

123: Enterprise Campus Wired Design Fundamentals

cisco Live !

An Overview of Cisco Validated Designs and Best Practices

Chris Courtelis

Networking Solutions Engineer – Public Sector Victoria & Tasmania

Based on the following sessions:

Cisco Live US – BRKENS1500 & BRKENS2500

Session ID: BRKENS-1501

Cisco Webex App


Questions?

Use Cisco Webex App to chat with the speaker after the session

How

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

Webex spaces will be moderated by the speaker until 14 November 2025.



<https://ciscolive.ciscoevents.com/ciscolivebot/#BRKENS-1501>

Who am I?



Chris Courtelis
Networking
Solutions Engineer
8 Years in Cisco

2012 UTS Electrical Engineering

2017 Cisco Sydney TAC Engineer

2020 Enterprise Network Specialist

2021 Account Solutions Engineer

2024 Networking Solutions Engineer



This session is about enabling you to **understand network design decisions**

and assisting you to **make those design choices
yourself**

**This session is not a deep dive
or Cisco product update**

Additional sessions for further learning provided



Session Agenda

- 01 Campus Design Fundamentals**
- 02 Campus Design Principles**
- 03 Campus Foundational Services**
- 04 Platform Design Considerations**
- 05 Campus Design Best Practices**
- 06 Campus Integrations**

Campus Design Fundamentals

Design Fundamentals



- **What is a Campus?**
- **Place in Network (PIN)**



What is a Campus?

A basic definition of a Campus is:

A group of *one or more buildings*, and surrounding grounds, where *people and their belongings* work together




Campus is focused on User Access

Campus = Geography





Buildings are spread out. Multiple floors per building.

Campus ≠ Data Centre





One or few large buildings nearby. Usually a single floor.

Campus Networks – Real Life



Campus Networks – Floorplans and Physical Layout





MDF = Main Distribution Framework
IDF = Intermediate Distribution Framework


Campus Design Fundamentals

Design Fundamentals

- **What is a Campus?**
- **Place in Network (PIN)**




Campus 'Place in Network' (PINs) & Topology




Campus Design Principles

Design Fundamentals

- **Multi-Layer Model**
 - Campus Multi-Layer
 - Hierarchical Design
- **Access Layer**
- **Distribution Layer**
- **Core Layer**



Campus Design Model – Multiple Layers





3 ‘Logical’ Layers

If you ‘collapse’ layers your device needs to support all ‘logical’ functions


Each layer has a specific set of requirements and provides a specific set of functions

Campus Design Principles – Scalability & Stability



Fault Domain


Hierarchical Network Design – Solid Foundation





Hierarchy – each **layer** has a **specific role**

Creates small fault domains with **clear demarcations** and isolation

Promotes **load balancing and redundancy**




Multilayer Architecture – Solid Foundation





Multilayer “Architecture” – Lack of Foundation

Not an architecture
to “design” or build




Multilayer Architecture – Solid Foundation



Campus Design Principles

Design Fundamentals

- **Multi-Layer Model**
- **Access Layer**
 - Campus Access (Baseline)
 - Extended Access (IOT/FTTX)
 - Routed Access
- **Distribution Layer**
- **Core Layer**




Campus Access (Baseline)


WHAT: Connecting Users & Devices to the Distribution layer

HOW: Connect users to the network


Considerations:




L2




L2



Multiple
L2
Features



Low to Med
L2 Scale





Extended Access (IOT/FTTX)

WHAT: Extension of the Access to connect more devices or device types

HOW: Extend the size and scale of the Access layer and connect more hosts

Considerations:




Routed Access


WHAT: Connecting Users & Devices to the Distribution layer via L3 and limits L2

HOW: Connect users to the network and leverage routing to reduce L2 challenges


Considerations:




L3




L2



Multiple
L2 & L3
Features




Low to Med
L2 & L3
Scale



Best Practice: Access Layer Oversubscription Ratios

4 x 10G Uplinks spread
across two switches in stack



4 x 48 port switches in stack
12 x 10mGig + 36 x 1Gig

Soft recommendation for
Access to Distribution $\leq 20:1$

Access Uplinks: **40 Gbps**

Potential Downlinks:
48 x 10 Gbps

+ **144 x 1 Gbps**


Result: **624 Gbps**

Oversubscription ratio:
~15.6 : 1

Campus Design Principles

Design Fundamentals

- **Multi-Layer Model**
- **Access Layer**
- **Distribution Layer**
 - Campus Distribution (Baseline)
 - Collapsed Core
- **Core Layer**






