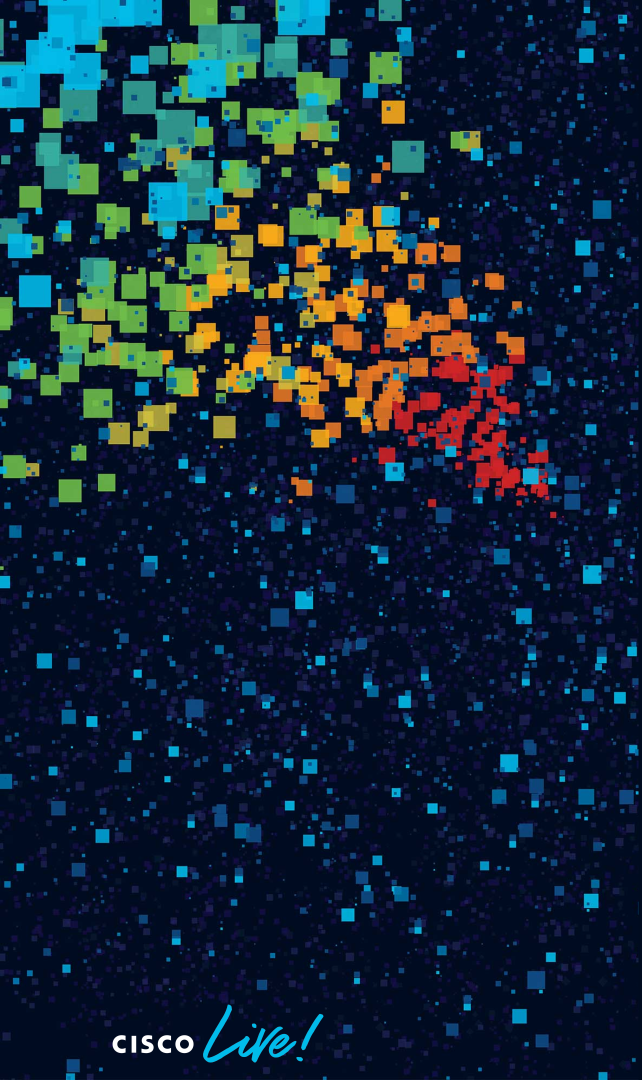

Agenda

- Introduction
- Migration philosophy and single site tools and strategies
- Multiple Cisco SD-Access sites
- Paper migration scenario
- Migration demo
- Conclusion

Introduction

Nomenclature

- For the remainder of this presentation
  - SDA = Cisco Software Defined Access
  - DNAC = Cisco DNA Center
  - ISE = Identity Service Engine
  - FMC = Firepower Management Center
  - FTD = Firepower Threat Defence
  - IBN = intent-based networking

A snapshot in time

- This solutions and strategies discussed in this presentation are evolving quickly. This is the current state as of May 2020. If you’re reviewing this well into the future, please check for the latest collateral at www.ciscolive.com

Level Setting

- Assumed fundamental SD-Access knowledge, as per session abstract:
  - BRKCRS-2810 (fundamentals)
  - BRKCRS-2811 (borders and external connectivity)
  - BRKCRS-2812 (introduction to migration strategies)
  - BRKCRS-2815 (connecting multiple fabrics)
- Ideally also familiar with:
  - BRKCRS-3810 (forwarding deep dive)
- We'll be moving at a fast pace. This will be recorded – review again later if needed
- Based on real-world migration experiences, no roadmap
- Correct R&S is fundamental to everything else
- Wireless migration out of scope, well covered in BRKCRS-2812

Recommended for Enthusiasts

- If you missed them this week, review all the IBN presentations from Cisco Live Barcelona (January 2020) at www.ciscolive.com , they all add value:

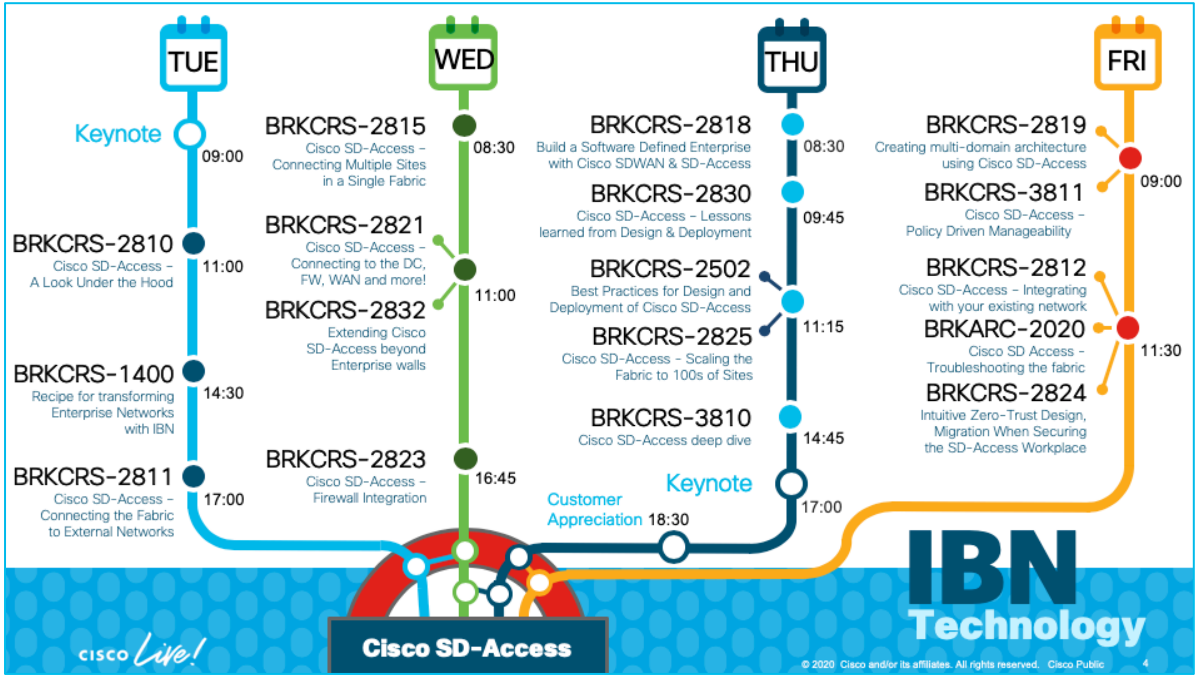

TUE

Keynote 09:00

BRKCRS-2810 Cisco SD-Access – A Look Under the Hood 11:00

BRKCRS-1400 Recipe for transforming Enterprise Networks with IBN 14:30

BRKCRS-2811 Cisco SD-Access – Connecting the Fabric to External Networks 17:00

WED

BRKCRS-2815 Cisco SD-Access – Connecting Multiple Sites in a Single Fabric 08:30

BRKCRS-2821 Cisco SD-Access – Connecting to the DC, FW, WAN and more! 11:00

BRKCRS-2832 Extending Cisco SD-Access beyond Enterprise walls 11:00

BRKCRS-2823 Cisco SD-Access – Firewall Integration 16:45

THU

BRKCRS-2818 Build a Software Defined Enterprise with Cisco SDWAN & SD-Access 08:30

BRKCRS-2830 Cisco SD-Access – Lessons learned from Design & Deployment 09:45

BRKCRS-2502 Best Practices for Design and Deployment of Cisco SD-Access 11:15

BRKCRS-2825 Cisco SD-Access – Scaling the Fabric to 100s of Sites 11:15

BRKCRS-3810 Cisco SD-Access deep dive 14:45

Keynote 17:00

Customer Appreciation 18:30

FRI

BRKCRS-2819 Creating multi-domain architecture using Cisco SD-Access 09:00

BRKCRS-3811 Cisco SD-Access – Policy Driven Manageability 09:00

BRKCRS-2812 Cisco SD-Access – Integrating with your existing network 11:30

BRKARC-2020 Cisco SD Access – Troubleshooting the fabric 11:30

BRKCRS-2824 Intuitive Zero-Trust Design, Migration When Securing the SD-Access Workplace 11:30

cisco Live!
Cisco SD-Access
IBN
Technology
© 2020 Cisco and/or its affiliates. All rights reserved. Cisco Public
4

Other Valuable Resources

- SDA compatibility matrix

http://cs.co/sda-compatibility-matrix

Please note the Cisco recommended versions for stability. Do not upgrade just because the option is available

- SDA design best practices

https://community.cisco.com/t5/networking-documents/cisco-sda-design-guidance-and-best-practices/ta-p/3865954

- Cisco Validated Designs

http://cisco.com/go/cvd/campus

Cisco's intent-based network
Delivered by Cisco Software Defined Access

RECAP

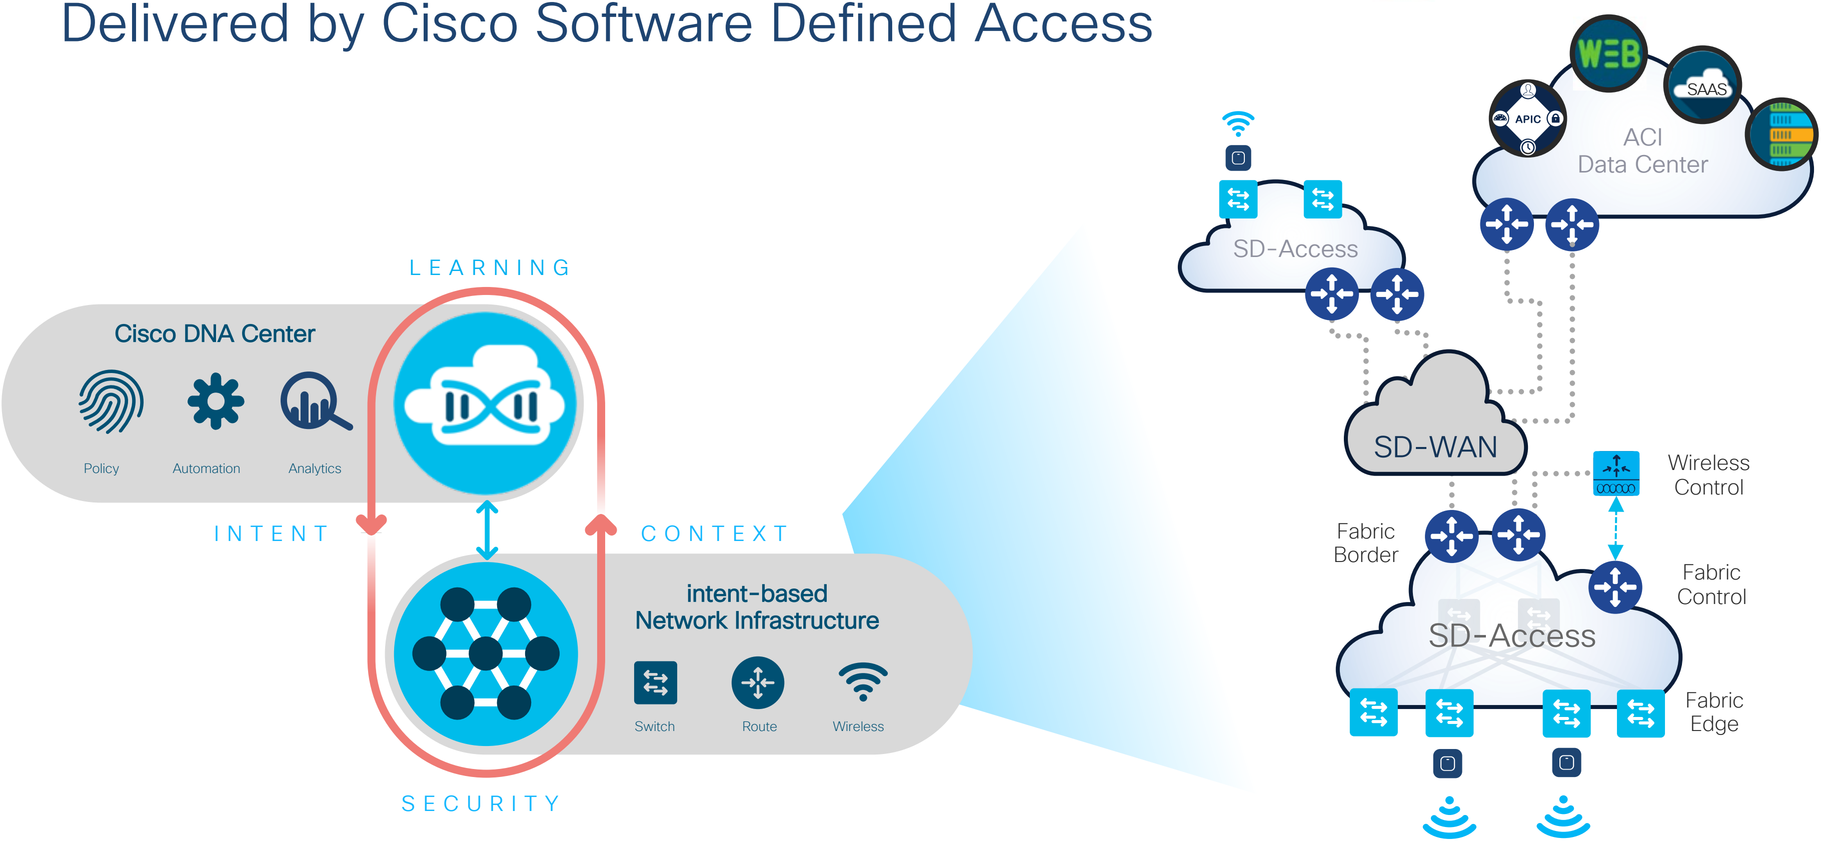

Cisco SD-Access

Fabric Roles and Terminology

RECAP

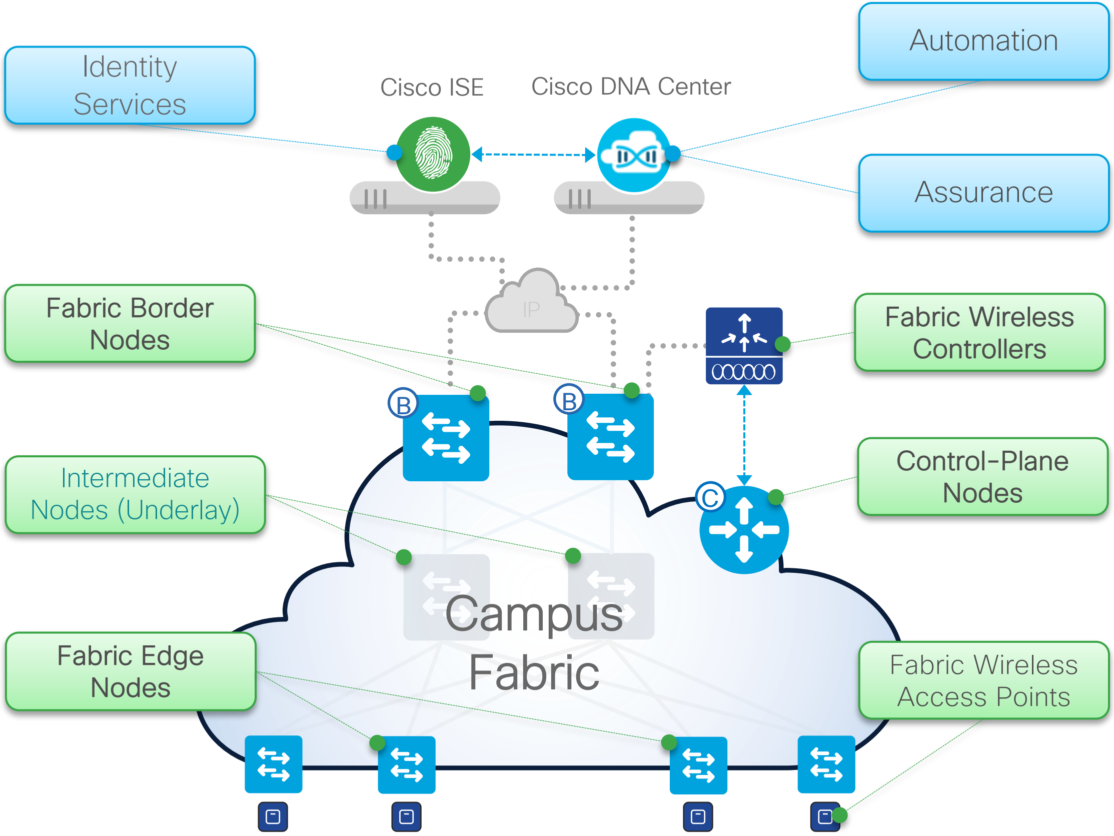

- **Network Automation** – Simple GUI and APIs for intent-based Automation of wired and wireless fabric devices
- **Network Assurance** – Data Collectors analyze Endpoint to Application flows and monitor fabric network status
- **Identity Services** – NAC & ID Services (e.g. ISE) for dynamic Endpoint to Group mapping and Policy definition
- **Control-Plane Nodes** – Map System that manages Endpoint to Device relationships
- **Fabric Border Nodes** – A fabric device (e.g. Core) that connects External L3 network(s) to the SD-Access fabric
- **Fabric Edge Nodes** – A fabric device (e.g. Access or Distribution) that connects Wired Endpoints to the SD-Access fabric
- **Fabric Wireless Controller** – A fabric device (WLC) that connects Fabric APs and Wireless Endpoints to the SD-Access fabric

Migration philosophy and single site tools and strategies

A Design and Migration Philosophy

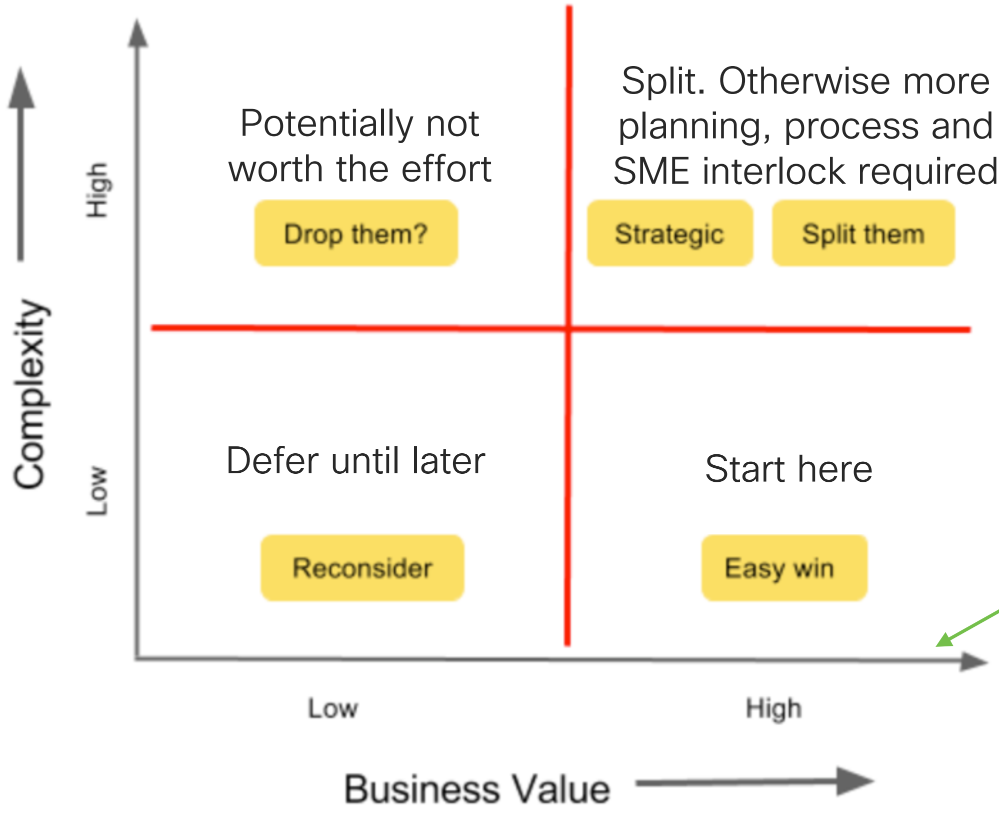

Potentially not worth the effort

Drop them?

Split. Otherwise more planning, process and SME interlock required

Strategic

Split them

Defer until later

Reconsider

Start here

Easy win

Find the happy starting place!

https://medium.com/@ank.mahajan/value-vs-complexity-a-prioritisation-tool-4a2a1ba08eda

A Design and Migration Philosophy

RECAP

Parallel IMPLEMENTATION RESOURCES	Incremental RESOURCES IMPLEMENTATION
Usually best for small scale deployments	Good for any size deployment
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

Lowest Complexity Brownfields Migration Scenario

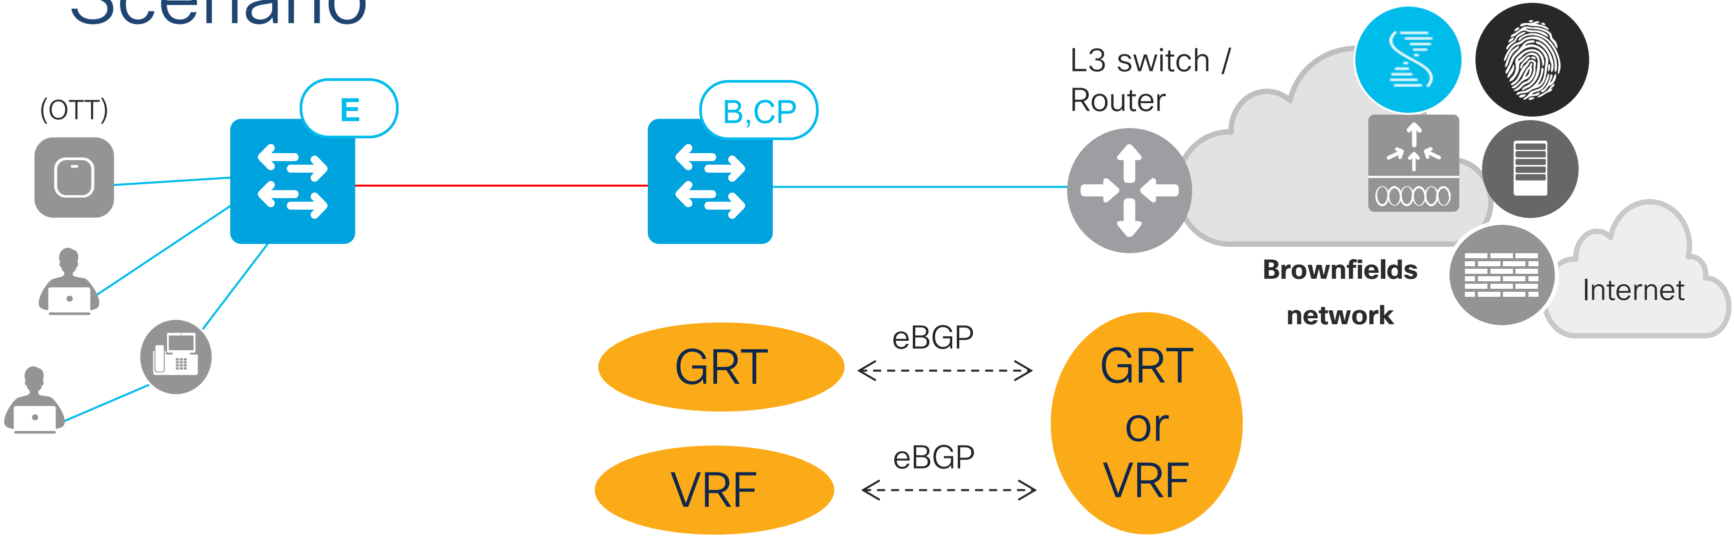

- SDA Compatibility matrix recommended code
- Single fabric site
- OTT local mode wifi
- < 20ms RTT (for wifi)
- Static FE port assignments
- No multicast

- Dedicated border node
- New FE switches
- No underlay over brownfields network
- One SDA VN
- No inter-VN security
- No SGACLs
- No SGTs

- No network transit over fabric site
- No redundancy
- No inter-border iBGP
- New IP ranges in SDA
- No L2 flooding in SDA
- No SDA L2 border
- Jumbo MTU everywhere

Then Layer on High Business Value Features

For example:

11. And so on.....

10. Fabric enabled wireless

09. Application health

08. TrustSec

07. Multicast over fabric

06. L2 border for static IP hosts

05. Inter-network security

04. Macro-segmentation with VN

03. Border redundancy / resiliency

02. SDA to brownfields L3 connectivity

01. Install DNAC and integrate with ISE

Integrating Cisco DNA Center with Existing ISE

- Recommended to integrate Cisco DNA Center with existing ISE
- ISE cluster should be compatible with SDA as per compatibility matrix
- At minimum take ISE operational backup before Cisco DNA Center integration
- If existing TrustSec/SGACL config in ISE
  - If multiple matrices or multiple SGACLs per cell, use ISE to manage TrustSec policy, not Cisco DNA Center
  - Prior to Group-Based Access Control (Cisco DNA Center <1.3.1)
    - After integration do not manage TrustSec policies in both in ISE and Cisco DNA Center. Cisco DNA Center can overwrite changes previously made directly in ISE GUI
      - Cisco DNA Center will learn policies, SGTs and contracts from ISE
      - Will not learn SGACLs with ICMP, log or source ports
  - As of Group-Based Access, Cisco DNA Center (1.3.1+) will pull in ISE TrustSec policies. This is a merge. After the merge, either Cisco DNA Center or ISE is the “master” for SGT policies

