

Cisco SDA Design and Best Practices

Imran Bashir

Technical Marketing Engineer, Enterprise Business

This content is based on
Cisco DNAC 1.2.10
Cisco ISE 2.4 Patch 6

Agenda

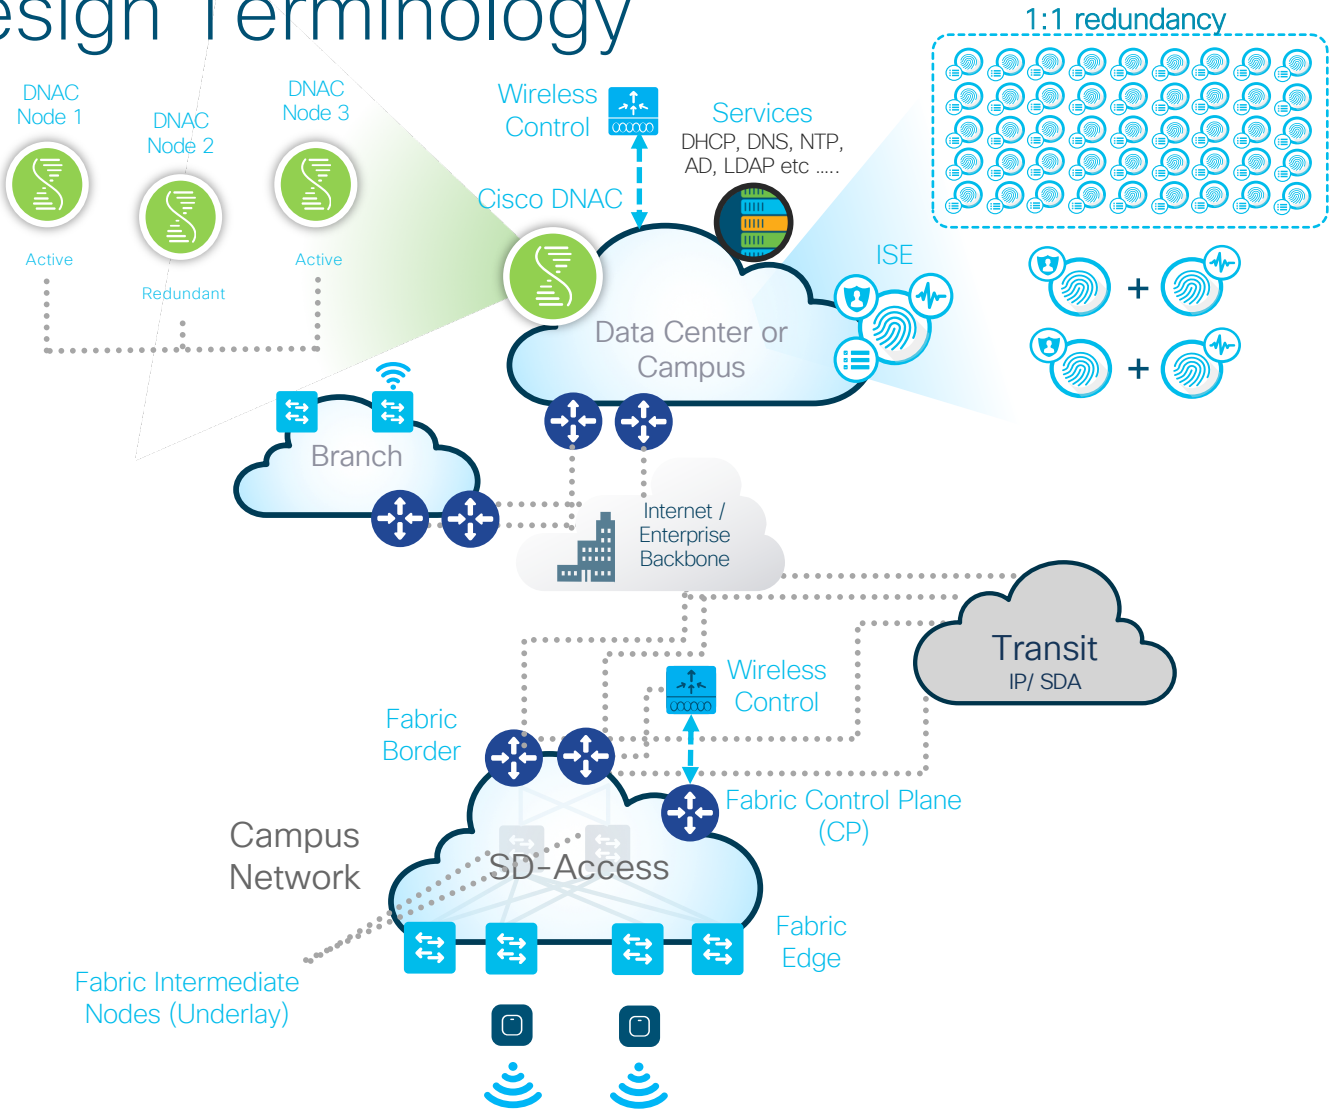


1. Design Terminology
2. Introduction to Cisco SDA design types
3. Designing Cisco DNA Center
4. Designing Cisco Identity Services Engine
5. Policy and Segmentation
- ***** BREAK *****
6. Scaling Cisco SDA Fabric
7. Think about Bandwidth and Latency
8. Best Practices
 - Cisco DNAC
 - Cisco ISE
 - SDA Fabric
9. Conclusion

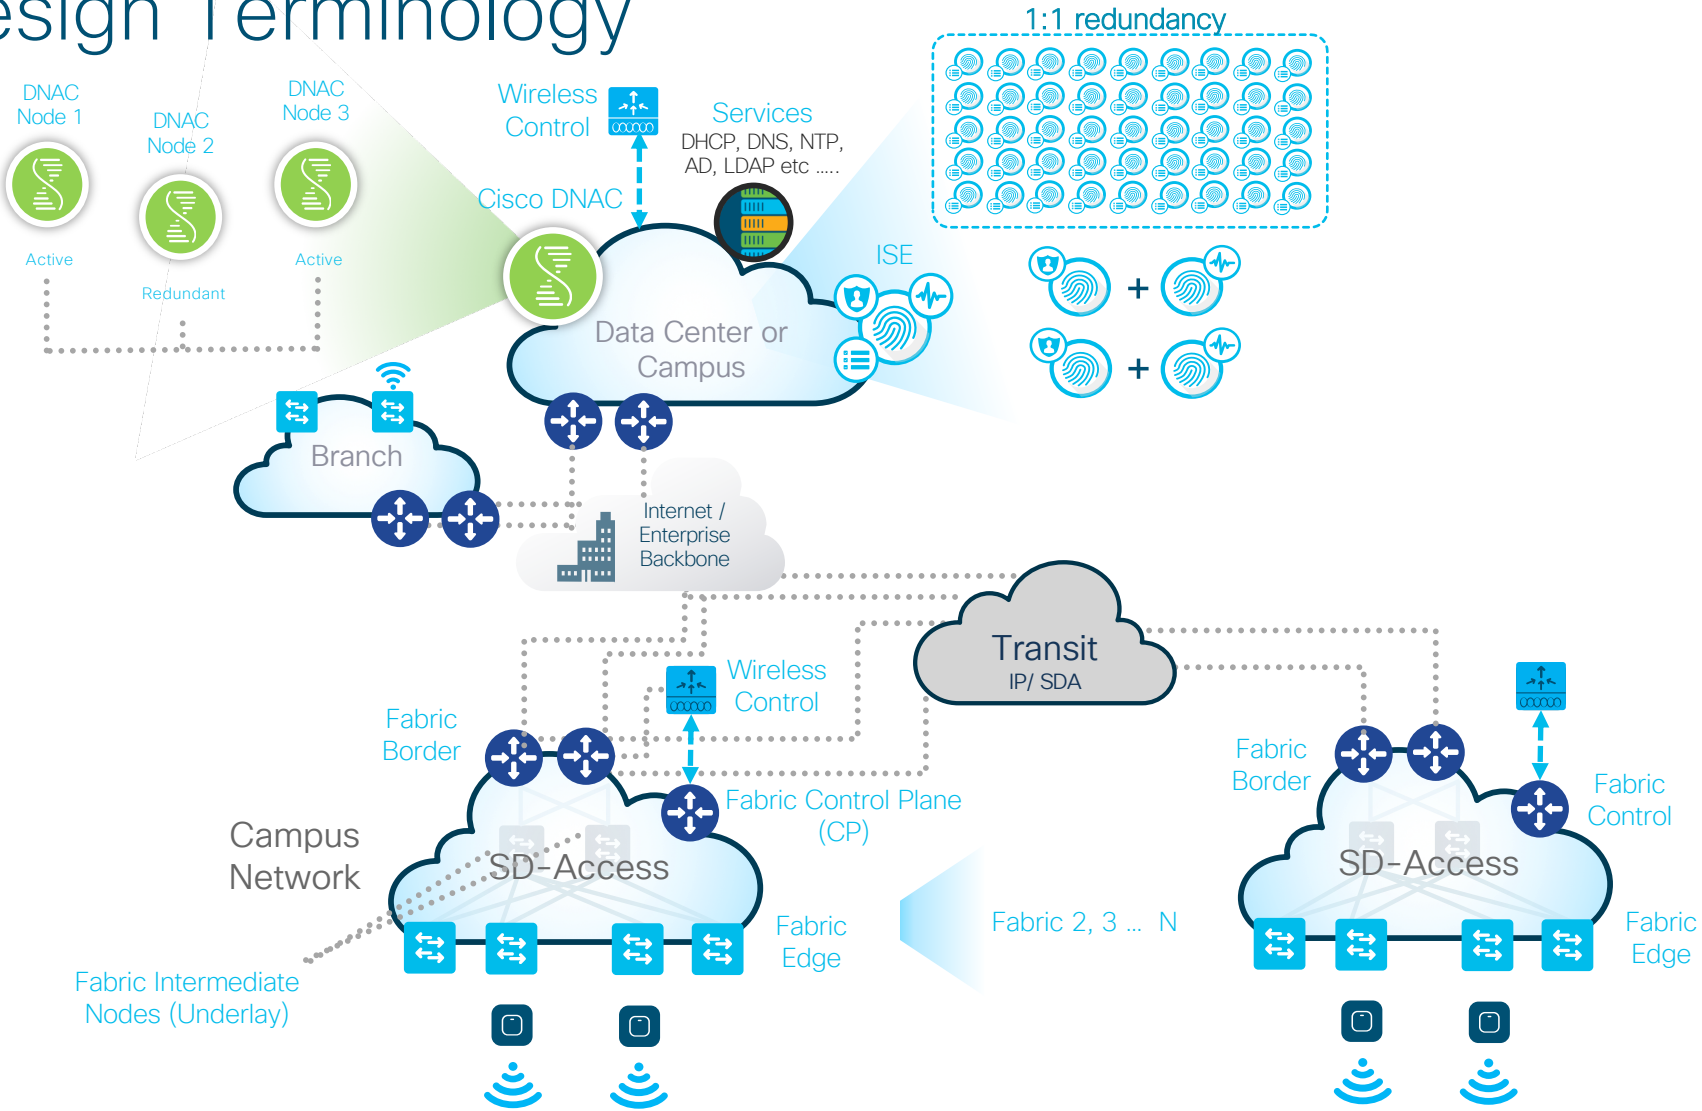
Design Terminology



Design Terminology



Design Terminology



Design Terminology

For Reference

Term	Description
Cisco DNA Center	Cisco DNA Center is a software solution that resides on the Cisco DNA Center appliance. The solution receives data in the form of streaming telemetry from every device (switch, router, access point, and wireless access controller) on the network.
ISE	AAA Policy Server -- Cisco Identity Services Engine (ISE) is a network administration product that enables the creation and enforcement of security and access policies for endpoint devices connected to the company's routers and switches. The purpose is to simplify identity management across diverse devices and applications.
Wireless LAN Controller	A wireless LAN (WLAN) controller is used in combination with the Lightweight Access Point Protocol (LWAPP) to manage light-weight access points in large quantities by the network administrator or network operations center. The wireless LAN controller is part of the Data Plane within the Cisco Wireless Model.
Fabric Domain	A logical (administrative) construct consisting of one or more Fabric or more Transits. Multiple independent Fabrics are connected to each other using a Transit.
Fabric Control Node	The SD-Access fabric control plane node is based on the LISP Map-Server (MS) and Map-Resolver (MR) functionality combined on the same node. The control plane database tracks all endpoints in the fabric site and associates the endpoints to fabric nodes, decoupling the endpoint IP address or MAC address from the location (closest router) in the network.
Fabric Border	The location where traffic exits the fabric as the default path to all other networks is an external border
Fabric Edge Nodes	The SD-Access fabric edge nodes are the equivalent of an access layer switch in a traditional campus LAN design. The edge nodes implement a Layer 3 access design
Fabric Intermediate Node	The fabric intermediate nodes are part of the Layer 3 network used for interconnections among the edge nodes to the border nodes

Introducing Cisco SDA Design types



Sample Network with Multiple Sites

SDA Design is driven by Customer requirements



Use Cases

- Mobility
- Survivability
- Scale
- Segmentation and Policy

Building/ Floor



Branch/ Campus



Metro Region



Very Small

Small

Medium

Large

Sample Network with Multiple Sites

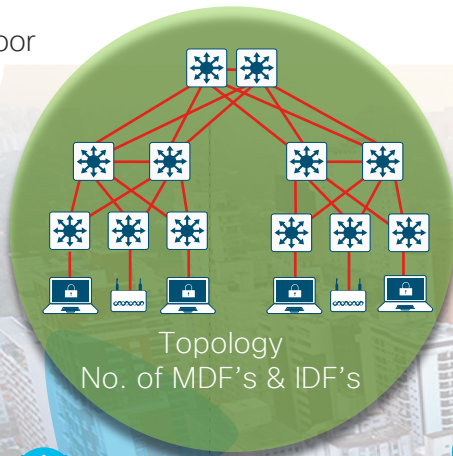
SDA Design is driven by Customer requirements



Use Cases

- Mobility
- Survivability
- Scale
- Segmentation and Policy

Building/ Floor



Branch/ Campus



Metro Region



Very Small

Small

Medium

Large

Types of SDA Designs

Fabric Design Categories

FIAB - Fabric In a Box

- Single wiring closet (MDF)
- Border, CP & E and Wireless in a single box
 - Limited Survivability
 - Limited Redundancy
- Stack supported (up to 8) with redundancy and survivability for Control plane
- Total endpoints < 2K (software limit)

Small Site

- Multiple wiring closets (MDF's)
- 2 x (collocated Border & CP) (in a single box)
 - Limited Survivability for Border & CP
 - Limited Redundancy for Border & CP
- Dedicated Edge (no stacking)
- Local WLC, Integrated
- Standalone ISE

Multiple Sites

- Multiple Sites is driven by customer design requirement
- Multiple Fabrics
- MAN or WAN Underlay
- Site Borders & Transit Area
- Distributed ISE

Very Small Site

Small Site

Large Site

Medium Site

Large Site

- 2 dedicated CPs (w SDA Wireless) – 6 with Wired ONLY. Up to 4 Border nodes
 - Full Survivability for Border & CP
 - Full Redundancy for Border & CP
- Local WLC + HA
- PAN - Local PSN

Medium Site

- Dedicated CP's for higher survivability (Site, building, floor)
OR
- 2 x collocated Border & CP (in a single box)
 - Full Survivability for CP
 - Limited Redundancy for Border
- Dedicated Edge (no stacking)
- Local WLC + HA
- ISE PAN - Local PSN

Solution Scale

Overall Solution Scale is Driven by Cisco DNAC

Cisco DNAC 1.2.10



Cisco DNAC

	Cisco DNAC (Overall Scale)	Cisco DNAC (Per Fabric Scale)
No. of Endpoints Max concurrent endpoints	*25,000 <small>(5K Wired + 20K Wireless)</small>	Same as overall
No. of Fabric Nodes Inc all managed devices Switches, Routers, WLC	500	500
Access Points No of AP's + Sensors	4K <small>(Max 200 Sensors)</small>	Same as overall
DNAC Sites No of Fabrics	*200	N/A
Virtual Networks No of VN's	64	Same as overall
IP Pools Max No. of IP Pools	N/A	100

Scale Numbers

* = Higher numbers with newer appliance



DN1-HW-APL
44 Core- UCS M4
Not Orderable



DN2-HW-APL
44 Core- UCS M5



DN2-HW-APL-L
56 Core- UCS M5

Solution Scale

Overall Solution Scale is Driven by Cisco DNAC Scale



Cisco DNAC

Overall Scale

Cisco DNAC

	Cisco DNAC (Overall Scale)	Cisco DNAC (Per Fabric Scale)
No. of Endpoints Max concurrent endpoints	*25,000 (5K Wired + 20K Wireless)	Same as overall
No. of Fabric Nodes Inc all managed devices Switches, Routers, ME + WLC	500	500
Access Points No of AP's + Sensors	4K (Max 200 Sensors)	Same as overall
DNAC Sites No of Fabrics	*200	N/A
Virtual Networks No of VN's	64	Same as overall
IP Pools Max No. of IP Pools	N/A	100
No. of Endpoints Max concurrent endpoints	*25,000 (5K Wired + 20K Wireless)	Same as overall


Scale Numbers



SDA

	SDA
Fabric Devices Max devices part of Fabric	N/A
AP's Max number of AP's	Switch Ports (Max 100 AP in 16.10)
Client Endpoints Scalable Groups	CP LISP Entries (3K - 200K)
IP Pools IP subnets (VLAN's)	N/A
Sites	N/A
Max Ports	No. of switch Ports
Max VN's	64 - 200,000


Scale Numbers



Cisco ISE


	ISE
Fabric Devices Max devices part of Fabric	N/A
AP's Max number of AP's	Max NAD devices
Client Endpoints Scalable Groups	10,000 - 2M
IP Pools IP subnets (VLAN's)	N/A
Sites	N/A
Max Ports	N/A Max NAD = 100K
Max VN's	N/A

Scale Numbers



Very Small Site

FIAB -- Fabric In A Box




Very Small



Small Design



Medium Design



Large Design

Overview

FIAB - Fabric In a Box

- Total **endpoints** < 2K (software limit)
- Border, CP & E and Wireless in a single box
 - Limited Survivability for CP and Border
- Single wiring closet (MDF)

Benefits

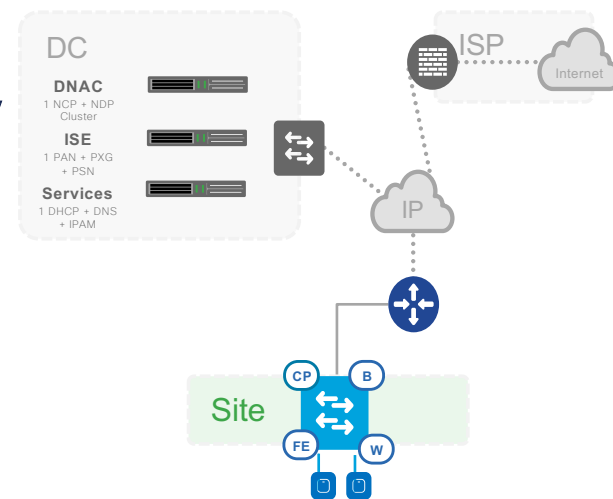
- Reduces cost to deploy SDA for very small sites
- FE + FB + CP on same C9K
- Supports 9800 & Embedded-Wireless in 1.2.10 (16.10.1e for C9300)

Border, Control and Edge

	9300
End Points/Hosts Max number of Endpoints	< 2K
Fabric Nodes	1
Virtual Networks Maximum number of VN's	< 8
IP Pools	< 8
Access Points	100
	B, CP & FE


Note: Platforms numbers can be higher but consider these solution numbers for design

Sample Topology



Very Small Site

Stacks of FIAB




Very Small



Small Design



Medium Design



Large Design

Overview

Stack of FIAB's

- Total **endpoints** < 2K (software limit)
- If a member of the Stack fails (with CP and Border), the next available member in the stack taker over the CP and Border functionality
 - **Limited Survivability for CP and Border**
- **Single wiring closet (MDF)**
- Max of 8 boxes can be in a Stack
- All the stack members must be the same platform


Benefits

- Get additional ports in a FIAB
- Still reduced cost to deploy SDA for very small sites
- FE + FB + CP on same C9K
- Supports 9800 & Embedded-Wireless in 1.2.10 (16.10.1e for C9300)

	Border, Control and Edge
	9300
End Points/Hosts Max number of Endpoints	< 2K
Fabric Nodes	1
Virtual Networks Maximum number of VN's	< 8
IP Pools	< 8
Access Points	100
	B, CP & FE

Note: Platforms numbers can be higher but consider these solution numbers for design

Sample Topology



Small Site

Overview

- Multiple wiring closets or even single.
- Border and CP are collocated in a single box
- Redundancy for Border or CP
- Limited Survivability
- Total endpoints < 10K (recommendation, but DNAC and platform scale can drive this number)

Benefits

- Small site design
- Tends to be Building or Office with < 10,000 endpoints and < 100 IP Pools/Groups
- 1-2 Collocated CP + External Border (Single Exit)
- Could be local WLC connected to Border (e.g. Stack) or Embedded WLC
- Looking at <1000 dynamic authentications and <250 group based policies.
- FB + CP + Wireless (9300)with distributed Fabric Edges
- Supports 9800 & Embedded-Wireless in 1.2.10 (16.10.1e for C9300). (maximum of 2 Embedded WLC) in N + 1 config.