Campus Distribution (Baseline)


WHAT: Connecting multiple Access switches and the Core Layer

HOW: Distribution connectivity from Core/WAN to Access


Considerations:



Multiple
L2 & L3
Features



Med to High
L2 & L3
Scale





Campus Collapsed Core


WHAT: Connecting multiple Access layers and the WAN/Edge layer

HOW: Collapse Core & Distribution layers into a single layer

Considerations:



Best Practice: Distribution Layer – Oversubscription Ratios



Soft recommendation for
Distribution to Core $\leq 4:1$

Distribution to Core: **200 Gbps**

Access to Distribution:
4 x 2 x 4 x 25 Gbps


Result: **800 Gbps**

Oversubscription ratio:
4 : 1

Campus Design Principles


Design Fundamentals

- **Multi-Layer Model**
- **Access Layer**
- **Distribution Layer**
- **Core Layer**
 - Campus Core (Baseline)
 - Campus Core + Interconnect
 - Campus Core + Edge



Do I need a Core Layer?

Second Building Block– 4 New Links




4th Building Block
• **12 New Links**
• **8 IGP Neighbours**

3rd Building Block
• **8 New Links**
• **5 IGP Neighbours**

Do I need a Core Layer?

Second Building Block– 4 New Links



- 4th Building Block
- 4 New Links
 - 3 IGP Neighbours


- 3rd Building Block
- 4 New Links
 - 3 IGP Neighbours

Campus Core (Baseline)


WHAT: Connects multiple Distribution layers to an Interconnect and/or other network domains

HOW: A simple, high-bandwidth, L3 transport between other network layers


Considerations:




L3




L3



Limited
Features



High
Scale




Campus Core Interconnect


WHAT: Extension of the Core, used to connect multiple Core Layers and/or other network domains

HOW: Distribute the bandwidth and density requirements of multiple Cores


Considerations:




L3




L3



Limited
Features



Higher
Scale





Campus Core + Edge (SP/WAN)

WHAT: Connects multiple Campus areas to remote domains and/or the Internet

HOW: Collapse Core & Edge Layers


Considerations:



Campus Foundational Services

Design Fundamentals

- **Layer 1 physical layer & links**
 - Copper vs Fibre Cabling
 - L1.5 with UDLD
 - L1.5 with Ether Channel
- **Layer 2 switching protocols**
- **Layer 3 routing protocols**



Copper vs Fibre Media


Category 5, 6 and 7

Unshielded (UTP) Shielded (STP)

RJ45 (Access to Endpoints)



Category	Frequency	Distance	Data Rate	Shielding
5E	100-350 MHz	100m	1000 Mbps	UTP or STP
6	250-550 MHz	1G - 100m 10G - 50m	1 Gbps 10 Gbps	UTP or STP
6A	500-550 MHz	100m	10 Gbps	UTP or STP
7	600 MHz	100m	10 Gbps	Shielded only



OM3, OM4 and OM5

Multi-Mode (MMF) Single-Mode (SMF) Wave-Division Multiplex (WDM)

SFP (Access and Distribution)


QSFP (Core and Edge)



- Multimode**

 - Short distance cable runs (less than 1000ft.)
 - High bandwidth support
 - Higher cable cost
 - Lower electronics cost
 - Easier to terminate due to larger core size
- Single Mode**

 - Long distance cable runs (greater than 1000ft.)
 - Highest bandwidth support
 - Lower cable cost
 - Higher electronics cost
 - Harder to terminate due to smaller core size




<https://tmgmatrix.cisco.com/>

Fibre Best Practices: Unidirectional Link Detection (UDLD)


WHAT: Protects against one-way communication generally on fibre links

HOW: Utilises “echos” for hellos between peers to determine if the link is unidirectional or bidirectional


Considerations:



Enable Global
Config



UDLD
Aggressive




Campus + EtherChannel


WHAT: Logical link that increases bandwidth and provides link redundancy while reducing peer complexity

HOW: Combines multiple physical links into a single logical link


Considerations:




L2/L3
Config



Tune
Load-
Balancing




Exponents
of Two



Campus Foundational Services

Design Fundamentals

- **Layer 1 physical layer & links**
- **Layer 2 switching protocols**
 - Layer 2 Design
 - Layer 2 Spanning-Tree
 - Layer 2 VLAN Trunks
- **Layer 3 routing protocols**