Why Integrate Cisco DNA Center with Existing ISE

- Supplicants on endpoints don’t need to change
- Profiling policies can be reused
- Existing AD integrations can be reused
- Existing posturing rules can be reused
- AuthZ rules can be reused, but with different results profile

Multi-VN Border Route Peering Scenarios

- Many permutations are possible. Most common 1:1

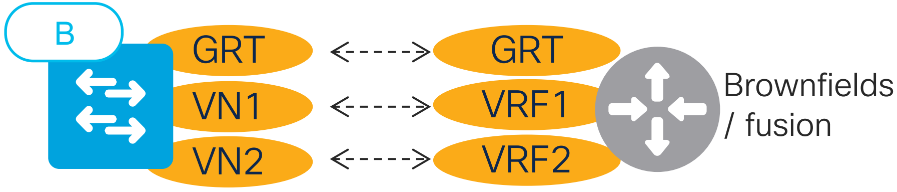

Brownfields / fusion

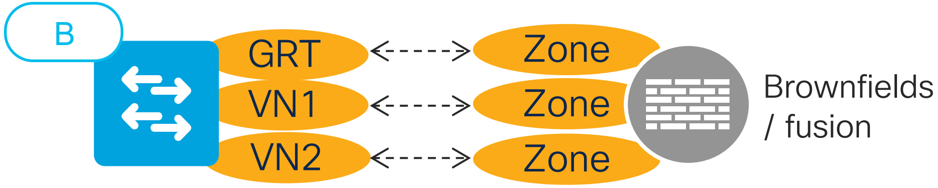

Brownfields / fusion

- But more generally, 1:n and n:1 are also possible

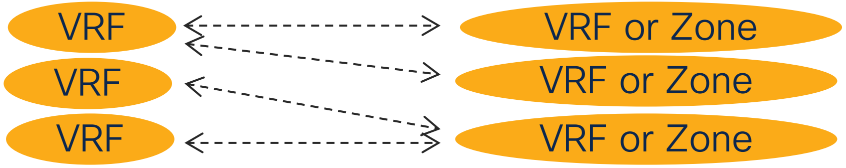

- Properly captured connectivity and security requirements dictate the 'right' answer. Define and agree on these early
- Easier to get right before start, harder to change afterwards

New Scopes and Routing

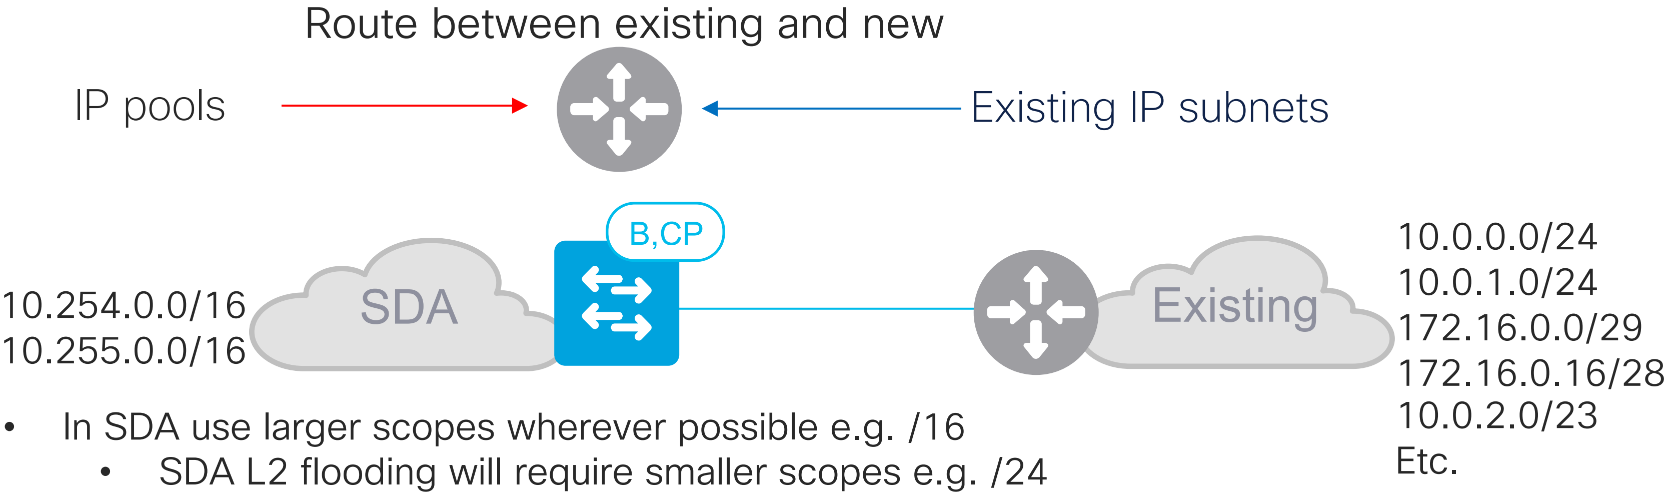

- In SDA use larger scopes wherever possible e.g. /16
  - SDA L2 flooding will require smaller scopes e.g. /24
- Discourage moving existing IP subnets into SDA due to IP pool limits:
  - Every IP pool in a fabric site is an SVI on an FE
  - Limits per fabric site
  - Limits per Cisco DNA Center cluster (virtual + physical interfaces), check the data sheet!
  - Leave room to grow!
- Moving 1000 existing /29s-23s may exceed Cisco DNA Center scale limits, and very manual
- Underlay will require new IP pool(s)
  - Do not manually assign any IP addresses from LAN automation pools

Application Visibility and Experience

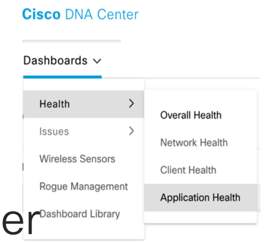

Cisco DNA Center

Dashboards

- Health
- Issues
- Wireless Sensors
- Rogue Management
- Dashboard Library

- Overall Health
- Network Health
- Client Health
- Application Health

- Quantitative metrics: Application name, usage, and throughput
- Qualitative metrics: Packet loss, network latency, client network latency, server network latency, application network latency and jitter
- Application Visibility = DNAC rendering of quantitative metrics
- Application Experience = Application Health = DNAC rendering of quantitative + qualitative metrics
- As of now, in an SDA fabric:

	Application Visibility	Application Experience
Catalyst 9000 switch	Yes, on access ports with no VXLAN encap, if ETA is not enabled, on capable supported C9K switching platforms – check latest Assurance configuration guide	No
Fabric enabled wireless + C9800 WLC	Not today. Roadmap	No
ISR or ASR border router	Yes, on ports with no VXLAN encap	Yes, on ports with no VXLAN encap

Application Visibility in SDA

- In an SDA network:
  - Application Visibility supported on Cat9K FE switches and Cat9K borders running 16.12.2T+ with DNAC 1.3.3+
    - Not supported on fabric APs, yet
  - Application Visibility NOT supported if:
    - Traffic is VXLAN encapsulated
    - Encrypted Traffic Analytics is enabled on the Cat9K switch
  - Application Visibility:
    - Enable on FE ports connected to endpoints, other than fabric APs or extended nodes
    - Enable using the “lan” interface description + maximal telemetry profile

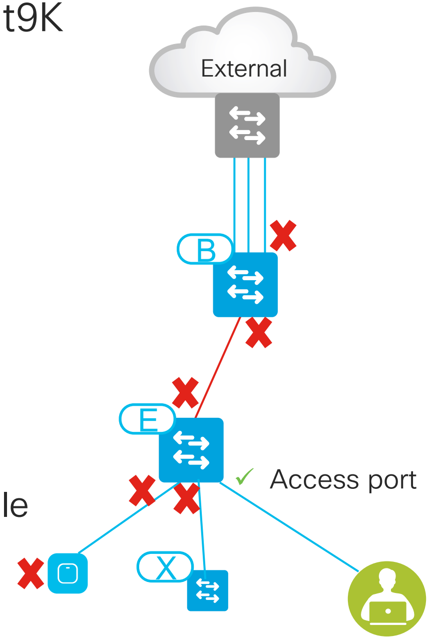

Application Experience in SDA

- In an SDA network:
  - Application Experience is supported on border router running 16.12.2T+ with DNAC 1.3.3+
  - Application Visibility NOT supported if:
    - Traffic is VXLAN encapsulated
  - Application Experience:
    - Enable on border VRF handoff ports
    - Enable using the “lan” interface description + maximal telemetry profile

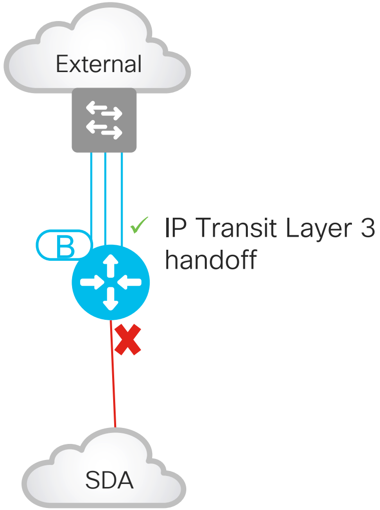

Border Platform Selection

Router B	Switch B
• Lower port density • Patch FEs to intermediate switches • Higher fabric scale. CP, SGT, SGACL, etc. • Future support for co-located SDW:SDA functionality • A site re-provision might be required • Qualitative data for Application Experience	• Higher port density • Patch FEs directly to border switch • L2 handoff support on C3K, C9K and C6K • Multi-chassis • Back-stack Catalyst 9300 • Stackwise Virtual support for 9500 and 9500H with Cisco DNA Center 1.3.3+

Border Switch Redundancy

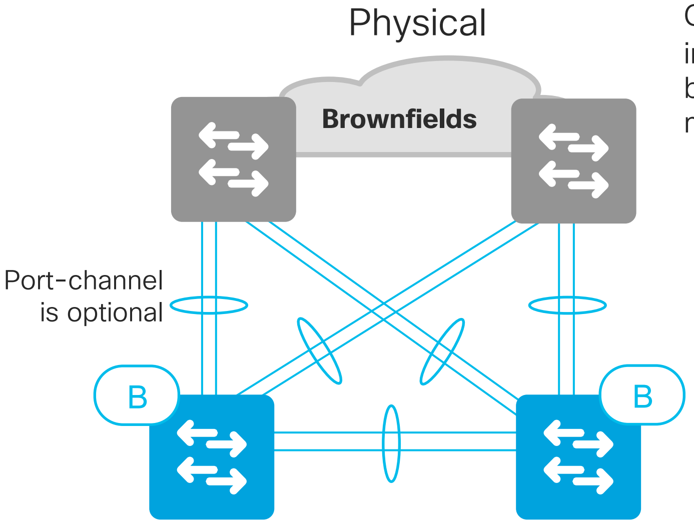

Can also be routed sub-interface, depending on brownfields switch model / router

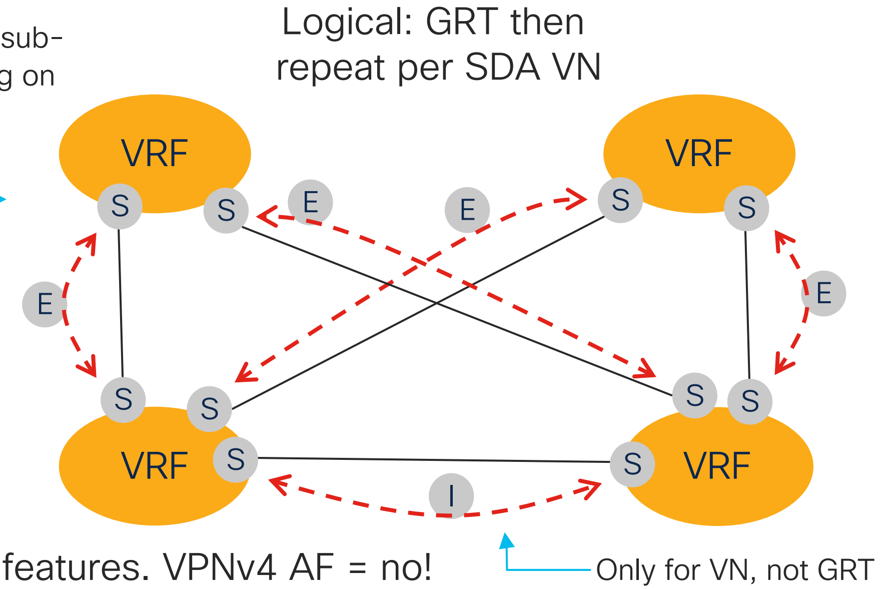

- Per-VRF BGP recommended. Best route control features. VPNv4 AF = no!
  - Other protocols allowed but not recommended
- BFD optional, but recommended if other side capable, SVI can stay up when physical link fails
  - Manually add on CLI or with template
- N-S and E-W manual port-channel optional/permitted. Reduces BGP peerings to 5x per VRF

S	SVI
E	eBGP+BFD
I	iBGP+BFD

Border Router Redundancy

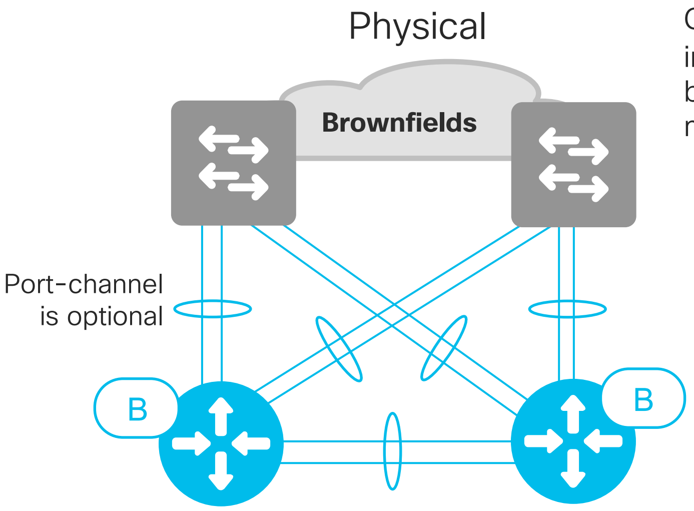

Can also be routed sub-interface, depending on brownfields switch model / router

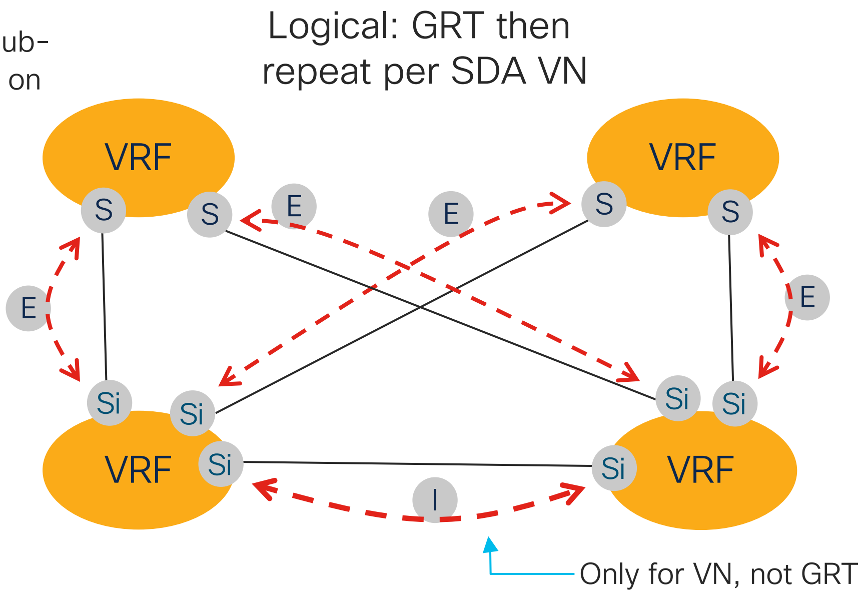

- Same as switches, but with sub-interfaces

S	SVI
E	eBGP+BFD
I	iBGP+BFD
Si	Sub-int

Border Node Redundancy

- North-South eBGP
  - Deploy manually or via Cisco DNA Center automation with IP Transit
    - If manual then on border set neighbour weight to 65535 as per CSCvg86018 and CSCvi29660

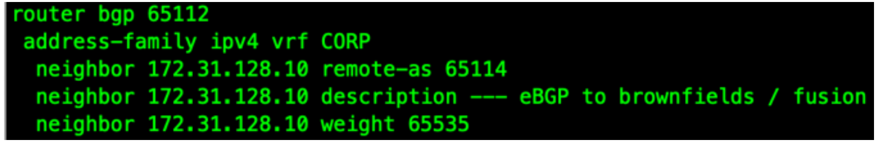

```
router bgp 65112
 address-family ipv4 vrf CORP
  neighbor 172.31.128.10 remote-as 65114
  neighbor 172.31.128.10 description --- eBGP to brownfields / fusion
  neighbor 172.31.128.10 weight 65535
```

    - Cisco DNA Centre L3 border handoff automation will automatically set weight to 65535
  - Automated L3 border handoff uses VLAN ranges 3001-3500
    - Prior to 1.3.3 VLANs are automatically selected by Cisco DNA Centre
      - If deleted, will not get same again
    - As of 1.3.3 we can manually choose the VLAN in 3001-3500 range, but not the IP subnet
    - No BFD automation

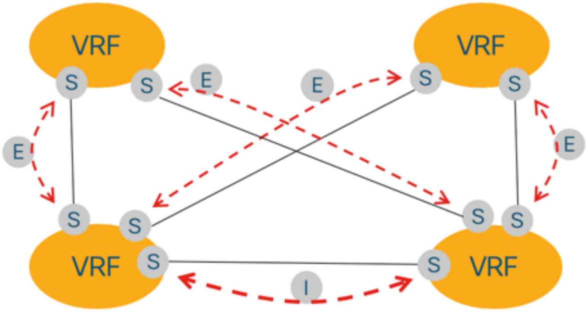

Border Node Redundancy

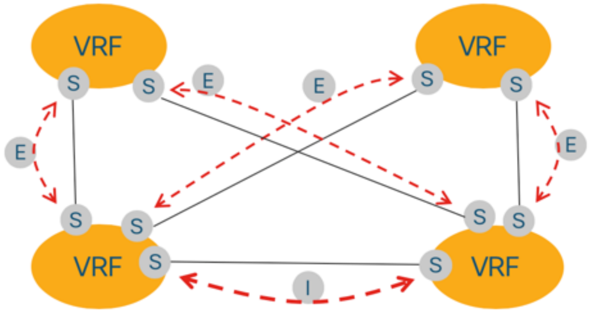

- East-West iBGP
  - Do NOT implement on internal-only border pairs
  - Not necessary if borders will never lose connectivity to underlay and outside world
  - Not necessary for underlay, only overlay
  - Solves routing convergence in external border failure scenarios
- Will never be automated by Cisco DNA Center, because the need for this will go away in future
- If internal + external borders then see CSCvm77399 for additional necessary filtering logic. Make sure on both borders
  - One possible solution on internal + external border:

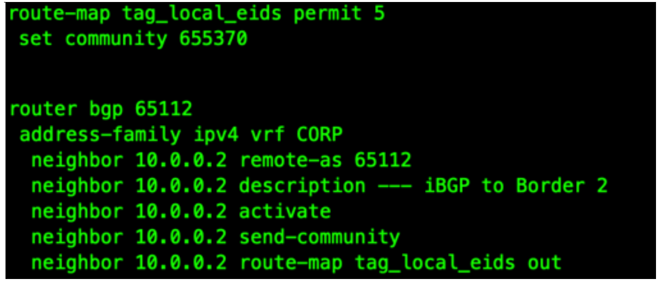

```
route-map tag_local_eids permit 5
 set community 655370


router bgp 65112
 address-family ipv4 vrf CORP
  neighbor 10.0.0.2 remote-as 65112
  neighbor 10.0.0.2 description --- iBGP to Border 2
  neighbor 10.0.0.2 activate
  neighbor 10.0.0.2 send-community
  neighbor 10.0.0.2 route-map tag_local_eids out
```

Border Node Provisioning

- Be clear on what border role is required. Review BRKCRS-2811
- Cannot easily change after provisioning
- If site borders will all have full connectivity to all external networks then usually:
  - There is no point importing external prefixes into fabric
  - Choose external border, simplifies fabric routing tables and East-West iBGP configuration

External border

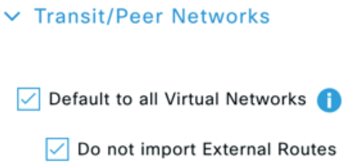

Transit/Peer Networks

☑ Default to all Virtual Networks

☑ Do not import External Routes

External+Internal border

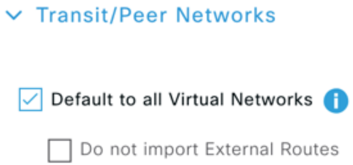

Transit/Peer Networks

☑ Default to all Virtual Networks

□ Do not import External Routes

Internal border

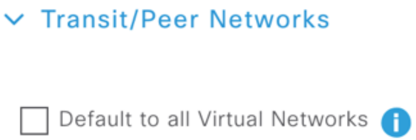

Transit/Peer Networks

□ Default to all Virtual Networks

Fusion Firewall

- Why use a fusion firewall?
  - Comprehensive inter-VN policy, stateful inspection, AVC
  - ASA: Source SGT to Destination SGT policy
  - FTD: NGIPS, AMP/TG integration, Source SGT to Destination SGT policy as of 6.5, TLS decryption, URL filtering
  - SGT can be derived from ACI EPG. Review latest BRKDCN-2489
  - Rich reporting in FMC
    - Top blocks, top malwares top hosts effected by malware, network risk, customized, etc.
  - Firewall grade logging. Potential data enrichment
    - IOS/IOS-XE SGACL logging not guaranteed

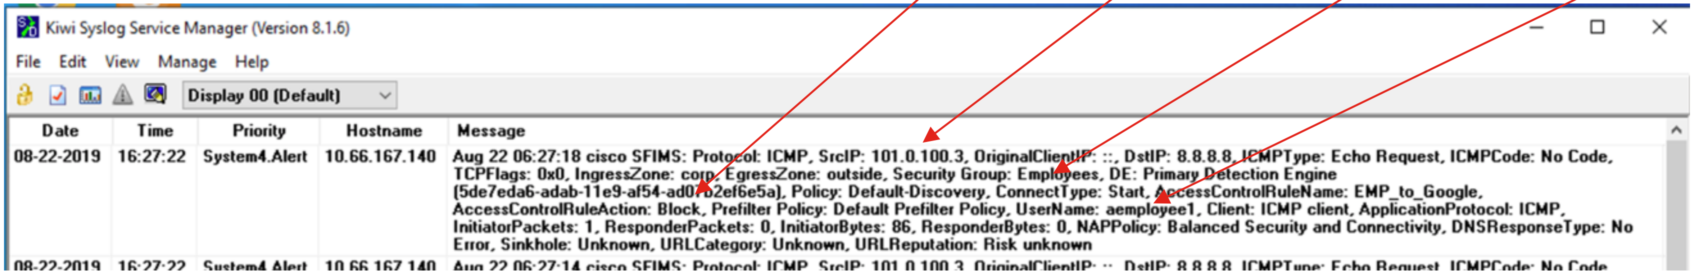
Kiwi Syslog Service Manager (Version 8.1.6)
File Edit View Manage Help
Display 00 (Default)

Action Source IP Source SGT Source username

Date	Time	Priority	Hostname	Message
08-22-2019	16:27:22	System4.Alert	10.66.167.140	Aug 22 06:27:18 cisco SFIMS: Protocol: ICMP, SrcIP: 101.0.100.3, OriginalClientIP: ::, DstIP: 8.8.8.8, ICMPType: Echo Request, ICMPCode: No Code, TCPFlags: 0x0, IngressZone: corp, EgressZone: outside, Security Group: Employees, DE: Primary Detection Engine (5de7eda6-adab-11e9-af54-ad07b2ef6e5a), Policy: Default-Discovery, ConnectType: Start, AccessControlRuleName: EMP_to_Google, AccessControlRuleAction: Block, Prefilter Policy: Default Prefilter Policy, UserName: aemployee1, Client: ICMP client, ApplicationProtocol: ICMP, InitiatorPackets: 1, ResponderPackets: 0, InitiatorBytes: 86, ResponderBytes: 0, NAPPolicy: Balanced Security and Connectivity, DNSResponseType: No Error, Sinkhole: Unknown, URLCategory: Unknown, URLReputation: Risk unknown
08-22-2019	16:27:22	System4.Alert	10.66.167.140	Aug 22 06:27:14 cisco SFIMS: Protocol: ICMP, SrcIP: 101.0.100.3, OriginalClientIP: ::, DstIP: 8.8.8.8, ICMPType: Echo Request, ICMPCode: No Code