Very Small



Small Design



Medium Design



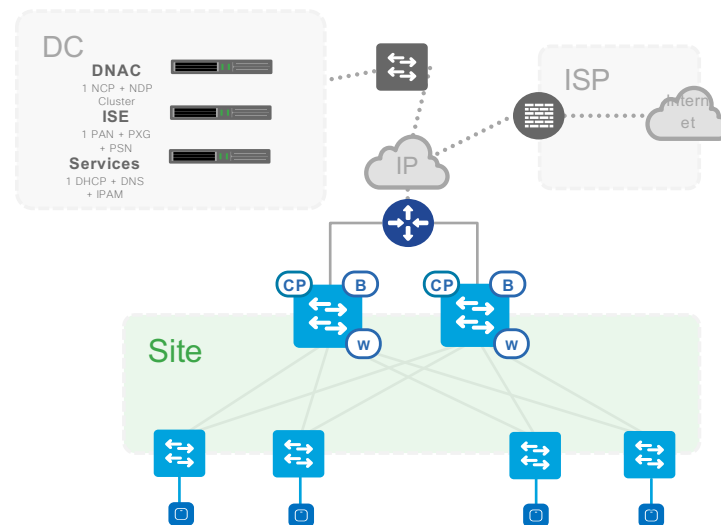
Large Design

● = Scale Numbers are currently being tested

	Border, Control		Fabric Edge	
	9300	9500	9200	9300
End Points/Hosts Max number of Endpoints	< 10K	< 10K	●	< 10K
Fabric Nodes Maximum number of VN's	2 (Collocated)	2 (Collocated)	●	< 25
Virtual Networks Maximum number of VN's	< 64	< 64	●	< 64
IP Pools	< 64	< 64	●	< 64
Access Points	200	200	●	200
	B, CP		FE	

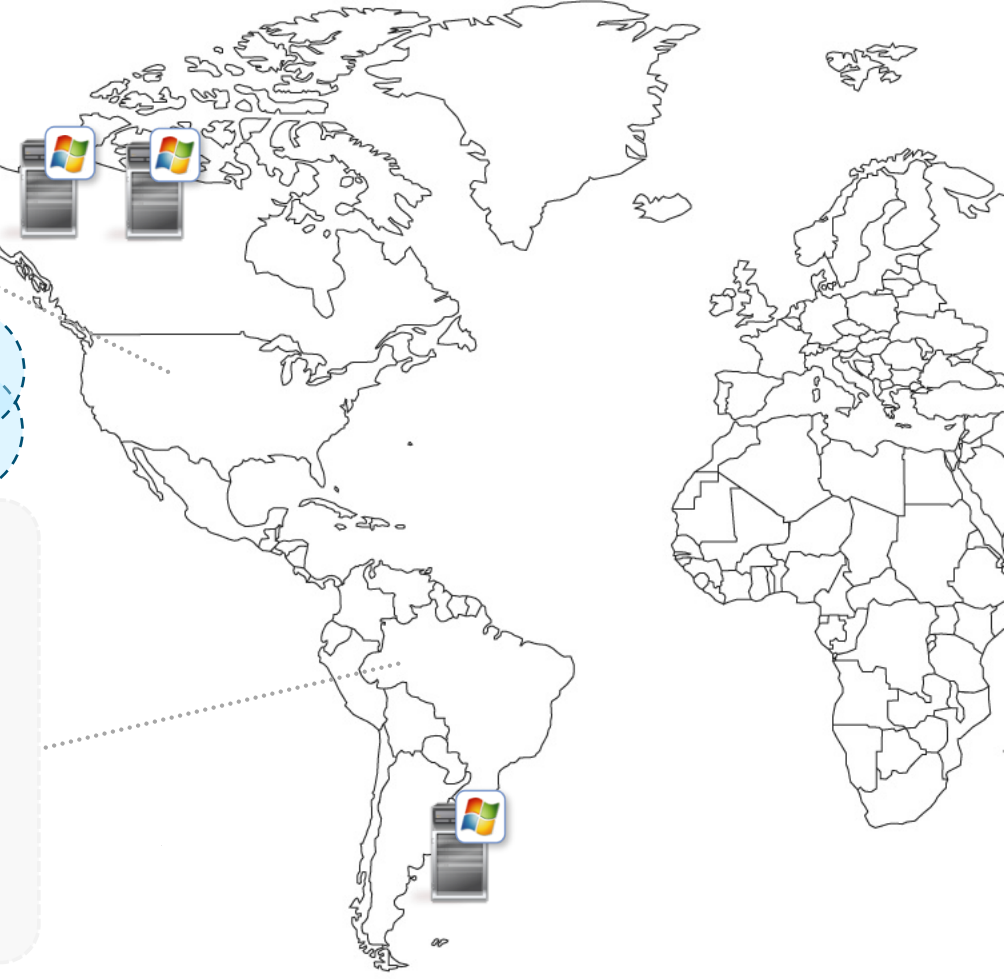
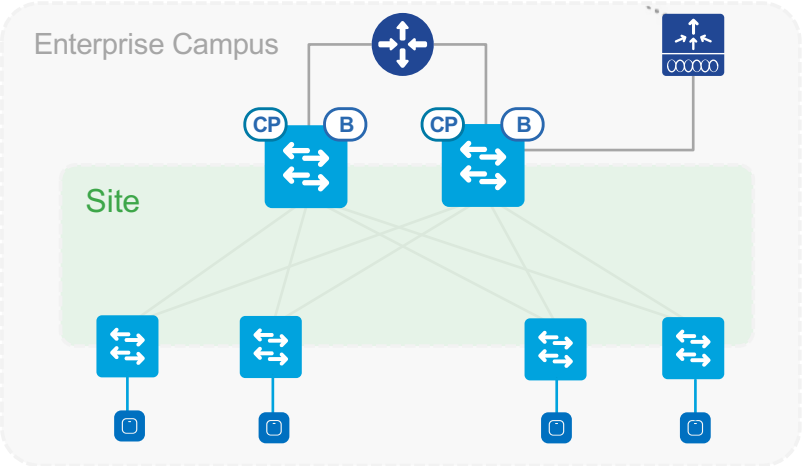

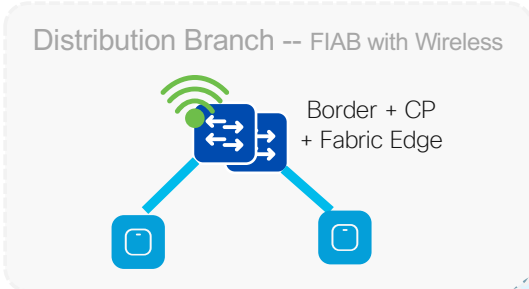
Note: Platforms numbers can be higher but consider these solution numbers for design

Sample Topology



Strategy for Cisco SD-Access in a small site

Design for a small site



Medium Site

Overview

Medium Site

- Multiple wiring closets or even single.
- Dedicated CP's for higher survivability (Site, building, floor)
- 2 x collocated Border & CP (in a single box)
 - Full Survivability for CP
 - Limited Redundancy for Border
- Dedicated Edge (no stacking)
- **Recommended** total endpoints < 10K (recommendation, but DNAC and platform scale can drive this number).

Benefits

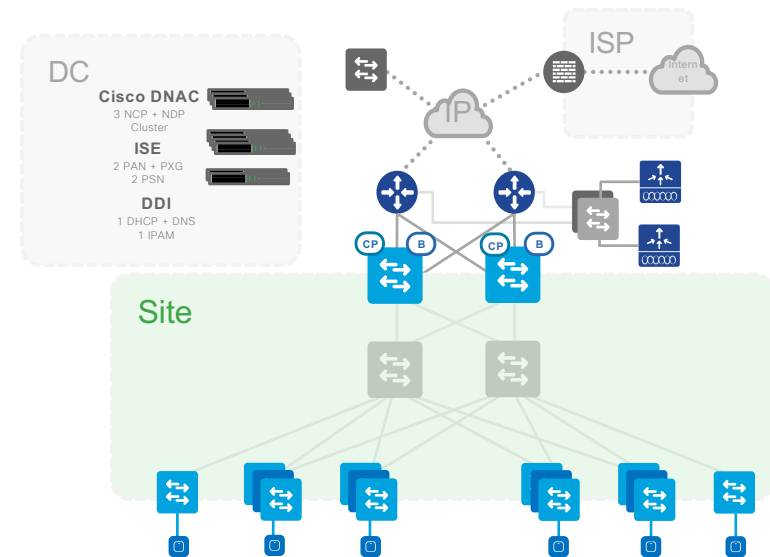
- Next level up to a small design.
- Max Control Plane nodes = 6 (Wired Only); 4 with Wireless (2 Enterprise and 2 Guest CP's).
- Tends to be Multiple Buildings with < 25,000 endpoints
- Most likely a 3 Tier design, recommendation is to use 9400 & 9500 as intermediate nodes.
- Can choose a Co-located or a Distributed/Dedicated CP + Border(Single Exit) design.
- Tends to be WLC + FEW via Services Block or a local Data Center
- Looking at < 25,000 dynamic authentications and < 1000 group based policies

● = Scale Numbers are currently being tested

	Border, Control		Fabric Edge	
	9500	9600	9300	9400
End Points/Hosts Max number of Endpoints	< 25K	< 25K	●	< 25K
Fabric Nodes	4 (4 CP, 2 B)	4 (4 CP, 2 B)	●	<250
Virtual Networks Maximum number of VN's	< 64	< 64	●	< 64
IP Pools	< 64	< 64	●	< 64
Access Points	200	200	●	200
	B, CP		FE	

Note: Platforms numbers can be higher but consider these solution numbers for design

Sample Topology



Large Site

Overview

Large Site

- Multiple wiring closets (most likely).
- Max Control Plane nodes = 6 (Wired Only); 4 with Wireless.
- Max Border nodes = 4
- Dedicated CP's for higher survivability (Site, building, floor)
- Dedicated Borders for site exits
 - Full Survivability for CP
 - Full Redundancy for Border
- Dedicated Edge (no stacking)
- **Recommended** total endpoints < 25K (recommendation, but DNAC and platform scale can drive this number).

Benefits

- Dedicated borders can provide multiple exits to different DC's or destinations.
- Tends to be Many Buildings with < 25,000 endpoints and < 500 IP Pools/Groups
- Most likely a 3 Tier design, recommendation is to use 9500 as intermediate nodes.
- Can choose a Co-located or a Distributed/Dedicated CP + 2-4 Borders (Multiple Exits)
- Looking at < 25,000 dynamic authentications and < 2000 group based policies




Very Small



Small Design



Medium Design




Large Design

● = Scale Numbers are currently being tested

	Border, Control		Fabric Edge	
	9500	9600	9300	9400
End Points/Hosts Max number of Endpoints	< 25K	< 25K	●	< 25K
Fabric Nodes	6 + 4 (6 CP, 4 B)	6 + 4 (6 CP, 4 B)	●	<1000
Virtual Networks Maximum number of VN's	< 64	< 64	●	< 64
IP Pools	< 64	< 64	●	< 64
Access Points	?	?	●	200
	B, CP		FE	

Note: Platforms numbers can be higher but consider these solution numbers for design

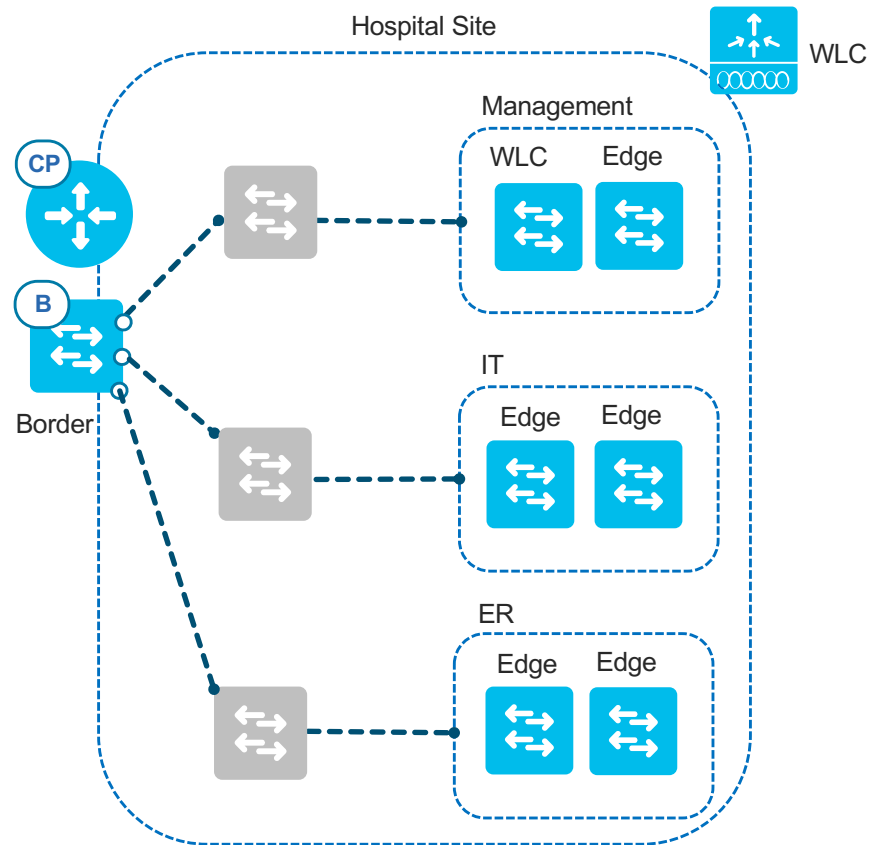



Why Multiple sites

Survivability or WAN separated networks

Use Case

- I need high survivability for my ER department

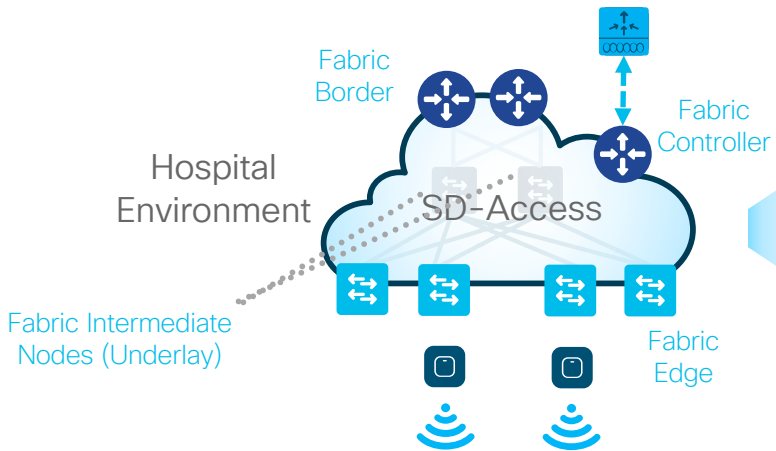


Why Multiple sites

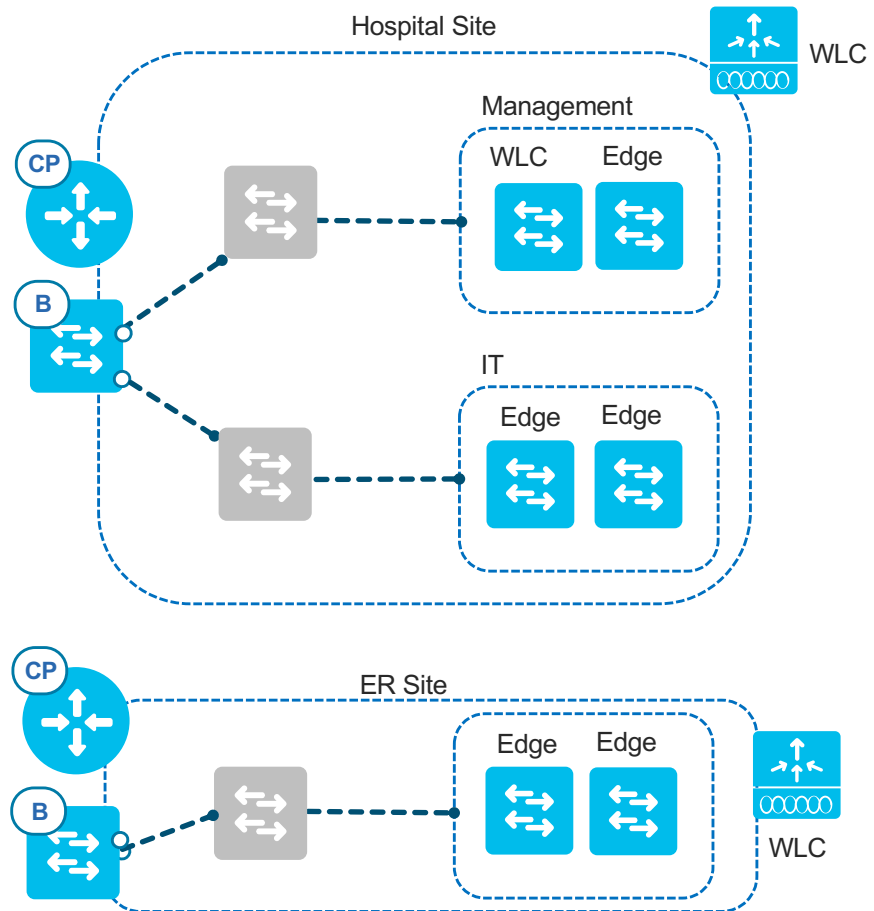
Survivability or WAN separated networks

Use Case

- I need high survivability for my ER department



ts affiliates. All rights reserved. Cisco Confidential



Branch or Campus

Roaming Wireless Endpoints

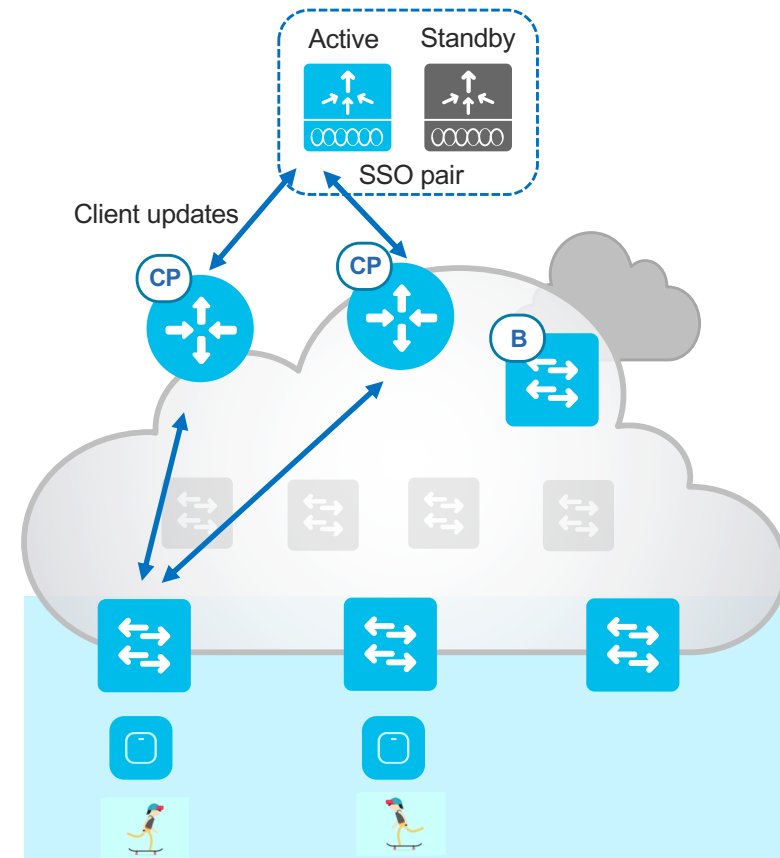
Use Case

- Wireless Endpoints roaming in Campus network

- Client roaming increases control plane (CP) events
- Deploy a dedicated CP with the correct platform

● = Scale Numbers are currently being tested

	Border, Control		Fabric Edge	
	9500	9600	9300	9400
End Points/Hosts Max number of Endpoints	< 25K	< 25K	●	< 25K
Fabric Nodes	4 (4 CP, 2 B))	4 (4 CP, 2 B))	●	<250
Virtual Networks Maximum number of VN's	< 64	< 64	●	< 64
IP Pools	< 64	< 64	●	< 64
Access Points	200	200	●	200
	B, CP		FE	



Sample Network with Multiple Sites

SDA Design is driven by Customer requirements



Use Cases

- Mobility
- Survivability
- Scale
- Segmentation and Policy


Building/ Floor



Branch/ Campus



Metro Region



Very Small


Small

Medium

Large

Types of Transit

Transit Design – IP vs SDA transit



Why IP Transit

Customers already using existing WAN or have adopted SD-WAN

Less than <1G circuits from Provider(s)

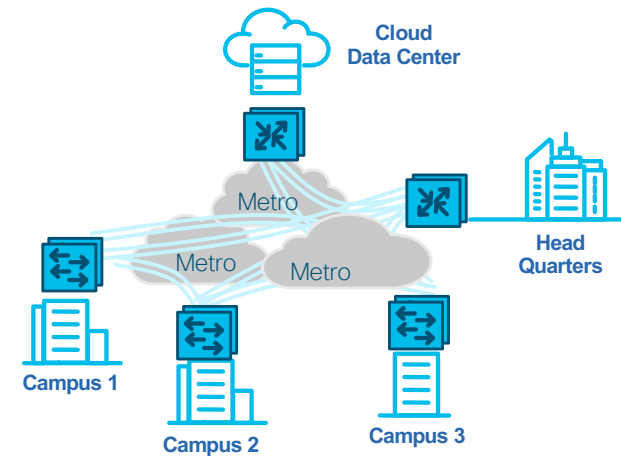
Higher latencies because sites are in different regions (many miles apart)

Use-cases

Internet Handoff
P2P IPSEC encryption

Policy Based Routing
WAN Accelerators

Traffic engineering
Mobile Backhaul LTE



Why SDA Transit

Smaller or isolated Failure Domains
Helps scaling number of Endpoints

DNAC provides Automation and Single View of entire system

VNs and SGTs gets pushed to all sites (consistent policy)
Local breakout at each Site for Direct Internet Access (DIA)

Use-cases

Consistent policy and end-to-end segmentation using VRFs and SGTs

Smaller and Isolated fault domains

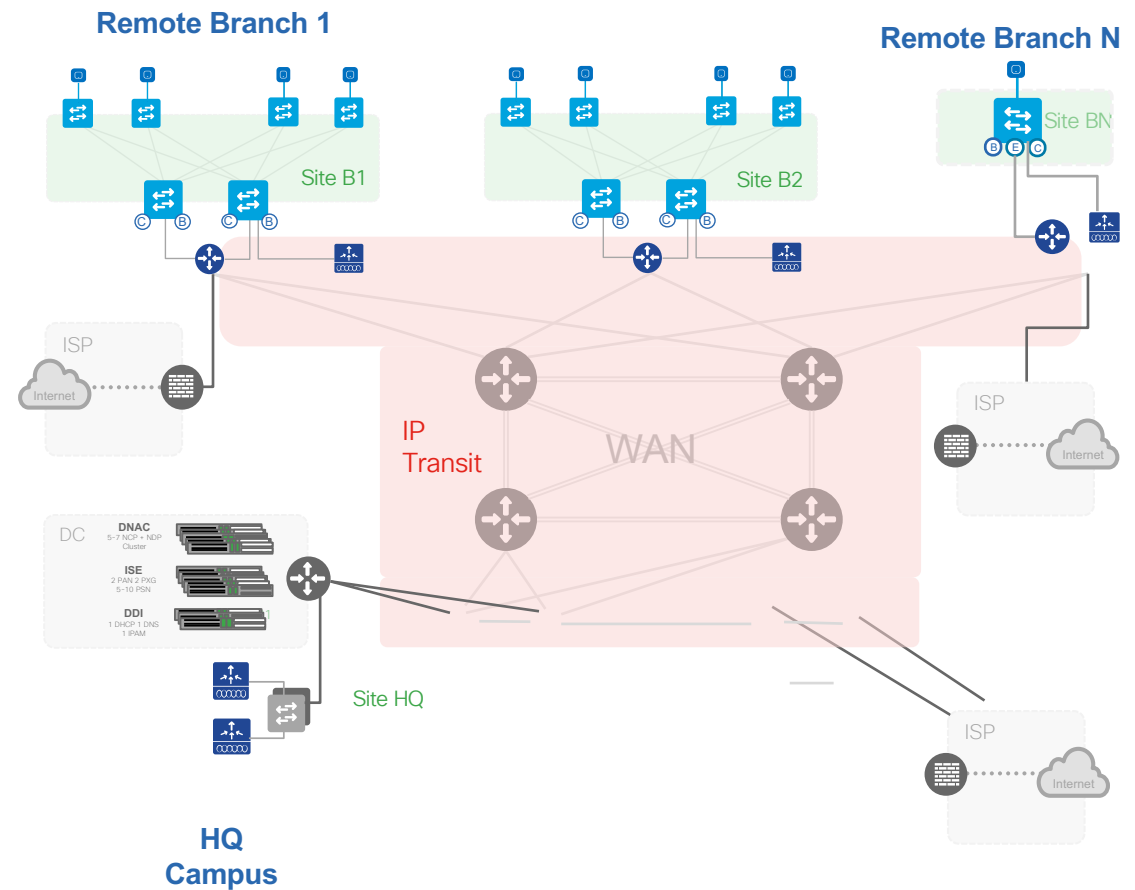
Resiliency and Scalability

IP Transit

Design for a multi site with IP Transit

Overview

- Tends to be many remote branch offices connected
- Customers already using existing WAN or have adopted SD-WAN
- Higher latencies because sites are in different regions (many miles apart)
- **Typical use cases**
 - Internet Handoff
 - P2P IPSEC encryption
 - Policy Based Routing
 - WAN Accelerators
 - Traffic engineering
 - Mobile Backhaul LTE