Layer 2 Design and Spanning Tree – Best Practices

WHAT: Protects against L2 loops

HOW: Determines correct port state and administers port state to prevent loops


Considerations:




Limit VLAN
Extension



RPVST+



MST




Layer 2 Design and Trunks - Best Practices


WHAT: Layer 2 interface type that supports multiple VLANs

HOW: Deployed on interconnection between Access and Distribution layers to provide multiple VLAN access


Considerations:



Prune unused VLANs




Change native VLAN



Campus Foundational Services

Design Fundamentals

- **Layer 1 physical layer & links**
- **Layer 2 switching protocols**
- **Layer 3 routing protocols**
 - Layer 3 Design
 - Layer 3 First Hop Routing
 - Layer 3 with BFD




Layer 3 Design and Routing Protocols – Best Practices


WHAT: Provides connectivity between different networks/subnets

HOW: Typically deployed on Distribution-to-Core and Core-to-Core interconnects providing L3 reachability


Considerations:



Build
Triangles



Redundant
Paths




First Hop Redundancy


WHAT: Pair of redundant L3 Routers function as a single virtual Router

HOW: One Router is elected as Active and performs packet forwarding for local hosts. The other router is the Standby.


Considerations:



HSRP or
VRRP



Timer
Tuning




GLBP

R1 – Active , Forwarding traffic
R2 – Hot Standby, Idle

IP: 10.0.0.254
MAC: 0000.0c12.3456
vIP: **10.0.0.1**
vMAC: **0000.0c07.ac00**

IP: 10.0.0.253
MAC: 0000.0c78.9abc
vIP: **10.0.0.1**
vMAC: **0000.0c07.ac00**



Redundancy & Protocol Interaction – L2 vs L3

L3 routed interface provides faster convergence than **L2 switch port** with an associated L3 SVI



1. Link Down
2. Interface Down
3. Routing Update

~ 8 msec loss

21:38:37.042 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet3/1, changed state to down

21:38:37.050 UTC: %LINK-3-UPDOWN: Interface GigabitEthernet3/1, changed state to down

21:38:37.050 UTC: IP-EIGRP(Default-IP-Routing-Table:100): Callback: route_adjust GigabitEthernet3/1



1. Link Down
2. Interface Down
3. Autostate
4. SVI Down
5. Routing Update

~ 150-200 msec loss

21:32:47.813 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet2/1, changed state to down

21:32:47.821 UTC: %LINK-3-UPDOWN: Interface GigabitEthernet2/1, changed state to down

21:32:48.069 UTC: %LINK-3-UPDOWN: Interface Vlan301, changed state to down


21:32:48.069 UTC: IP-EIGRP(Default-IP-Routing-Table:100): Callback: route_adjust Vlan301

Bidirectional Forwarding Detection (BFD)

WHAT: Detects faults between 2 routers very quickly (in milliseconds)

HOW: Signals to upper-layer routing protocols that the link is down, faster than routing protocol dead timers


Considerations:



250ms x 3
Physical
Interfaces



750ms x 3
SVI



```
interface Gig1/0/1
 ip address 1.1.1.1 255.255.255.0
 bfd interval 250 min_rx 250 multiplier 3
 ip ospf 1 area 0


router ospf 1
 bfd all-interfaces
```

Layer 3 Design - Summarise Routes


WHAT: Utilises an overarching route that encompasses multiple more specific routes

HOW: Utilise a wider subnet mask route or a less specific route for advertisement


Considerations:




Faster
Reconvergence



More HW
Efficient



Sub-
Optimal
Routing



Campus Foundational Services - Best Practices Summary

Core


Distribution

Access

Layer 3: Summarise routes towards the Core for optimal network reconvergence

Layer 2: Tune L2 Protocols

Layer 1: Use EtherChannels/UDLD for link resiliency




Platform Design Considerations

Design Considerations




- **Chassis Considerations**
 - Catalyst 9000 Portfolio
 - Modular vs. Fixed Chassis
- **Cabling Considerations**
- **Feature Considerations**



What's the **right switch** for the job?