Example: Firepower logging. FMC + ISE integration + AD integration

Fusion Firewall

- Other notes from the field:
  - TrustSec policies not downloaded from ISE to firewall
    - Usually not an issue:
      - East-West SGT polices in fabric
      - North-South SGT policies in fusion firewall
  - Size appropriately: max throughput, max connections per second, etc. Factor in unit failures

Fusion Firewall

- FTD and ASA use GR for HA during failover
  - GR maintains route peering session while the newly elected master re-establishes the session
  - BFD session will drop on active firewall master failover, taking down BGP peering, causing an outage
- On firewalls, if BGP session is expected to seamlessly move between units, do not use BFD, it counterfeits GR

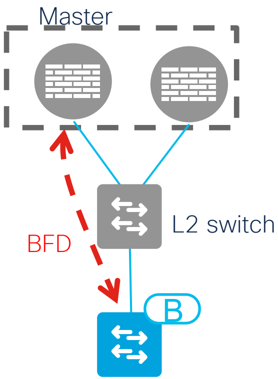

Fusion Firewall

Example scenarios: Switch cluster between border and firewall

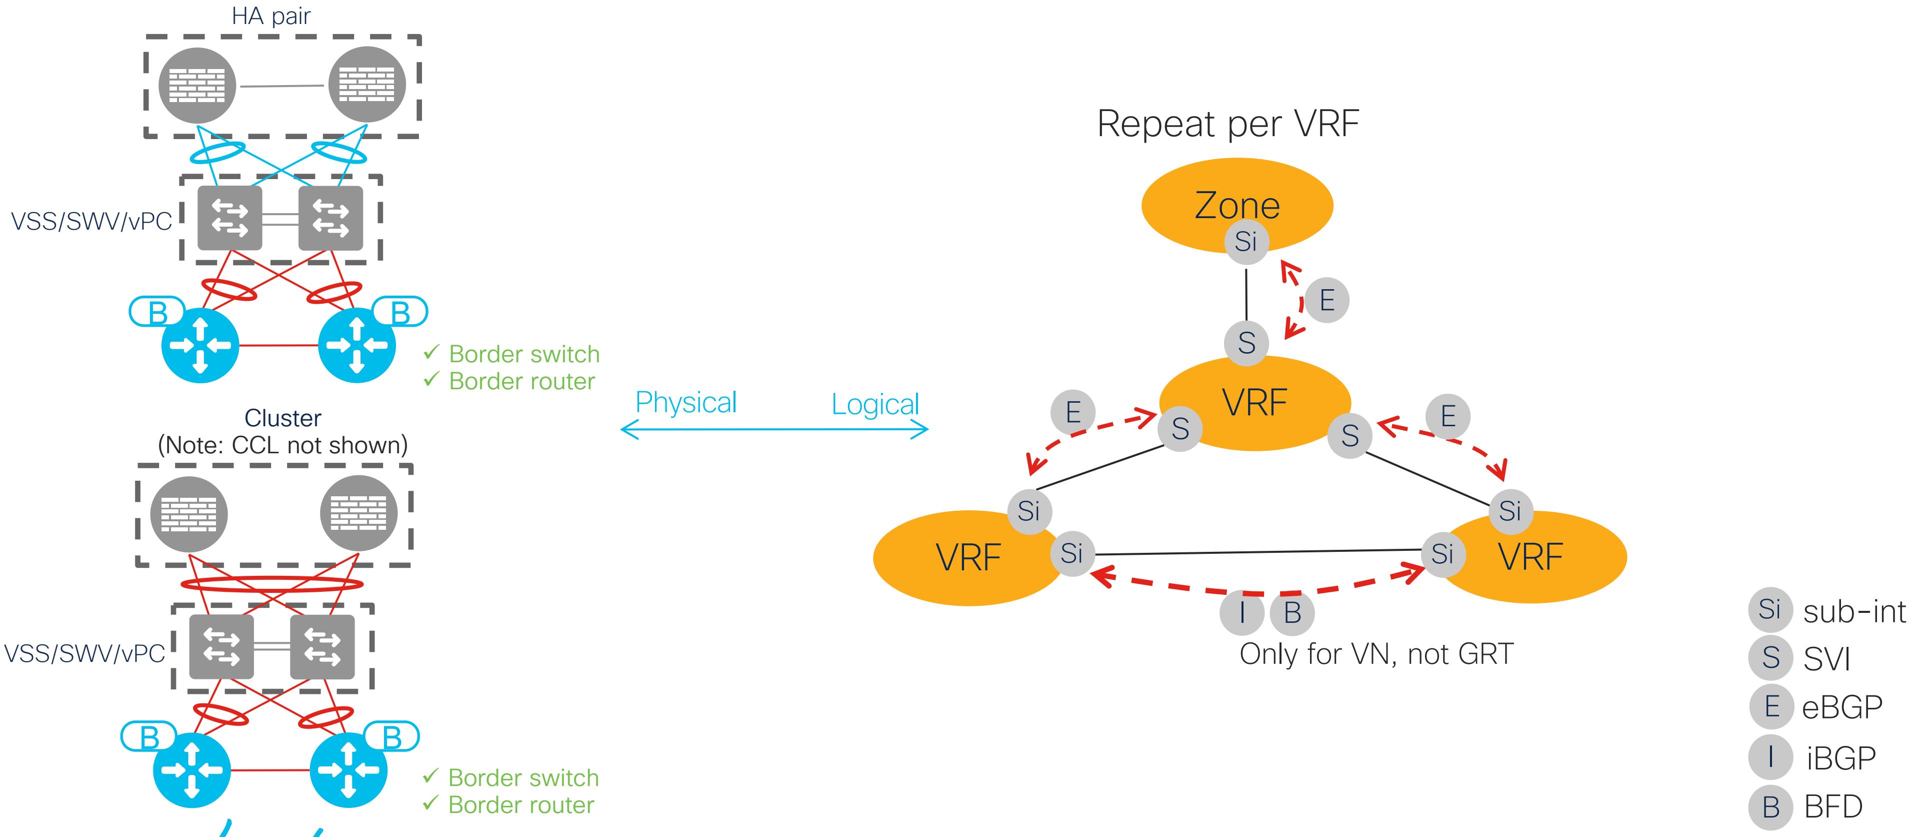

✓ Border switch
✓ Border router

✓ Border switch
✓ Border router

Si sub-int
S SVI
E eBGP
I iBGP
B BFD

Fusion Firewall

Example scenarios: Firewall attached directly to border

HA pair

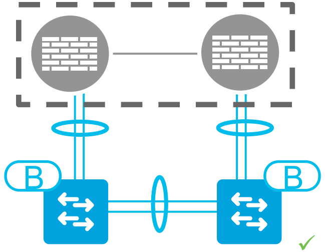

✓ Border switch
X Border router

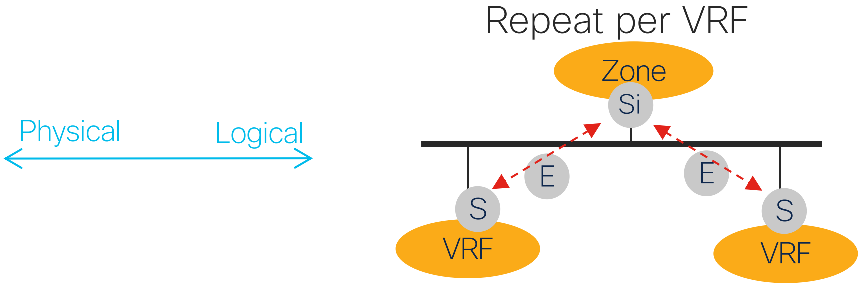

Cluster
(Note: CCL not shown)

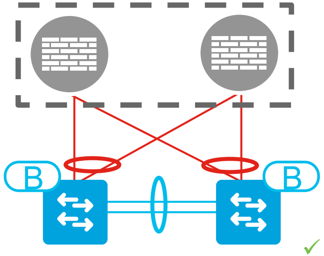

✓ Border switch
~ Border router, depends on CCL placement

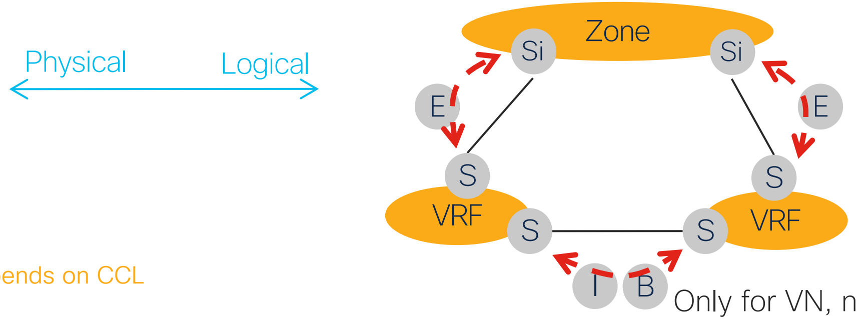

Only for VN, not GRT

Si	sub-int
S	SVI
E	eBGP
I	iBGP
B	BFD

Fusion Firewall

Example scenarios: Firewall attached directly to border

Stackwise Virtual border

- DNAC 1.3.3+
- 9500 or 9500H with 16.12.2t+

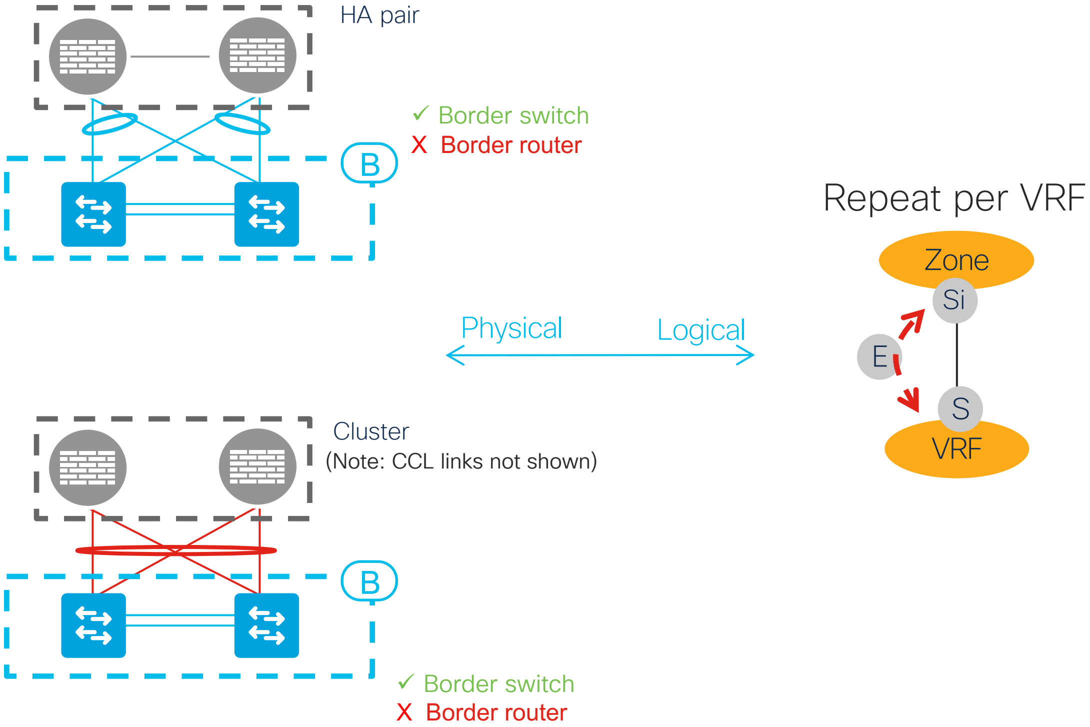

Si	sub-int
S	SVI
E	eBGP
I	iBGP
B	BFD

Test Your Configuration Before Critical Apps

- Routing convergence is also dependent on correct brownfields routing configuration
- Building 1-2 fabric edges does not impact existing network and can be used for testing of endpoints
- Before adding critical production traffic, test failure scenarios:
  - Fabric endpoint PING to hosts external to fabric
    - ECMP. Run multiple parallel PINGs
  - Fail links
  - Reload borders
  - Fix anomalies
  - Repeat tests
- Get a sampling of all the exotic endpoints and test on this fabric edge – if it works here, it will work 99.999% elsewhere
- Once the borders are right, the whole fabric can leverage this. Correct border routing is the same for fabric of 1x FE or 250x FE

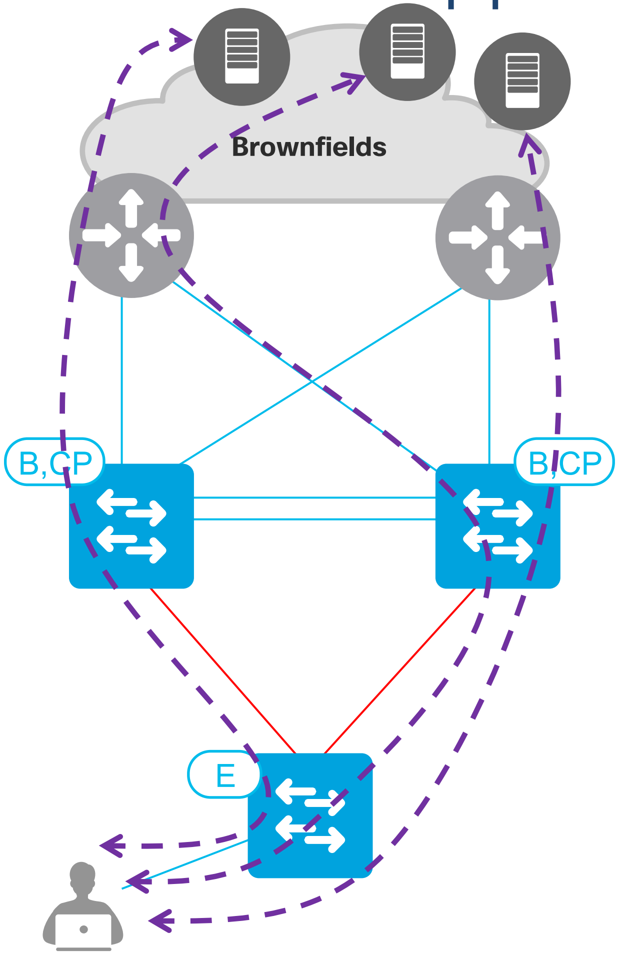

SDA and PXE Boot

- Successful strategies from the field:
  1. Set DHCP options to redirect the PXE client to a TFTP PXE server e.g. https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc732351(v=ws.10)?redirectedfrom=MSDN
  2. Use an PXE implementation that supports reflection of option 82 and add the PXE/SCCM server into the IP helpers list in SDA fabric
     - Beware, not all implementations support option 82 reflection
     - SDA FE sets option 82. On the return packet, SDA border requires option 82 to be same value as set by the FE. See BRKCRS-3810 for details of how option 82 is used with IP helper

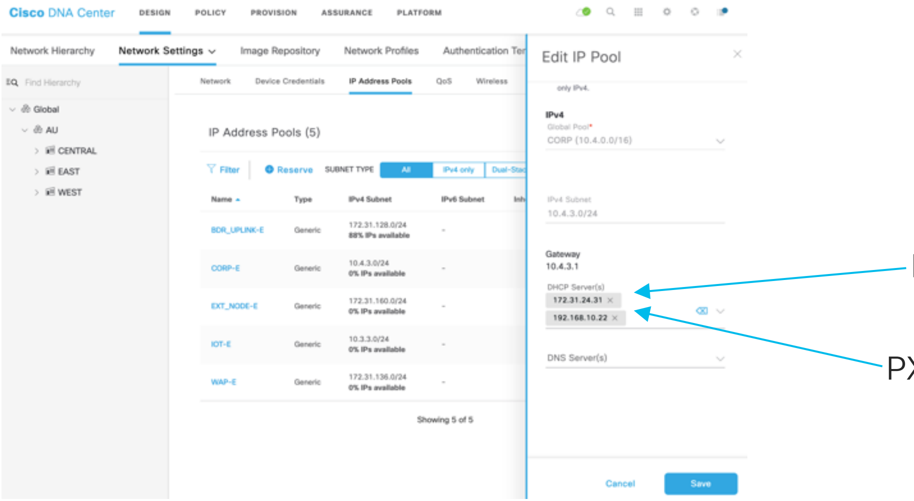

IP Address Pools (5)

Name	Type	IPv4 Subnet	IPv6 Subnet	Inh
BDR_UPLINK-E	Generic	172.31.128.0/24 88% IPs available	-	
CORP-E	Generic	10.4.3.0/24 0% IPs available	-	
EXT_NODE-E	Generic	172.31.160.0/24 0% IPs available	-	
IOT-E	Generic	10.3.3.0/24 0% IPs available	-	
WAP-E	Generic	172.31.136.0/24 0% IPs available	-	

Showing 5 of 5

Edit IP Pool

only IPv4.

IPv4

Global Pool*: CORP (10.4.0.0/16)

IPv4 Subnet: 10.4.3.0/24

Gateway: 10.4.3.1

DHCP Server(s): 172.31.24.31, 192.168.10.22

DNS Server(s)

DHCP server

PXE/SCCM server

Disclaimer: These are observations from successful real-world SDA migrations. All PXE implementations use 3rd party code, thus results may differ depending on the PXE implementation

3rd Party IP Phone

- SDA FE access port supports separate authentication for IP phone and PC behind IP phone
  - 3rd party phone and PC behind phone can be assigned to different VLANs
- Voice VLAN can be communicated to 3rd party phone through DHCP option, LLDP or CDP (exact method implementation dependent)
- 3rd party IP phone
  - Some types do not support 802.1x and will require MAB
    - If MAB, timers may need to be accelerated for some 3rd party phones
    - Same is true regardless of SDA or traditional network
  - Should be under support from the vendor
    - Some issues observed where unsupported IP phone has problem that will not be fixed by vendor due to no support
  - Should support 802.1x passthrough for PC behind phone

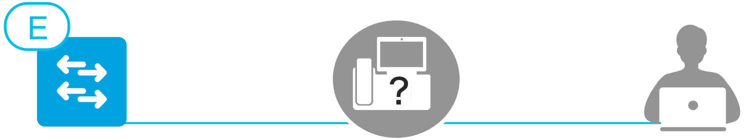

FE Access Port with Unintelligent Switch

- Use extended node if possible
- An "intelligent" switch consumes EAPoL, which breaks 802.1X between FE and endpoints
- "Unintelligent switch" is one that does not consume EAPoL
  - Finding a switch that behaves this way is responsibility of partner / customer
- Unintelligent switch tradeoffs:
  - Microsegmentation between endpoints physically connected to unintelligent switch is lost
  - Assurance and automation not possible for unintelligent switch
  - Cisco TAC do not support the unintelligent switch
- Unintelligent switch connected to FE is supported
  - Make sure tradeoffs are clearly understood
  - Each endpoint can dynamically authorised into a different network and SGT
- FE edge port supports maximum 10 IP addresses. 11th IP address is dropped by FE
  - IPv6 hosts will have multiple IP address, typically 3-4 per host

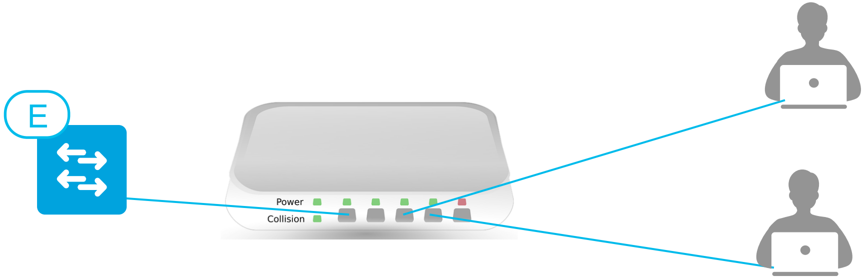

Why MTU Matters

RECAP

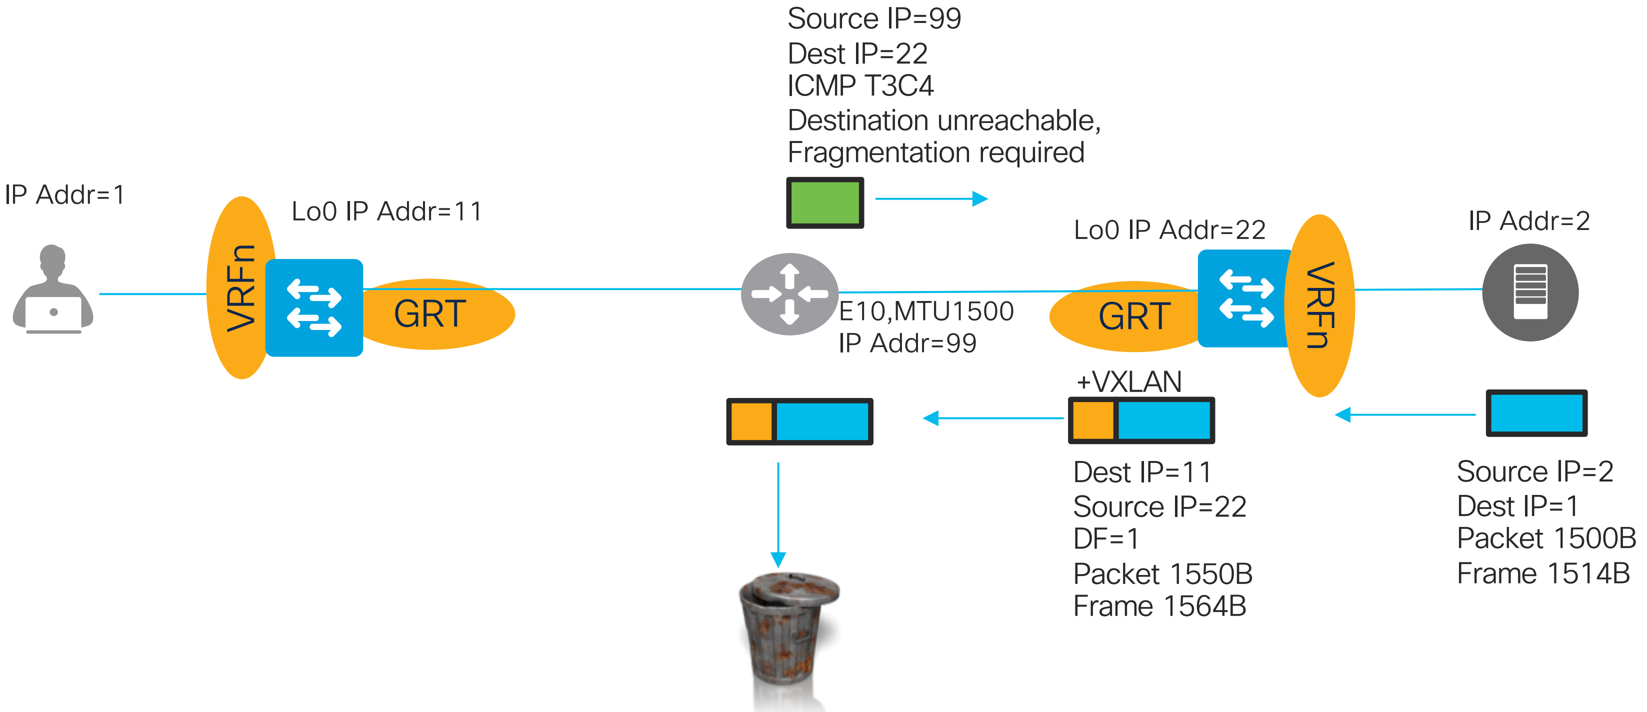

- ICMP T3C4 sent in GRT, thus never received by offending endpoint
- Exceeding MTU in underlay = black hole
- Use TCP adjust-MSS in overlay, but this won’t help large UDP packets (if they exist?)

