Modernize your data center

Cisco’s blueprint for success

December 3, 2024

Raul Arias – DC Solutions Engineer

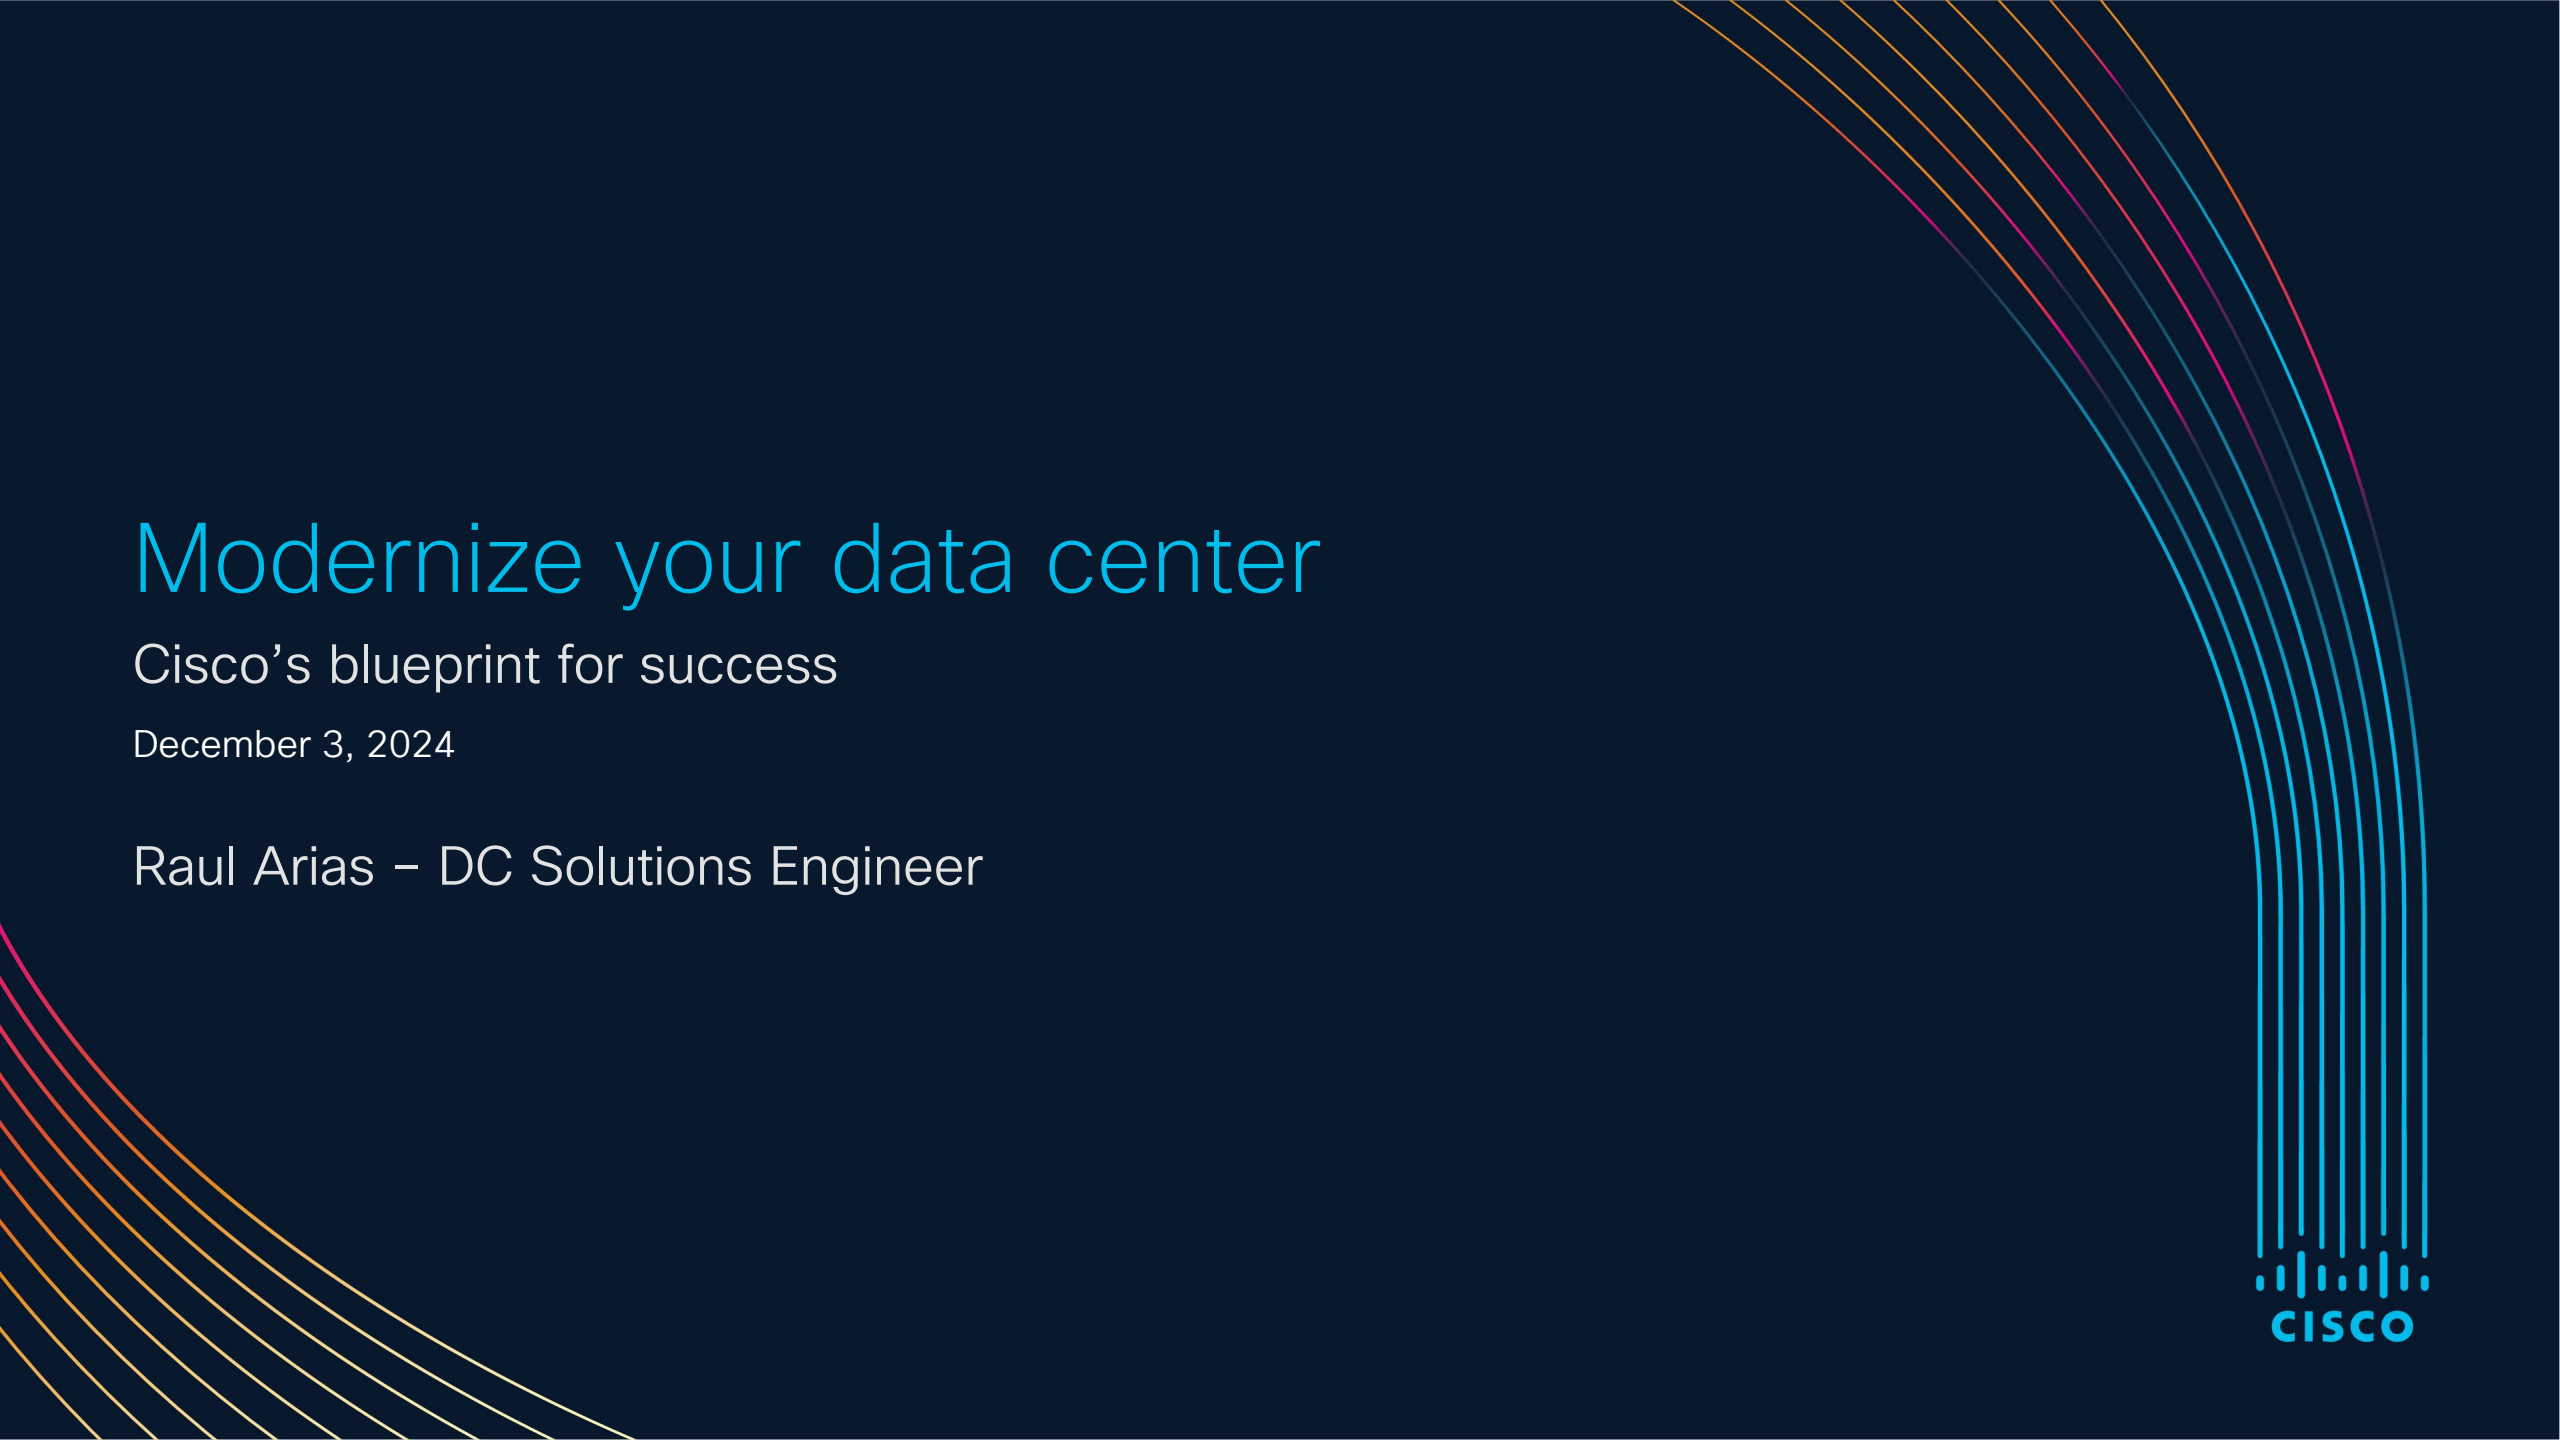

Cisco Data Center Networking Fabric Options

- Application Centric Infrastructure(ACI) fabrics
- Nexus Dashboard(NDFC) fabrics
- Nexus Hyperfabric

Application-Centric Infrastructure (ACI)

What Is ACI?

- Hardware-integrated Software-Defined Networking (SDN) solution
- Configuration is based entirely around reusable policies, not individual config commands
- Security policy is tightly integrated down to the data plane - and required, not optional
- Policy model is based around multi-tier application structure, not traditional networking
- Based on a Spine-Leaf IP Fabric with VXLAN overlay