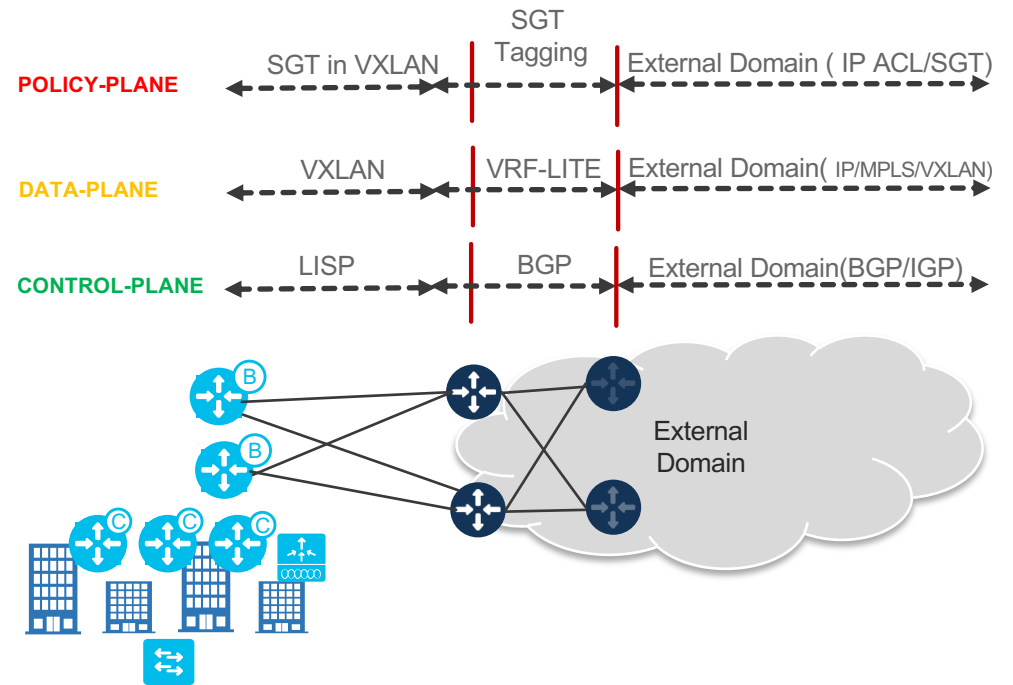


IP Transit

Border Hand-off

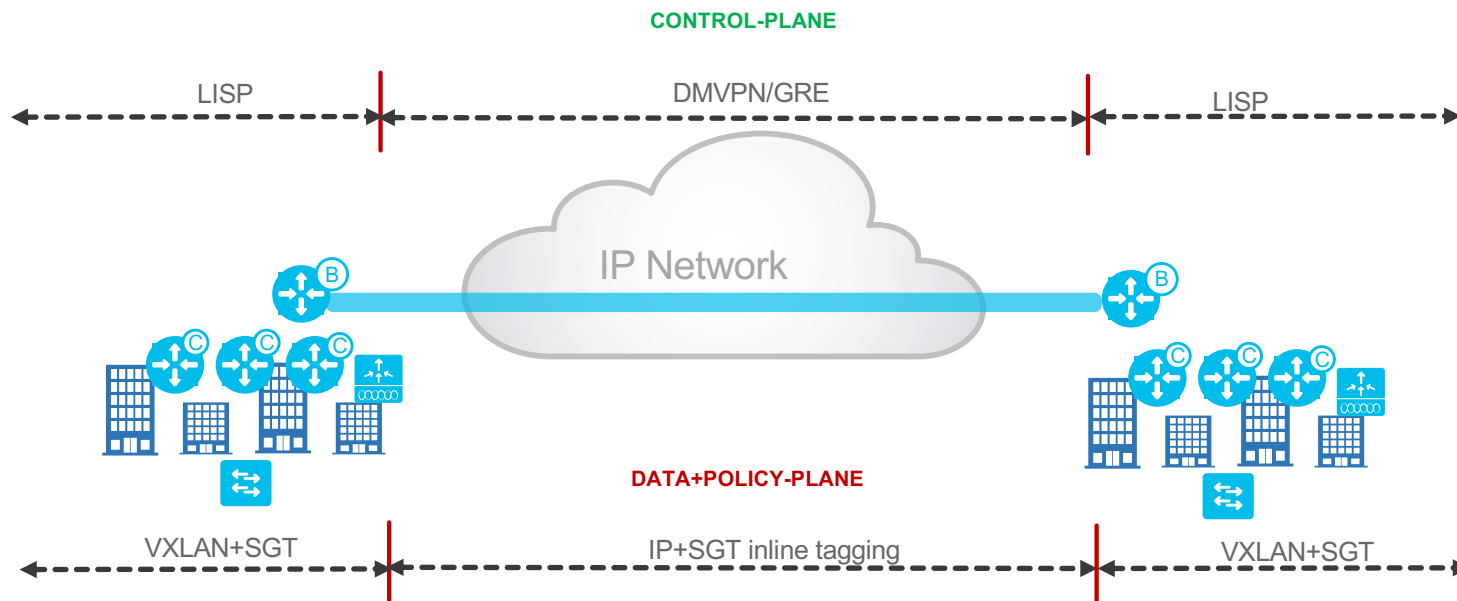
Overview

- Traffic hand-off from Fabric to outside domains

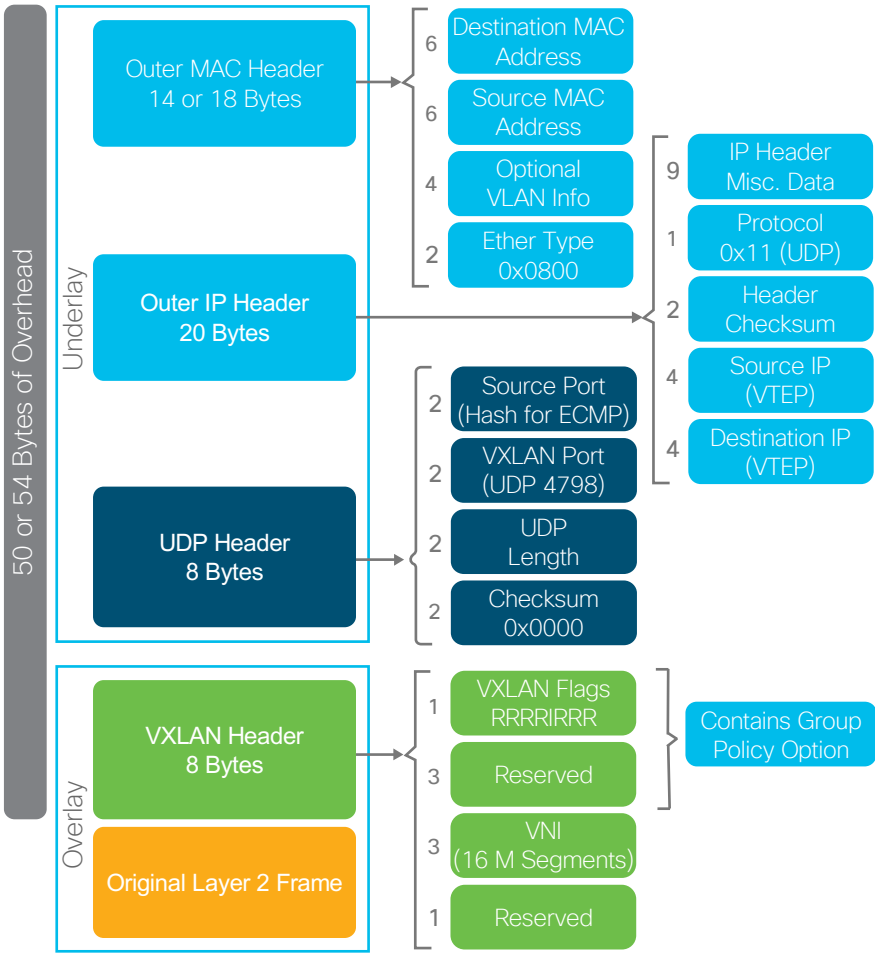


Inter-Connecting Fabric Domains/Sites

Multiple Fabric Domains



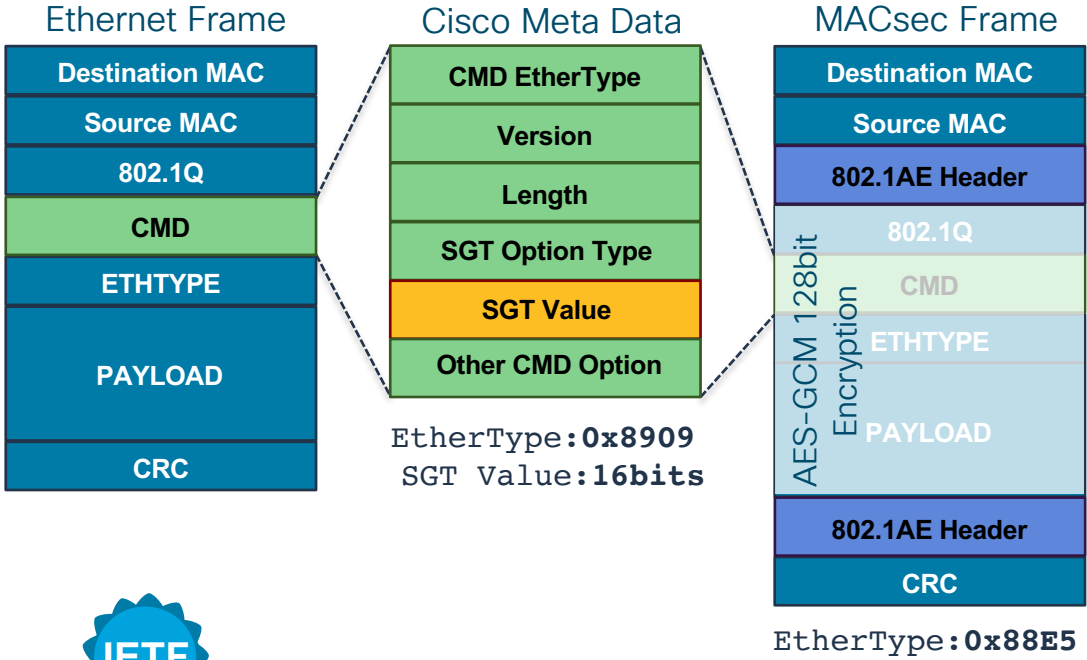
Packet Walks



Underlay Network
 Routing ID (RLOC) - IP address of the LISP router facing ISP

Overlay Network
 Endpoint Identifier (EID) - IP address of a host VRF
 Instance Id
 Dynamic EID
 VLAN

Ethernet inline tagging



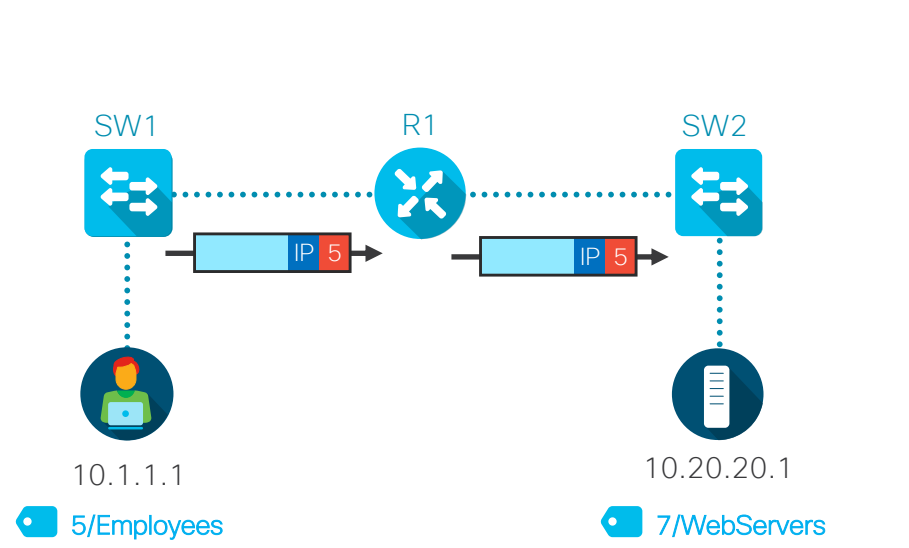
- Faster, and most scalable way to propagate SGT within LAN or DC
- SGT embedded within Cisco Meta Data (CMD) in Layer 2 frame
- Capable switches understands and process SGT in line-rate
- Optionally protect CMD with MACsec (IEEE802.1AE)
- No impact to QoS, IP MTP/Fragmentation
- L2 Frame Impact: ~20 bytes
- 16 bits field ~ 64,000 tag space
- **Non-capable device drops frame with unknown Ethertype**



<http://tinyurl.com/sgt-draft>

Two ways to propagate


DATA PLANE PROPOGATION



SGT carried inline in the data traffic. Methods include, SGT over:

- Ethernet
- MACSec
- LISP/VxLAN
- IPSec
- DMVPN
- GETVPN

CONTROL PLANE PROPOGATION



IP-to-SGT data shared over control protocol. No SGT in the data plane. Methods include, IP-to-SGT exchange over:

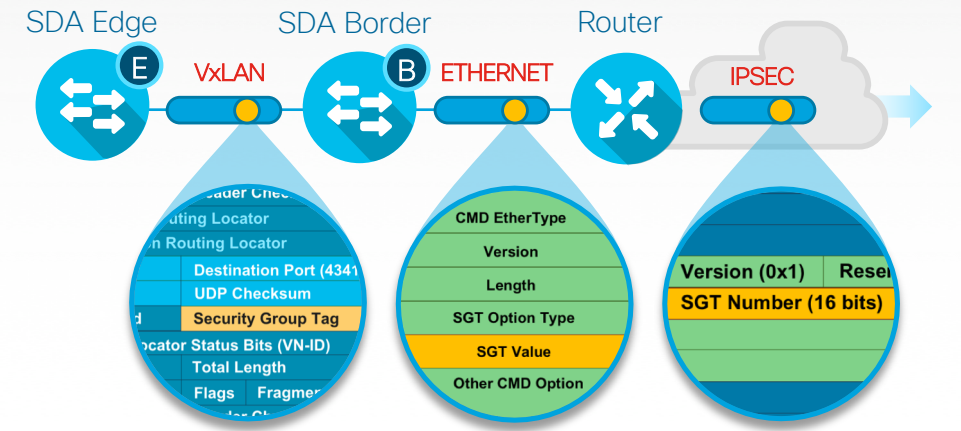
- SXP
- pxGrid



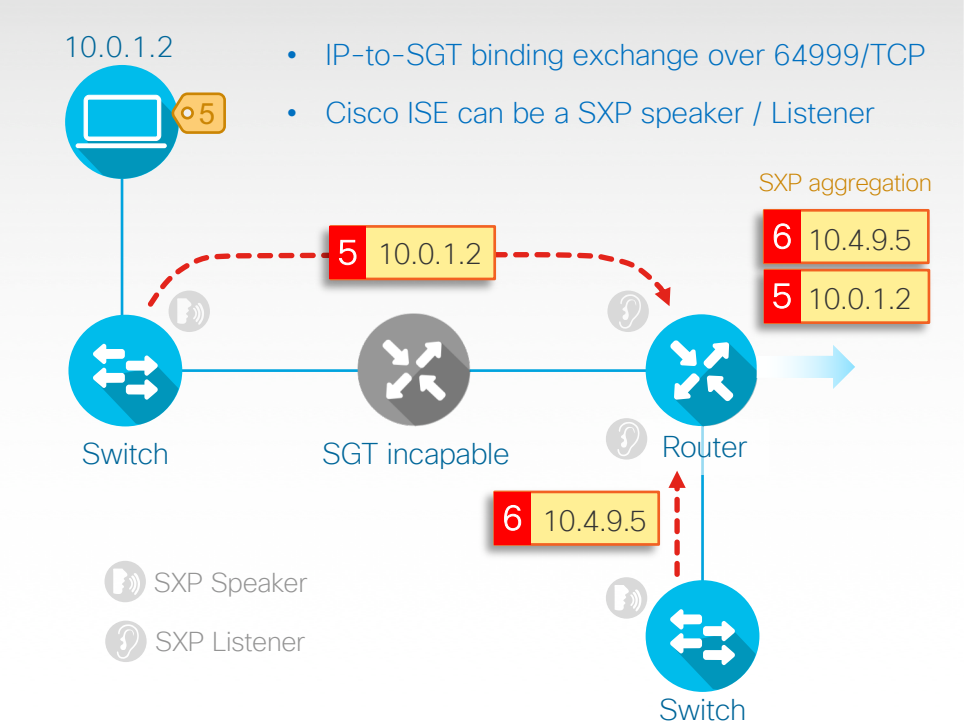
Propagation examples

Inline Methods

- Ethernet** **Ethernet Inline Tagging:** (EtherType:0x8909) 16-Bit SGT encapsulated within Cisco Meta Data (CMD) payload.
- IPSec** **L3 Crypto:** Cisco Meta Data (CMD) uses protocol 99, and is inserted to the beginning of the ESP/AH payload.
- VxLAN** SGT (16 bit) insertion in the Nonce field (24 bit)

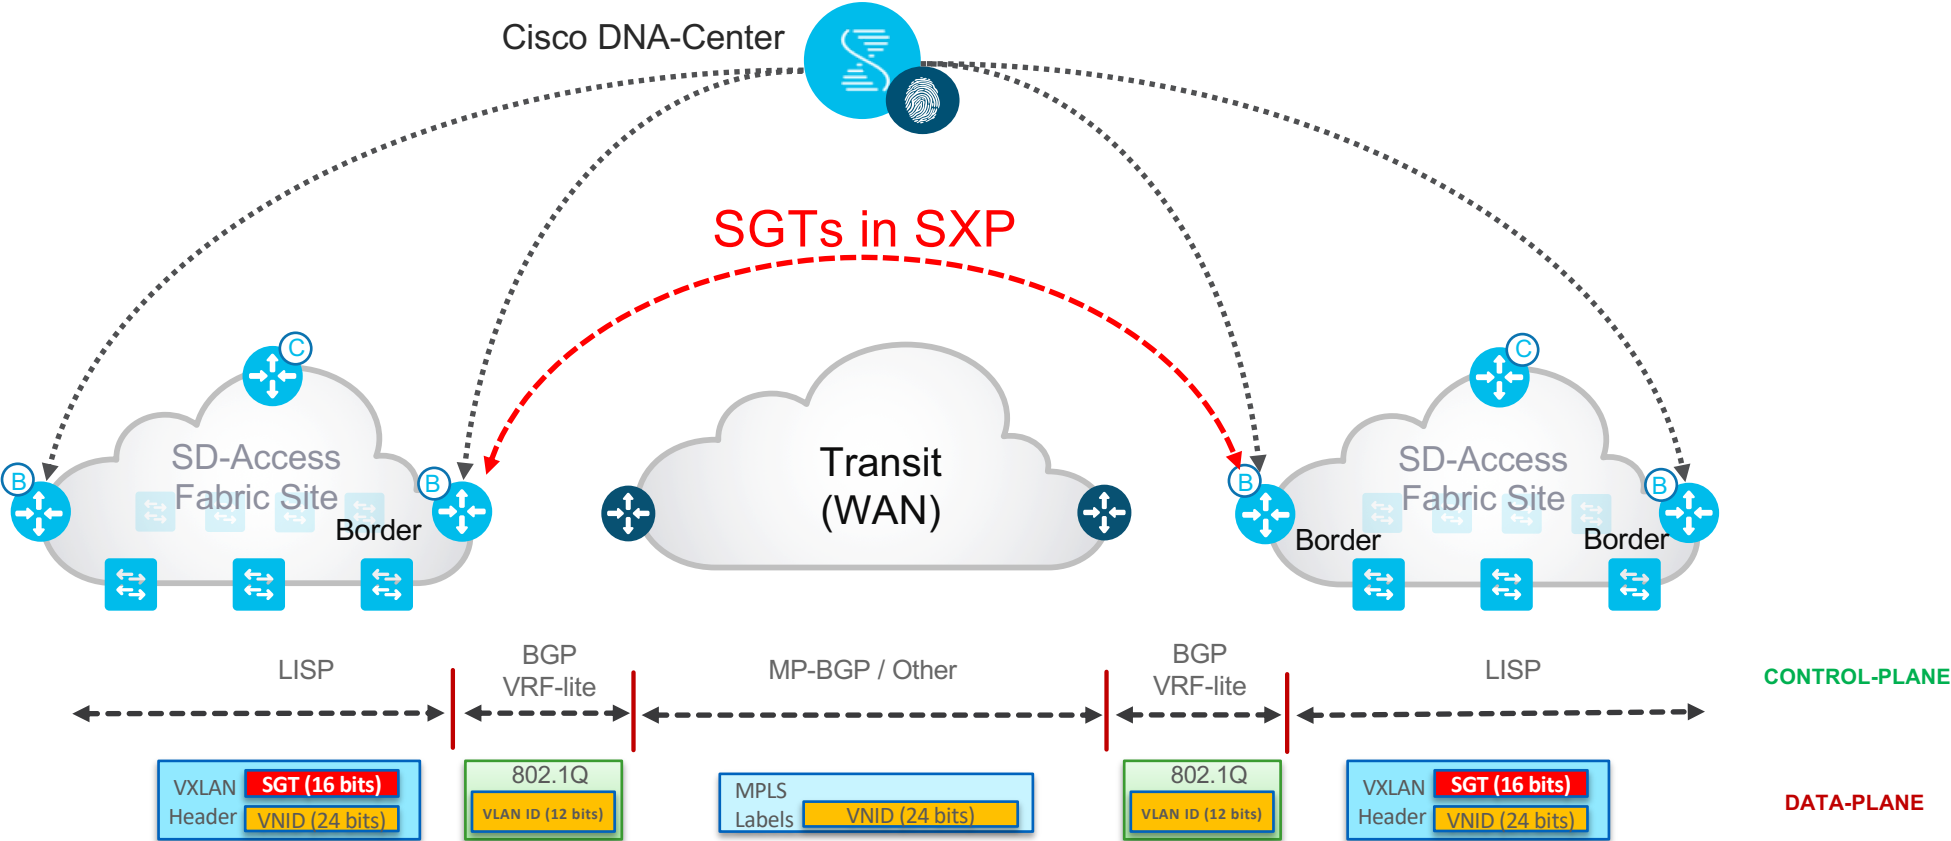


SGT Exchange Protocol (SXP)