Cisco Catalyst 9000 Switching Portfolio




Cisco Catalyst 9000 Switching Portfolio


To Deep Dive on Catalyst 9K Series please refer to:

BRKARC-2092: Cisco Catalyst 9000 Series, Cisco Silicon One, and IOS XE
Architecture and Innovation

Tuesday 1:45pm, Room 203



Switching Chassis: Modular vs. Fixed Platforms



Modular

PROs

- **More Flexible**
- Longer Life-Cycle
- Higher Port Density
- More Power/Cooling
- Redundant Processors

CONs

- **More Complex**
- BW limit by Chassis
- Slow(er) Dev & Test
- Lower MTBF
- Higher COGs

Fixed

PROs

- **Less Complex**
- Swap Chassis for BW
- Faster Dev & Test
- Higher MTBF
- Lower COGs


CONs

- **Less Flexible**
- Shorter Life-Cycle
- Lower Port Density
- Less Power/Cooling
- Single Processor

Platform Design Considerations

Design Considerations

- **Chassis Considerations**
- **Cabling Considerations**
 - Multigigabit 1G to 10G
 - Power over Ethernet Up to 90W
 - Next Gen Connectivity
- **Feature Considerations**



Multigigabit: Beyond 1G Access Connectivity



Category 5, 6 & 7

Unshielded (UTP)

Shielded (STP)

RJ45 (Access to Endpoints)



Cat6A
(Offset Wires)

Cat5E
(Flush Wires)



Category	Frequency	Distance	Data Rate	Shielding
5E	100-350 MHz	100m	1000 Mbps	UTP or STP
6	250-550 MHz	1G - 100m 10G - 50m	1 Gbps 10 Gbps	UTP or STP
6A	500-550 MHz	100m	10 Gbps	UTP or STP
7	600 MHz	100m	10 Gbps	Shielded only



The Ethernet Evolution

10GBASE-T 2006

10G




5GBASE-T 2016

5G




2.5GBASE-T 2016

2.5G



1000BASE-T 1999

1G



100BASE-T

1995

100M



10BASE-T

1990

10M




Multigigabit: The Need for Speed

3 clients on 2.4GHz @ 20MHz
15 clients on 5 GHz @ 80MHz


Wi-Fi 6 Performance with Gigabit switch

THROUGHPUT **872.14 Mbps**




Wi-Fi 6 Performance with mGig switch

THROUGHPUT **1.537 Gbps**



Wi-Fi 6E/7 is Driving Multi-Gigabit & UPOE



Cisco 90W
UPOE[®]+
(low voltage)

Multi-Gigabit
Ethernet
(up to 10Gbps)




2.4 GHz

5 GHz


6 GHz




Power over Ethernet: Evolving Requirements



Current Campus 3-Tier Design





Next Generation 3-Tier Campus Design



Platform Design Considerations

Design Considerations

- Chassis Considerations
- Cabling Considerations
- Feature Considerations
 - Unicast (IPv4/IPv6)
 - Multicast (IPv4/IPv6)
 - Quality of Service



Campus Networks: L2/L3 Unicast Technologies

IPv4 Unicast

- MP-BGP, VPNv4
- Internet (v4), NAT, PBR
- MPLS-VPN, VRF-Lite
- IPv4 SSO, NSF/NSR, GIR


Core

- EIGRP, OSPFv2, ISIS, RIP
- SVI, HSRP/VRRP
- ARP, DHCP Relay
- IPDT/SISF, DAI
- BFD, Echo
- IPv4 SSO, NSF/NSR, GIR

Distribution

- PVST, MST, REP/RENN
- 802.1Q, DTP
- VLANs, VTP
- DHCP Snooping
- MAC Leaning
- L2 SSO

Access



IPv6 Unicast

- MP-BGP, VPNv6
- Internet2 (v6), NAT64, PBR
- MPLS-VPN, VRF-Lite
- IPv6 SSO, NSF/NSR, GIR

Core

- EIGRPv6, OSPFv3, ISISv6, RIPng
- SVI, HSRPv6/VRRPv6
- NDP, DHCPv6 Relay
- SISF (v4/v6), RA Guard
- BFDv6, Echo
- IPv6 SSO, NSF/NSR, GIR

Distribution

- PVST, MST, REP/RENN
- 802.1Q, DTP
- VLANs, VTP
- DHCPv6 Snooping
- MAC Leaning
- L2 SSO