Fragmentation

- Traditional non-fabric IOS-XE routers and switches
  - Fragment packets when IP MTU set on an SVI or routed interface
- SD-Access VXLAN:
  - No fragmentation support within fabric site
  - No fragmentation support within SD-Access transit
  - E2E MTU must be sufficient for VXLAN overhead

Unknown Unicast in Traditional Networking

RECAP

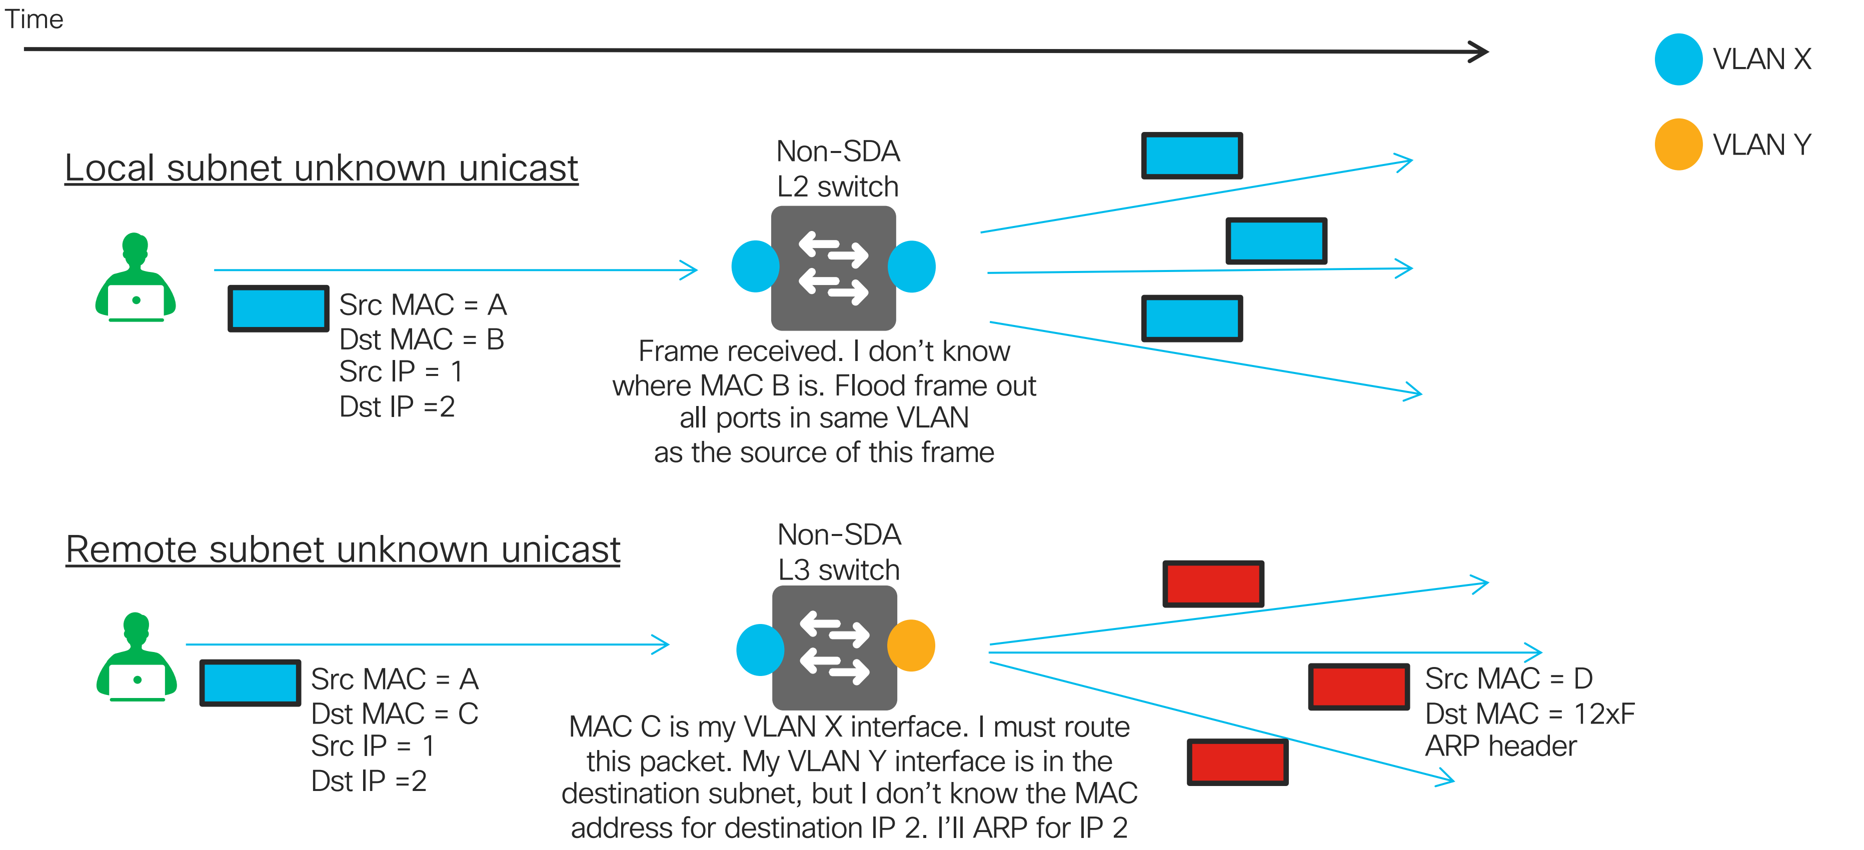

Time

● VLAN X

● VLAN Y

Local subnet unknown unicast

Non-SDA
L2 switch

Src MAC = A
Dst MAC = B
Src IP = 1
Dst IP =2

Frame received. I don't know where MAC B is. Flood frame out all ports in same VLAN as the source of this frame

Remote subnet unknown unicast

Non-SDA
L3 switch

Src MAC = A
Dst MAC = C
Src IP = 1
Dst IP =2

MAC C is my VLAN X interface. I must route this packet. My VLAN Y interface is in the destination subnet, but I don't know the MAC address for destination IP 2. I'll ARP for IP 2

Src MAC = D
Dst MAC = 12xF
ARP header

L2 Flooding

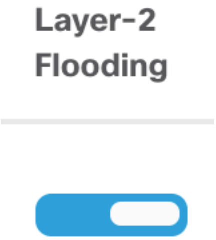

Layer-2 Flooding

- Disabled by default. Enable selectively and sparingly
- Floods Ethernet broadcast and link local multicast in overlay
- Does not flood unknown unicast
  - Beware of the silent host
    - Rumours of endpoints that don’t use/respond to ARP, no tangible evidence ….. yet

Is that you Bigfoot?

    - Can hardcode IP/MAC into IPDT on FE switch CLI. A not scalable work around

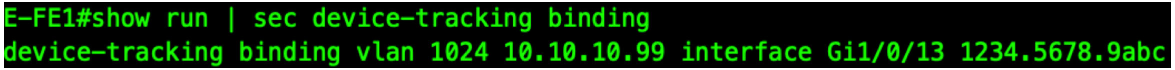

```
E-FE1#show run | sec device-tracking binding
device-tracking binding vlan 1024 10.10.10.99 interface Gi1/0/13 1234.5678.9abc
```

- Wake on LAN where source and destination on same subnet is supported – Ethernet broadcast
- Wake on LAN where source is remote does not work well yet – directed broadcast
  - There is some workarounds, please consult with your technical presales team
  - Permanent and automated fix is on roadmap for later this year

L2 Flooding – Local Unknown Unicast

- May help solve local subnet unknown unicast, because end hosts should ARP first

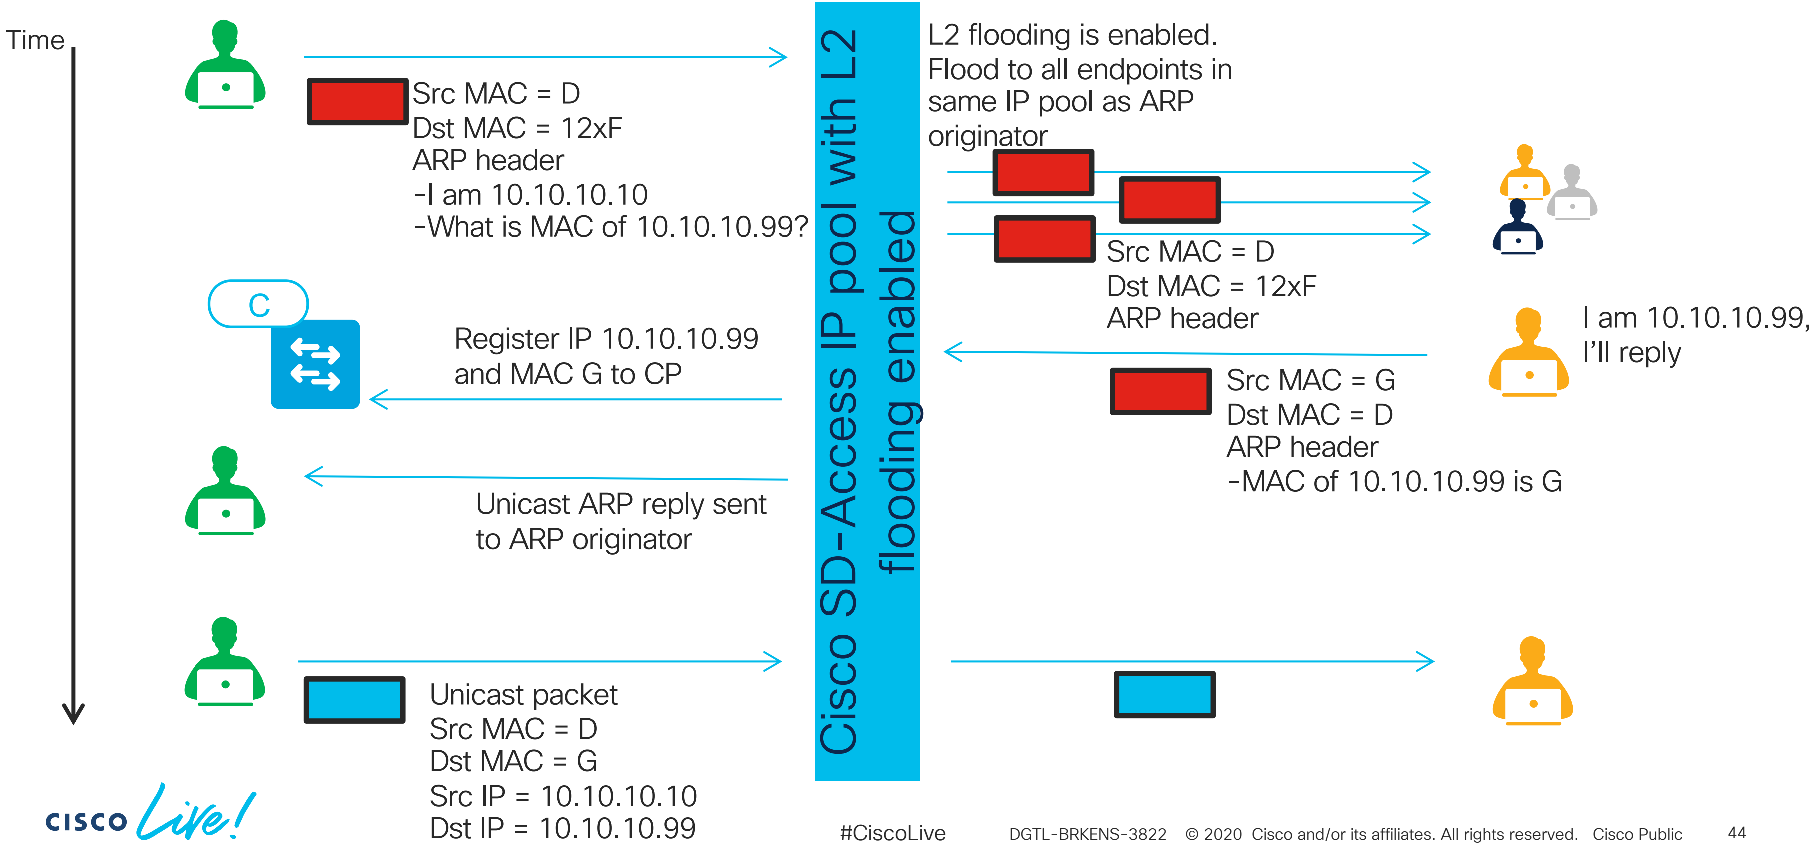

Src MAC = D
Dst MAC = 12xF
ARP header
-I am 10.10.10.10
-What is MAC of 10.10.10.99?

L2 flooding is enabled. Flood to all endpoints in same IP pool as ARP originator

Src MAC = D
Dst MAC = 12xF
ARP header

Register IP 10.10.10.99 and MAC G to CP

I am 10.10.10.99, I'll reply

Src MAC = G
Dst MAC = D
ARP header
-MAC of 10.10.10.99 is G

Unicast ARP reply sent to ARP originator

Unicast packet
Src MAC = D
Dst MAC = G
Src IP = 10.10.10.10
Dst IP = 10.10.10.99

L2 Flooding – Remote Unknown Unicast

- **Cannot** solve remote subnet unknown unicast
- Cisco SD-Access border does not ARP. Use IPDT workaround described previously, or have something else in same IP pool ARP, thus populating CP

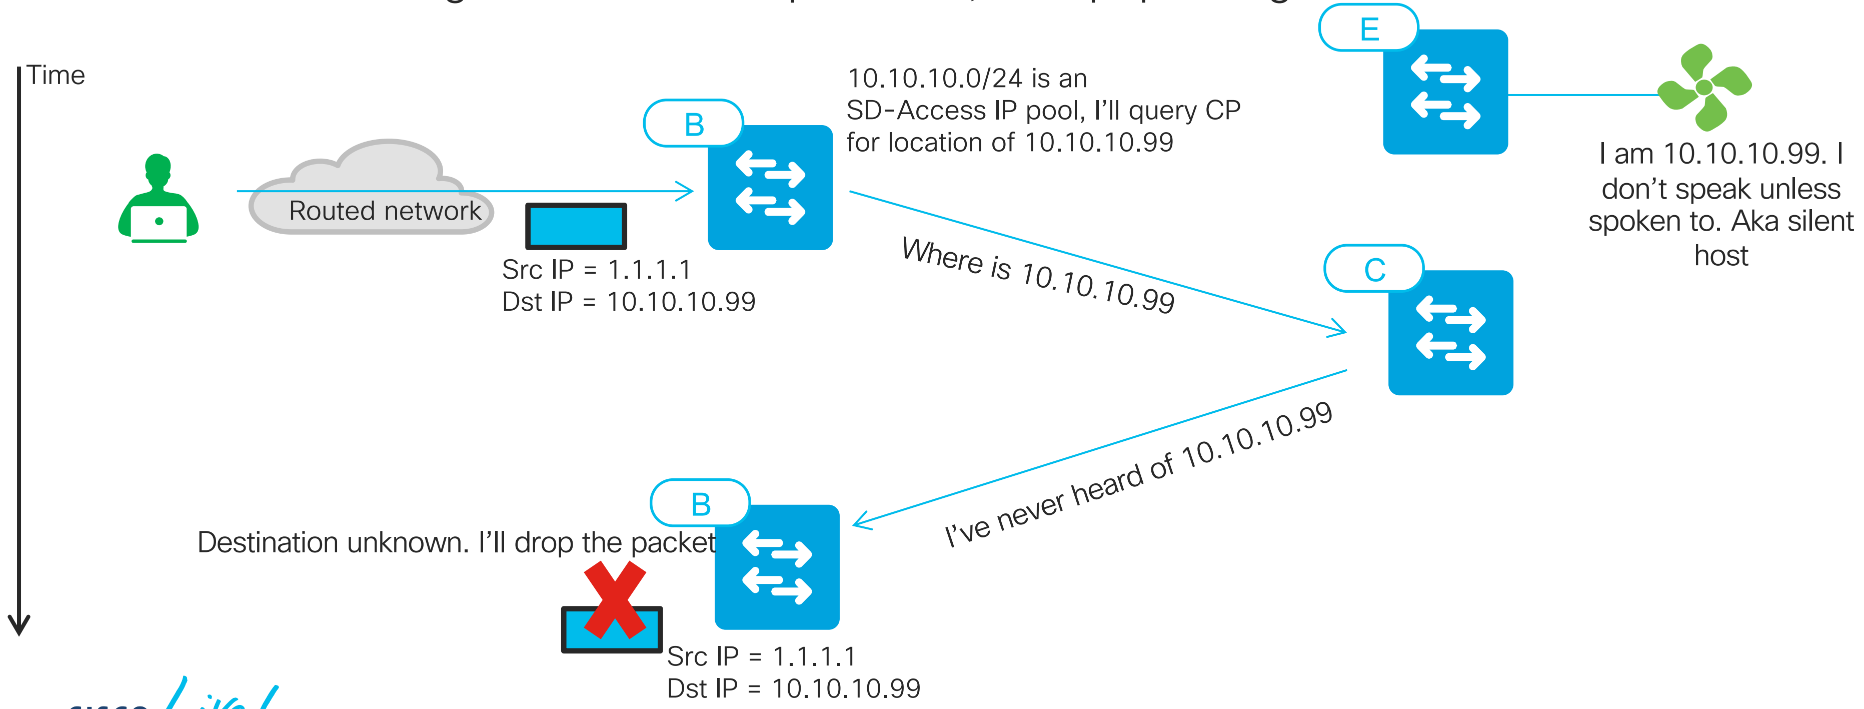

Time
Routed network
B
Src IP = 1.1.1.1
Dst IP = 10.10.10.99
10.10.10.0/24 is an SD-Access IP pool, I'll query CP for location of 10.10.10.99
E
I am 10.10.10.99. I don't speak unless spoken to. Aka silent host
Where is 10.10.10.99
C
I've never heard of 10.10.10.99
B
Destination unknown. I'll drop the packet
Src IP = 1.1.1.1
Dst IP = 10.10.10.99

Underlay Over Brownfields

- Can increase risk. Brownfields routing domain must be stable / reliable
- P2P routed links on all paths between fabric nodes
- All underlay IGPs supported
- All brownfields link MTU must be able to accommodate VXLAN overhead
- If SDA native multicast or L2 flooding is required, then brownfields must reliably support PIM SSM and/or ASM:
  - Pre 1.3.3: 239.0.0.1-239.0.1.246 for L2 flooding – PIM ASM, requires RP
  - 1.3.3+: Single unique PIM ASM multicast group per fabric site, starting with 239.0.17.1
  - Can use Borders as RP for L2 flooding group ranges, to not rely on existing RP for flooding function
  - 232.0.0.1-232.0.3.232 for native multicast – PIM SSM, does not require RP
  - These groups should not be in use in the brownfields network

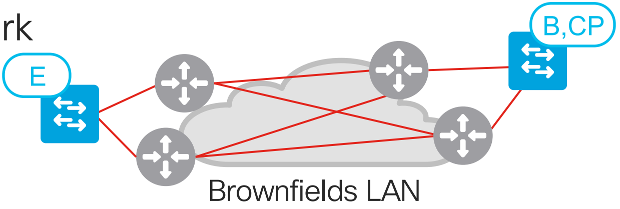

L2 Border – Deployment Model

Same VLAN on two borders not supported

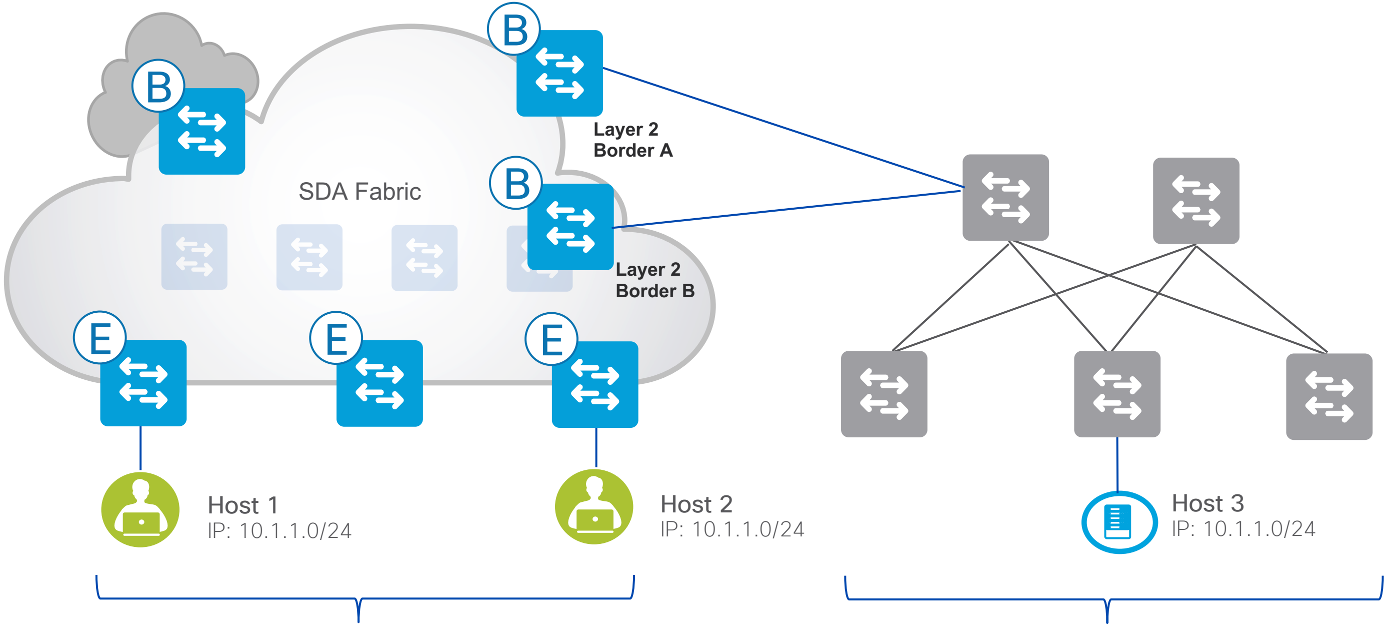

Hosts attached to SDA Fabric
Edge nodes in Address Pool (1024)

Hosts attached to traditional
Access switches in VLAN (300)

L2 Border – Deployment Model

Supported

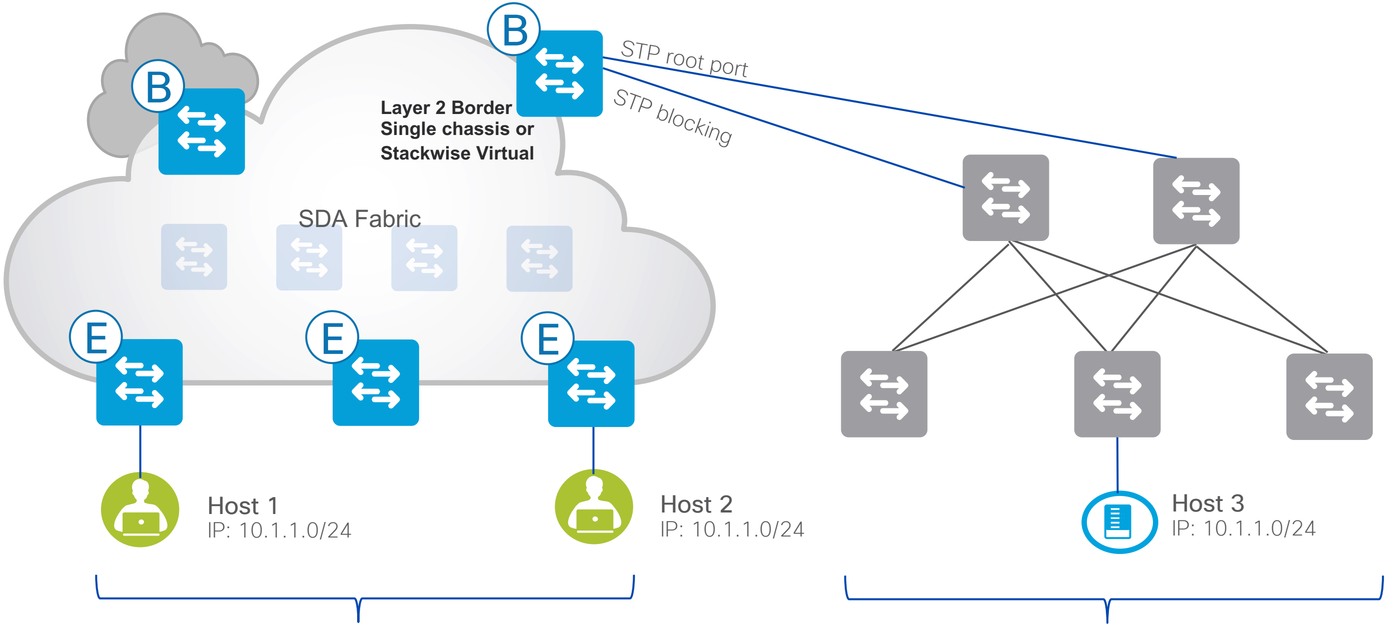

Hosts attached to SDA Fabric
Edge nodes in Address Pool (1024)

Hosts attached to traditional
Access switches in VLAN (300)