Cisco SD-Access for Distributed Campus IP Based WAN Transit

Management and Policy

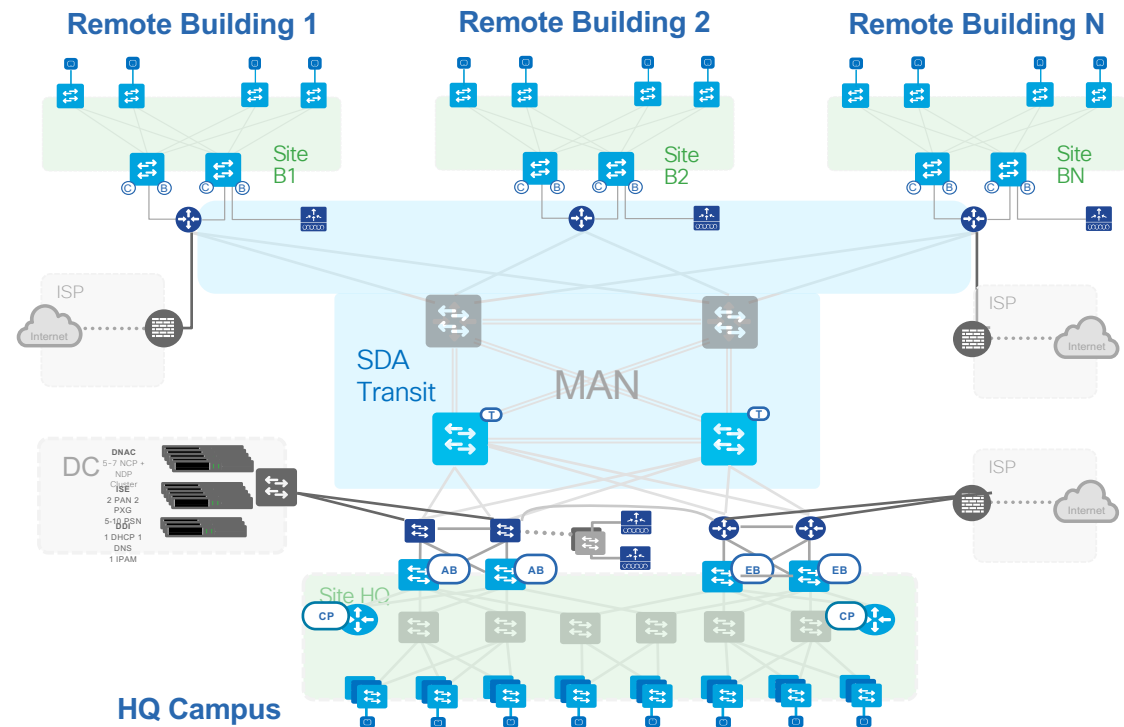


SDA Transit

Design for a multi site with SDA Transit

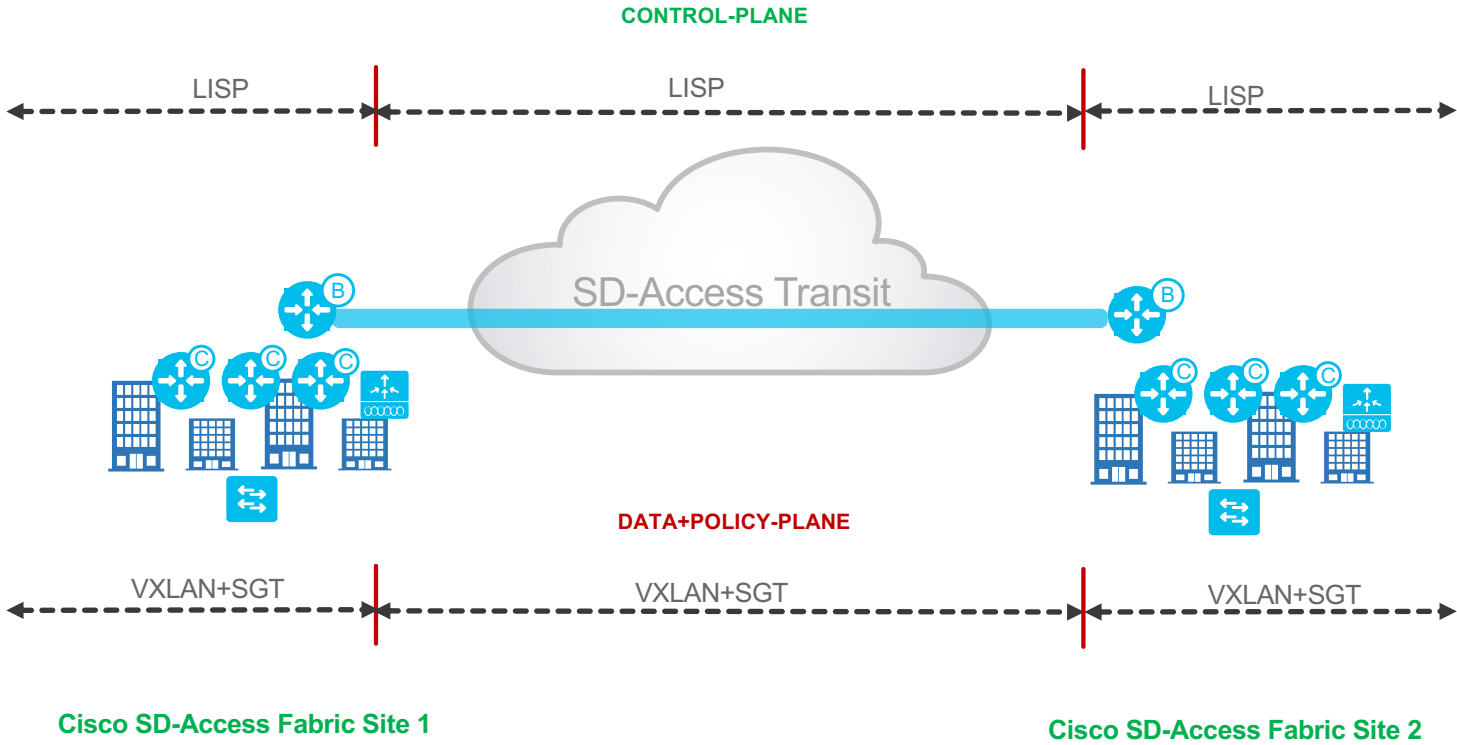
Overview

- Customers have multiple sites connect via “Dark Fiber” links or DWDM links
- Sites are in same Metropolitan area (a few hundred miles apart)
- Typical use cases
 - Consistent policy and end-to-end segmentation using VRFs and SGTs
 - Smaller and Isolated fault domains
 - Resiliency and Scalability



Cisco SD-Access Transit

Multiple SD-Access Fabric Sites

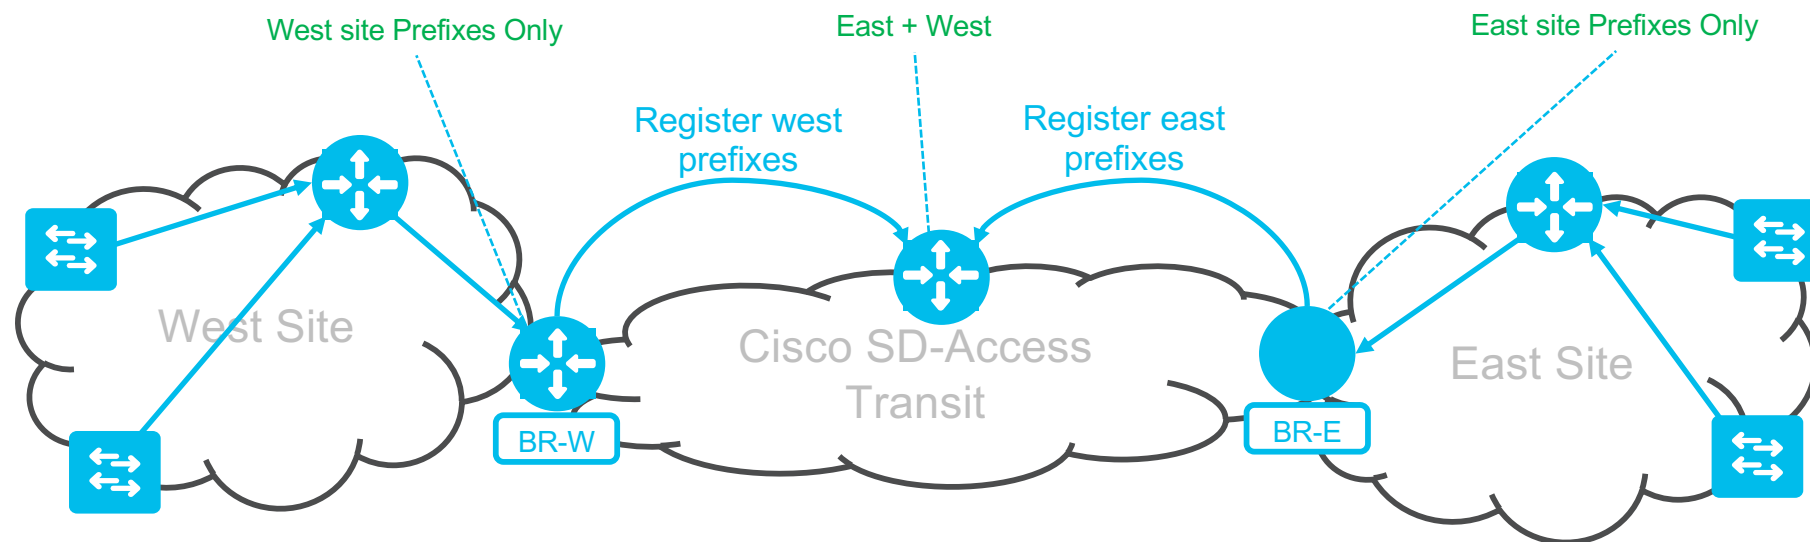


Cisco SD-Access Distributed Site Control Plane for Global Scale

Multiple SD-Access Fabric Sites

Use Case

- Each site only maintains state for in-site end-points.
- Off site traffic follows default to transit.
- Survivability, each site is a fully autonomous resiliency domain
- Each Site has its own unique subnets

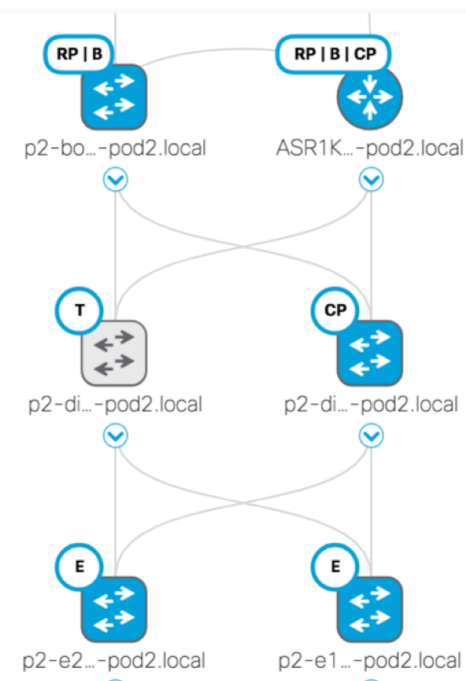


Cisco SD-Access Distributed Campus

Transit Control Plane

Best Practices for Site Transit Control Plane

- **Transit Control Plane** nodes **today** receive aggregate routes from **site borders** at each fabric site using LISP
- It is recommended that these nodes must be **dedicated**. Do **NOT** collocate them as they are critical for **inter-site** communications
- Deploy **2** Transit Control Plane nodes for **redundancy and load balancing**



Cisco SD-Access Distributed Campus

Fabric Border Support Matrix

Best Practices for Fabric Border Selection

- Consider the following

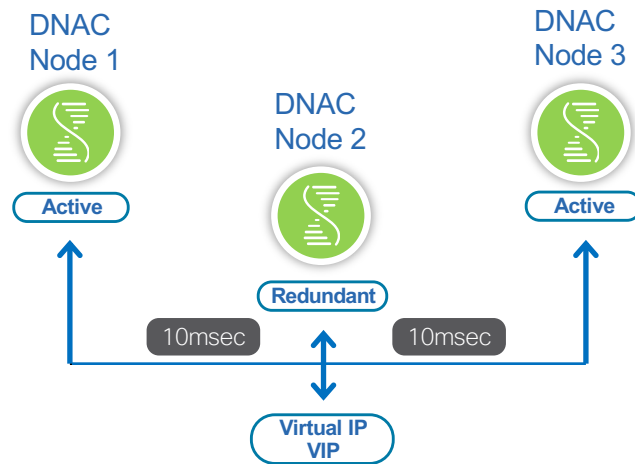
Cisco SD-Access Border Node	Cisco SD-Access Transit	IP Transit
c9K	●	●
ASR1K/ISR4K	●	●
C6K	●	●
N7K	●	●

Cisco DNAC Design Considerations



Cisco DNAC

Cisco DNA Center Design- Three Node High Availability



Important Call-outs

- **DNAC HA placement:** All DNAC should be in same DC, since the latency is 10ms (1-hop away)
- **Disaster Recovery:** Today DR offering is to restore the last known configuration to the DR site.
- **Automation vs Assurance:** Automation can be Active-Active, where as Assurance is Active-Passive.
- **DNAC deployment (single-node):** DNAC now supports VIP (required), if there is a single server make it cluster-ready.
- **DNAC deployment (2-nodes):** When running DNAC with 2 nodes, HA is not supported but the servers can be deployed in a cold stand-by mode.

Cisco DNAC Platform Support

Compatibility Information

Plug and Play (PnP)

Software Image Management (SWIM)
Upgrade OS on switches and WLC

Automation
Configuration on Underlay, Overlay and Policy

Device Type	Device Family	Recommended Software Version	Min. Supported Software Version ¹	Essentials	Advantage	Inventory	Topology	SWIM	PnP	Assurance	ENFV & Routing	EasyQoS	Patching (SMU)	IWAN	SDA
Switch	CAT2K	IOS 15.2(2)E8	IOS 15.2(2)E3	Y	N	Y	Y	Y	Y	Y	NA	Y	NA	N	NA
Switch	CAT3K	IOS-XE 16.6.1	IOS-XE 3.6.5E	Y	Y	Y	Y	Y	Y	Y	N	Y	N	N	Y
Switch	CAT4K	IOS-XE 3.10E	IOS-XE 3.6.5E	Y	Y	Y	Y	Y	Y	Y	NA	Y	NA	NA	Y
Switch	CAT6K	IOS 15.5.1 SY	IOS 15.5.1 SY	Y	Y	Y	Y	Y	Y	Y	NA	Y	NA	NA	Y
Switch	CAT9K	IOS-XE 16.6.3	IOS-XE 16.6.2	Y	Y	Y	Y	Y	Y	Y	NA	Y	Y	NA	Y
IOT Switch	IE 2K	IOS 15.2(6)E1	IOS 15.2(6)E1	Y	NA	Y	Y	Y	Y	Y	NA	NA	NA	NA	NA
IOT Switch	IE 3K	IOS 15.2(6)E1	IOS 15.2(6)E1	Y	NA	Y	Y	Y	Y	Y	NA	NA	NA	NA	NA
IOT Switch	IE 4K	IOS 15.2(6)E1	IOS 15.2(6)E1	Y	NA	Y	Y	Y	Y	Y	NA	NA	NA	NA	Y
IOT Switch	IE 5K	IOS 15.2(6)E1	IOS 15.2(6)E1	Y	NA	Y	Y	Y	Y	Y	NA	NA	NA	NA	Y
DC Switch	N7K	NX-OS 7.3.2	NX-OS 7.3.1	Y	Y	Y	Y	Y	NA	Y	N	N	N	N	Y
Router	ASR 1K	IOS-XE 16.3.5	IOS-XE 16.3.1	Y	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Router	ISR 11XX	IOS-XE 16.7.1	IOS-XE 16.6.1	Y	Y	Y	Y	Y	Y	Y	NA	N	NA	Y	N
Router	ISR 4K	IOS-XE 16.6.3	IOS-XE 3.16	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Wireless	AP (11n)	AireOS 8.5.130	AireOS 8.5.120	Y	Y	Y	Y	Y	Y	Y	N	Y	N	N	N
Wireless	AP (Outdoor)	AireOS 8.5.130	AireOS 8.5.120	Y	Y	Y	Y	Y	Y	Y	NA	NA	NA	NA	Y
Wireless	AP (Wave1)	AireOS 8.5.130, 8.7.106	AireOS 8.5.120	Y	Y	Y	Y	Y	Y	Y	N	Y	N	N	Y
Wireless	AP (Wave2)	AireOS 8.5.130, 8.7.106	AireOS 8.5.120	Y	Y	Y	Y	Y	Y	Y	NA	NA	NA	NA	Y
Wireless	Controller	AireOS 8.5.130, 8.7.106	AireOS 8.5.120	Y	Y	Y	Y	Y	NA	Y	N	Y	N	N	Y
Virtual	Varies	Varies	Varies	N	Y	Y	N	Y	Y	Y	Y	N	N	N	N

<https://www.cisco.com/c/en/us/support/cloud-systems-management/dna-center/products-device-support-tables-list.html>

Device Compatibility

<https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/compatibility-matrix.html>

SD-Access 1.2.x Hardware and Software Compatibility Matrix

SDA compatibility is supported only for the specific software versions listed in the table below:

Features	Hardware	SDA 1.2.8 ³	SDA 1.2.6 ²	SDA 1.2.5 ¹	SDA 1.2.4	SDA 1.2.3	SDA 1.2.2
Management	Cisco DNA Center	Cisco DNA Center 1.2.8	Cisco DNA Center 1.2.6	Cisco DNA Center 1.2.5	Cisco DNA Center 1.2.4	Cisco DNA Center 1.2.3	Cisco DNA Center 1.2.2
Identity	Identity Services Engine	ISE 2.4 Patch 2, Patch 5 ISE 2.3 Patch 1, Patch 2, Patch 4, Patch 5	ISE 2.4 Patch 2, ISE 2.3 Patch 1, Patch 2, Patch 4, Patch 5	ISE 2.4 Patch 2, ISE 2.3 Patch 1, Patch 2, Patch 4, Patch 5	ISE 2.4 Patch 2, ISE 2.3 Patch 1, Patch 2, Patch 4	ISE 2.4 Patch 2, ISE 2.3 Patch 1, Patch 2, Patch 4	ISE 2.4 Patch 1, ISE 2.3 Patch 1, Patch 2, Patch 4
Fabric Edge	Cisco Catalyst 9200 Series Switches	IOS XE 16.10.1s (9200L), IOS XE 16.10.1s (9200)					
	Cisco Catalyst 9300 Series Switches	IOS XE 16.10.1s, IOS XE 16.9.2s, IOS XE 16.9.1s, IOS XE 16.6.5, IOS XE 16.6.4s, IOS XE 16.6.4a***	IOS XE 16.9.2, IOS XE 16.6.5, IOS XE 16.6.4a,*** IOS XE 16.6.4s, IOS XE 16.9.1s	IOS XE 16.6.4s, IOS XE 16.9.1s	IOS XE 16.6.4s, IOS XE 16.9.1	IOS XE 16.6.4s	IOS XE 16.6.4
	Cisco Catalyst 9400 Series Switches	IOS XE 16.10.1s, IOS XE 16.9.2s, IOS XE 16.9.1s, IOS XE 16.6.5, IOS XE 16.6.4s, IOS XE 16.6.4a***	IOS XE 16.9.2, IOS XE 16.6.5, IOS XE 16.6.4a,*** IOS XE 16.6.4s, IOS XE 16.9.1s	IOS XE 16.6.4s, IOS XE 16.9.1s	IOS XE 16.6.4s, IOS XE 16.9.1	IOS XE 16.6.4s	IOS XE 16.6.4
		IOS XE 16.10.1s,	IOS XE 16.9.2				

Cisco ISE Design Considerations



Cisco Identity Services Engine design

Standalone and Distributed

- Applies to both physical and virtual deployment
- Compatible with load balancers



Lab and Evaluation

100 Endpoints





Small HA Deployment
2 x (PAN+MNT+PSN)

20,000 Endpoints



Small Multi-node Deployment
2 x (PAN+MNT), <= 5 PSN



Large Deployment
2 PAN, 2 MNT, <=50 PSN

500,000 Endpoints

ISE 2.4 is the recommendation

Long-term (LTR) “suggested release”

- <https://community.cisco.com/t5/security-blogs/announcing-the-quot-suggested-release-quot-status-of-ise-2-4/ba-p/3775587>
- <https://www.cisco.com/c/en/us/products/collateral/security/identity-services-engine/bulletin-c25-740738.html>

Announcing the "Suggested Release" status of ISE 2.4

Policy and Access

4942 VIEWS 40 HELPFUL 16 COMMENTS



yshchory Cisco Employee

01-08-2019 04:21 PM

Happy New 2019!

About a year ago, we have started a journey to make ISE even more the robust solution our customers expect it to be.

This journey is a journey everyone subscribed with - our Engineering team have and are investing a huge amount of resources to ensure that ISE's code is simply better, in terms of robustness and quality, our testing environments are better and continually improving, our processes are better in terms of maintaining high quality, and today we are announcing another milestone in this journey. ISE 2.4, our latest release, had made it to the "Suggested Release" milestone!



MENU



Products & Services / Security / Network Visibility and Segmentation / Cisco Identity Services Engine / Bulletins /

Cisco Identity Services Engine Software Release Lifecycle Product Bulletin

Download Print

Updated: May 24, 2018 Document ID: 1526605485745274

The Cisco® Identity Services Engine (ISE) plays a critical role in enforcing access policies and limiting exposure to a continuously evolving threat landscape. This landscape drives the need for constant innovation and a rapid release cadence. Delivering multiple releases in a short timeframe can be challenging to organizations that require long-term stability and predictability when planning deployments and upgrades. To address these needs, the Cisco ISE team is striving to implement a predictable release lifecycle, as described in this document.

Cisco ISE software release timelines


Cisco plans to release a new ISE software version approximately every 6 months: one in March or April ("spring release") and one in September or October ("fall release"). Each release will continue to be characterized by feature richness and software quality that address market requirements.

The March-April release will be designated a **Long-Term Release (LTR)**, and the September-October release will be designated a **Short-Term Release (STR)**.

The LTR will typically be even numbered, for example, 2.0, 2.2, 2.4, and so on.

The STR will typically be odd numbered, for example, 2.1, 2.3, and so on.

SNS-36xx appliances



What are we solving?

- Increased endpoint capacity per appliance and deployment
- [UCS M4](#) Feb 2019 End Of Sale

How do we solve it?