Access

Campus Networks: L2/L3 Multicast Technologies

IPv4 Multicast

- PIM-SM, SSM and Bidir
- AutoRP, BSR RP, MSDP
- MVPN, Multicast VRF-Lite
- Multicast load splitting
- IPv4 multicast HA


Core

- Dual-stack IPv4 / IPv6
- PIM-SM, SSM and Bidir
- IGMPv2,v3 snooping
- Stub multicast routing
- PIM BFD
- IPv4 multicast HA

Distribution

- IGMP v1,v2,v3 snooping
- IPv4 multicast QoS & ACL
- IGMP v1,v2 filtering

Access



IPv6 Multicast

- PIM-SM and SSM
- IPv6 BSR RP
- IPv6 embedded RP
- IPv6 multicast HA

Core

- Dual-stack IPv4 / IPv6
- PIM-SM and SSM
- MLDv1,v2 snooping
- HW register and RPF
- HSRP-aware PIM
- IPv6 multicast HA

Distribution

- MLD v1,v2 snooping
- IPv6 multicast QoS & ACL
- MLD v1,v2 filtering

Access


Understanding L2 Scale – MAC Address Scale

Soft recommendation for
Access to Distribution $\leq 20:1$

- Each unique **Endpoint** will have 1x **MAC address**
- **Access:** No. of Hosts = No. of MACs
- All **MACs** are also learned on Distribution
- **Distro:** Sum of No. of Access

1000–2000 x 20


Result: **20–40K MACs**



Understanding L3 Scale – IP Route Scale

Soft recommendation for
Access to Distribution $\leq 20:1$
Distribution to Core $\leq 4:1$

- Each **Endpoint** will have 1 x **ARP** (and/or 3+ **NDP**)
- All **ARP/NDP** resolve on Distribution (L3 SVI)
 - **Distro**: Sum of No. of Access = **1-3K x 20 = 20-60K**
 - **VLANs**: 5-10 per Access = **4-5 x 20 = 100-200**
- All L3 + WAN/DC (x VRF) Subnets on Core
 - **Core**: Sum of No. of Distro = **10-20 x 200 = 2K-4K**
 - **WAN/DC**: Sum of No. of Sites = **10-20 x 2K = 20K-40K**
 - **Internet**: Q3 2025 = **~1M IPv4**, **~230K IPv6**



Understanding ACL Types Storage and Scale



PACL Port Based ACL

Applied to Physical Ports



```
access-list extended PACL permit 36.48.0.3
access-list extended PACL deny 36.48.0.0 0.0.255.255
access-list extended PACL permit 36.0.0.0 0.255.255.255
!
interface gigabitethernet 2/0/1
  switchport
  ip access-group PACL in
```

Applied to L2 VLAN domains

VACL VLAN Based ACL



```
ip access-list extended SERVER_ACL
  permit ip 10.1.2.0 0.0.0.255 host 10.1.1.100
  permit ip host 10.1.1.4 host 10.1.1.100
vlan access-map VACL
  match ip address SERVER_ACL
  action drop
!
vlan filter VACL vlan-list 10
```



RACL Router (IP) ACL

Applied to L3 Ports & SVIs



```
access-list extended RACL permit 36.48.0.3
access-list extended RACL deny 36.48.0.0 0.0.255.255
access-list extended RACL permit 36.0.0.0 0.255.255.255
!
interface vlan 36
  ip address 36.64.0.1 255.255.255.0
  ip access-group RACL in
```

OGACL Object-Group ACL

Maps IP to Object



```
object-group network NetOps
  11.96.179.0
  11.96.180.0
object-group network Servers
  10.10.172.12
  10.10.179.210
ip access-list extended OGACL
  permit object-group NetOps Servers any
```




The (main) reason for Campus LAN QoS


WHAT: Quality of Service (QoS) primary role is to manage packet loss

HOW: Utilise traffic classes and packet queues to prioritise critical or sensitive applications. i.e. audio/video traffic


Considerations:





QoS Consistency
Wired & Wireless



Timing
(ms)





ASIC
Differences



Campus Design Best Practices


Design Considerations

- **LAN High Availability**
 - SSO/NSF
 - StackWise and StackWise Virtual
 - Campus Architecture Resiliency
- **LAN Security**
- **Virtual Networking (Fabrics)**



How much resiliency is enough?