L2 Border – Deployment Model

Supported

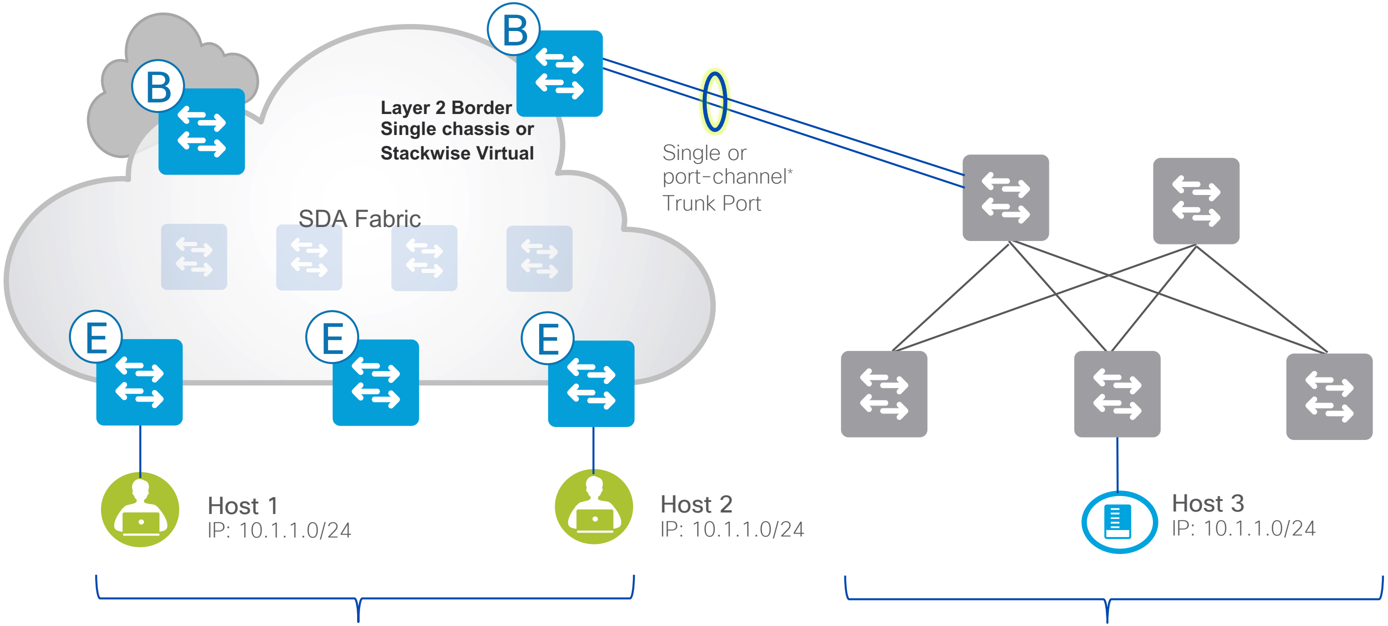

Hosts attached to SDA Fabric Edge nodes in Address Pool (1024)

Hosts attached to traditional Access switches in VLAN (300)

L2 Border – Deployment Model

Stackwise Virtual 9500 and 9500H supported in 1.3.3+ with 16.12.2t+

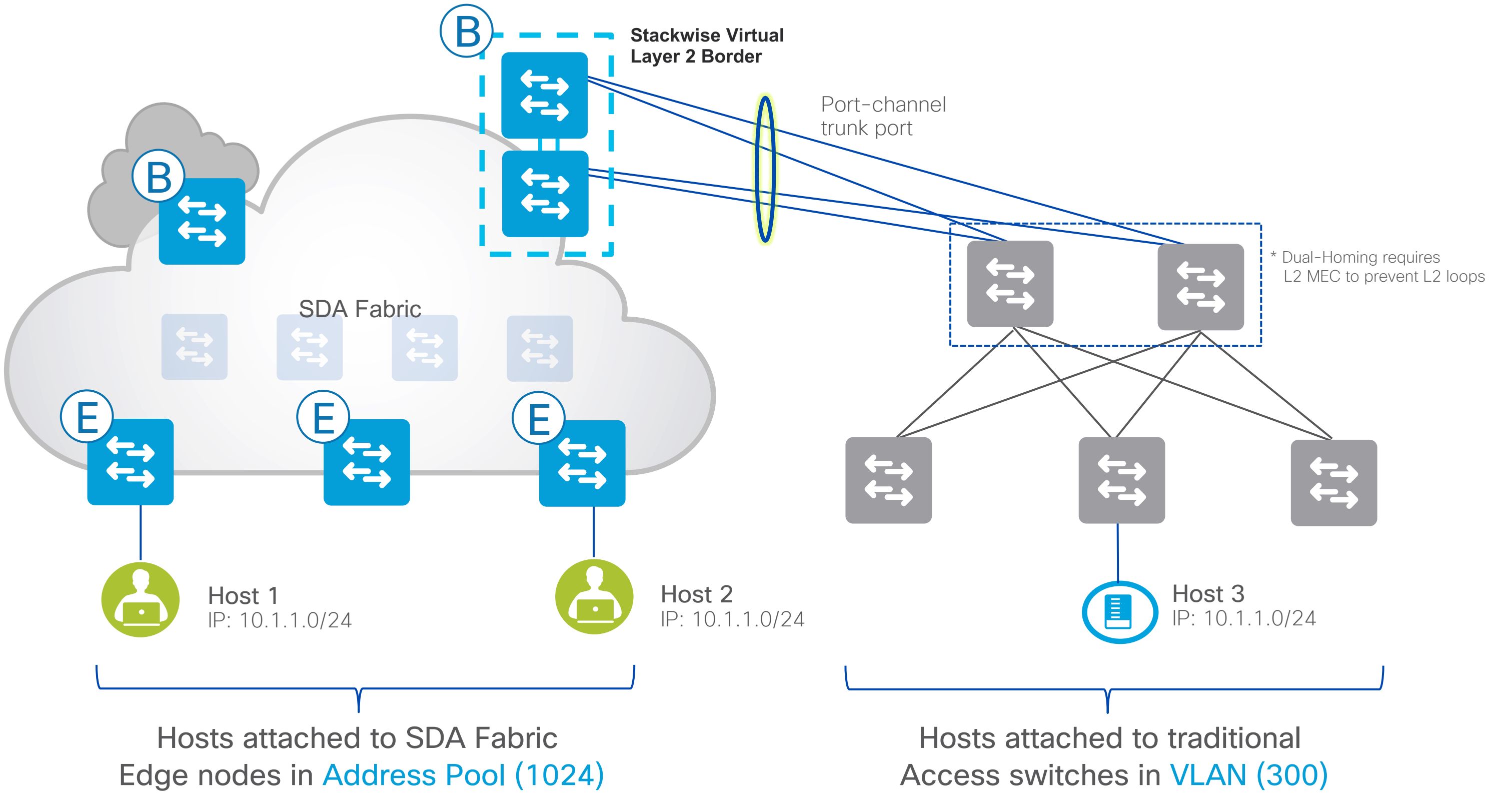

L2 Border for Traditional Network Migration

- Fabric SVIs replaces legacy network d/g for IP range(s) connected to L2 border
  - IP subnets will be advertised from SDA borders
- Supported on C3K, C6K and C9K
- Max EID supported depends on platform
- Caution. L2 storm in legacy network = L2 storm in fabric = Potential border outage
  - Recommend using dedicated L2 border, reduce risk
- VLAN IDs in traditional network cannot overlap with Cisco DNA Center/fabric reserved VLANs
  - Reserved VLANs:

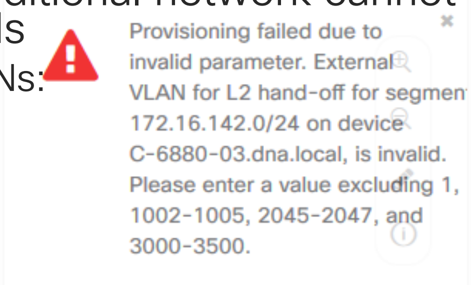

Provisioning failed due to invalid parameter. External VLAN for L2 hand-off for segmen 172.16.142.0/24 on device C-6880-03.dna.local, is invalid. Please enter a value excluding 1, 1002-1005, 2045-2047, and 3000-3500.

  - Manual VLAN translation is solution, on non-fabric switch, not on L2 border
    - Not supported on C3K
    - In C9K family requires 9500H or 9600

L2 Border for Gateway Outside of Fabric

- Always try and use anycast SVI on fabric edge switches instead of external default gateway
  - FE SVI = routing efficiency, no hairpin on external gateway
  - FE SVI = E-W SGT more easily preserved
- If not possible then solve requirement using L2 border. Please consult with your technical presales team on suitability and design before proceeding
- Proper / final DNAC workflow to solve this scenario is roadmap

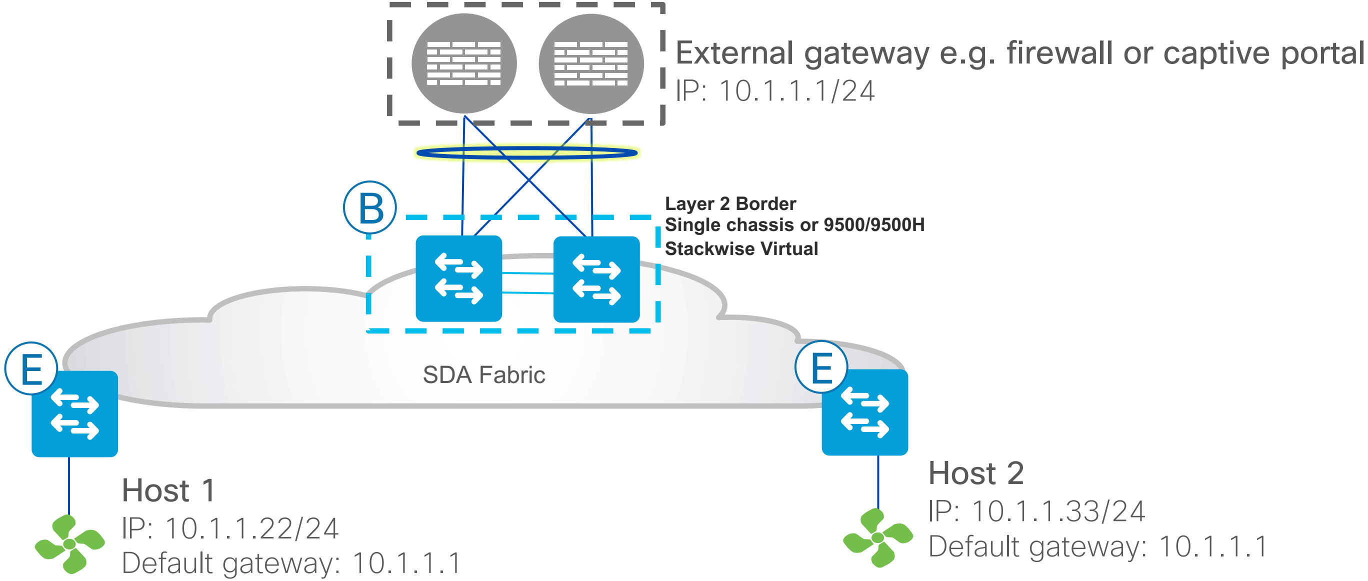
External gateway e.g. firewall or captive portal
IP: 10.1.1.1/24
B
Layer 2 Border
Single chassis or 9500/9500H
Stackwise Virtual
SDA Fabric
E
E
Host 1
IP: 10.1.1.22/24
Default gateway: 10.1.1.1
Host 2
IP: 10.1.1.33/24
Default gateway: 10.1.1.1

Multicast in Fabric

RECAP

- Multicast can be enabled on a per-VN basis through Cisco DNA Center workflow
  - PIM-ASM or SSM can be running in the overlay
  - ASM uses a fabric RP. <1.3.3 RP must be in fabric. 1.3.3+ RP can be inside or outside of fabric
- Replication
  - Head end: RP or ingress fabric node replicates multicast into unicast VXLAN tunnels to receivers. Not scalable. Low bandwidth only. Negates need for multicast forwarding in underlay
  - Native: Reduced latency and increased scale. Uses underlay to do replication
    - Incoming Multicast traffic for a given VN is encapsulated in VXLAN, and then sent with {Source IP = ingress node RLOC, Destination IP = Underlay Multicast Group} as the outer IP addresses
    - PIM SSM is used in the underlay for multicast transport
    - The state created in the underlay will include (S, G) entries for all the groups , in essence (RLOC,G) entries
    - Underlay ranges 232.0.0.1-232.0.3.232 for native multicast
    - Underlay scale of 1000 (RLOC,G) for all overlay VNs

Multicast with Fabric RP + RP Outside the Fabric

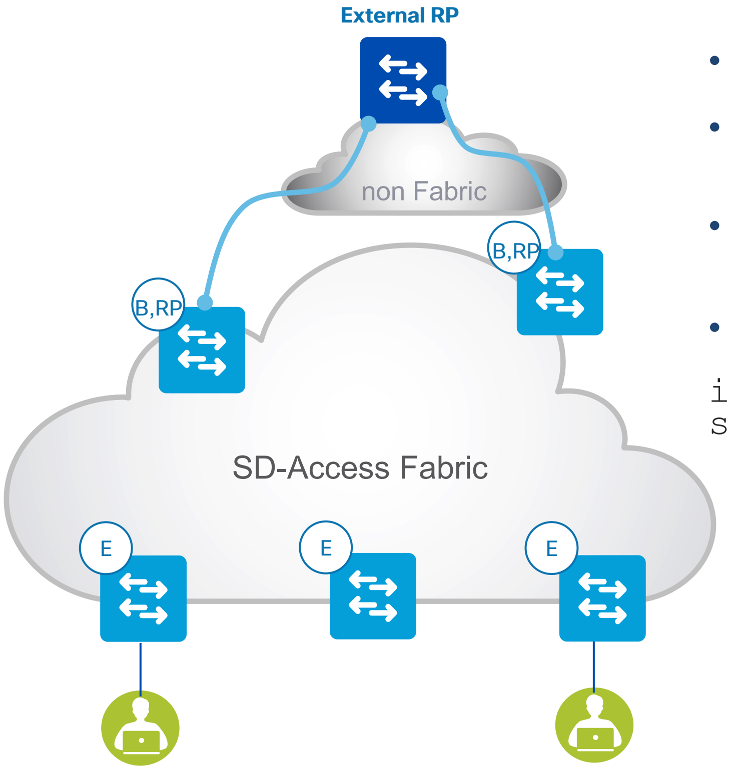

- Prior to Cisco DNA Center 1.3.3 RP per VN is mandatory
- Cisco DNA Centre today automates the MSDP peering between the Fabric RPs
- As of 1.2.5 we support peering with external RPs via template editor
- Command to add via template editor to fabric RPs :

```
ip msdp vrf <vrf_name> peer External_RP_ip connect-source Loopback<Lo_created_for_multicast_VN>
```

Multicast with RP Outside the Fabric

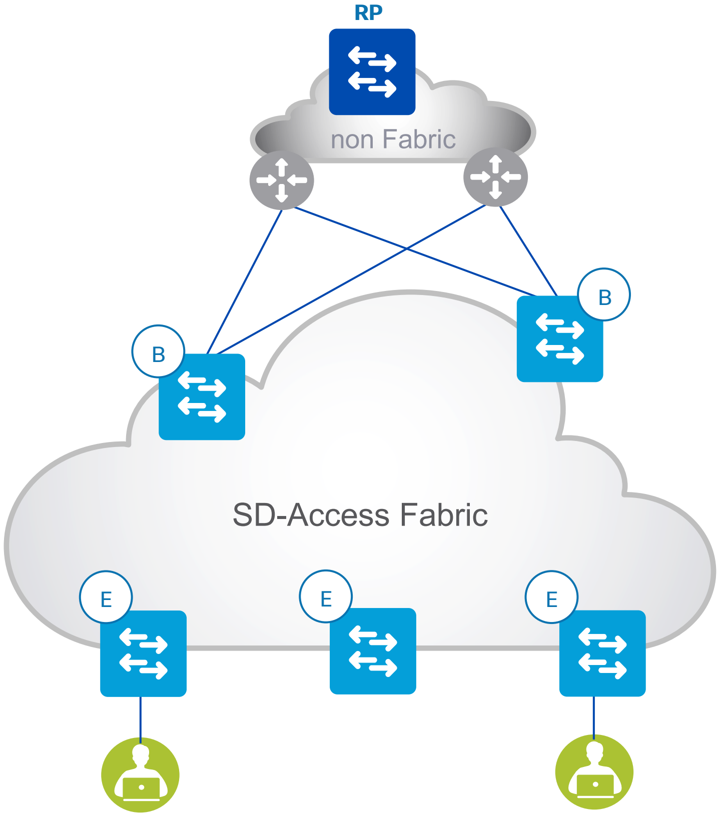

- As of Cisco DNA Center 1.3.3
  - All previous multicast scenarios are supported, and
  - VN RP can exist solely outside of fabric. Per-VN RP in fabric is no longer mandatory

Converting Brownfields Access Switch to FE

Rebuild the switch:

1. IOS-XE version complies with SD-Access compatibility matrix
2. License level / subscription level sufficient
3. Factory reset the switch as per LAN Automation deployment guide (wr erase insufficient)

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/tech_notes/b_dnac_sda_lan_automation_deployment.html

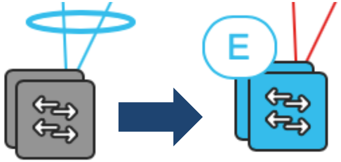

LAN automation

OR

Manual conversion

4. Execute LAN automation (ISIS IGP)
5. Provision, add to fabric site as FE
6. Provision edge ports in host onboarding

4. Replace startup configuration with tailored startup configuration and reload the switch:
   - Routed p2p uplinks, Loopback0
   - MTU accommodates VXLAN overhead
   - Multicast routing and PIN, if required
   - SSH login creds, SNMP
5. Modify distribution layer to have routed downlinks, or, repatch switch to new distribution
6. Discover just-reloaded switch in Cisco DNA Center (SSH,NETCONF,SNMP)
7. Provision, add to fabric site as FE
8. Provision edge ports in host onboarding, if required

~40K ports migrated

Multiple Cisco SD-Access sites

Border Nodes – One Box and Two Box

RECAP

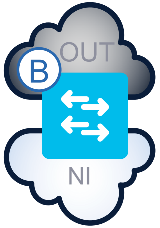

One box

- Internal and external domain routing is on the same device
- Simple design, without any extra configurations between the border and outside routers
- The Border device will advertise routes to and from the local fabric domain to the external domain

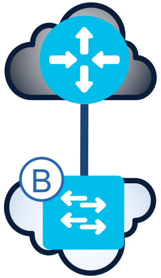

Two box

- Internal and external domain routing are on different devices
- Requires two devices with BGP in between to exchange connectivity and reachability information
- This model is chosen if border does not support some functionality (This can due to hardware, or software, or solution validation) to run the external domain on the same device e.g. border switch cannot terminate IPsec

Multi-site Transits

- 1-box + DMVPN is supported
- 1-box + IPsec is supported

Requires suitable router!

- 1-box + MPLS PE not supported
- 1-box + manual GRE tunnels not supported

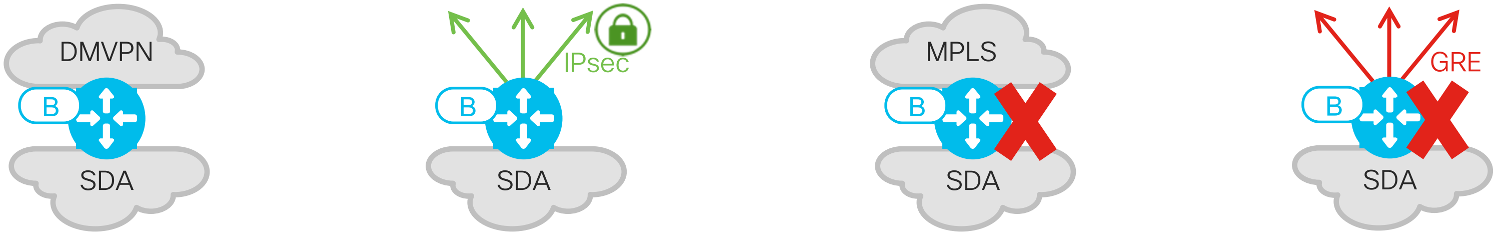

- 1-box warnings:
  - Cannot create DMVPN instance or static VTI until after Cisco DNA Center automation creates VRF/VN. Watch for race condition
  - Must run IPsec/DMVPN on router code versions from SD-Access combability matrix

Multi-site Transits

- SDA Transit
  - No 2-box option
  - 1-box only
    - Cisco DNA Center automated
    - SGTs and VNs are preserved
    - Cannot tolerate tunnelled frame MTU exceeding transit MTU
      - TCP Adjust MSS on FEs can solve TCP packet sizes, but not large UDP
    - No possibility of service insertion e.g. firewall / IPS / content accelerator

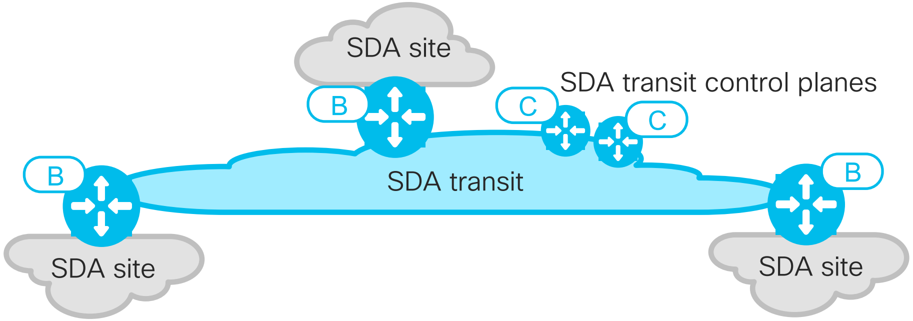

Multicast over SDA Transit

- Not supported if RPs are located within the SDA fabric sites. This is roadmap
- Supported today if all of these are true:
  1. VN at all fabric sites has same external RP IP address (min DNAC 1.3.3)
  2. External RP is reachable in VN on all borders connected to SDA Transit
  3. All fabric sites are set to either head-end replication or native multicast, mix of head-end and native not possible
- If all sites use native multicast, then underlay between fabric sites must support PIM-SSM
- All the usual SDA Transit design rules apply too e.g. no service insertion, jumbo MTU, Transit CPs mandatory, etc.