- New appliances based on UCS M5

Prerequisites

- Must be running ISE 2.6
- <http://cs.co/ise-feedback>

Cisco Identity Services Engine design

Many reference for ISE design

- Design Guide
- Cisco Live Design Session -- BRKSEC-3432

TECSEC-3416

Walking on solid ISE: advanced use cases and deployment best practices

DEVNET-2334

How to Operationalise Security with ISE and APIs

BRKSEC-2430

ISE Deployment Staging and Planning

LTRSEC-1655

Configuring ISE (Identity Service Engine) PIC (Passive Identity Connector)

BRKCLD-2412

Consistent Group-based Policy for On-premise, Hybrid & Multi-cloud with Cisco DNA Intent-based Networking

BRKSEC-3432

Advanced ISE - Architect, Design and Scale ISE for your production networks

Imran Bashir

Thursday 08:30-10:30

LTRSEC-2502

The Art of ISE posture, configuration and troubleshooting

BRKSEC-2059

Deploying ISE in a Dynamic Environment

BRKSEC-3229

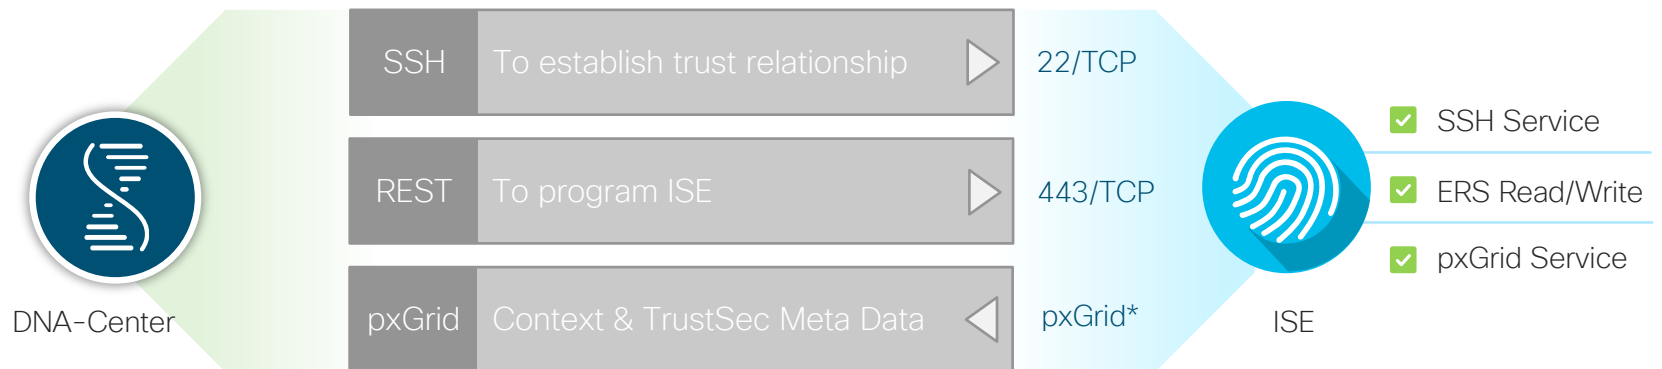
ISE under magnifying glass. How to troubleshoot ISE

Segmentation and Policy



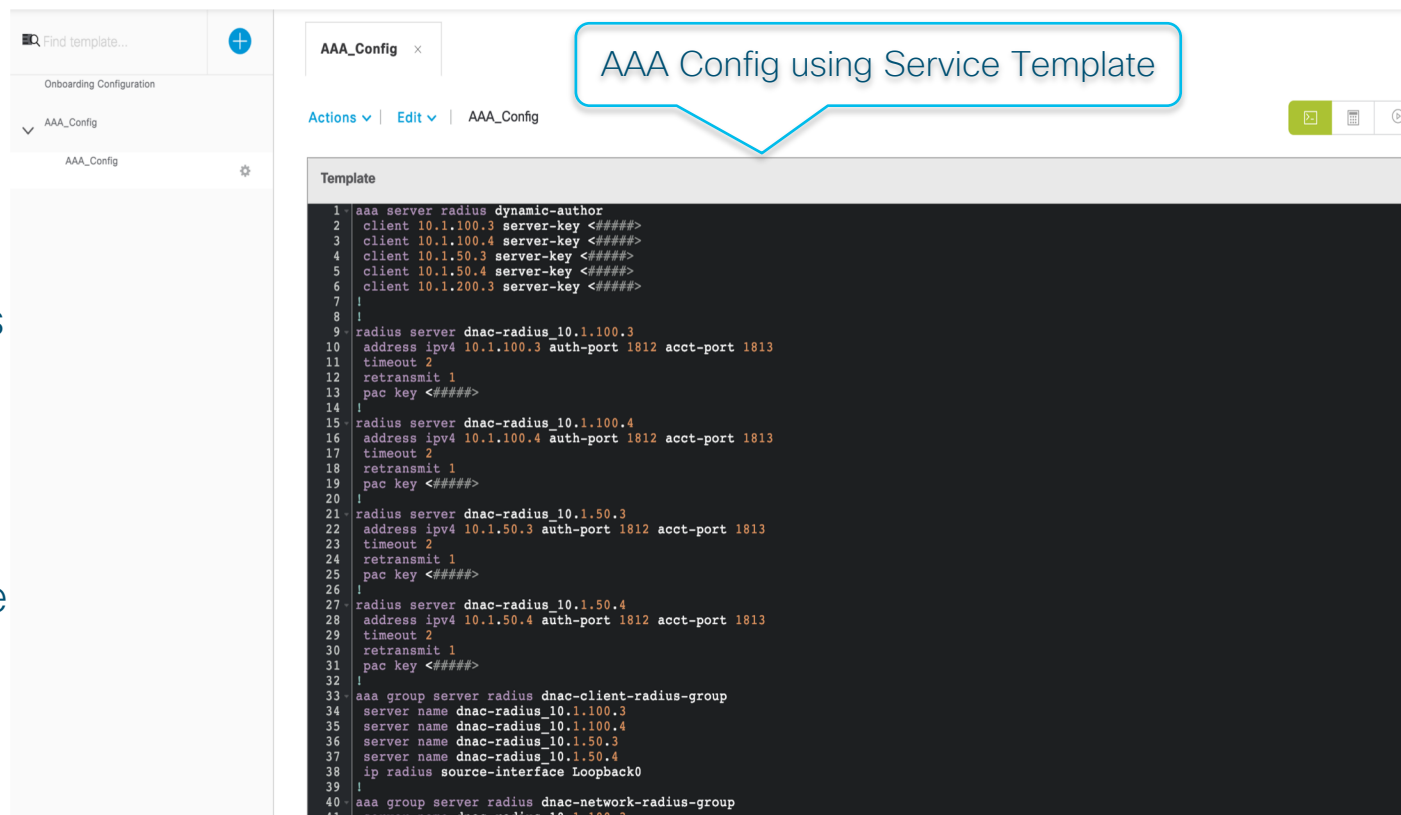
DNAC & ISE

Communication channels for integration



Additional option without Load Balancer

- Optionally use a Service Template and configure the necessary AAA configuration
- DNAC won't override this configuration during re-syncs
- Useful for deployments without an LB and for site survivability



The screenshot displays a network configuration interface. On the left, a sidebar shows a tree view with 'Onboarding Configuration' expanded to 'AAA_Config'. The main area shows a configuration editor for 'AAA_Config'. A callout box with a blue border and white background points to the title 'AAA Config using Service Template'. Below the title, a 'Template' section contains the following configuration:

```
1 aaa server radius dynamic-author
2 client 10.1.100.3 server-key <####>
3 client 10.1.100.4 server-key <####>
4 client 10.1.50.3 server-key <####>
5 client 10.1.50.4 server-key <####>
6 client 10.1.200.3 server-key <####>
7 !
8 !
9 radius server dnac-radius_10.1.100.3
10 address ipv4 10.1.100.3 auth-port 1812 acct-port 1813
11 timeout 2
12 retransmit 1
13 pac key <####>
14 !
15 radius server dnac-radius_10.1.100.4
16 address ipv4 10.1.100.4 auth-port 1812 acct-port 1813
17 timeout 2
18 retransmit 1
19 pac key <####>
20 !
21 radius server dnac-radius_10.1.50.3
22 address ipv4 10.1.50.3 auth-port 1812 acct-port 1813
23 timeout 2
24 retransmit 1
25 pac key <####>
26 !
27 radius server dnac-radius_10.1.50.4
28 address ipv4 10.1.50.4 auth-port 1812 acct-port 1813
29 timeout 2
30 retransmit 1
31 pac key <####>
32 !
33 aaa group server radius dnac-client-radius-group
34 server name dnac-radius_10.1.100.3
35 server name dnac-radius_10.1.100.4
36 server name dnac-radius_10.1.50.3
37 server name dnac-radius_10.1.50.4
38 ip radius source-interface Loopback0
39 !
40 aaa group server radius dnac-network-radius-group
41 server name dnac-radius_10.1.100.3
```

Assigning Endpoints to the Correct VN

Authorization Profiles > **iot_authz**

Authorization Profile

* Name

Description

* Access Type

Network Device Profile

Service Template

Track Movement

Passive Identity Tracking

Common Tasks

DAACL Name

ACL (Filter-ID)

Security Group Virtual Network:

VLAN

Map SGT to VN

Virtual Networks

- IoT
- DEFAULT_VN

Advanced Attributes Settings

SGT, VN and Address Pool Association


Security Groups List > IoT_Sensors

Security Groups

* Name

IoT_Sensors

* Icon



Description

Propagate to ACI

Tag Value

60

(Enter value between 2 and 65519)

Generation Id: 0

SGT+VN+IP Pool
Association

Associated Virtual Networks and Subnet/IP Address Pools

Virtual Network Name	Subnet/IP Address Pool Name	Type	Is Default	Max Value
IoT			false	
▼ DEFAULT_VN			true	
	10_202_2_0-DEFAULT_VN	Data	false	256

802.1X Switch Provisioning by Cisco DNAC

The screenshot displays the Cisco DNA Center interface for configuring Open Authentication. The top navigation bar includes 'Cisco DNA Center' and tabs for 'DESIGN', 'POLICY', 'PROVISION', and 'ASSURANCE'. A secondary navigation bar shows 'Network Hierarchy', 'Network Settings', 'Image Repository', and 'Network Policy'. The main content area is titled 'Open Authentication' and contains the following settings:

- Deployment Mode: Open
- First Authentication Order: 802.1x MAC Auth Bypass(MAB)
- 802.1x to MAB Fallback: A slider set to 21, with a range from 3 to 120.
- Wake on LAN: Yes No
- Number of Hosts: Single Unlimited

At the bottom of the configuration panel are 'Cancel' and 'Submit' buttons. On the left side, a sidebar titled 'AuthTemplate Method' lists three options: 'Closed Authentication', 'Easy Connect', and 'Open Authentication', with 'Open Authentication' selected.

Advanced Policy Options Using ISE

Advanced Policy options using ISE

Group-Based Access Control Policies Scalable Groups Access Contract

Last updated: 12:47 am Refresh **Advanced Options** Add Policy

Filter Edit Delete Deploy

<input type="checkbox"/> Policy Name ▲	Status	Description
<input type="checkbox"/> Cont2Cont	DEPLOYED	
<input type="checkbox"/> Emp2Dev	DEPLOYED	Emp to Dev policy to Deny WebAccess
<input type="checkbox"/> Emp2Emp	DEPLOYED	Emp to Emp policy to block malware
<input type="checkbox"/> GuestToWeb	DEPLOYED	Guest accessing the DMZ Web Server
<input type="checkbox"/> Prod2Emp	DEPLOYED	Production Server to Employee policy to Allow Web Access

Show 10 entries Showing 1 - 5 of 5 Previous 1 Next

Advanced Policy Concepts Using ISE

- Out of the box - Blacklist policy model
 - Traffic permitted unless specifically blocked
- Easy to move to whitelist model when ready
 - Populate matrix as needed
 - Change the Default Egress Policy to Deny

The screenshot displays the 'Production Matrix' interface in Cisco ISE. The interface shows a grid of source and destination policies. The source policies listed are Contractors, ACI_Production_..., ACI_Development..., Employees, Internet, and Unknown. The destination policies listed are ACI_Production_..., Contractors, and ACI_Development... The grid cells are colored based on the policy status: green for 'Enabled' and red for 'Disabled'. The 'Default' policy is shown as 'Enabled' with a description of 'Default egress rule'.

Source	ACI_Production_...	Contractors	ACI_Development...
Contractors	Enabled	Enabled	Enabled
ACI_Production_...	Enabled	Enabled	Enabled
ACI_Development...	Enabled	Enabled	Enabled
Employees	Enabled	Disabled	Enabled
Internet	Enabled	Enabled	Enabled
Unknown	Enabled	Enabled	Enabled

Default Enabled SGACLs : Permit IP Description : Default egress rule

SGACL Policy in ISE with Source SGT

- DNAC today has no Source SGT option in the policy definitions within the contracts
- Use ISE instead of DNAC to manage the SGACL policies to write Source SGT

Advanced Policy in ISE with Source SGT

Security Groups
IP SGT Static Mapping
Security Group ACLs
Network Devices
Trustsec AAA Servers

Security Groups ACLs List > Block_Malware

Security Group ACLs

* Name: Block_Malware Generation ID: 11

Description:

IP Version: IPv4 IPv6 Agnostic

* Security Group ACL content

```
deny tcp src dst eq telnet
deny udp src dst eq domain
deny tcp src dst eq 2280
deny tcp src dst eq 1433
deny tcp src dst eq 1521
deny tcp src dst eq 445 log
deny tcp src dst eq 137
deny tcp src dst eq 138
deny tcp src dst eq 139
deny udp src dst eq snmp
deny tcp src dst eq www
deny tcp src dst eq 443
deny tcp src dst eq 22
deny tcp src dst eq pop3
deny tcp src dst eq 123
deny icmp log
```

SGACL Policy in ISE with ICMP

- DNAC today has no ICMP option in the policy definitions within the contracts
- Use ISE instead of DNAC to manage the SGACL policies to allow or deny ICMP traffic

SGACL Policy in ISE to Filter ICMP Traffic

Security Groups ACLs List > Deny_ICMP

Security Group ACLs Generation ID: 1

* Name

Description

IP Version IPv4 IPv6 Agnostic

* Security Group ACL content

```
deny icmp log
```

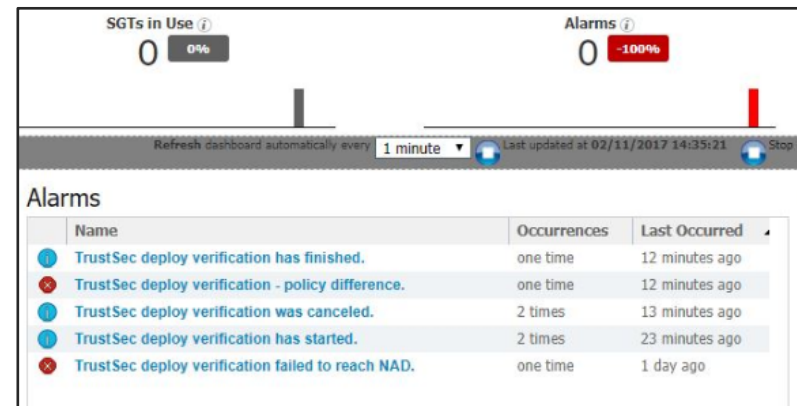

Deployment Verification

- Admin can see the status for a NAD when a new configuration change is being pushed.
- In IOS verifies SGTs/GEN-IDs and SGACLs/GEN-IDs along with ACEs
- In NX-OS verifies SGTs/Names and SGACLs along with ACE entries.
- Verification can be done automatically with new policy changes or manually.

Automatic verification after every deploy [?](#)

Time after deploy process minutes (10-60) [?](#)

[Verify Now](#)



Policy Deployment Validation

The screenshot displays the Cisco Identity Services Engine (ISE) interface. The top navigation bar includes 'Home', 'Context Visibility', 'Operations', 'Policy', 'Administration', and 'Work Centers'. A secondary navigation bar shows 'TrustSec' with sub-menus for 'BYOD', 'Profiler', 'Posture', 'Device Administration', and 'PassiveID'. A notification banner at the top right reads: 'Click here to do wireless setup and visibility setup Do not show this again.' The main content area is titled 'Trustsec Deployment Verification' and shows a report from 2018-10-30 00:00:00.0 to 2018-10-30 14:34:33.0. A table lists the verification results, including columns for 'Logged At', 'Verification ID', 'Message Code : Status', and 'Details'. A 'Click for Details Report' button is visible next to one of the entries. The bottom of the interface shows 'Rows/Page' set to 1 and '11 Total Rows'.

Logged At	Verification ID	Message Code : Status	Details
Today	Verification ID		
2018-10-30 14:34:29.204	d6715c0e-a51c-4faf-97c1-0fb06cfc8c90	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 14:22:18.803	f4462393-3cd3-4628-a2a7-ad69f39f2f52	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 14:21:53.772	6d0cc718-6b46-438d-af99-8eee79924746	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 03:12:25.23	696f5b99-1b65-467b-99b6-049992236e18	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 03:11:53.202	d9c72622-538d-40bd-89e0-aded14f902ce	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 03:11:19.02	e281b66c-2b80-4f8a-90b3-52716896ff47	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 02:29:15.867	618b87a8-8cd2-4784-b342-674d0068dff9	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 02:28:33.426	16f0ced5-aa6a-4955-bc62-ef151e320506	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 02:27:23.564	415b347e-ede1-4415-94fa-187e5f39b2cc	61029 : TrustSec deploy verification has finished.	Click for Details Report
2018-10-30 02:27:21.091	20932da6-3cc1-45d6-8521-ec465054d02d	61030 : TrustSec deploy verification was canceled.	Click for Details Report
2018-10-30 02:26:33.889	96545e06-f0ca-49e2-821e-1663785fe6a4	61029 : TrustSec deploy verification has finished.	Click for Details Report

CoA Push from PSN

- From ISE 2.4+ network administrator can push (CoA) changes from PSN
- Provides an option to pick the PSN from which the network device can receive the updates.
- Improves the performance in large scale deployments

▼ TrustSec Notifications and Updates

* Download environment data every

* Download peer authorization policy every

* Reauthentication every ⓘ

* Download SGACL lists every

Other TrustSec devices to trust this device

Send configuration changes to device Using CoA CLI (SSH)

Send from

Ssh Key

CTS Server List for SGACL Download

- Server List needed to be defined in ISE in case of multiple PSNs.
- Switch requests the policy from the first server (PSN) for the SGT it protects.
- Falls back to the next server when the first one goes down.
- Default server list will only have Primary PAN name and address.

The screenshot displays the Cisco Identity Services Engine (ISE) Administration interface. The top navigation bar includes 'Home', 'Context Visibility', 'Operations', 'Policy', 'Administration', and 'Work Centers'. The 'TrustSec' menu is expanded, showing 'BYOD', 'Profiler', and 'Posture'. The 'TrustSec Policy' menu is also expanded, showing 'Overview', 'Components', 'TrustSec Policy', 'Authentication Policy', 'Authorization Policy', 'SXP', 'Troubleshoot', 'Reports', and 'Settings'. The 'Components' menu is selected, and the 'TrustSec AAA Servers' option is highlighted in the left sidebar. The main content area is titled 'AAA Servers' and features a table with columns for 'Name', 'Description', and 'IP Address'. The table contains two entries: 'ise21-psn1' with IP address '10.200.100.95' and 'ise21-psn3' with IP address '10.200.100.94'. Above the table are action buttons for 'Edit', 'Add', 'Move Up', 'Move Down', 'Delete', and 'Push'. A 'Show' button is located in the top right corner of the table area.

<input type="checkbox"/>	Name	Description	IP Address
<input type="checkbox"/>	ise21-psn1		10.200.100.95
<input type="checkbox"/>	ise21-psn3		10.200.100.94

Verify SGACL Policy on IOS Switch

```
Switch#show cts role-based permissions
IPv4 Role-based permissions default:
  Permit IP-00
IPv4 Role-based permissions from group 3 to group 5:
  Deny IP-00
IPv4 Role-based permissions from group 4 to group 5:
  ALLOW HTTP HTTPS-20
IPv4 Role-based permissions from group 3 to group 20:
  Deny IP-00
IPv4 Role-based permissions from group 4 to group 6:
  Deny IP-00
IPv4 Role-based permissions from group 3 to group 7:
  Deny IP-00
IPv4 Role-based permissions from group 4 to group 7:
  Permit IP-00
```

SGACL policies could be statically defined on NAD

SGACL Mapping Policy should match to one on ISE

SGACL policies coming from ISE have precedence over static

Source Tree Destination Tree Matrix

Egress Policy (Source Tree View)

Edit + Add X Clear Mapping ⚙ Configure + Push ☹ Monitor All - Off

Source Security Group

BYOD (3/0003)

Source Inner Table

Status	Destination Security Group	Security Group ACLs	Description
<input checked="" type="checkbox"/> Enabled	Data_Center	Deny IP	

SGACL Monitoring – Best Effort Syslog

```
Switch#show cts role-based permissions
```

```
IPv4 Role-based permissions from group 8:EMPLOYEE_FULL to group 8:EMPLOYEE_FULL:  
Malware_Prevention-11
```

```
Switch#show ip access-list
```

```
Role-based IP access list Deny IP-00 (downloaded)
```

```
10 deny ip
```

```
Role-based IP access list Malware_Prevention-11 (downloaded)
```

```
10 deny icmp log (51 matches)
```

```
20 deny udp dst eq 445 log
```

```
30 deny tcp dst range 1 100 log
```

```
40 deny udp dst eq domain log
```

```
*May 24 04:50:06.090: %SEC-6-IPACCESSLOGDP: list Malware_Prevention-11 denied icmp  
10.10.18.101 (GigabitEthernet1/1 ) -> 10.10.11.100 (8/0), 119 packets
```

Verifying SGACL Drops

- Use show cts role-based counter to show traffic drop by SGACL

```
Switch#show cts role-based counters
Role-based IPv4 counters
```

From	To	SW-Denied	HW-Denied	SW-Permitted	HW Permitted
*	*	0	0	48002	369314
3	20	53499	53471	0	0
4	5	0	0	0	3777
3	6	0	0	0	53350
4	6	3773	3773	0	0
3	7	0	0	0	0
4	7	0	0	0	0

From * to * means Default Rule

- This show command displays the content stats of RBACL enforcement. Separate counters are displayed for both HW and SW switched packets. The user can specify the source SGT using the “from” clause and the destination SGT using the “to” clause.
- Mostly SGACL filtering is done in HW. Only if the packet needs to be punted to SW (e.g. TCAM is full, marked to be logged) , SW counter increments

Validating the SGT scale on Cat 3K/9K

```
3850#show platform hardware fed switch active fwd-asic resource tcam utilization
CAM Utilization for ASIC# 0
Table                                     Max Values          Used Values
-----
Unicast MAC addresses                    32768/512           56/23
Directly or indirectly connected routes  16384/7168        3107/96
L2 Multicast groups                      8192/512            0/7
L3 Multicast groups                      8192/512            0/9
QoS Access Control Entries               2816                 52
Security Access Control Entries         3072              211
Netflow ACEs                             768                  15
Input Microflow policer ACEs             256                   7
Output Microflow policer ACEs            256                   7
Flow SPAN ACEs                           512                   13
Control Plane Entries                    512                   272
Policy Based Routing ACEs                1024                   9
Tunnels                                  256                   13
Input Security Associations               256                   4
SPD                                       256                   2
Output Security Associations and Policies 256                   9
SGT_DGT                                4096/512         4060/512
CLIENT_LE                               4096/64              1/0
INPUT_GROUP_LE                           6144                  0
OUTPUT_GROUP_LE                           6144                  0
```

- IP/SGT Counter – 12K limit officially
- ACE Counter – ACEs are shared with like SGT/DGT
- SGT/DGT Hash table – Cells from the ISE Matrix

ISE 2.4 SXP Scaling Numbers

- Max ISE SXP nodes = 8 (four pairs in HA)
- Max ISE SXP peers = 200 for a SXP PSN
 - = 800 per ISE deployment (four SXP HA pairs)
- Max ISE SXP Binding = 350K per SXP PSN
 - = 1.4Million per ISE deployment (four SXP HA pairs)
- * HA here is Active - Active and they don't sync the mappings between the SXP nodes
- * These are the numbers for dedicated SXP nodes in ISE.