Resiliency and High Availability are designed at every level



Architecture

Physical Stacking

Virtual Stacking

Maintenance Mode

Operating System

Hot and Cold Patching

Minimal or no traffic loss upgrades

Minimised downtime stack upgrades

Platform


Redundant Control & Data Plane

Redundant Power & Fans


Power pooling

High-Availability – SSO & NSF

Modular Switch with Redundant Supervisors



StackWise Stack or StackWise Virtual Pair



Stateful Switchover (SSO) synchronizes the **active process state** and **running-config**, between Active & Standby supervisors or switches in a stack


Non-Stop Forwarding (NSF) allows for **graceful restart** of L3 routing protocols

StackWise Access – Physical Stacking


WHAT: Combines multiple Access switches into a single virtual switch to expand port density of Access layer

HOW: Stateful Switchover (SSO)


Considerations:




NSF/NSR
Protocols



L2 MEC



Extended Fast
Software
Upgrade
(xFSU)




StackWise Virtual Core/Distribution – Virtual Stacking


WHAT: Combines Core and/or Distribution into a single virtual switch to simplify and expand those layers

HOW: Stateful Switchover (SSO)


Considerations:




NSF/NSR
Protocols



L2 & L3
MEC




In-Service
Software
Upgrade
(ISSU)



Campus Architectures – Platform Resiliency

1

ECMP




Node Level Redundancy

Protocol Load-Balancing

More Neighbors =
Requires Protocol Tuning

2

EtherChannel (L2/L3 LAG)




Node & Link Redundancy

EtherChannel Load-Balancing

More Neighbors =
Requires Protocol Tuning

3

StackWise (L2/L3 MEC)




Layer Level Redundancy

Multi-chassis EtherChannel

Fewer Neighbors = No
Protocol Tuning

Campus Network utilising StackWise

L2/L3 Equal Cost Multi-Path




19 Nodes to manage

More IGP and STP Peers

Failover based on Protocol & CEF (sec)

L2/L3 with StackWise & MEC



8 Nodes to manage

Less IGP and STP Peers

Failover based on MEC (ms)

Campus Network utilising StackWise

L2/L3 Equal Cost Multi-Path

L2/L3 with StackWise & MEC

To Deep Dive on Highly Available Networks please refer to:

BRKENS-2095: Designing Highly Available Networks using Catalyst 9000 Series Switches

Thursday 10:30am, Room 210

19 Nodes to manage

More IGP and STP Peers

Failover based on Protocol & CEF (sec)

8 Nodes to manage



Less IGP and STP Peers

Failover based on MEC (ms)

Campus Design Best Practices

Design Considerations

- **LAN High Availability**
- **LAN Security**
 - End to End Security Approach
 - User Authentication
 - Segmentation
- **Virtual Networking (Fabrics)**



When is a network **secure**?

Campus Security – End to End occurs at different layers


WHAT: Implementing multiple security technologies at different layers of the network for a holistic approach

HOW: Utilising User Authentication, Protocol Authentication, Data Encryption, and Segmentation


Considerations:




Zero Trust



NetOps
/SecOps



Beyond
the LAN



LAN Security Best Practices

Cisco Umbrella

- uses DNS as a security tool to identify and block threats

802.1x User Authentication

- forces users to authenticate before allowing them on network

IP Source Guard / v6 RA Guard

- prevents IP/MAC Spoofing and IPv6 Man-in-the-Middle attacks

Dynamic ARP Inspection


- prevents current ARP attacks

DHCP Snooping

- prevents Rogue DHCP Server attacks

Port Security

- prevents CAM attacks and DHCP Starvation attacks





End User Authentication

WHAT: Provide security at the access layer to authenticate endpoints/users

HOW: Utilising the 802.1x standard for wired and wireless ports, including EAP, MAB, and WebAuth


Considerations:



Network Security: Visibility “How and Who/What”

HOW

Authenticate



WHO/WHAT

Identify

IP add: 192.168.2.101

ID

Bob (Employee)

Unknown

WHAT

Apple iPad/IOS/11.0.1

Unknown

WHEN

10:30 AM

Unknown

WHERE

Floor-1, Amsterdam, Building 1

Unknown

Conn TYPE

Wireless

Unknown

APPS

Firefox, MS Word, Anyconnect

Unknown


HW

Serial No, CPU, memory


Profiling

Network Security: Enforcement

Endpoints Authenticated and Profiled



Basic Segmentation Options you have:



Limitations

IP based Segmentation

Limited Flexibility & Complex Maintenance

Challenging to replication and retain consistency

Which policy is simpler to understand?

```
Switch-1# show ip access-list CorPolicy
```

```
Extended IP access list CorPolicy
```

```
10 permit tcp 10.1.100.0 0.0.0.255 172.16.100.0 0.0.0.255 eq 80
20 permit tcp 10.1.100.0 0.0.0.255 172.16.100.0 0.0.0.255 eq 443
30 permit tcp 10.1.101.0 0.0.0.255 172.16.100.0 0.0.0.255 eq 80
40 permit tcp 10.1.101.0 0.0.0.255 172.16.100.0 0.0.0.255 eq 443
50 permit tcp 10.1.100.0 0.0.0.255 172.16.101.0 0.0.0.255 eq 80
```

ACL
Policy

```
Switch-1# show cts role-based permissions
```

```
IPv4 Role-based permissions default:
```

```
Deny IP-00
```


```
IPv4 Role-based permissions from group 10:Employee_SGT to group 100:ProdServer_SGT:
```

```
Web_Only-10
```

SGACL
Policy




Cisco TrustSec (CTS) – Simplifying Segmentation

Traditional Access Control



Security Policy based on addresses
introduces complexity and lacks flexibility

Security Group Tags





Security Policy based on tags
provides topology flexibility, and granular security

Campus Design Best Practices


Design Considerations

- **Lan High Availability**
- **LAN Security**
- **Virtual Networking (Fabrics)**
 - Value of Overlays
 - BGP EVPN
 - SDA LISP




Layer 2 Demarcation – Understanding the History

Layer 2 Core




Large L2 Domains

Routed Access



Confined VLANs


Routed Distribution



Balanced


Is there a better way? **Fabrics?**

Fabrics: What is an Underlay?



Underlay Network = **Physical Infrastructure** to provide **IP reachability** with redundancy and **resiliency**.


Fabrics: What is an Overlay?



Overlay Network = Logical topology used to virtually connect devices to provide additional services, not delivered by the Underlay.

Examples: GRE, CAPWAP, IPSEC, LISP, EVPN, VXLAN


Fabrics: Why an Overlay?



- **Overlay (Fabric):** Flexible, Scalable and Extensible. Easy to add/modify services. Optimizes mobility events.

- **Underlay:** Build and forget! Reliable, manageable, and simple.

Fabrics: Virtual and Physical Topologies



Overlay

- Device
- Connectivity
- Authentication
- Authorisation
- Policy
- Segmentation


Underlay

- Predictable
- Low Touch
- Stable
- Simple

Campus Fabric Solutions & Designs

1

MPLS (L2/L3VPN)




L3 VRF Segmentation

WAN/Edge + VPN Services

MPLS/VPLS, LDP

2

EVPN (L2/L3VNI)




L2/L3 VNI Segmentation

Common WAN/LAN Services

MP-BGP/EVPN, VXLAN

3

SDA (L2/L3VNI + SGT)



L2/L3 VNI + SGT Segments

LAN + Group-Based Policy





LISP, VXLAN, SGT

EVPN Border and Spine

WHAT: Border and Spine connects EVPN Fabric to other network domains.

HOW: Border connects fabric to other network domains. Spine reflects the L2/L3 prefixes providing neighbor peering, learning and distribution point.

Considerations:






EVPN Leaf


WHAT: Leaf node connects Wired endpoints to EVPN Fabric domain.

HOW: An origination and termination point of VXLAN enabled overlay network to connect wired endpoints.


Considerations:



Overlay
Features



Med-High
L2/L3 Scale





SDA Border and Control Plane Node

WHAT: Border connects SDA Fabric to other network domains. Control Plane provides forwarding lookup/table.

HOW: Border connects fabric to north bound networks, 'ends' fabric domain. Control is LISP database for lookups.

Considerations:




SDA Edge Node


WHAT: Edge connects Wired/Wireless endpoints to SDA Fabric domain/network

HOW: Utilises an Anycast Gateway, and begins SDA fabric encapsulation for Wired and Wireless.