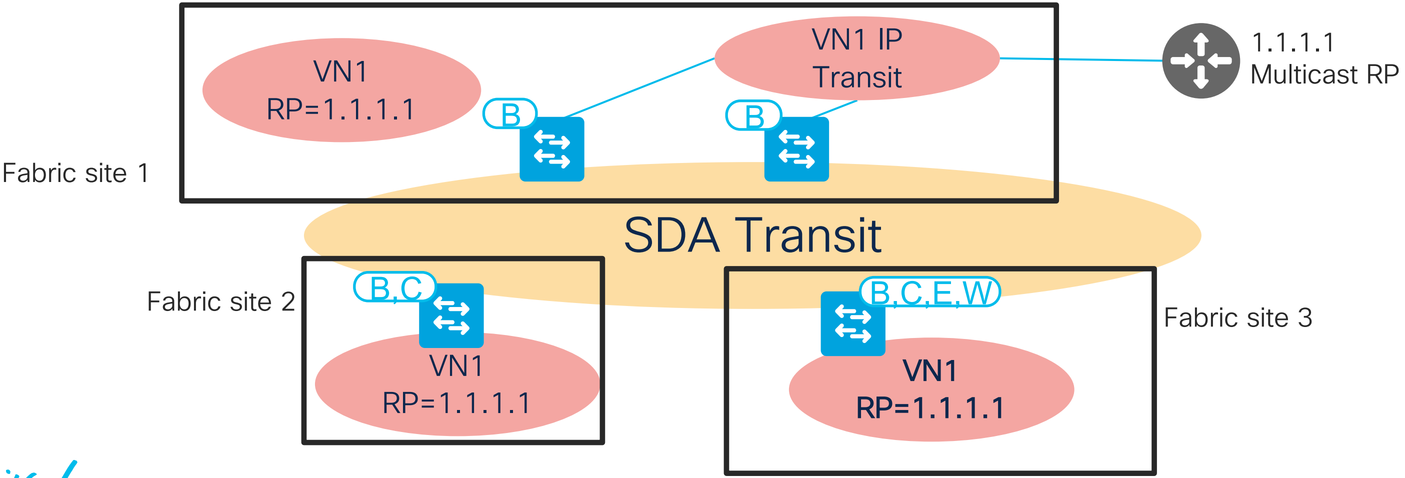

Multi-site Transits

- GRE tunnels, 2-box
  - Manual configuration on dedicated GRE routers
  - VNs can be preserved: one GRE tunnel per SDA VN/VRF
  - Must be router for SGT inline. SGTs are preserved in CMD in GRE
  - Tolerates overlay tunnelled frame MTU + GRE overhead exceeding transit MTU
    - TCP Adjust MSS can solve TCP packet sizes
    - ‘ip mtu’ on tunnel interface fragment packets before GRE encap, or sends ICMP T3C4

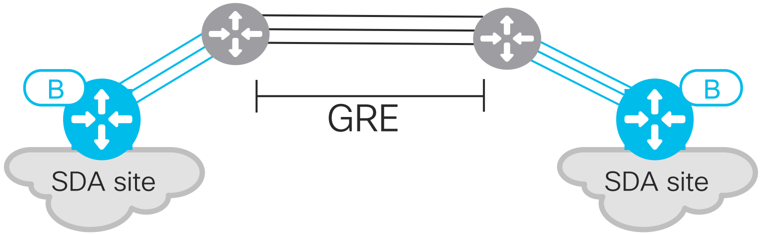

Multi-site Transits

- DMVPN, 2-box
  - Requires router that supports DMVPN
  - Manual configuration on DMVPN routers
  - SGTs can be preserved in CMD in GRE, or in CMD in IPsec
  - VNs can be preserved: one DMVPN cloud per SDA VN/VRF
  - Can handle WAN MTU limitations:
    - TCP adjust MSS can solve TCP packet sizes
    - ‘ip mtu’ on tunnel interface fragment packets before GRE encap, or sends ICMP T3C4

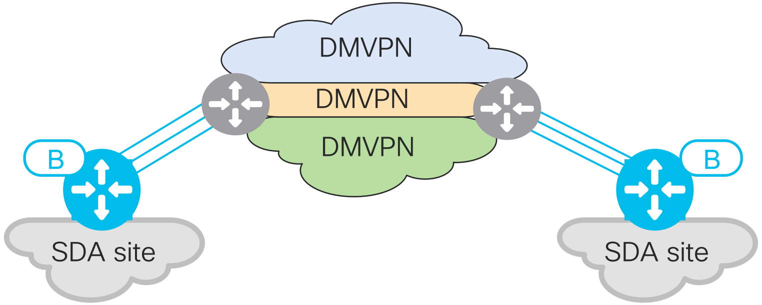

Multi-site Transits

- GRE in IPsec, 2-box
  - Requires suitable router
  - Manual configuration on GRE/IPsec routers
  - SGTs preserved in CMD in GRE
  - VNs can be preserved: one tunnel per SDA VN/VRF
  - Can tolerate tunnelled frame MTU exceeding transit MTU
    - TCP adjust MSS can solve TCP packet sizes
    - ‘ip mtu’ on tunnel interface fragment packets before GRE encap, or sends ICMP T3C4

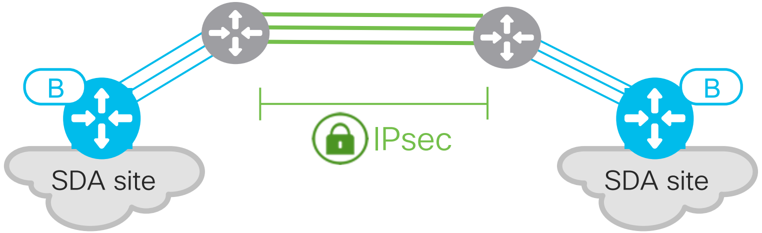

Multi-site Transits

- IP Transit
  - Border-side can be automated by Cisco DNA Center
  - SGTs can be preserved in Ethernet CMD
    - If SGT in Ethernet CMD then all external transit nodes hardware/software to be SGT capable
  - Or, SGTs can be preserved in SXP between borders and ISE
    - SXP is not VRF aware, must peer SXP per VRF from border to ISE
      - Can cause scaling problems / exceed platform memory. Recommend avoiding if possible
      - If SXP must be used, be very clear on all platform scaling limits. Don't assume. SXP table can be large and exceed a memory limit of lower resourced devices
    - If static FE ports SGT is lost. ISE doesn't know about static ports
  - VNs can be preserved: IP transit network needs to support VRFs
  - IP transit network should fragment and send ICMP T3C4 as required

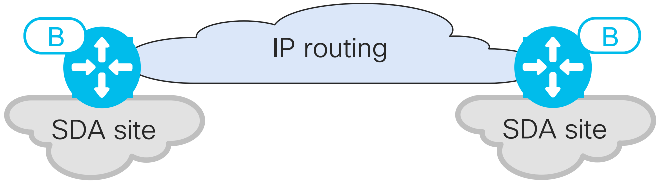

Multi-site Transits

- No roadmap used in this presentation, but, for completeness:
  - There is a plan to add SD-WAN as a Cisco DNA Center automated transit option for SDA multi-site scenarios
    - SDA and SD-WAN interworking will offer the benefits of SD-WAN (application aware routing, aggregation of multi transports, crypto, hierarchical WAN topologies) and VN + SGT preservation between SDA fabric sites
    - More details about SDA and SD-WAN interworking can be found in the most recent BRKCRS-2818

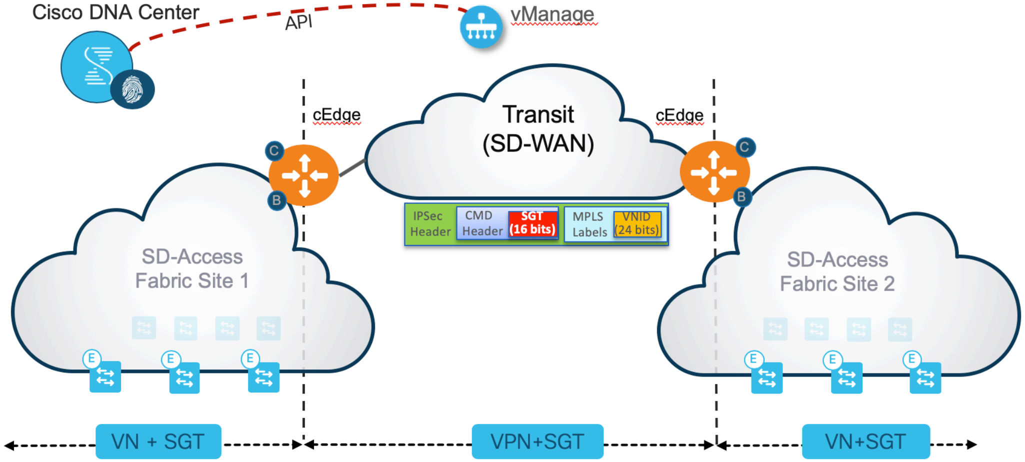

Paper migration scenario

Migration Example

Requirements:

- Move traditional network to SD-Access fabric
  - Will be a large fabric site at project completion. 400x FEs, 80x IP pools, 16x VNs, 2000x fabric APs (future), 40K endpoints on fabric
- BMS devices cannot change IP address
- BMS devices require Ethernet broadcast
- Corporate endpoints consume 3 ASM groups. Source in brownfields DC
- No cabling changes to existing brownfields network
- OTT wireless for now, future project for fabric wireless
- Corporate and IOT macro-segmentation
- Granular stateful AVC and guaranteed logging of inter-VN traffic
- Connect Corporate VN to second already-established fabric site with data privacy, VN and SGT preservation over 1500B MTU WAN
- Qualitative application metrics required for Application Experience in Assurance for traffic between fabric sites

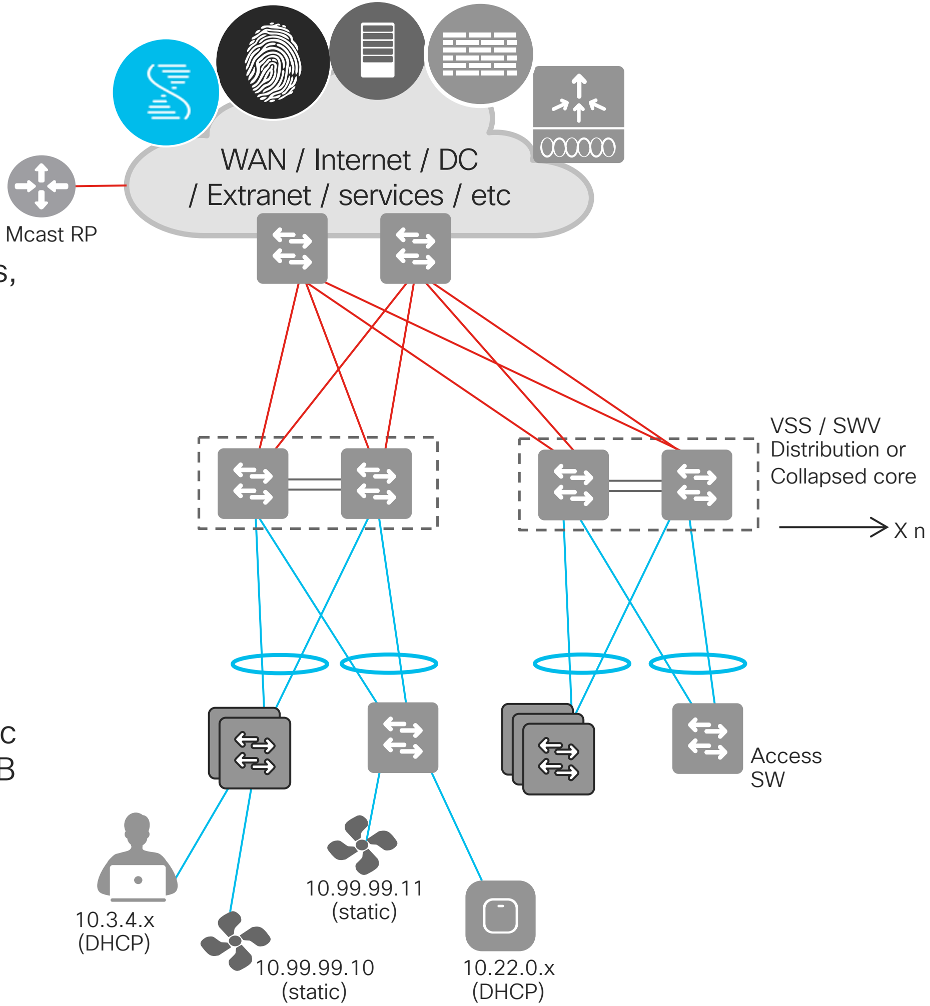

Migration Example - Logical

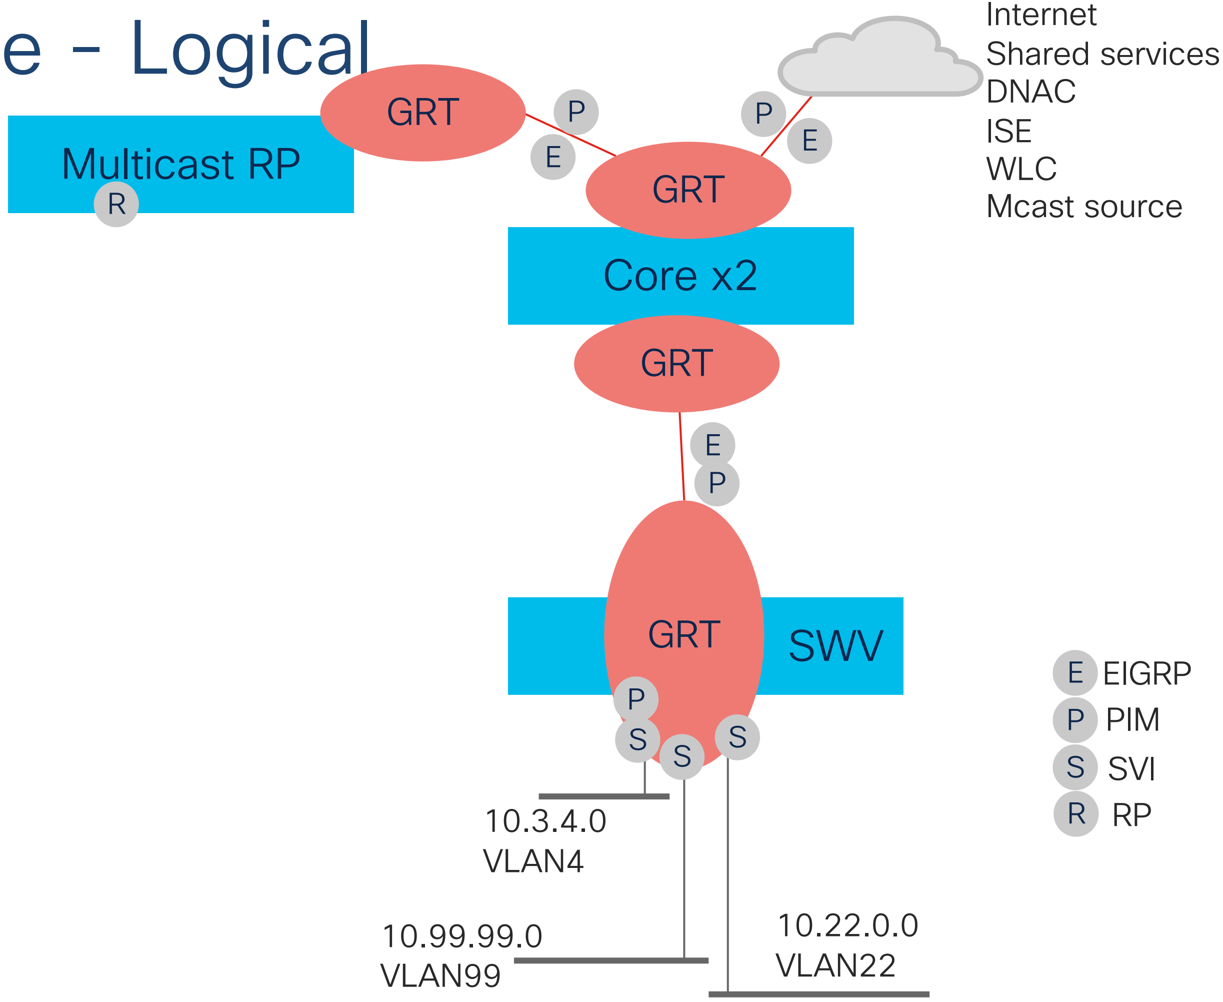

Migration Example - Physical

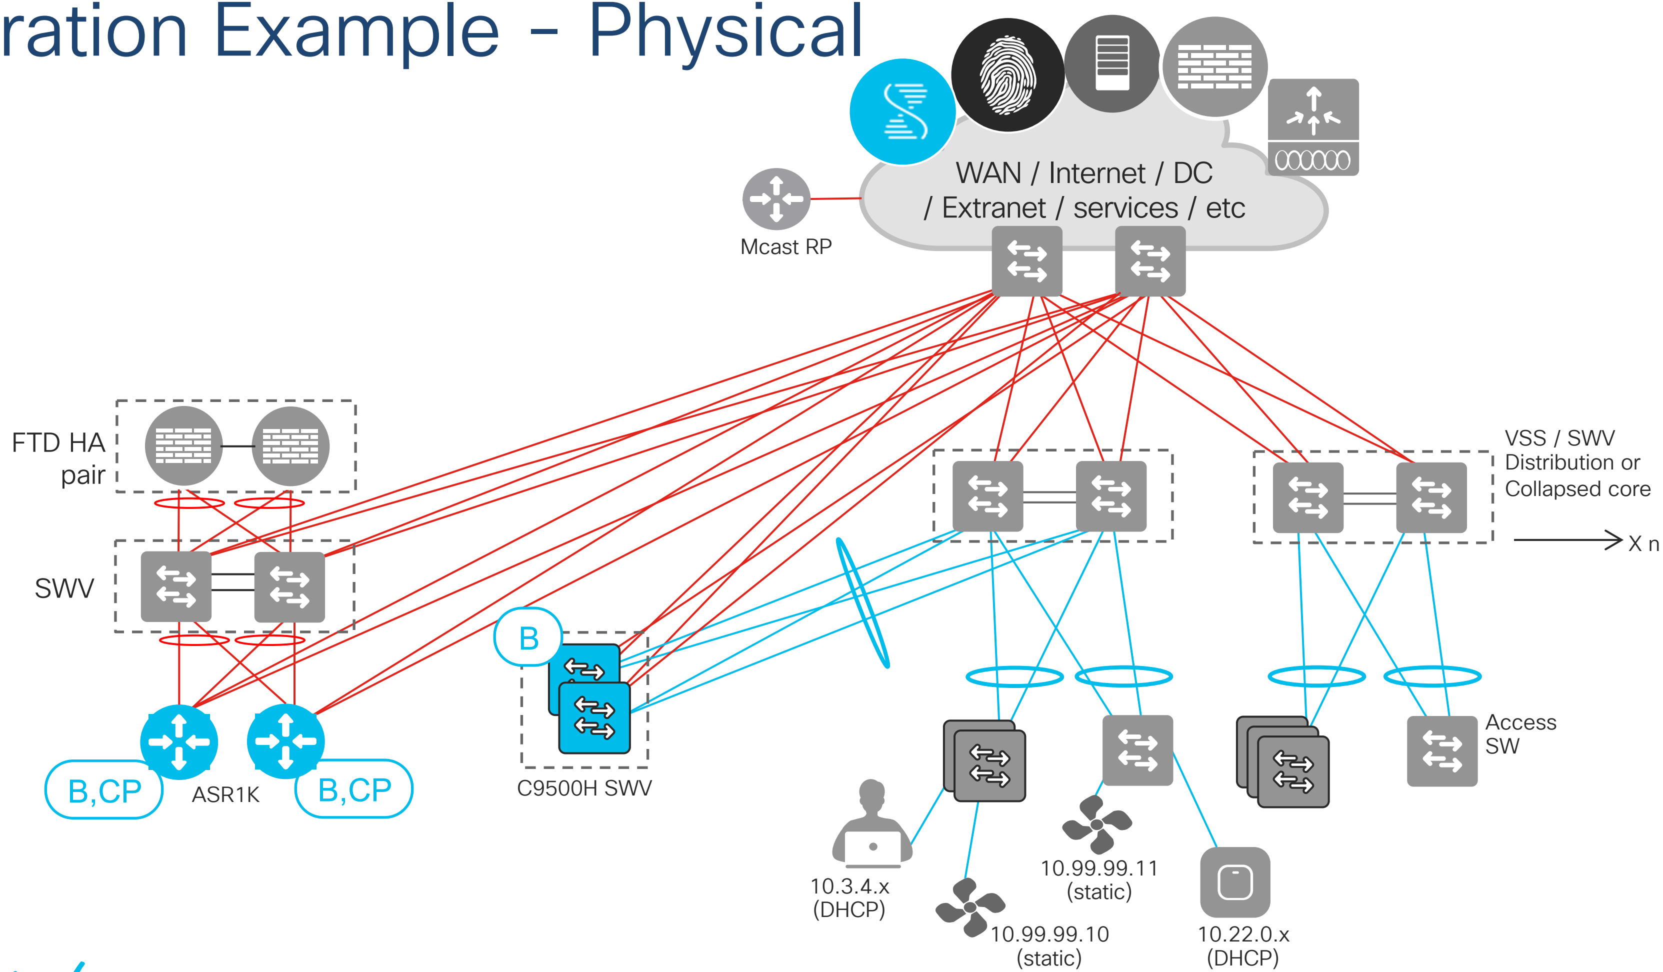

Migration Example - Logical

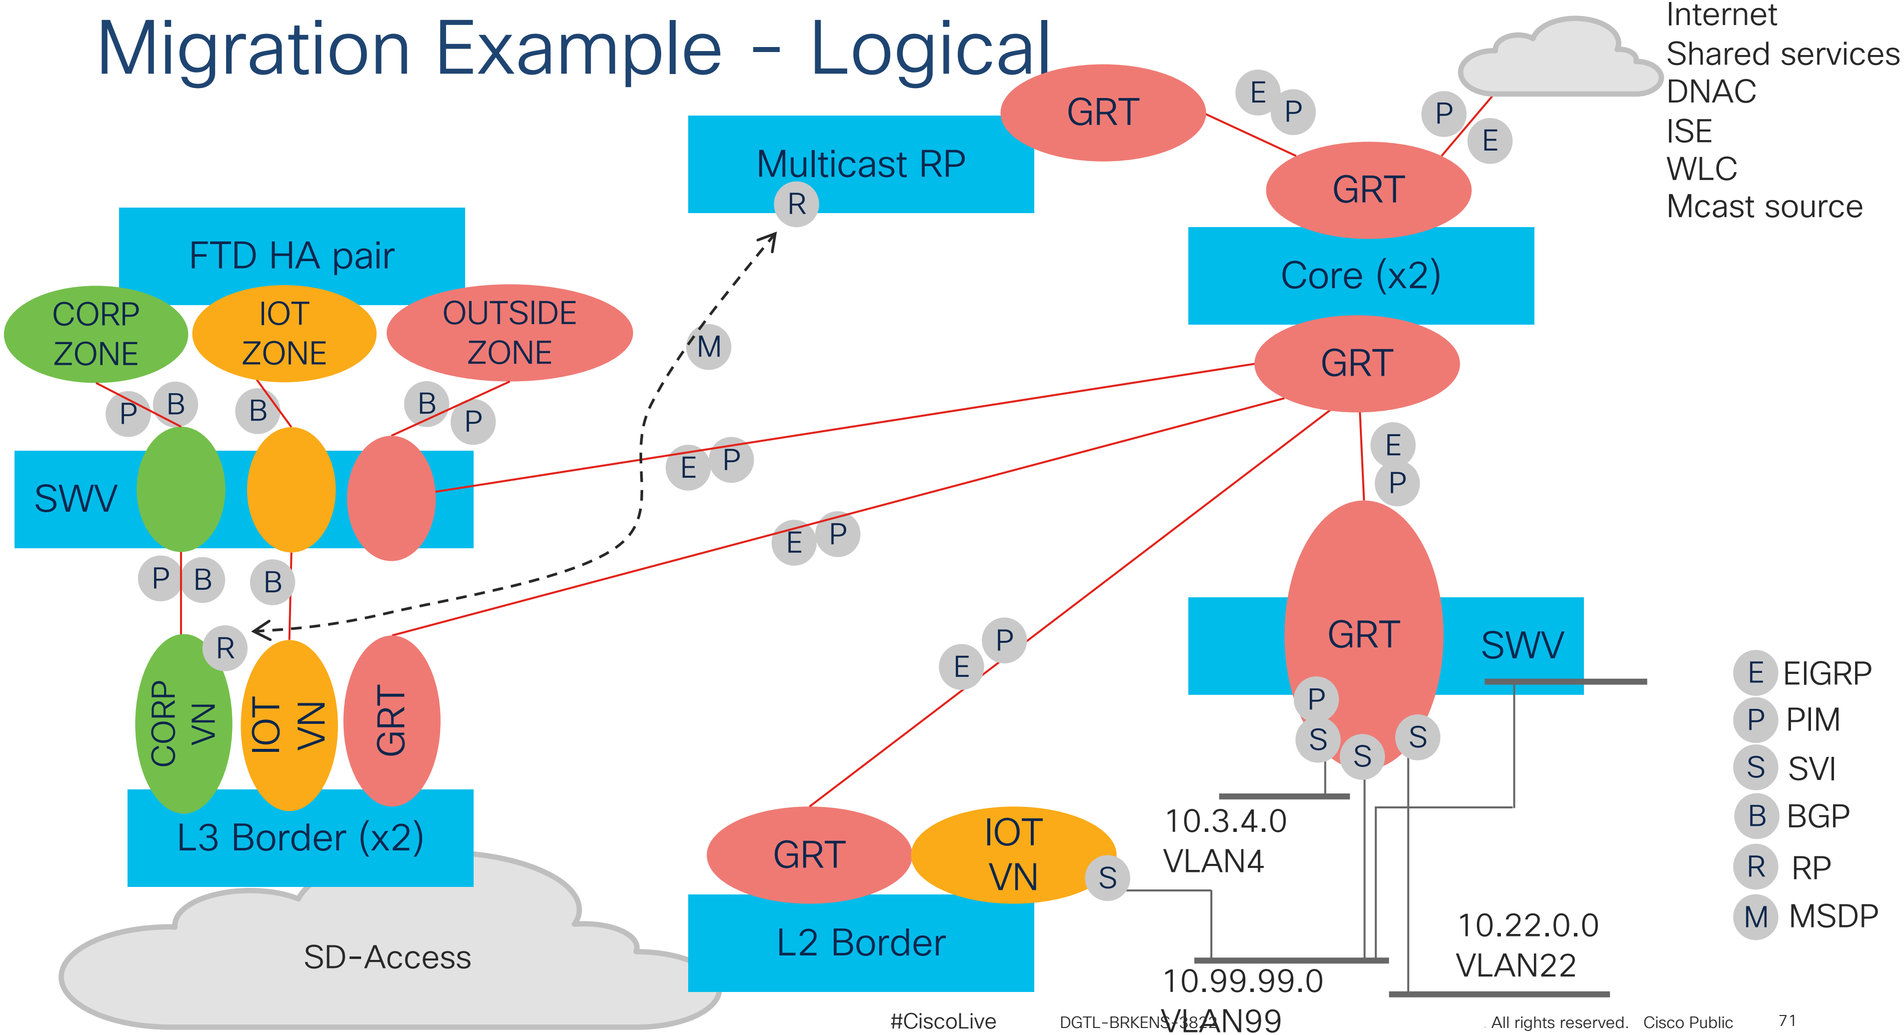

Migration Example

1. Connect L2 border to legacy L2 domain, shut d/g in legacy domain and create d/g on L2 border (as per logical slide)
2. Confirm correct licenses and SDA certified IOS-XE version on access stack
3. Replace startup config on access stack and reload
4. Convert SWV downlinks to P2P L3
5. Discover and provision as FE
6. Assign new FE ports to VNs/Pools