Cisco SDA Platform Support

Cisco SDA Platform Support

<https://www.cisco.com/c/en/us/support/cloud-systems-management/dna-center/products-device-support-tables-list.html>

Considerations for Bandwidth and Latency



Cisco Network Requirements

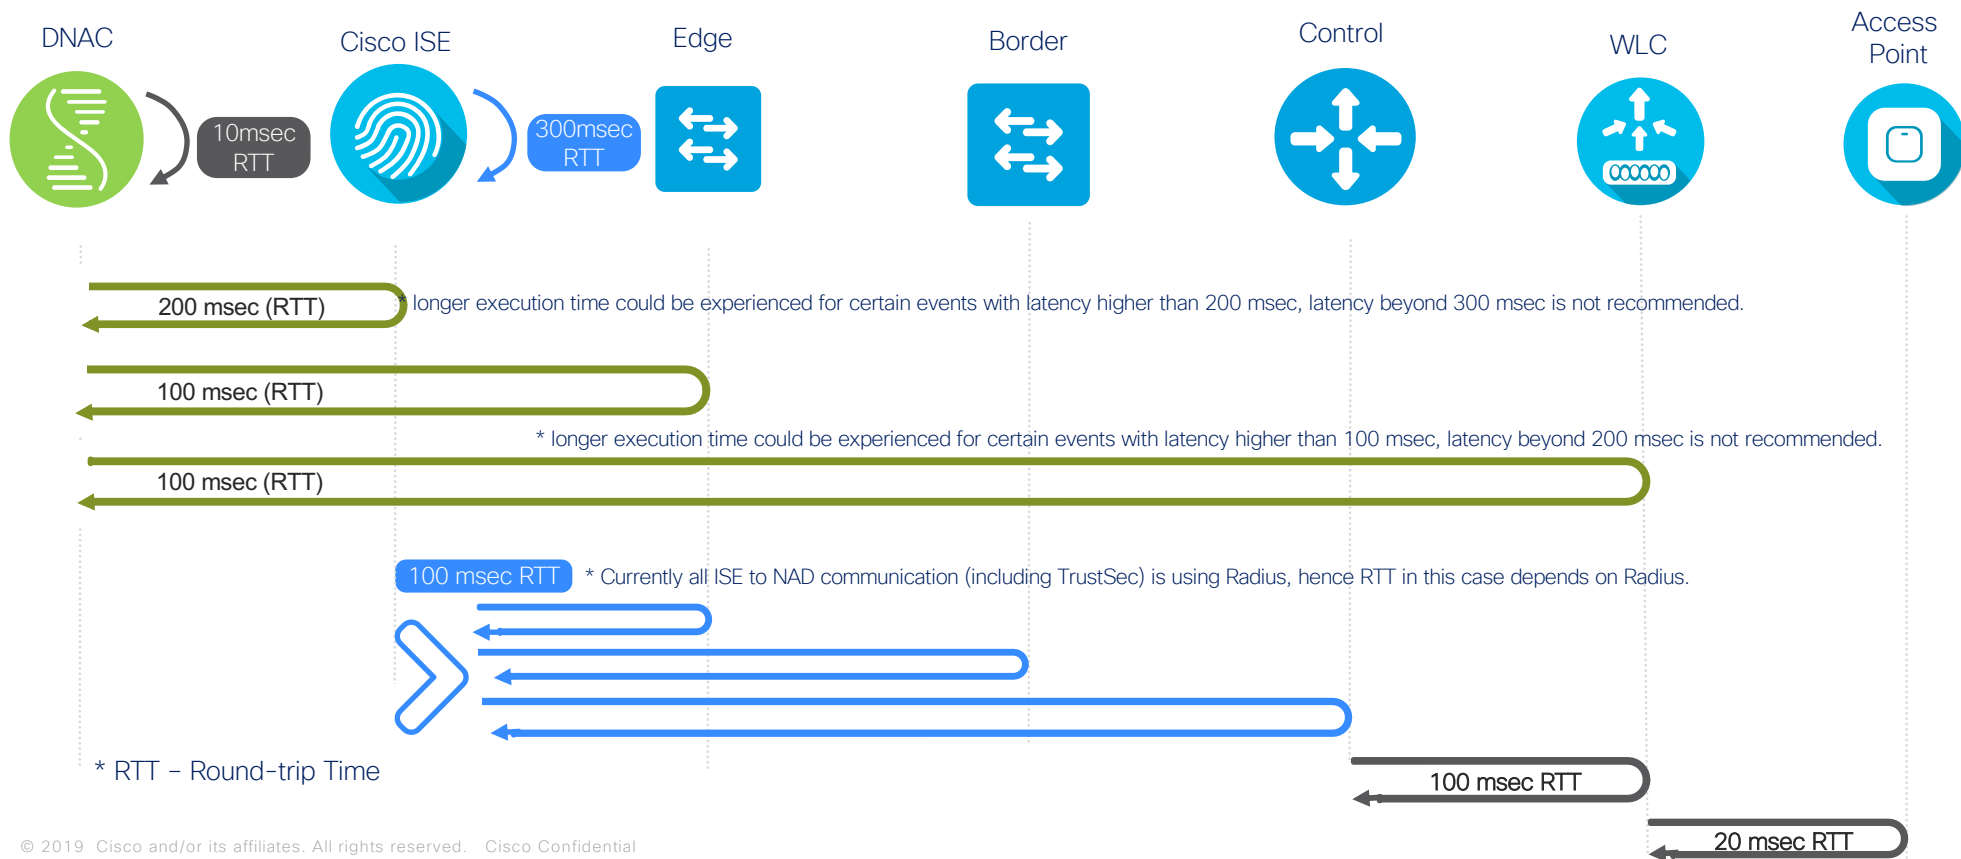
Latency Requirements (RTT)

In **Summary**, device latency should be around 100 msec RTT, you can go up to 200 msec RTT but there could be a performance hit. Anything beyond 200 msec is not recommended by Cisco at this time

The RTT (round-trip time) between Cisco DNA Center and network devices should be taken into consideration. The optimal RTT should be less than 100 milliseconds to achieve optimal performance for base automation and assurance. When RTT is between 100 milliseconds and 200 milliseconds, longer execution time could be experienced for certain events including Inventory Collection, Fabric Provision and Image Update, ranging from a few minutes to tens of minutes. Cisco does not recommend RTT more than 200 milliseconds.

Latency Requirements (RTT)

Cisco SD-Access Network Requirements



Cisco DNA Center Ports

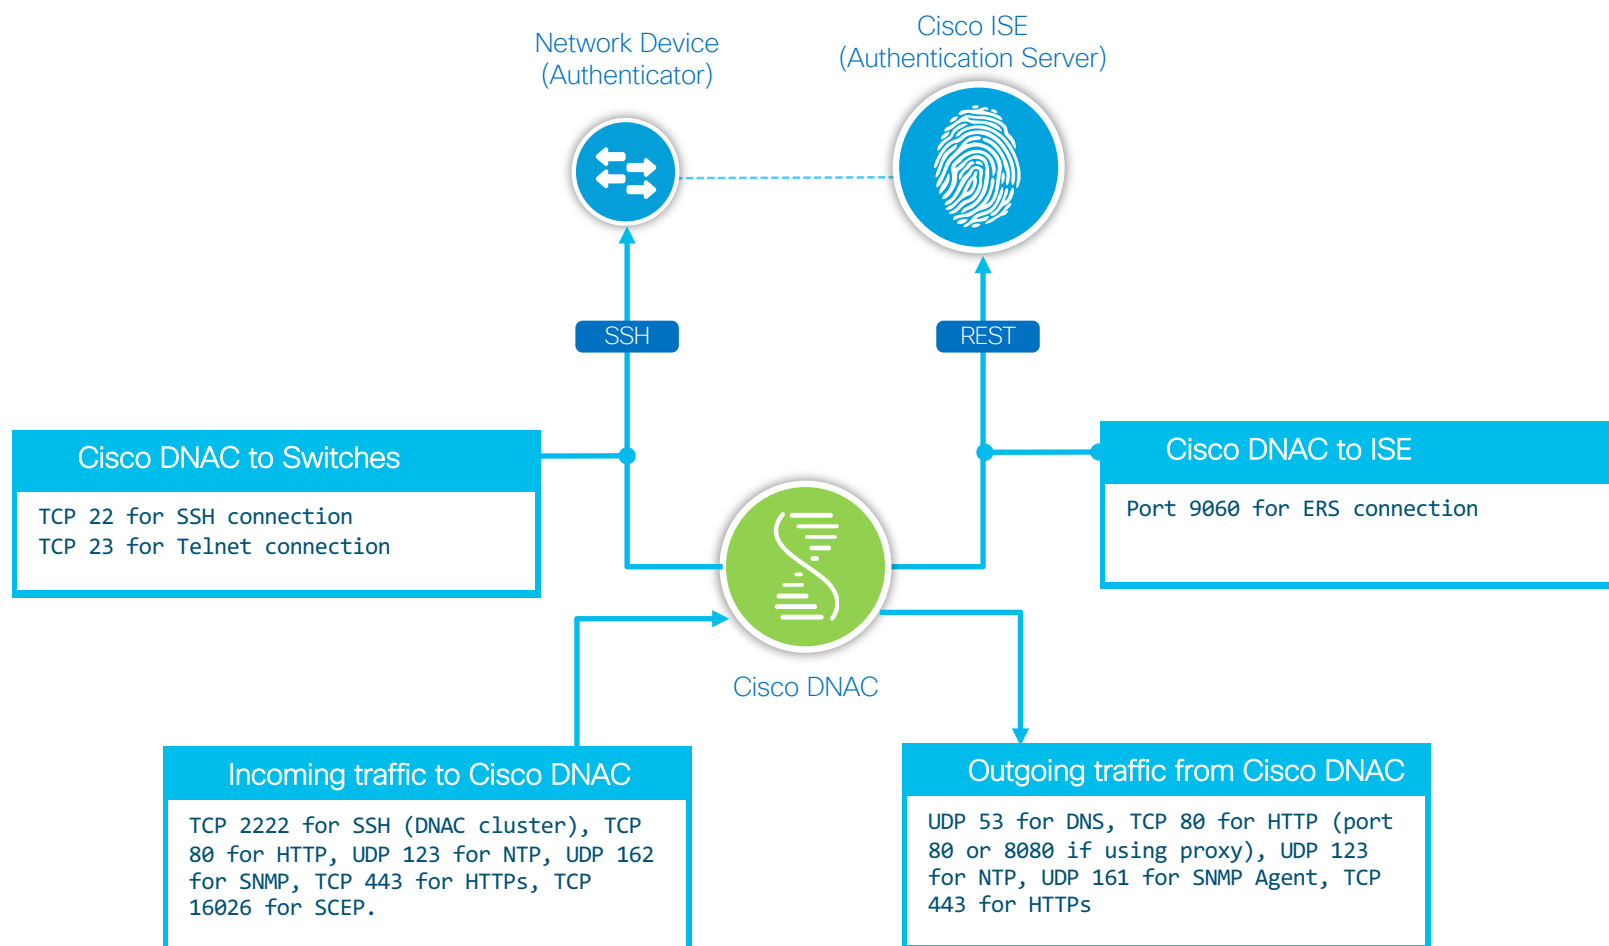
Is the appliance is behind Firewall

Cisco DNA-Center needs access to below URLs & FQDNs

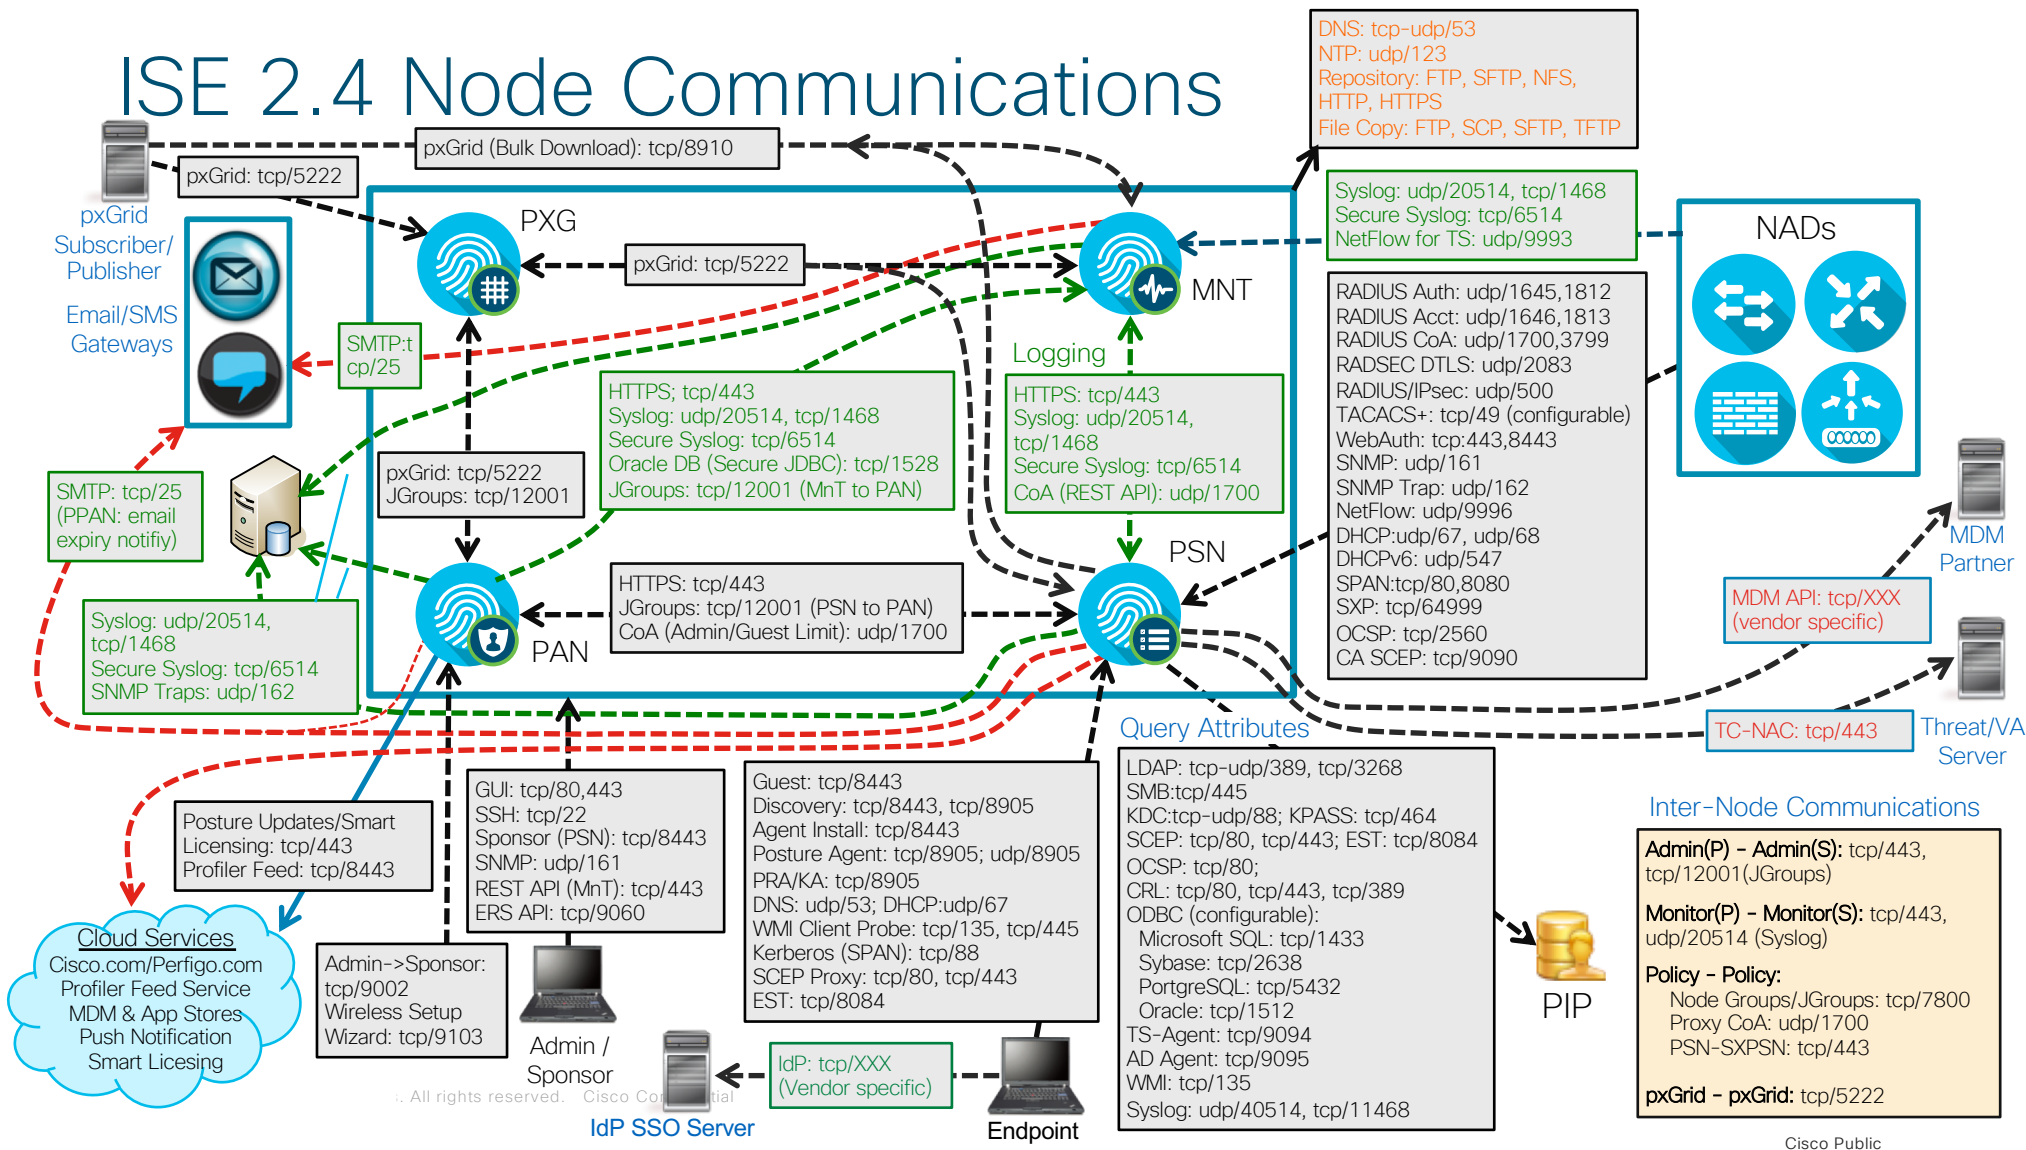
download System & Application package software	*.ciscoconnectdna.com:443
Integrate with cisco.com and Cisco Smart Licensing	*.cisco.com:443
Integrate with Cisco Meraki	*.meraki.com:443
Render accurate information in site & location maps	www.mapbox.com *.tiles.mapbox.com/* :443.