Considerations:




L3




L2



Overlay
Features



Higher
L3 Scale




SDA Edge Node

WHAT: Edge connects Wired/Wireless endpoints to SDA Fabric domain/network

To Deep Dive on Fabric Technologies please refer to:

HOW: EVPN: BRKENS-2041: EVPN Campus: Design and Implementation
Wednesday 8:30am, Room 210



SDA: BRKENS-2502: Cisco SD-Access LISP VXLAN Fabric Best Practices:
Considerations: Design and Deployment
Tuesday 1:45PM, Room 210



Campus Integrations

Design Considerations

- **Wireless LAN**
- **Firewalls, VRFs, and ACLs**




Wireless LAN


WHAT: Connecting APs centrally to one or multiple WLC. Thus, connecting Wireless Endpoints to the network

HOW: Tunnel L2 Overlay (CAPWAP)


Considerations:




Pros/Cons
of Central



SSO or
N+1 HA



Higher
L2/L3
Scale




Wireless: Local Mode Design Considerations


WHAT: Simplified central wireless configuration, troubleshooting, and roaming at scale.

HOW: WLC/APs are located at the same site, & centrally managed.


Considerations:



Large Site > 100 APs



Latency $\geq$ 100ms



Wireless: FlexConnect Design Considerations

WHAT: Enables control of remote site APs without deploying a WLC at each site

HOW: Control traffic is tunneled centrally, and data traffic is either centrally or locally switched


Considerations:




Many Small Sites



< 50 APs



Latency ≤ 100ms




SDA Fabric Enabled Wireless

WHAT: Connecting APs as part of a fabric network (SDA) for seamless L2 roaming and common policy.

HOW: Control plane traffic is managed via WLC, and data plane traffic is passed via the fabric directly.


Considerations:




SDA



WLC/AP
Support





<20ms
Latency



Campus Integrations

Design Fundamentals

- **Central Wireless LAN**
- **Firewalls, VRFs, and ACLs**




Firewalls, VRFs, and ACLs


WHAT: Controlling access into or out of network areas. To prevent unauthorized access to different network segments.

HOW: Segmentation and Access Control as well as Inspection


Considerations:




Inside &
Outside




L2/L3 &
L4-7
features





Med to High
L2/L3 Scale



Remember: Campus PINs & Topology




Remember: Hierarchical Campus




Remember: Campus Design Fundamentals and Principles


Collapse or Expand Layers?




Modular or Fixed Platforms?




Fiber or Copper Links?




ECMP, PortChannel or Stacking?



L2/L3 or MPLS or VXLAN?



Wireless or Security Included?



Keep Learning – Deeper Dive Sessions at Cisco Live

Switching

Cisco Catalyst 9000 Series, Cisco Silicon One, and IOS XE Architecture and Innovation [BRKARC-2092]

Tues Nov 11, 1:45PM – 3:15PM

Campus Switching Innovations for Future Proofed Workspaces [BRKENS-2609]

Wed Nov 12, 1:00PM – 2:30PM

Designing Highly Available Networks using Catalyst 9000 Switches [BRKENS-2095]

Thur Nov 13, 10:30AM – 12:00PM

Wireless

Wi-Fi 7 is Here – Are you Ready? [BRKEWN-2025]

Wed Nov 12, 8:30AM – 10:00AM

Design and Deploy Cisco Cloud Managed Enterprise Wireless Networks [BRKEWN-2046]

Wed Nov 12, 8:30AM – 10:00AM

Understanding Wireless Security and the Implications for Secure Wireless Network Design [BRKEWN-2104]

Wed Nov 12 1:00PM – 2:30PM

Fabric

Cisco SD-Access LISP VXLAN Fabric Best Practices: Design and Deployment [BRKENS-2502]

Tues Nov 11, 1:45PM – 3:45PM






EVPN Campus: Design and Implementation [BRKENS-2041]

Wed Nov 12, 8:30AM–10:00AM

What's New in SD-Access? [BRKENS-2568]

Wed Nov 12 8:30AM – 10:00AM

Keep Learning – Cisco Validated Designs



cs.co/en-cvds

Complete your session evaluations



Complete a minimum of 4 session surveys and the Overall Event Survey to claim a Cisco Live T-Shirt.



Earn up to 800 points by completing all surveys and climb the Cisco Live Challenge leaderboard.



Level up and earn exclusive prizes!



Complete your surveys in the Cisco Live Events app.

Continue your education



Visit the Cisco Stand for related demos



Book your one-on-one Meet the Expert meeting



Attend the interactive education with Capture the Flag, and Walk-in Labs



Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

Contact me at: ccourtel@cisco.com or via Webex Teams

Thank you

CISCO Live !