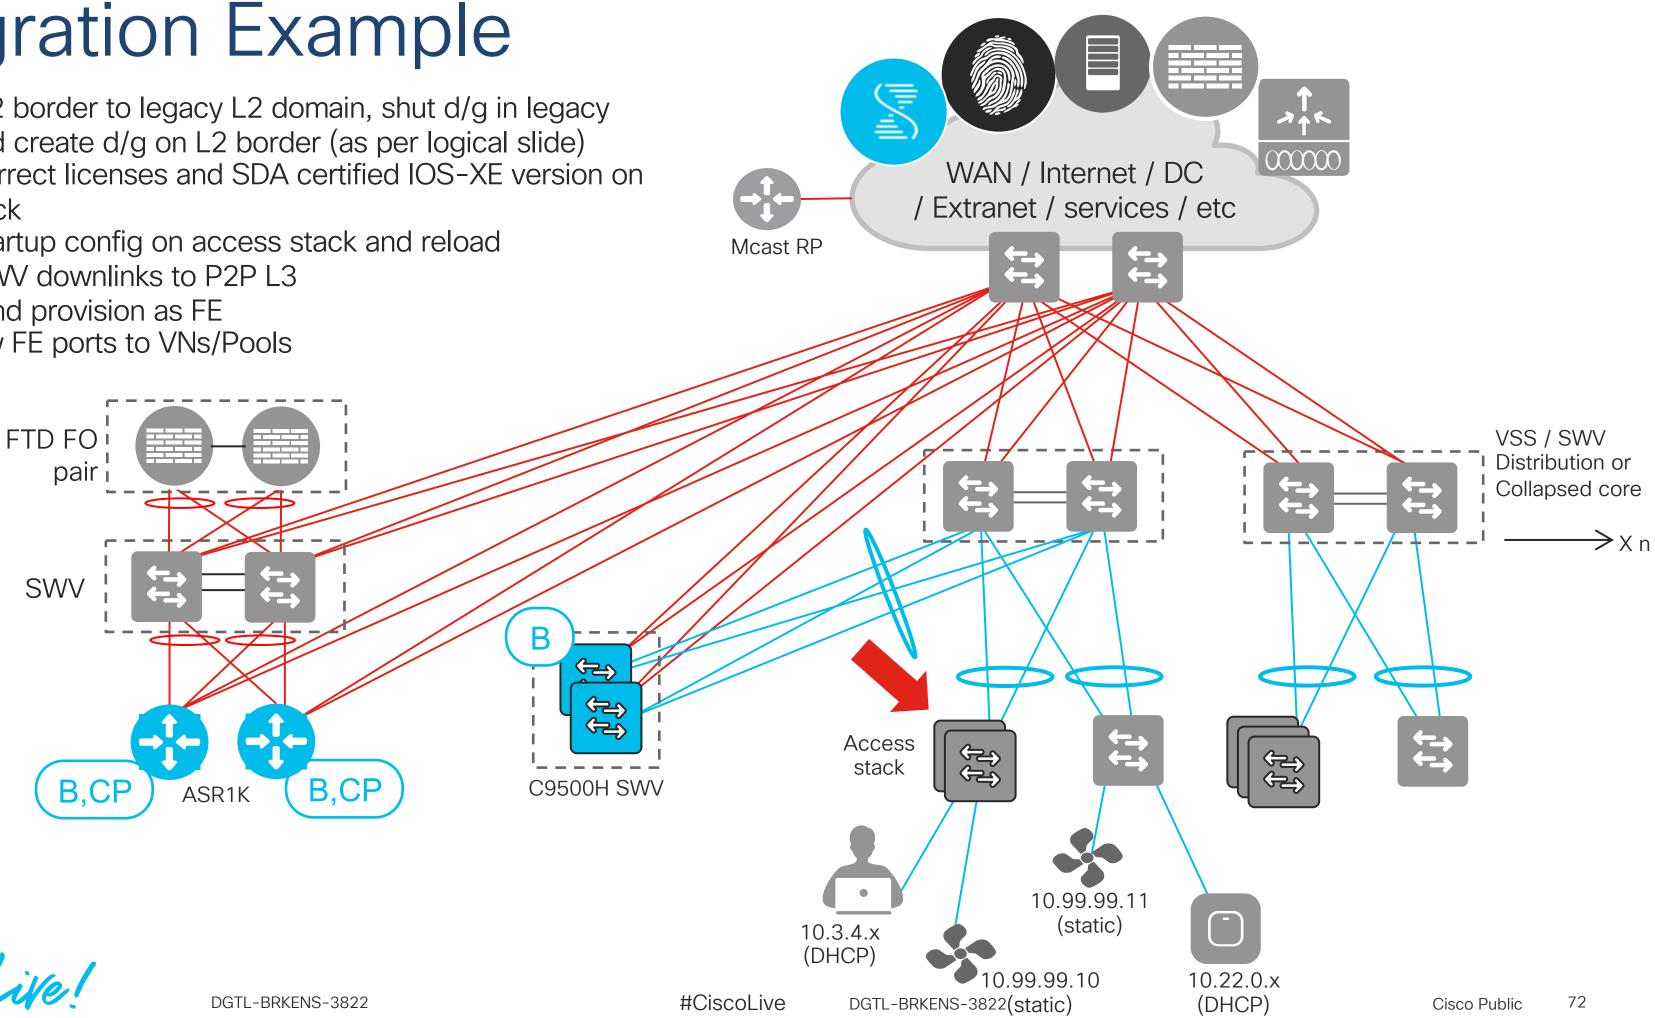

Migration Example

7. Confirm correct licenses and SDA certified IOS-XE version on access stack
8. Replace startup config on access stack and reload
9. Convert SWV downlinks to P2P L3
10. Discover and provision as FE
11. Assign new FE ports to VNs/Pools

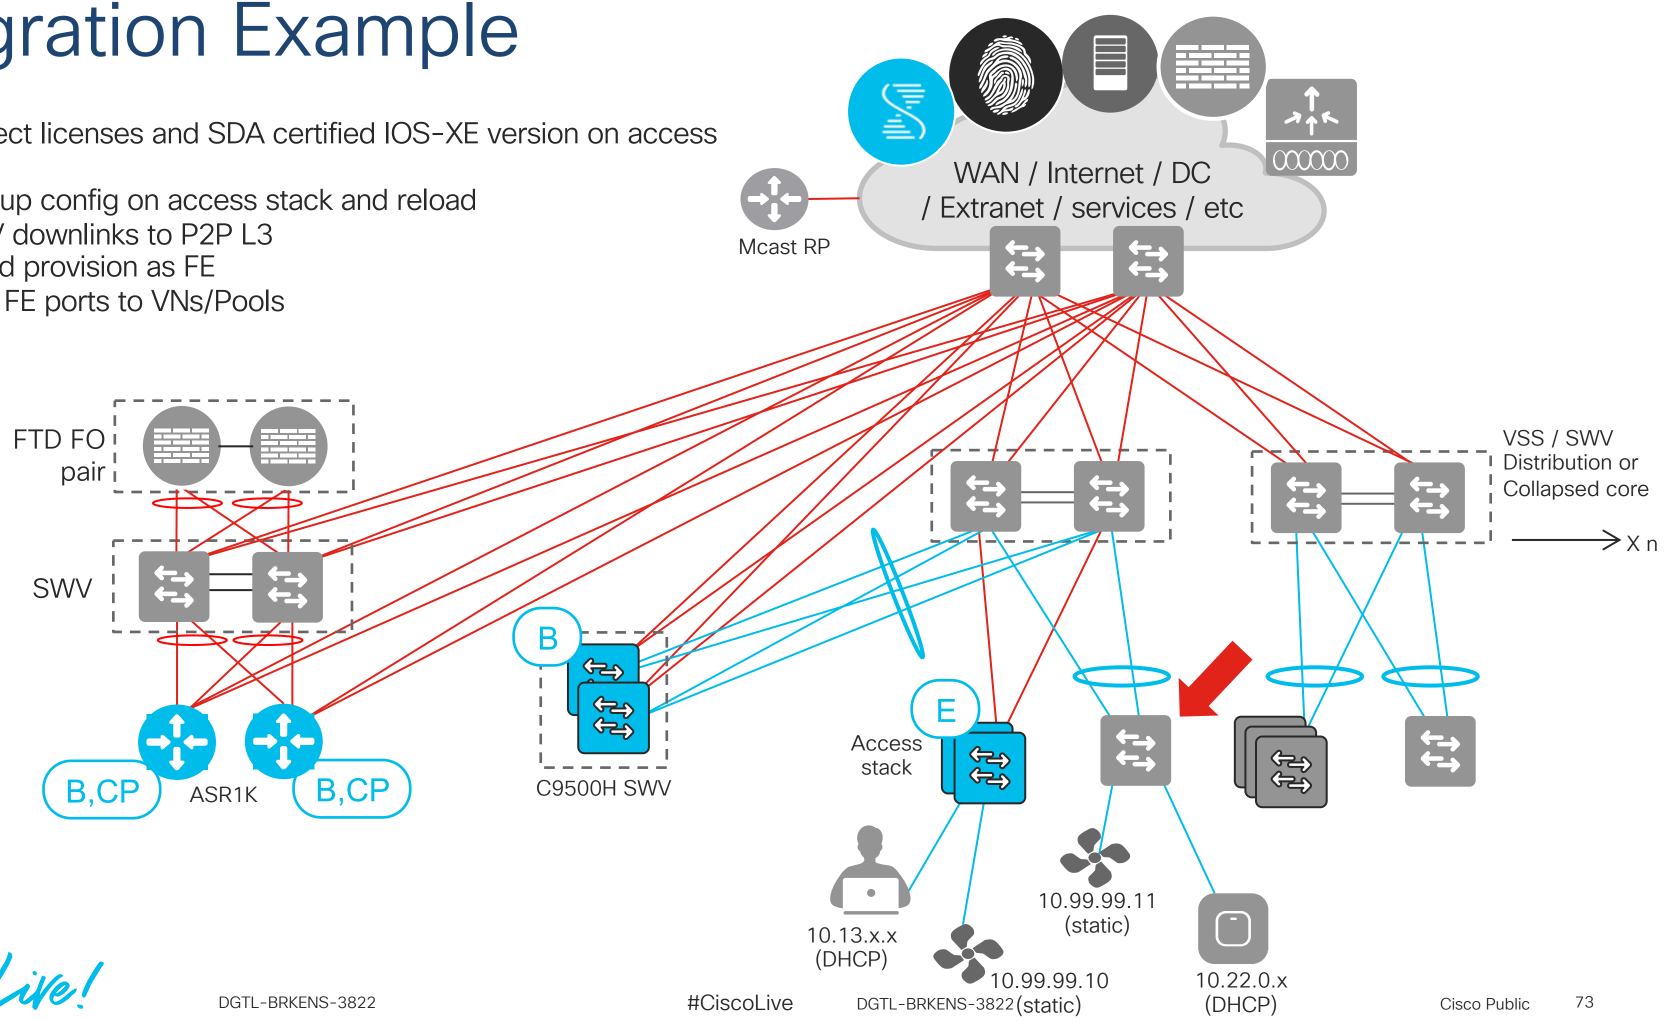

Migration Example

12. Remove L2 handoff from L2 border
13. Disconnect L2 border from legacy L2 domain

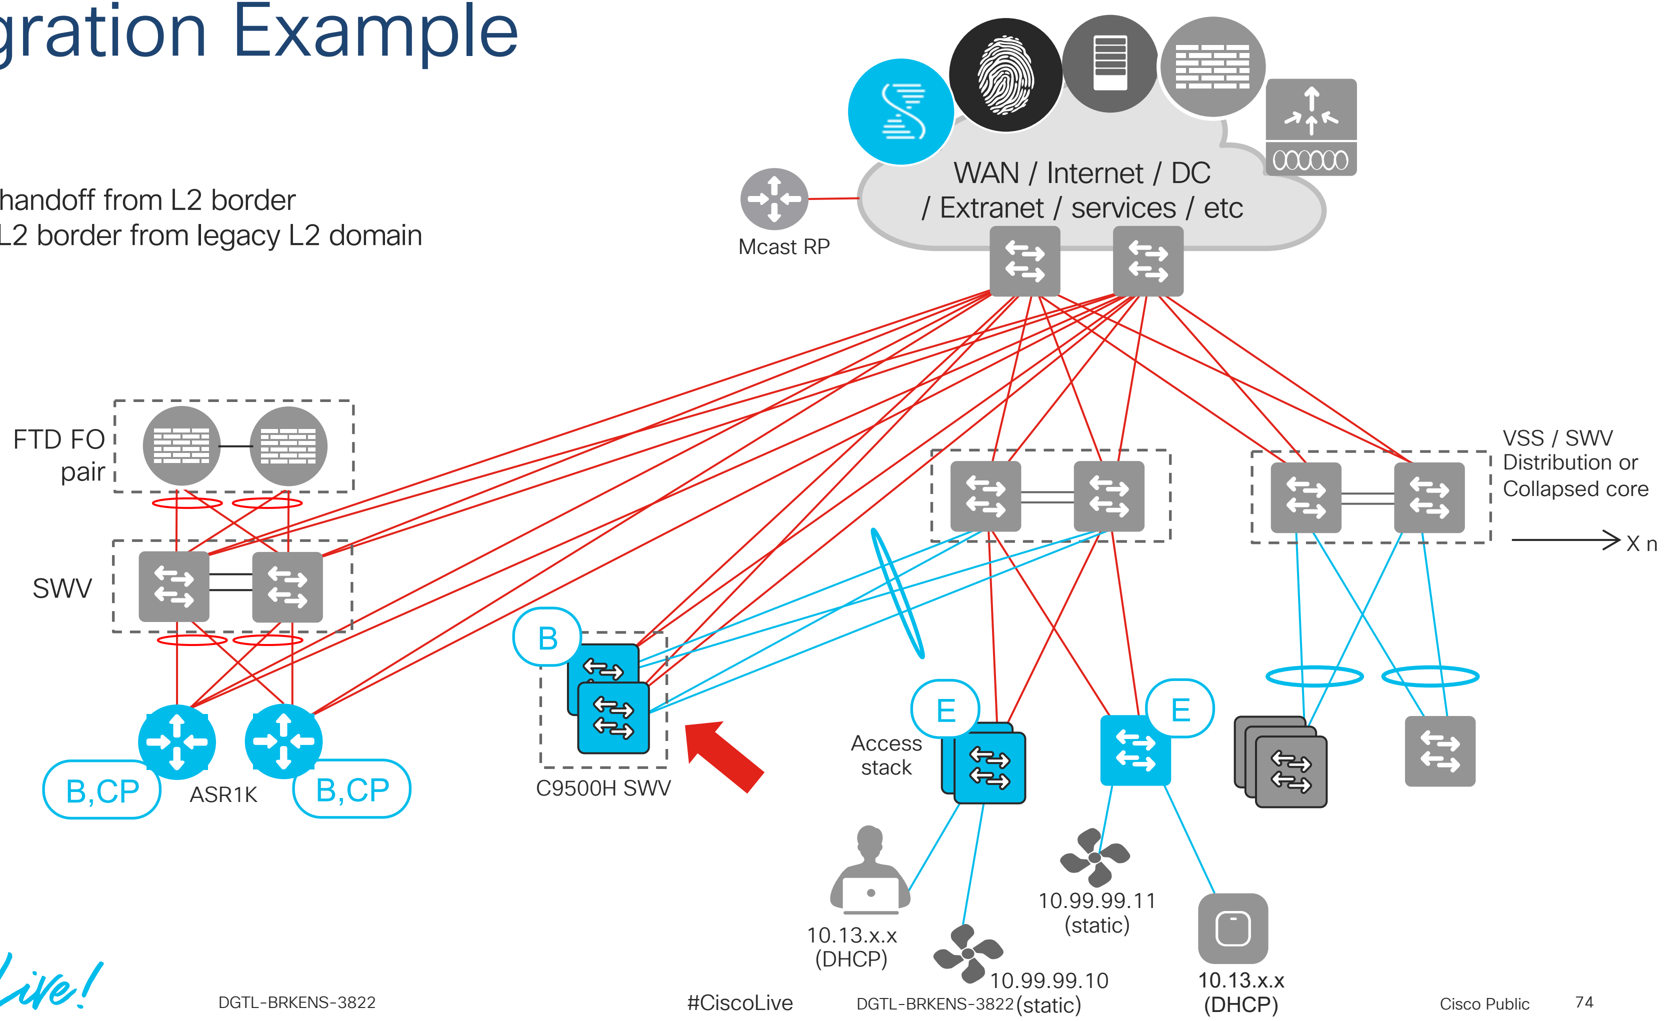

Migration Example

(For completeness: Repeat procedure per L2 domain, or if ambitious, convert entire L2 domain in one change window)

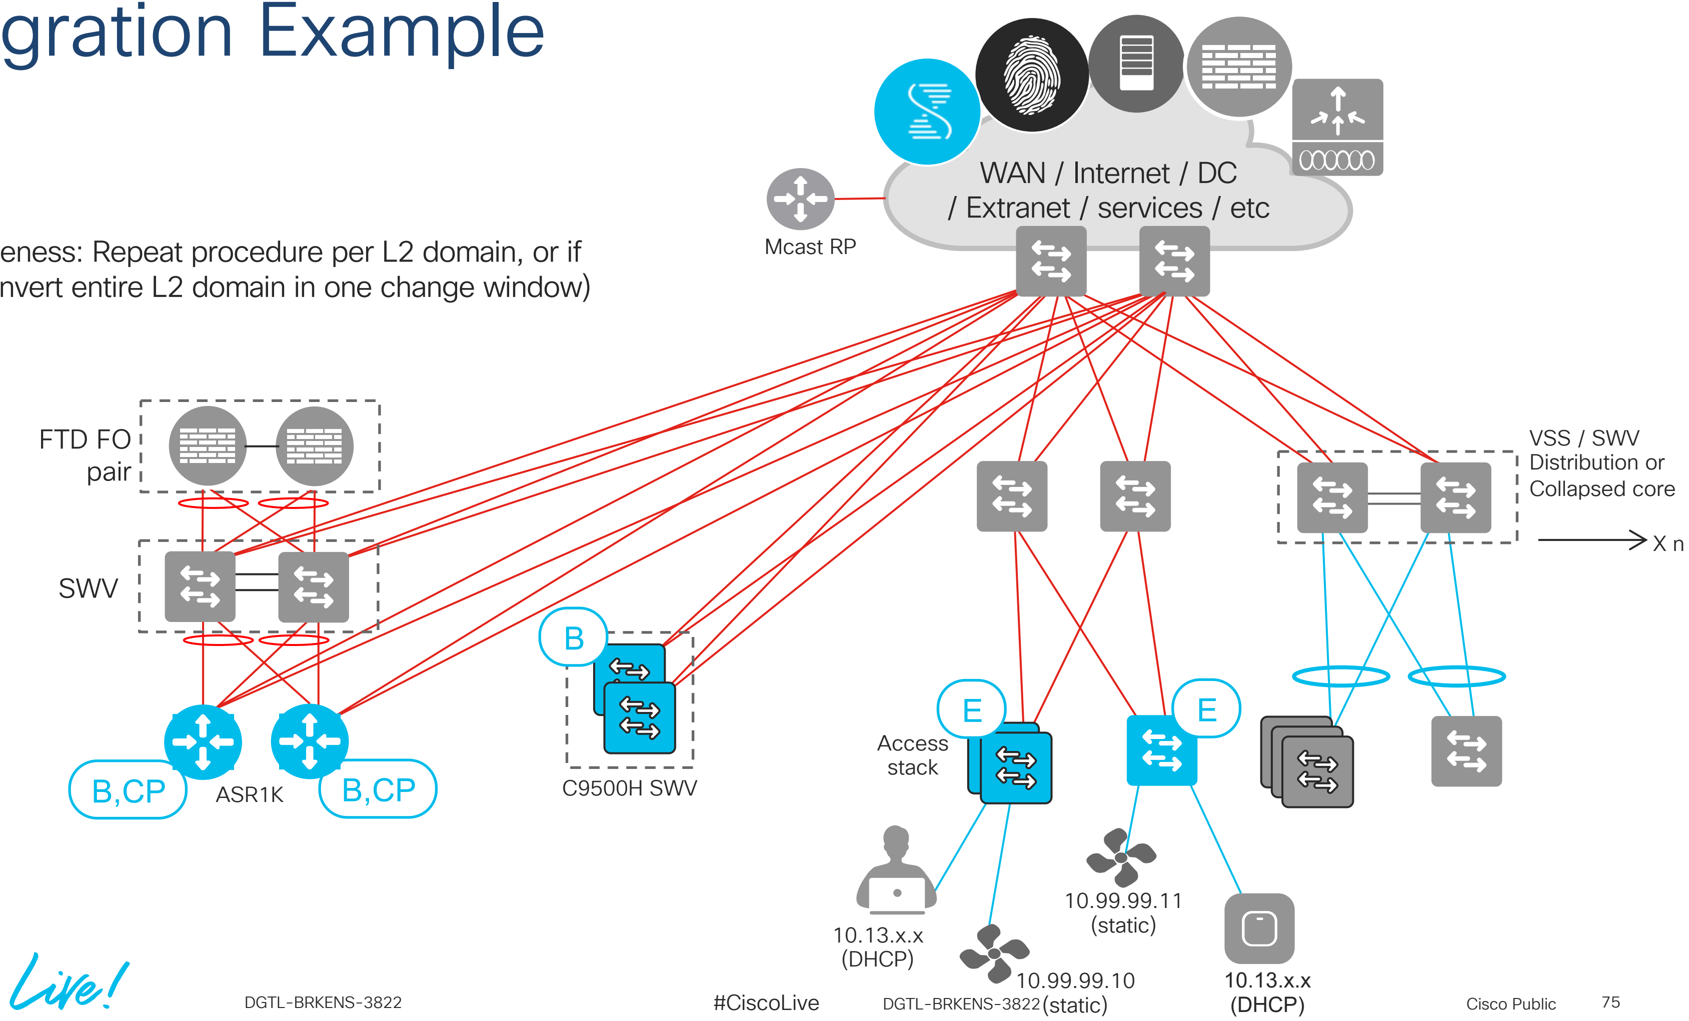

Migration Example

(For completeness: Repeat procedure per L2 domain, or if ambitious, convert entire L2 domain in one change window)