Note: Refer to the Cisco DNA Installation guide for more specific details

Cisco DNAC Node Communications



ISE 2.4 Node Communications

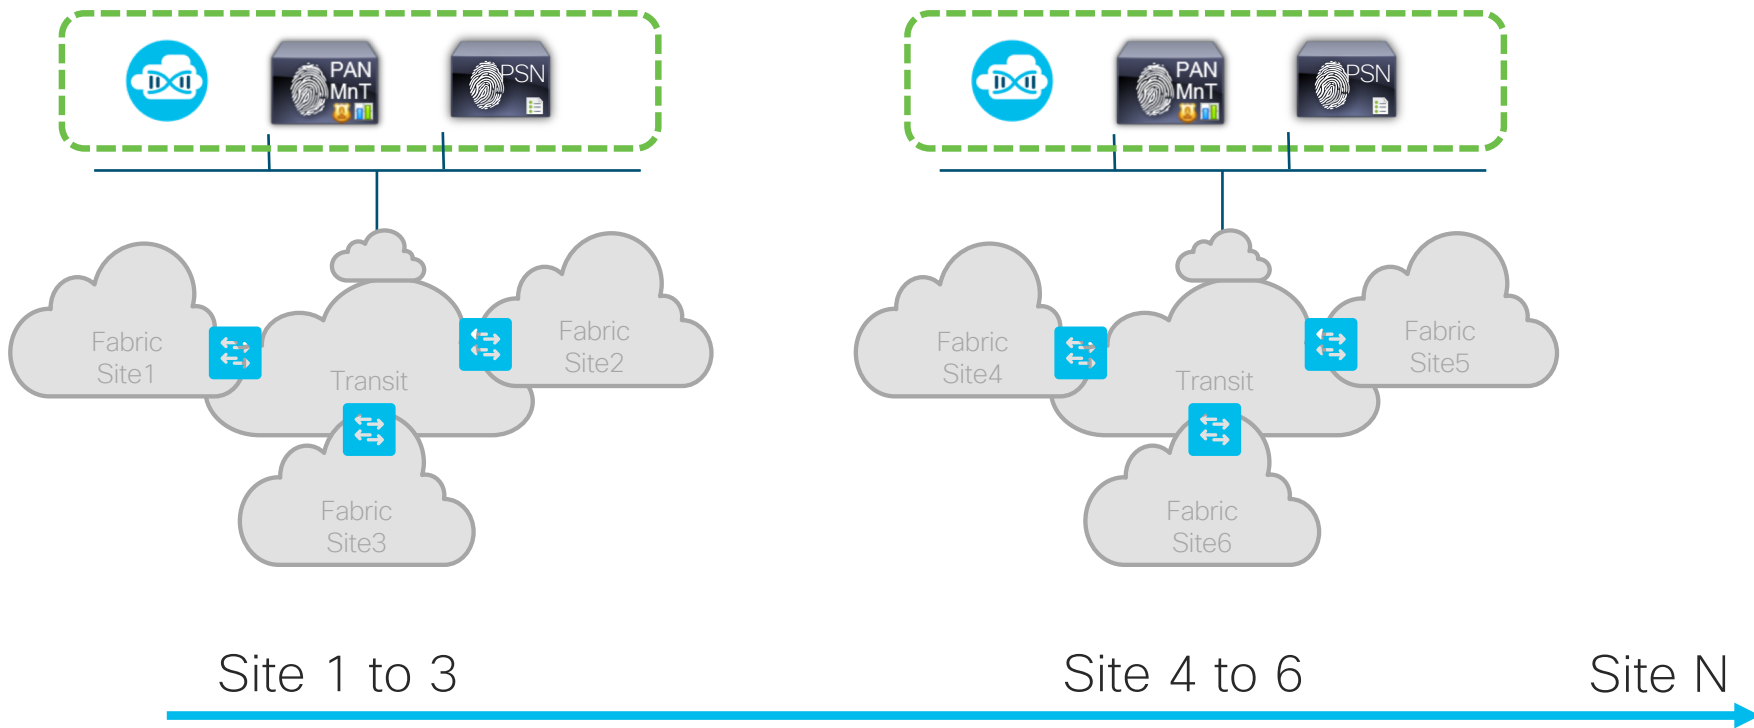


Cisco DNAC Best Practices



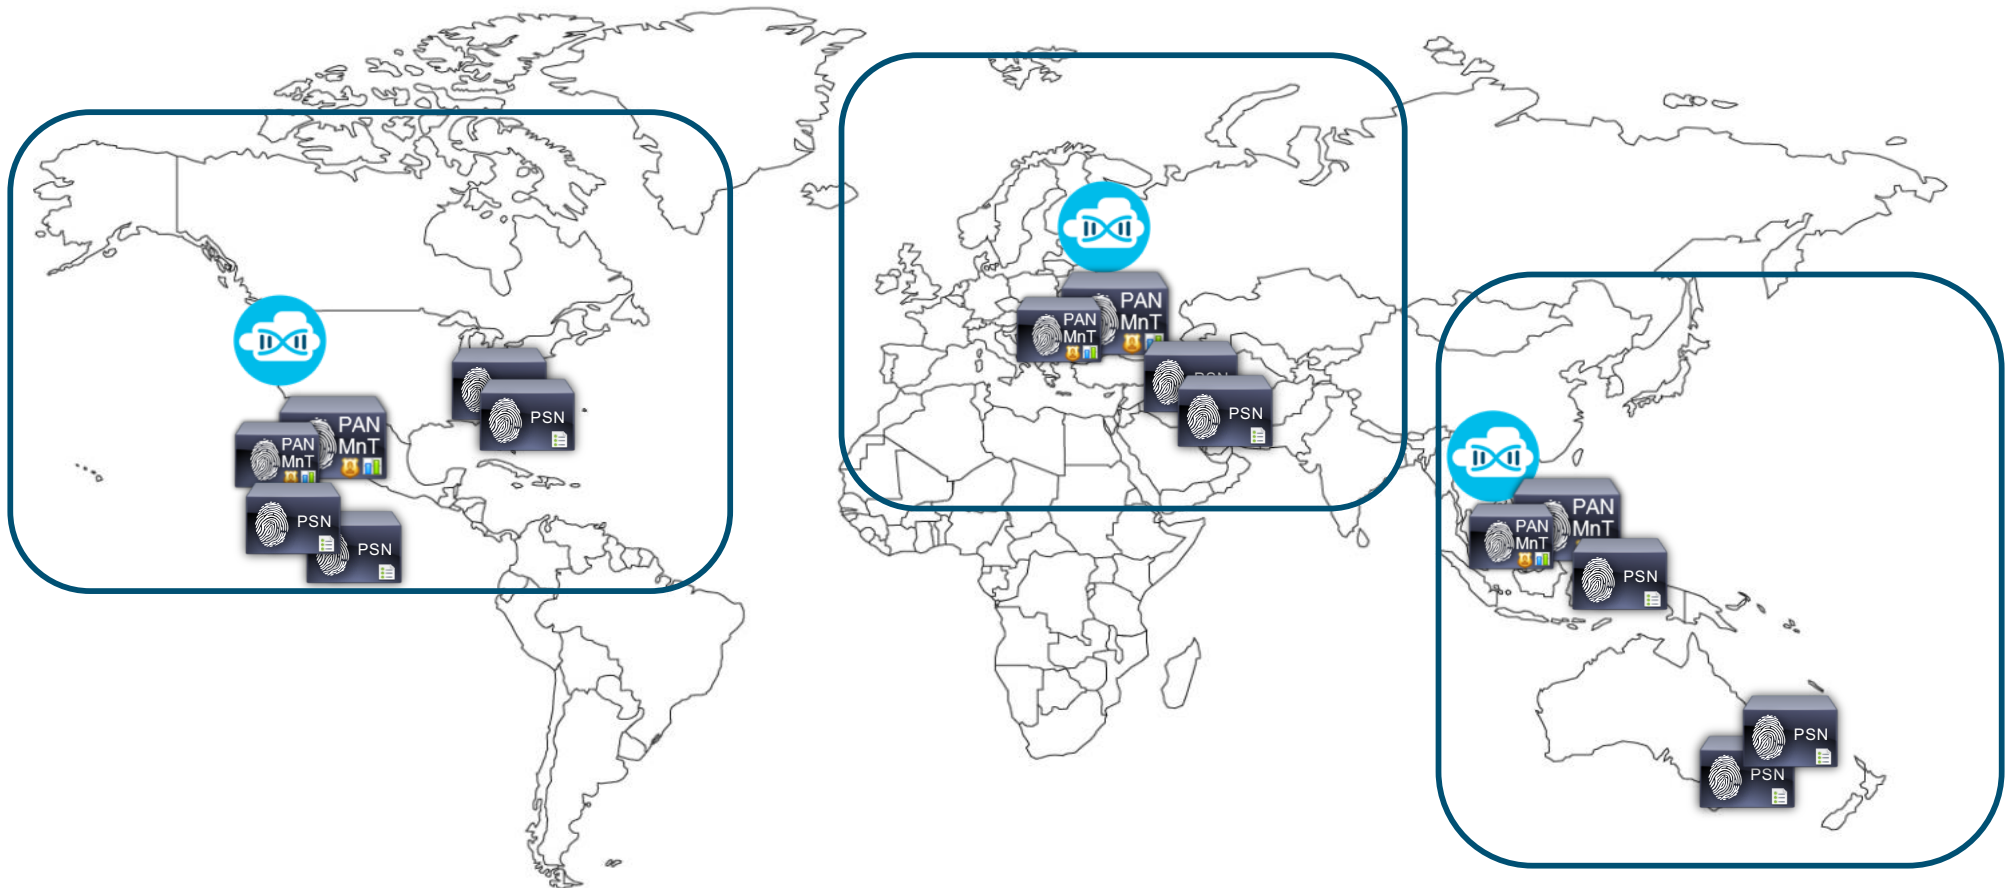
Scaling Strategy across Multiple Sites

Cisco DNAC Scale



Scaling Strategy across Multiple Sites

Cisco DNAC Scale



Recommendation for Cisco DNAC

Cisco DNA Center Design- Three Node High Availability

Users can choose to deploy DNA center as a single node or 3-node cluster.

- 3-node cluster deployment is for redundancy and to mitigate the split-brain problem.

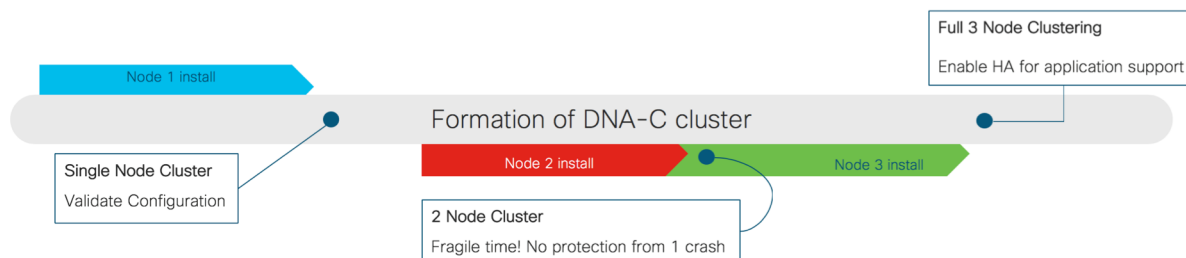
1. Bring up 1st DNAC node

- Complete the installation (Virtual IP, Intra-Cluster link) and let the services come up...

2. Bring up the 2nd DNAC node

- Let the installation complete

3. Bring up the 3rd DNAC node



Things to Remember:

- 2-node DNAC cluster cannot withstand a node failure
- A one node crash will lead to a stall of the other node

Recommendation for Cisco DNAC

Remove a Cisco DNAC node from Cluster

- If a node in a one of the node in cluster is in failed state and is not recovering after several hours, users should remove it from the cluster by running CLI : `$ maglev node remove <node_ip>`

Gracefully removing a node

- If for any reason, customer want to remove one of the active nodes in cluster, use the following steps:
- Move services on the given host another node by issuing:
`$ maglev node drain <node_ip>`
- Once all services are up and running, power down the node and remove it from the cluster: `$ maglev node remove <node_ip>`

Recommendation for Cisco DNAC

HA Command Cheat Sheet

HA commands:

- `maglev service nodescale status`
- `maglev service nodescale refresh`
- `maglev service nodescale progress`
- `maglev service nodescale history`
- `maglev node remove <node_ip>`
- `maglev node allow <node_ip>`
- `maglev cluster node display`

```
$ maglev cluster node display
```

```
maglev-1 [main - https://kong-frontend.maglev-system.svc.cluster.local:443]
```

ID	ADDRESS	PLATFORM
2ceaf148-1c73-11e9-ac26-380e4d37f009	9.9.10.102	DN1-HW-APL
3fd8700a-1d02-11e9-b858-40017afe6886	9.9.10.101	DN1-HW-APL
7fb3ec03-1d04-11e9-85ac-70df2ff7930c	9.9.10.100	DN1-HW-APL

Check All 3 nodes available

```
$ maglev service nodescale status
```

```
maglev-1 [main - https://kong-frontend.maglev-system.svc.cluster.local:443]
```

APPSTACK	SERVICE	CLUSTERED	ERROR
fusion	postgres	3/3	
maglev-system	cassandra	3/3	
maglev-system	elasticsearch	3/3	
maglev-system	glusterfs	3/3	
maglev-system	influxdb	3/3	
maglev-system	minio	3/3	
maglev-system	mongodb	3/3	
maglev-system	rabbitmq	3/3	
maglev-system	zookeeper	3/3	


Installing Cisco DNA Center

Important Considerations

- **Mandatory: NTP and DNS must be reachable** from the IP addresses used for DNA Center

(Note: Temporary Loopback can be used for DNS, but a real DNS server will be required after install.)

- Setup a single DNA Center node, as a Cluster node.
- With regard to Network Connectivity, DNA Center is simply a multi-homed appliance. (don't over complicate it) 😊



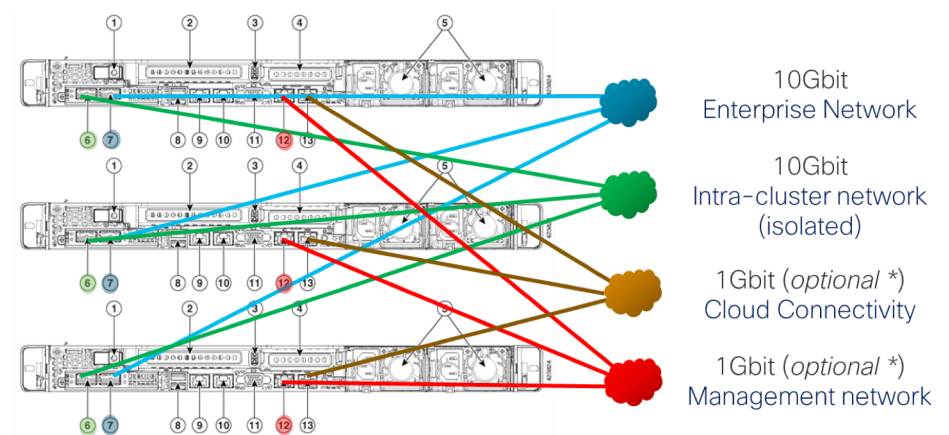
Config Wizard enumerates interfaces in alphabetical order

Network Patch Requirements for each DNAC Appliance

- Appliance Management port ["M" port] - CIMC (recommended)
- 10Gbit port [enp9s0] - Enterprise Network
- 10Gbit port [enp10s0] - Intra Cluster Link
- 1Gbit port [enp1s0f0] - Management (optional *)
- 1Gbit port [enp1s0f1] - Cloud Update (optional *)

(Interface #4 in Config Wizard)
 (Interface #1 in Config Wizard)
 (Interface #2 in Config Wizard)
 (Interface #3 in Config Wizard)

- Make sure that "Intra-Cluster Links" are connected to different switches !!!




* Required only if the Management network and/or the Cloud Update server is not reachable via the Enterprise Network

Installing Cisco DNA Center

Important Considerations

- As of DNA Center 1.2.5 VIPs are required on every interface



After upgrade, maglev config wizard will now prompt for users to configure a VIP for each IP'd interface

```
[Wed Oct 17 15:51:31 UTC] maglev@128.107.90.120 (maglev-master-1) ~  
$ sudo maglev-config update
```



Note: If the Cluster link is not up, the VIPs will not be enabled.

- DNA Center IP addressing (physical and VIPs) can be changed by “maglev-config update”
- DNA Center **will take time** to install and frequently needs to be immediately updated to get to the current release:

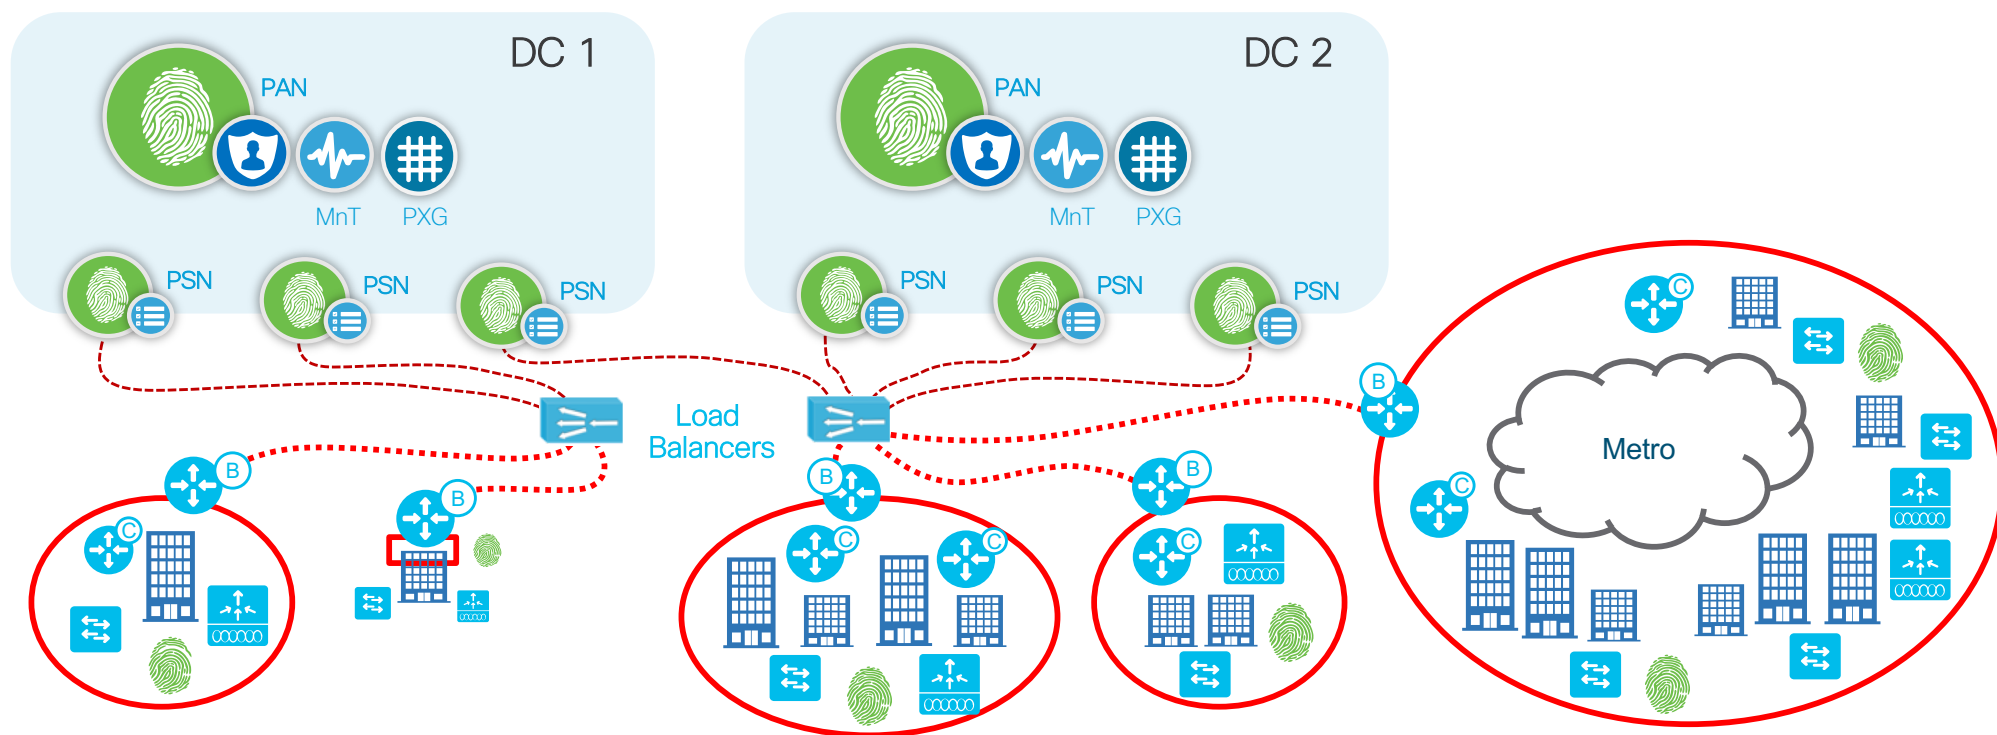
Best case 4-6 hours Worst case: up to 16 hours due to multiple upgrades required

Cisco ISE Best Practices



Scaling Strategy across Multiple Sites

Cisco ISE Scale Design



- PSN's are behind a dedicated Load Balancer
- DNAC site settings point to Load Balancer IP

ISE Integration for Automation

Important Considerations

ISE integration for Automation is mostly seamless

Under the following circumstance DNA Center Automation integration with ISE will fail:

- DNA Center or ISE IP Address has been changed since initial install
- DNA Center or ISE FQDNs have changed since initial install
- ISE VM has been cloned or restored

The fix: Re-gen and replace the ISE Root CA certificate and restart ISE

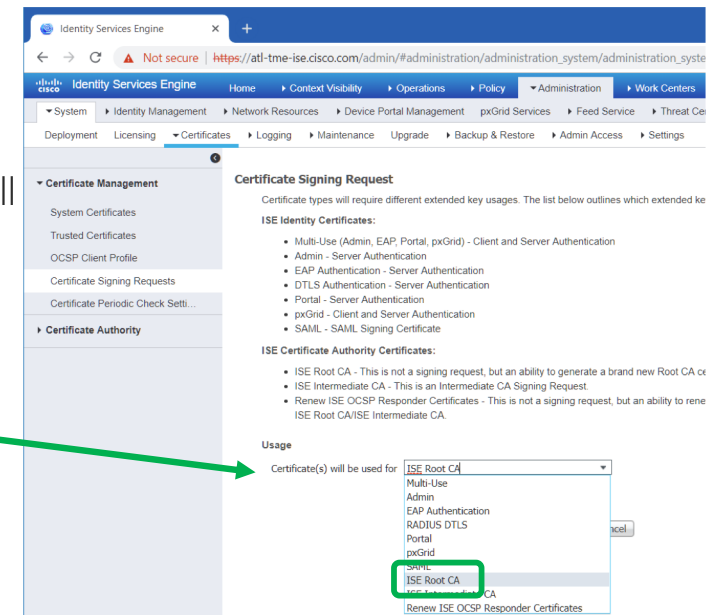
- ISE admin user has been changed since initial install

The fix: Ensure the ISE admin user is the same name for CLI and GUI and restart ISE
Verify APIs can be accessed using ERS SDK URL:
<https://<ISE-IP-Address>:9060/ers/sdk>

- ISE is not sync'd with NTP

The fix: Ensure the ISE, DNAC, AD, and other control devices are sync'd with NTP

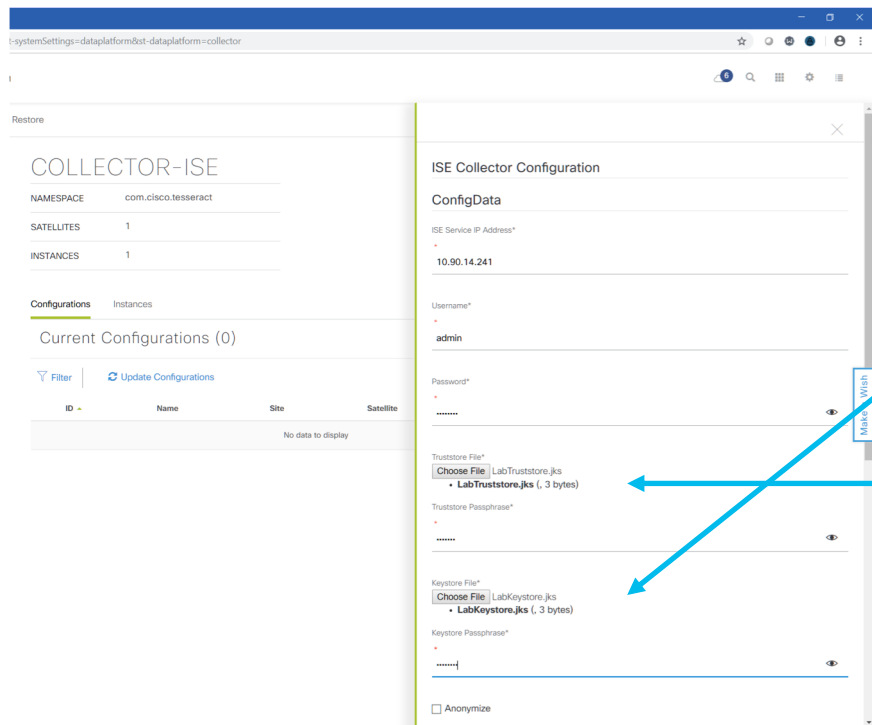
Best practice win POC ensure devices, DNAC, ISE, AD, etc use the same timezone



ISE Integration for Assurance

Important Considerations

ISE integration for Assurance is a manual process:
System Settings → Data Platform → ISE Collector



© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

It requires exporting ISE certificate, converting them from PEM to PKCS using OpenSSL and then to JKS format using Java Keytool

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center-assurance/1-2-5/b_dnac_assurance_1_2_5/dnac_assurance_1_2_chapter_01.html?bookSearch=true#id_71310

Export PxGrid Cert from ISE

Create KeyStore file in JKS format

a. Extract the Alias Name:
\$ keytool -v -list -storetype pkcs12 -keystore DNAC.cisco.com_192.168.0.1.p12 -storepass Cisco123 | grep -i alias
name: dnac.cisco.com_192.168.0.1

```
$ keytool -importkeystore -srckeystore DNAC.cisco.com_192.168.0.1.p12 -srcstoretype pkcs12 -srcalias dnac.cisco.com_192.168.0.1 -destkeystore keystore.jks -deststoretype jks -deststorepass Cisco123 -destalias Keystore
```

Create Truststore file in JKS format

```
$ keytool -importcert -file CertificateServicesEndpointSubCA-dnac-ise-01_cer -keystore truststore.jks -alias CertificateServicesEndpointSubCA-dnac-ise-01_
```

Enter keystore password:

Re-enter new password:

Owner: CN=Certificate Services Endpoint Sub CA - dnac-ise-01

Issuer: CN=Certificate Services Node CA - dnac-ise-01

Serial number: 9a91659bf1546c19e8ccd43fb4b6b62

Valid from: Sat Sep 16 11:59:43 PDT 2017 until: Fri Sep 17 11:59:40 PDT 2027

Certificate fingerprints:

SHA1: 6D:F5:B6:8F:E2:21:D5:91:44:23:28:7B:59:71:34:23:03:8F:F2:99

SHA256: 1A:8D:54:73:48:E6:2B:40:8A:64:AB:04:98:40:C9:C0:EB:07:28:54:C4:0C:4F:DD:7D:66:FA:5B:EB:C6:54:ED

Signature algorithm name: SHA256withRSA

Subject Public Key Algorithm: 4096-bit RSA k

Version: 3

<...snip...>

Trust this certificate? [no]: yes

Certificate was added to keystore

SDA Fabric Best Practices



What to know about IPAM Integration

- DNA Center IPAM integration supports both Infoblox and BlueCat
- The integration is very straight forward
- Attributes are not exchanged between DNA Center and IPAMs at this time.