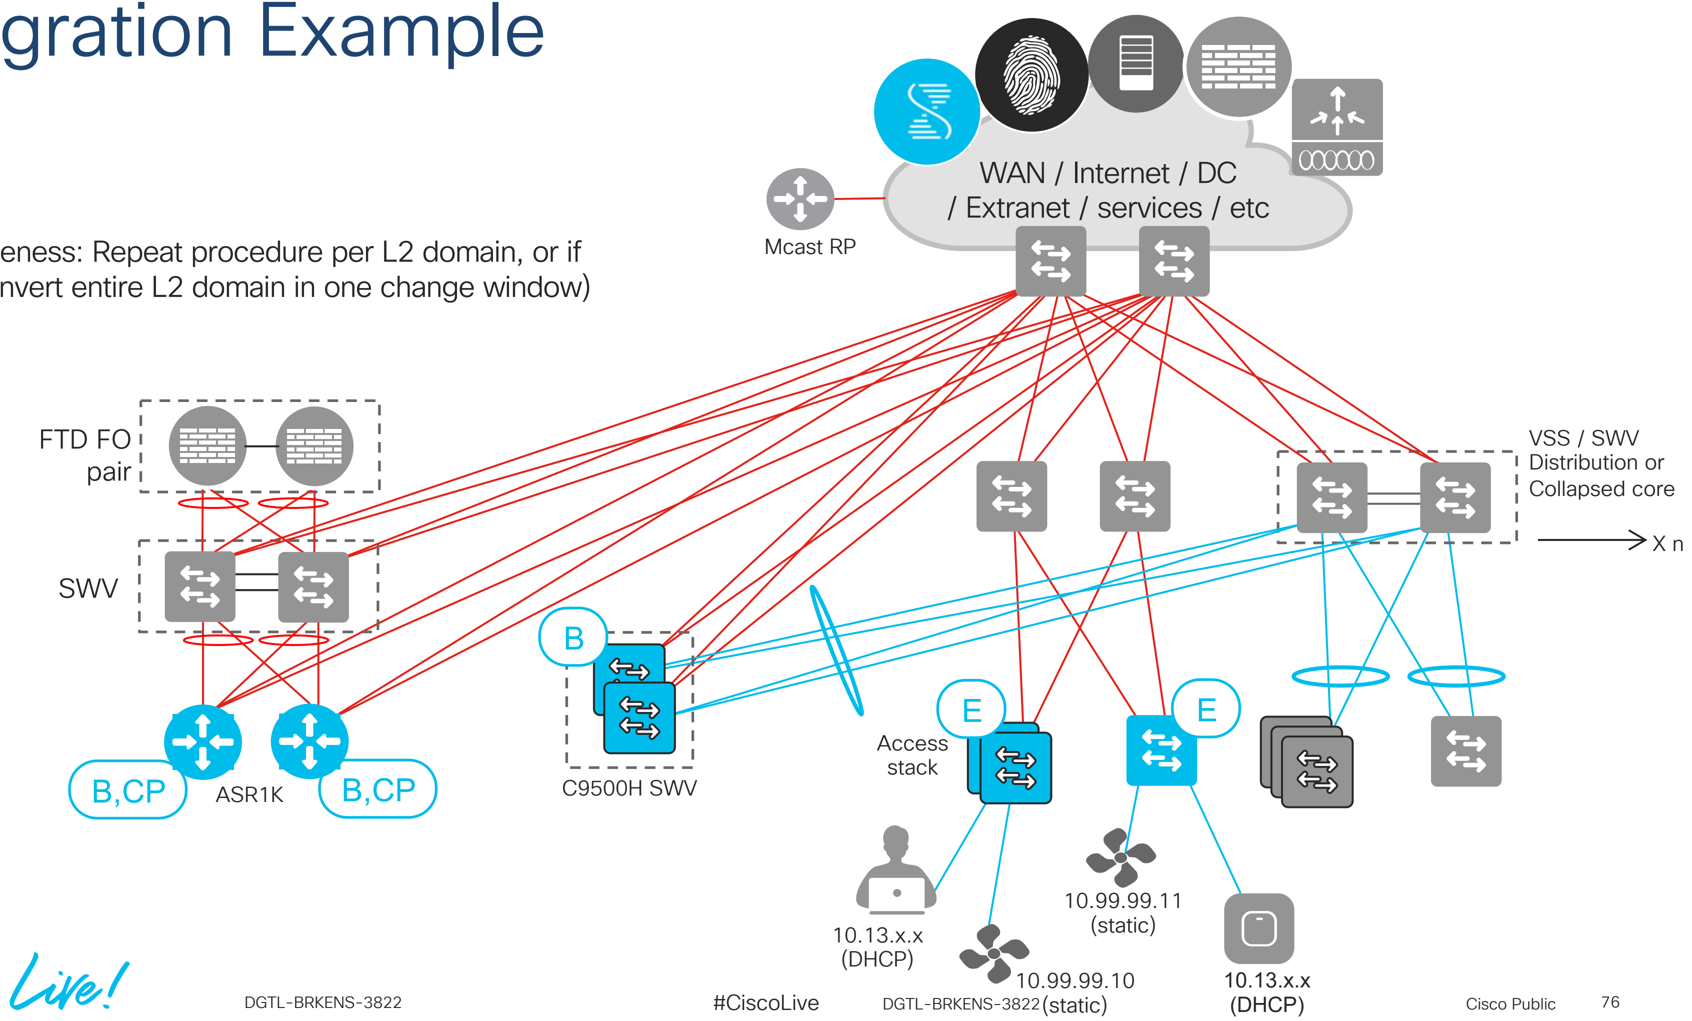

Migration Example

(For completeness: Repeat procedure per L2 domain, or if ambitious, convert entire L2 domain in one change window)

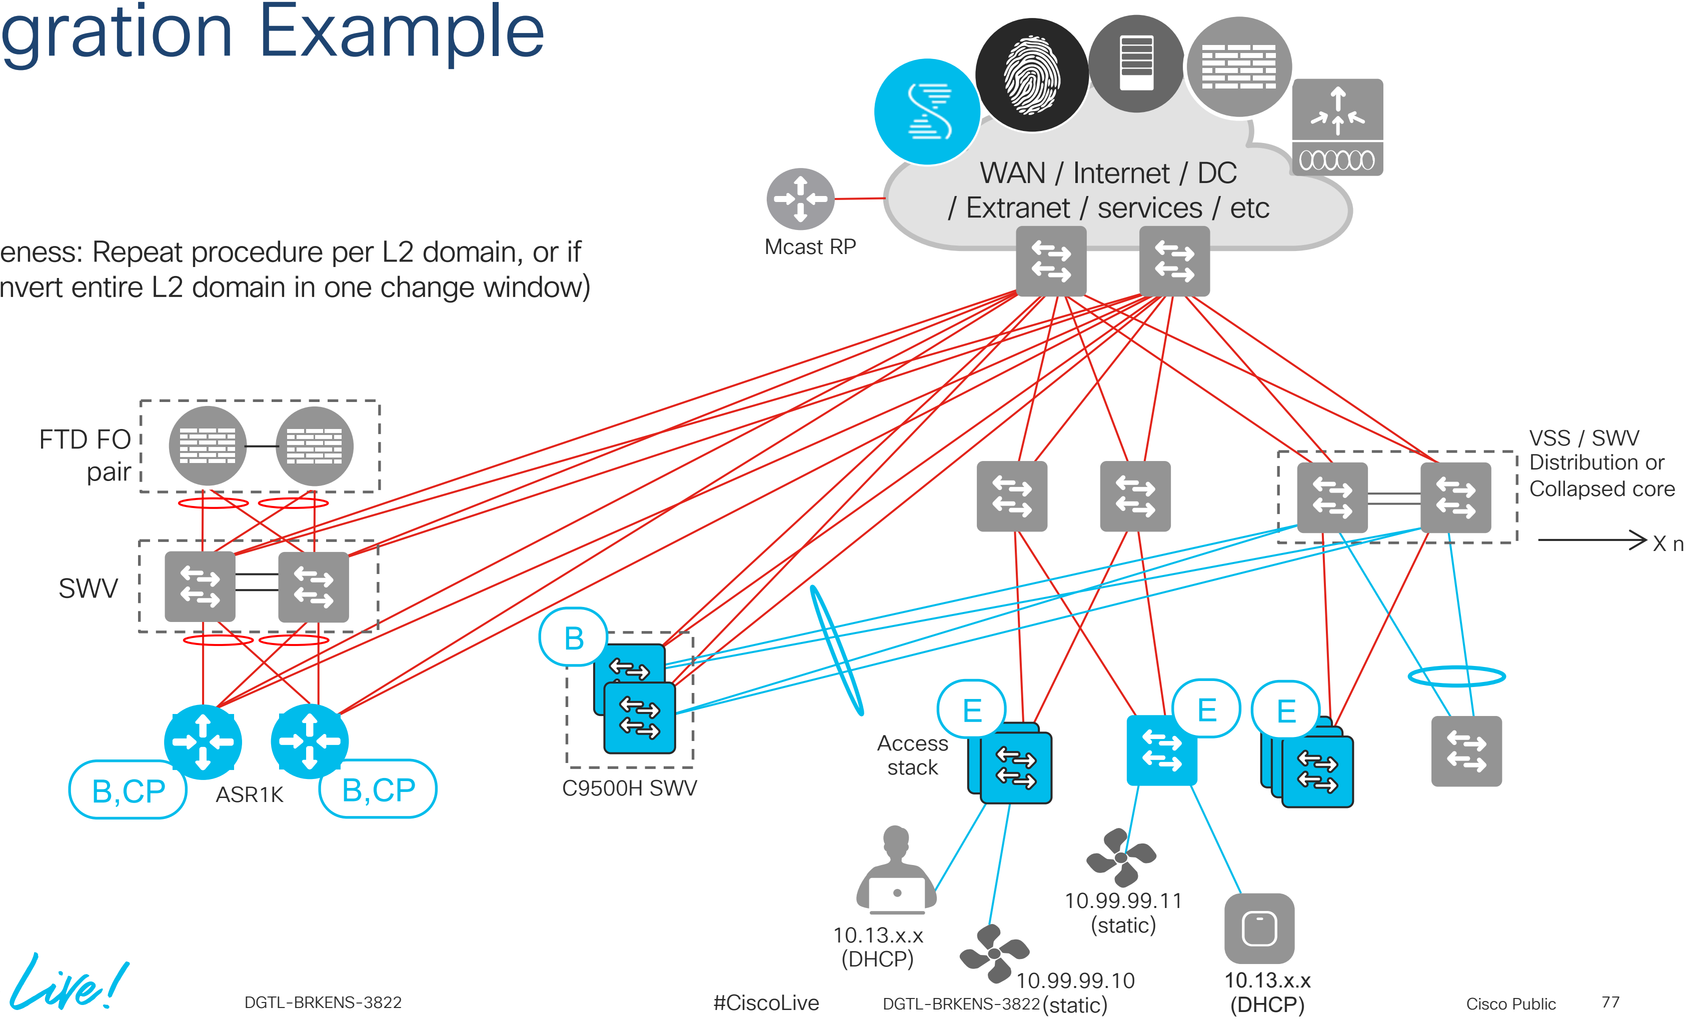

Migration Example

(For completeness: Repeat procedure per L2 domain, or if ambitious, convert entire L2 domain in one change window)

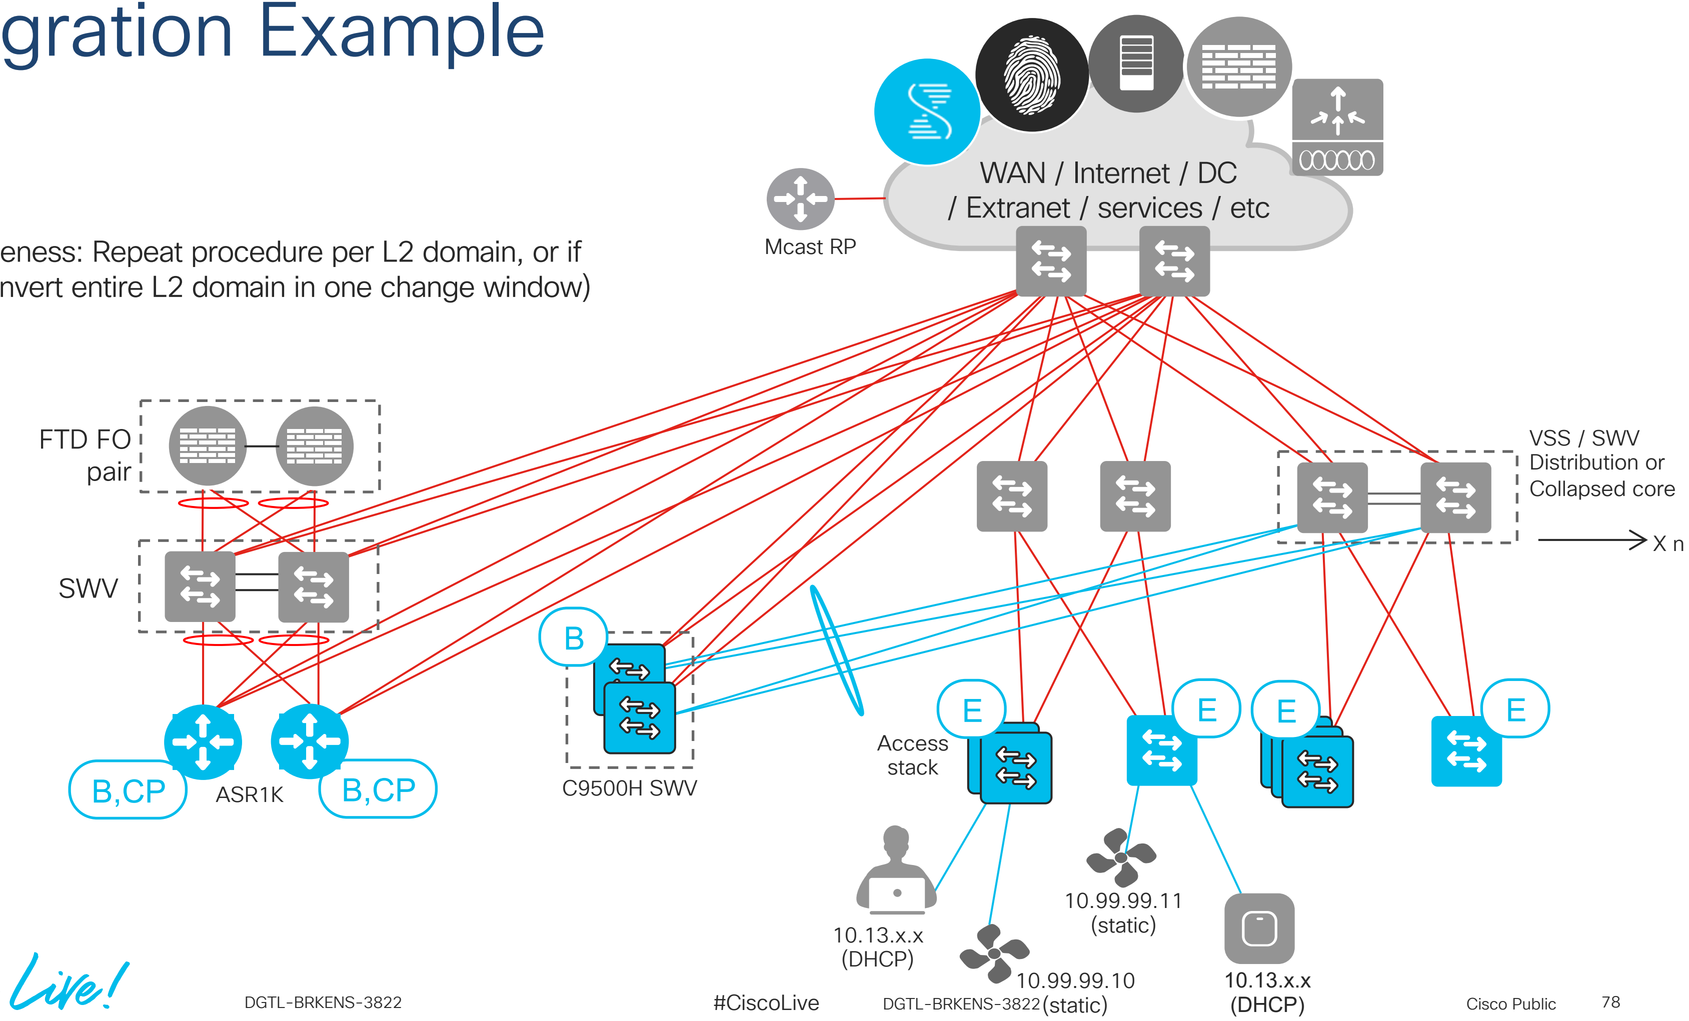

Migration Example

(For completeness: Repeat procedure per L2 domain, or if ambitious, convert entire L2 domain in one change window)

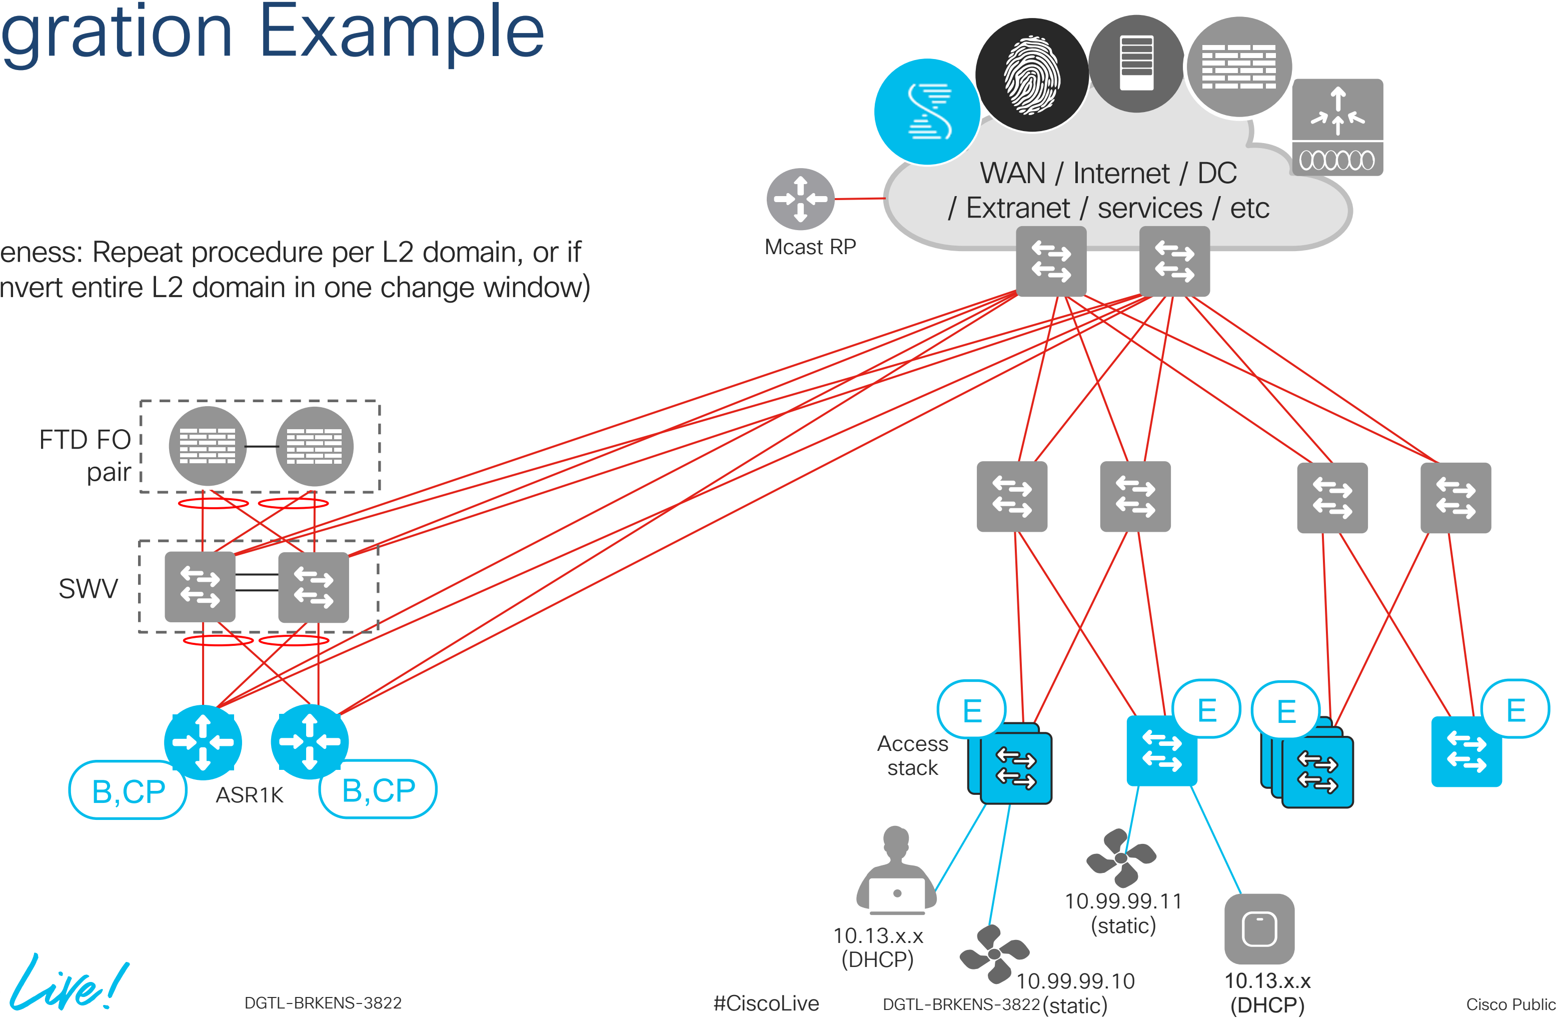

Migration Example

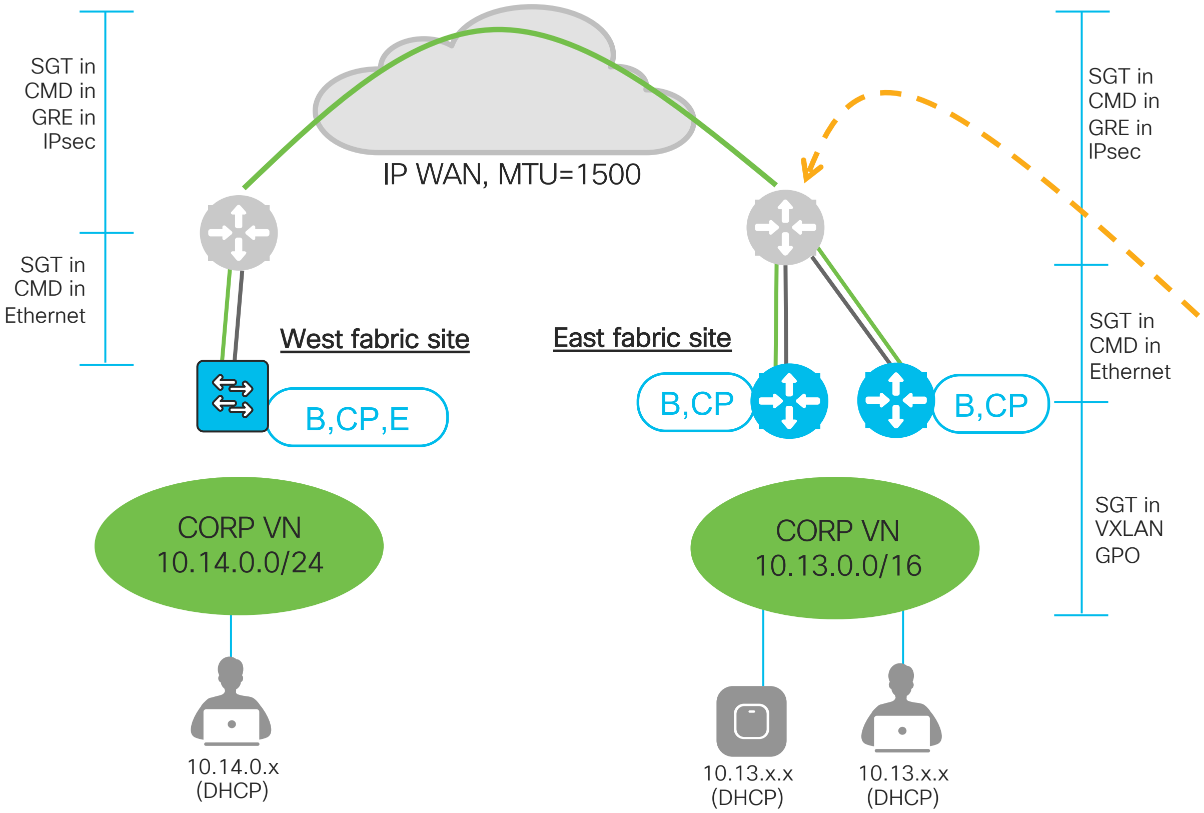

1. On routers create P2P static tunnel
2. Configure IP MTU and TCP adjust MSS in tunnels
3. Setup routing over tunnels e.g. BGP

```
interface Tunnel1002
 description --- To W-RT CORP VRF
 tunnel mode gre ip
 vrf forwarding CORP
 ip address 172.29.1.130 255.255.255.254
 ip mtu 1400
 ip tcp adjust-mss 1360
 cts sgt inline
 tunnel source Loopback0
 tunnel destination 192.168.8.29
 tunnel key 1002
 tunnel protection ipsec profile P1 shared
```

Underlay GRT

Overlay CORP

Migration demo

Conclusion

Summary

- Review recommended reference material, if needed: ciscolive.com !
- Understand the value proposition of SDA for the customer. Easiest high value features first
- Integrate Cisco DNA Center with ISE
- Add borders and pilot FEs
- Test
  - Failover scenarios
  - Exotic endpoints
  - Key use cases
- Deploy fabric
- Relax ☺

Thank you

cisco Live!

cisco

Possibilities

#CiscoLive