Applies to:

- Gateways
- DHCP
- DNS
- Reserved scopes

The screenshot shows the 'IP Address Manager' configuration page in the Cisco DNA Center interface. The page is titled 'IP Address Manager' and includes a search bar and a navigation menu. The main content area contains a form for configuring IPAM settings. The form includes the following fields:

- Server Name***: ATL-TME-InfoBlox
- Server Uri***: https://10.90.14.249
- Username***: admin
- Password***: [Redacted]
- Provider***: INFOBLOX
- View***: default

At the bottom of the form, there are two buttons: 'Apply' and 'Delete'.

What to know about Discovery

- DNA Center Discovery defaults to CDP with a level of 16.
- Could easily discover an entire Enterprise by mistake.

The screenshot shows the Cisco DNA Center Discovery interface. On the left, a sidebar shows a search for '135 Reachable Devices' under 'CDP 192.0.4.1'. The main area displays a 'DEVICES STATUS' summary with a green circle around '135 Devices'. Below this, 'Discovery Details' are shown, including 'CDP Level: 16'. A red box highlights the '135 Devices' count, and a green box contains the text 'Should be 8!'. At the bottom, a green callout box states: 'Best practice use a Discovery Range and manually enter lab IP ranges'. The interface also shows a table of discovered devices with columns for IP Address, Device Name, Status, and various protocols.

New Discovery

Discovery Name*

IP Address/Range*

Discovery Type *i*

CDP Range LLDP

IP Address* *i*

Subnet Filters *i* +

CDP Level *i*

16

Preferred Management IP *i*
None

What to know about Device Credentials

- As of DNA Center 1.2.5 cisco /cisco are no longer supported for SD-Access Automation!
- The new password should not be “cisco”, “ocsic”, or any variant obtained by changing the capitalization of letters therein, or by substituting “1”, “|”, or “!” for i, and substituting “0” for “o”, and substituting “\$” for “s”.
- LAN Automation will not accept user cisco or passwords with common cisco permutations
- Enable password must be set
- **To Test** from DNA Center
 - SSH to device & Authenticate
 - Verify the user can access exec mode


Best practice use
“service password-encryption”

Network Infrastructure – Underlay

SD-Access underlay options

Manual Underlay

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



Automated Underlay

- Discover Seed Device
- Input IP Address Pool
- Start LAN Automation
 - ✓ Discover the network device
 - ✓ Onboard the network device
 - ✓ Upgrade software
- Stop LAN Automation
 - ✓ Complete Configuration (L3 interface, IS-IS)
 - ✓ Manage Device in Cisco DNAC-Center

Border connectivity Best-Practice

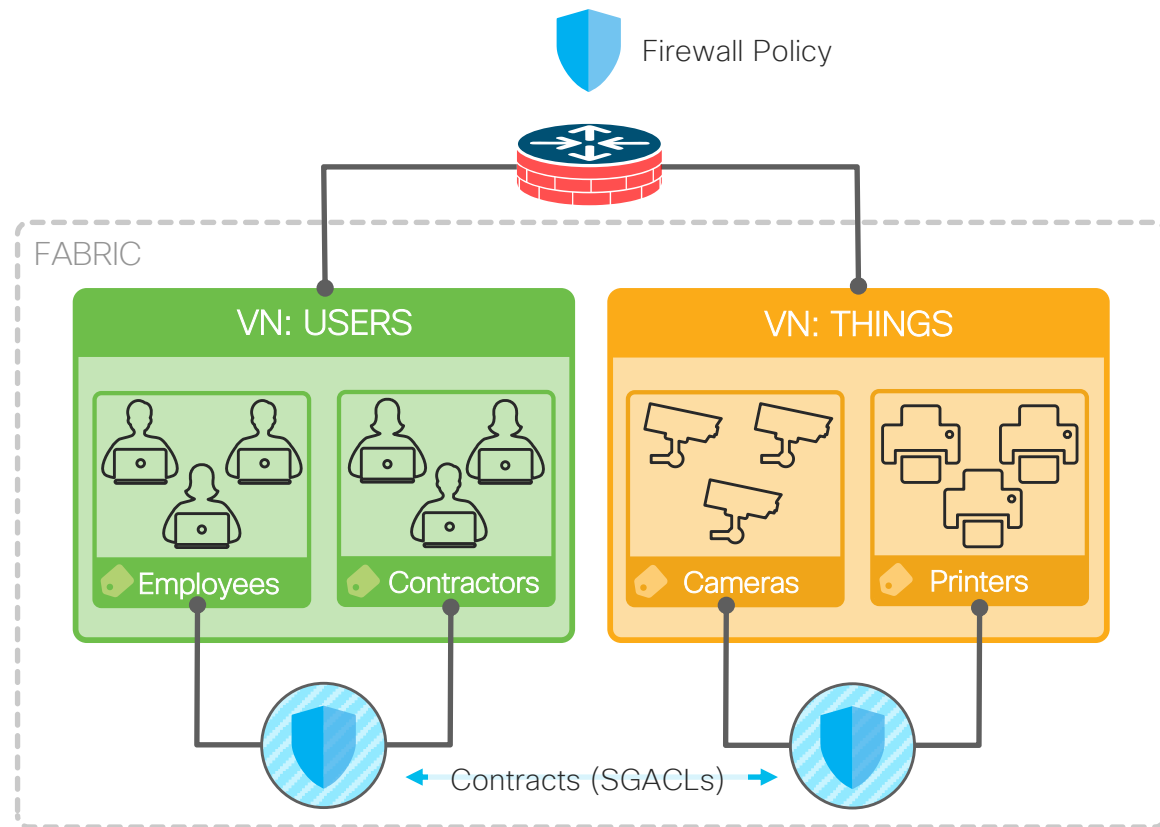
SDA enables Macro and Micro-segmentation

Inter-VN routing and policy enforcement on 'Fusion Router'

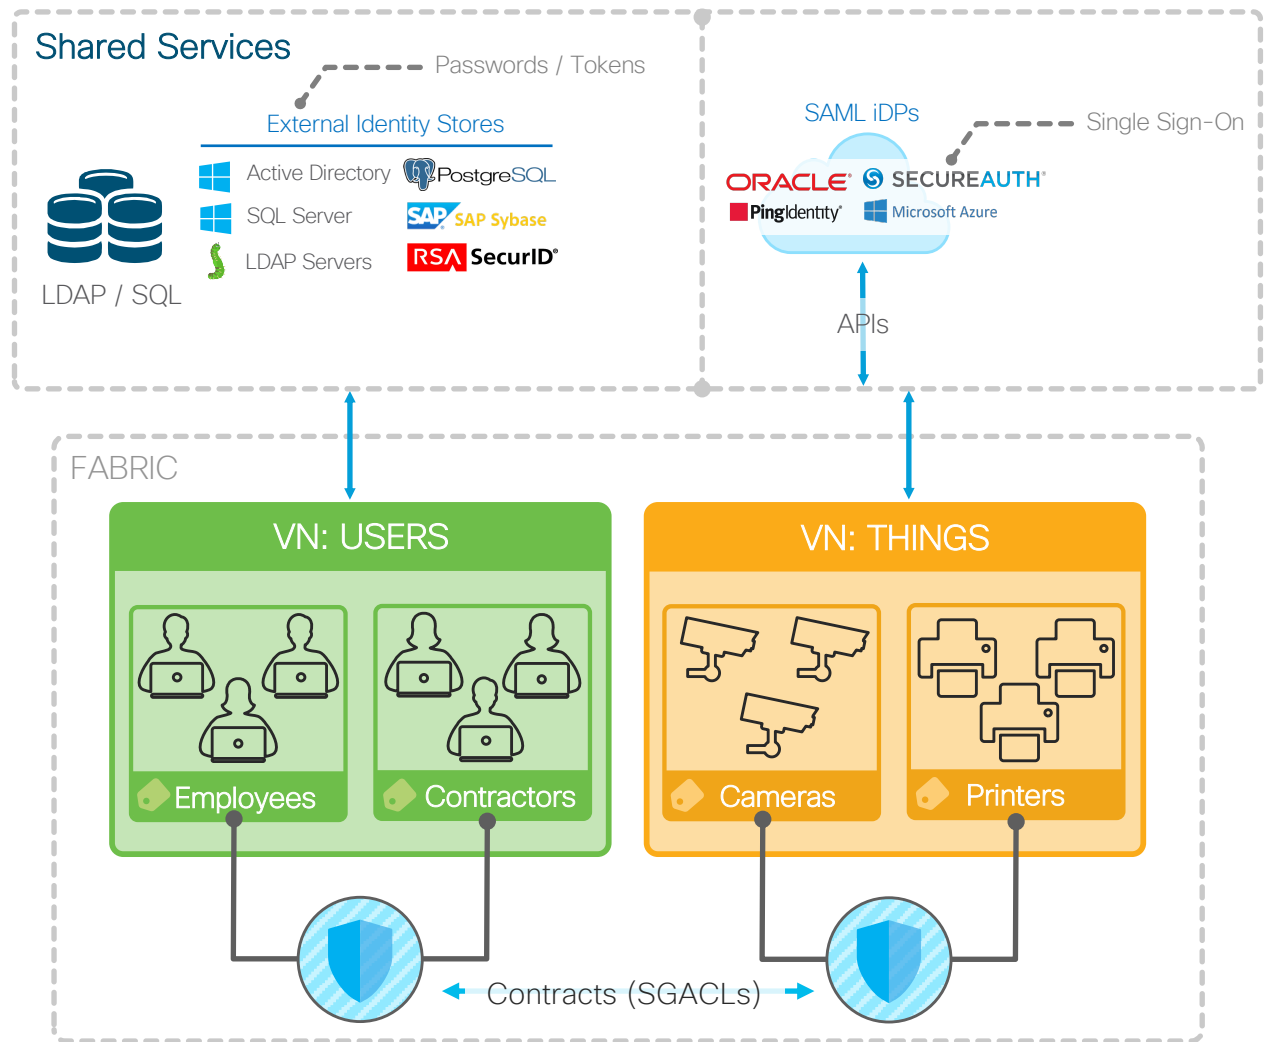
Macro segmentation with 'Virtual Networks'

Micro segmentation with 'Scalable Groups'

Contracts control access between SGTs



Objective:
Need access to Shared Services




Fusion Configuration

Connecting Fabric to Traditional Infrastructure


Extend

- Configure VRF
- Interfaces for each VN matching Border configuration



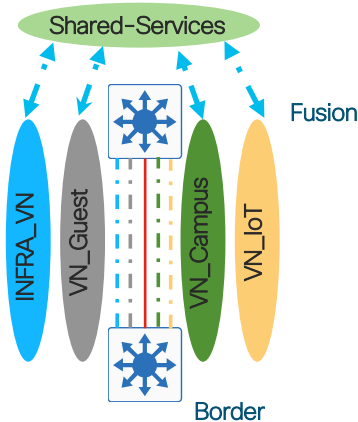
eBGP

- eBGP neighbors for each VN between Fusion and Border




Route Leak

- Route-leak shared-services subnets to each VN
- Route-leak VN subnets into Global



iBGP

- iBGP neighbors for each VN between Border nodes



- If Border / Fusion network device is Routing platform, L3 sub-interfaces will be used to extend Virtual Networks
- If Border / Fusion network device is Switching platform, VLANs & Trunk will be used to extend Virtual Networks


Fusion Configuration

Step 1: Extend

- Examine the below configs on the Fabric Border(s)
 - show running-config | section vrf definition
 - show running-config | section interface Vlan
 - show running-config | section interface <interface>

(OR)

- Navigate to DNAC --> Provision --> Fabric --> Fabric Site.
Select Border Node -> View Device Info option and drill down on the interface information.
- On the Fusion Device
 - Configure vrf matching Border Configuration.
 - Configure sub-interface(s) / Vlan(s) matching Border Configuration.



Fusion Configuration

Step 1: Extend – Fusion Node Configuration


• Step 1.1 – configure VRF

```
vrf definition Campus_VN
  rd 1:4099
  !
  address-family ipv4
    route-target export 1:4099
    route-target import 1:4099
  exit-address-family
!
vrf definition Guest_VN
  rd 1:4100
  !
  address-family ipv4
    route-target export 1:4100
    route-target import 1:4100
  exit-address-family
!!
vrf definition IoT
  rd 1:4101
  !
  address-family ipv4
    route-target export 1:4101
    route-target import 1:4101
  exit-address-family
```

• Step 1.2 – configure interface

```
!
interface gig 0/0/2.30xx
  description vrf interface to Border1-9500
  vrf forwarding Campus_VN
  encapsulation dot1q 30xx
  ip address 172.16.15.xx 255.255.255.252
  no ip redirects
  ip route-cache same-interface
  no shut
!
interface gig 0/0/2.30xx
  description vrf interface to Border1-9500
  vrf forwarding Shared_Services
  encapsulation dot1q 300x
  ip address 172.16.15.xx 255.255.255.252
  no ip redirects
  ip route-cache same-interface
  no shut
!
interface gig 0/0/2.30xx
  description vrf interface to Border1-9500
  vrf forwarding IoT
  encapsulation dot1q 30xx
  ip address 172.16.15.xx 255.255.255.252
  no ip redirects
  ip route-cache same-interface
  no shut
!
interface gig 0/0/2.30xx
  description vrf interface to Border1-9500
  vrf forwarding Guest_VN
  encapsulation dot1q 30xx
  ip address 172.16.15.xx 255.255.255.252
  no ip redirects
  ip route-cache same-interface
  no shut
!
```

Note: INFRA_VN on Border node maps to Shared_Services on Fusion node




Fusion Configuration

Step 2: eBGP – Fusion Node Configuration

- Configure BGP configuration to from eBGP neighbor with Border.

```
router bgp 65000
!
address-family ipv4 vrf Shared_Services
neighbor 172.16.15.xx remote-as 65001
neighbor 172.16.15.xx update-source gig0/0/2.30xx
neighbor 172.16.15.xx activate
neighbor 172.16.15.xx remote-as 65001
neighbor 172.16.15.xx update-source gig0/0/3.30xx
neighbor 172.16.15.xx activate
network 172.16.15.xx mask 255.255.255.252
network 172.16.15.xx mask 255.255.255.252
maximum-paths 2
exit-address-family
!
address-family ipv4 vrf Campus_VN
neighbor 172.16.15.xx remote-as 65001
neighbor 172.16.15.xx update-source gig0/0/2.30xx
neighbor 172.16.15.xx activate
neighbor 172.16.15.xx remote-as 65001
neighbor 172.16.15.xx update-source gig0/0/3.30xx
neighbor 172.16.15.xx activate
network 172.16.15.0 mask 255.255.255.252
network 172.16.15.40 mask 255.255.255.252
maximum-paths 2
exit-address-family
!
address-family ipv4 vrf Guest_VN
neighbor 172.16.15.xx remote-as 65001
neighbor 172.16.15.xx update-source gig0/0/2.30xx
neighbor 172.16.15.xx activate
neighbor 172.16.15.xx remote-as 65001
neighbor 172.16.15.xx update-source gig0/0/3.30xx
neighbor 172.16.15.xx activate
network 172.16.15.xx mask 255.255.255.252
network 172.16.15.xx mask 255.255.255.252
maximum-paths 2
exit-address-family
!
```

Note: INFRA_VN on Border node maps to Shared_Services on Fusion node

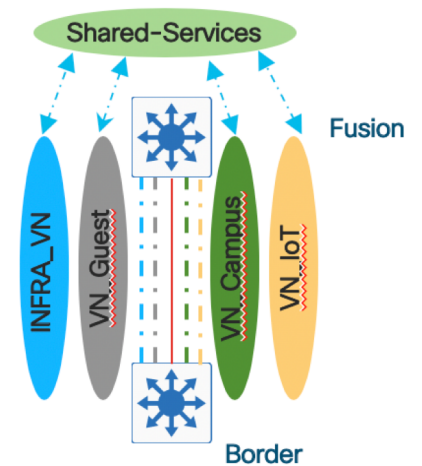


Fusion Configuration

Step 3: Route Leak - Fusion Node Configuration

- Controlled Route-leak between Global / Shared_Services-vrf and Fabric-vrf
 - Redistribute VN routes to Global / Shared_Services-vrf.
 - Redistribute Shared_Services / Global to VN.

```
!
ip prefix-list SHARED_SERVICES_NETS seq 5 permit 10.172.3.0/24
!
$$ SHARED_SERVICES_NETS - 10.172.3.0/24 contains ISE, DHCP, DNS in this subnet $$
!
route-map SHARED_SERVICES_NETS permit 10
match ip address prefix-list SHARED_SERVICES_NETS
!
vrf definition Campus_VN
rd 1:4099
!
address-family ipv4
import map SHARED_SERVICES_NETS
route-target export 1:4099
route-target import 1:4099
route-target import 100:100
exit-address-family
!
vrf definition Guest_VN
rd 1:4100
!
address-family ipv4
import map SHARED_SERVICES_NETS
route-target export 1:4100
route-target import 1:4100
route-target import 100:100
exit-address-family
!
vrf definition Shared_Services
rd 100:100
!
address-family ipv4
route-target export 100:100
route-target import 100:100
route-target import 1:4099
route-target import 1:4100
route-target import 1:4101
exit-address-family
```




Note: INFRVA_VN on Border node maps to Shared_Services on Fusion node

Fusion Configuration

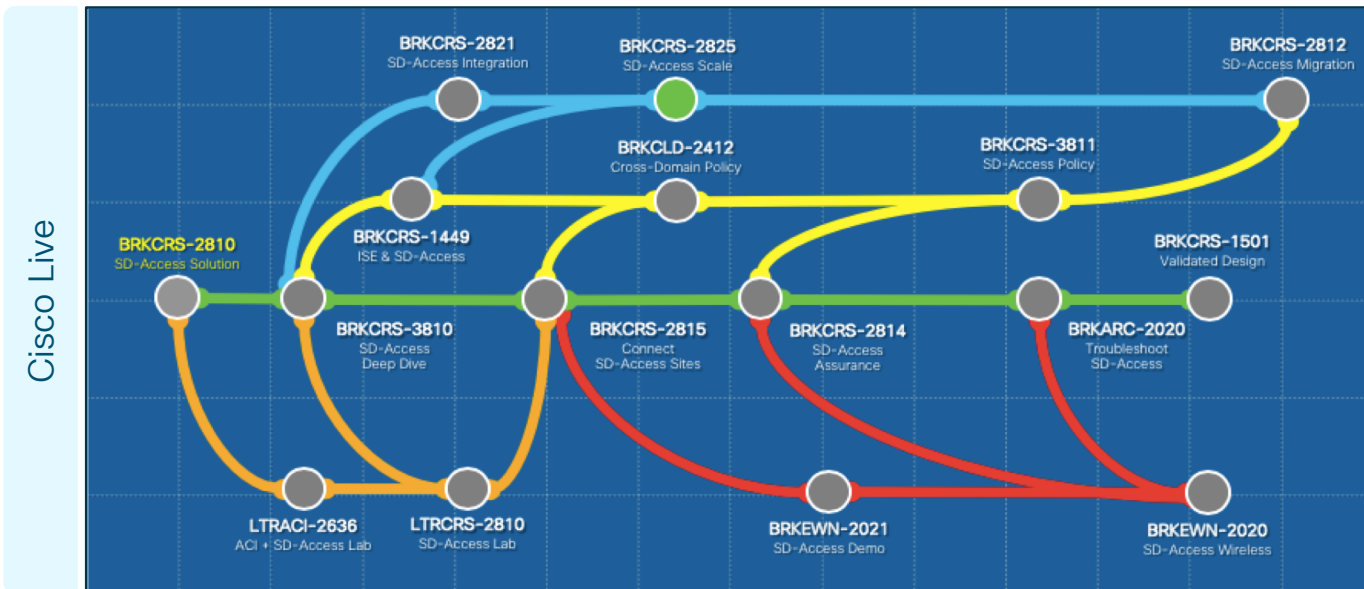
Step 4: iBGP – Border(s) Node Configuration

- Create iBGP session for every VN between Border nodes
 - Create Interface (vlan / sub-interface)
 - Configure iBGP session between Border Node

```
!
int vlan 103
description vrf interface to Border2-9500
vrf forwarding IoT
ip address 172.16.16.9 255.255.255.252
no ip redirects
ip route-cache same-interface
no shut
exit
!
int vlan 104
description vrf interface to Border2-9500
vrf forwarding Guest_VN
ip address 172.16.16.13 255.255.255.252
no ip redirects
ip route-cache same-interface
no shut
exit
!
router bgp 65001
neighbor 172.16.16.2 remote-as 65001
neighbor 172.16.16.2 update-source Vlan101
!
address-family ipv4
neighbor 172.16.16.2 activate
neighbor 172.16.16.2 weight 65535
neighbor 172.16.16.2 advertisement-interval 0
exit-address-family
!
address-family ipv4 vrf Campus_VN
neighbor 172.16.16.6 remote-as 65001
neighbor 172.16.16.6 update-source Vlan102
neighbor 172.16.16.6 activate
exit-address-family
!
address-family ipv4 vrf Guest_VN
neighbor 172.16.16.14 remote-as 65001
neighbor 172.16.16.14 update-source Vlan104
neighbor 172.16.16.14 activate
exit-address-family
!
address-family ipv4 vrf IoT
neighbor 172.16.16.10 remote-as 65001
neighbor 172.16.16.10 update-source Vlan103
neighbor 172.16.16.10 activate
exit-address-family
!
```



Additional Resources



Playbook

SDA Design Playbook

- Design guidance
- References
- Etc ..

TDM, BDM

- Latest collateral on outbound marketing