- Only uses specific Nexus 9000 Series switch models
- Centralized, clustered controller, called APIC, manages the whole fabric as one entity
- Integrates with virtual switching at hypervisor and container level

What Makes ACI Different?

- Designed around public cloud networking concepts, e.g. Bridge Domains, Endpoint Groups, and contracts vs. VLANs and SVIs
  - Explicit-allow access model (like Fibre Channel) vs. open-by-default like Ethernet…aka Whitelist model.
- Entire fabric managed as a whole from centralized controller cluster vs. switch-by-switch
  - APIC cluster acts as management plane, while control plane is distributed across fabric switches
  - Cluster does not participate in the control or data plane
  - Allows for automatic network/security policy consistency between physical and virtual

ACI Components

- Nexus 9000 ACI Spines (1 minimum, 2 for production use)
  - Nexus 9500/9400/9300 models are supported as modular or fixed ACI Spines
- Nexus 9000 ACI Leafs (1 minimum)
  - Nexus 9400/9300 models are supported as ACI Leafs
- APIC Controller Node / Cluster (Physical or Virtual)
  - There are Medium and Large physical APICs. Virtual APIC options include small/medium/large. AWS option available
  - Medium virtual and physical APICs support up to 1200 edge ports, Large physical/Virtual APIC supports beyond this
  - Currently (v6.0) on 4th generation of physical nodes, based on UCS C225 M6 server
  - Minimum of 3 nodes required for HA and supported by TAC for production use, scales up to 7 per fabric

Typical ACI network

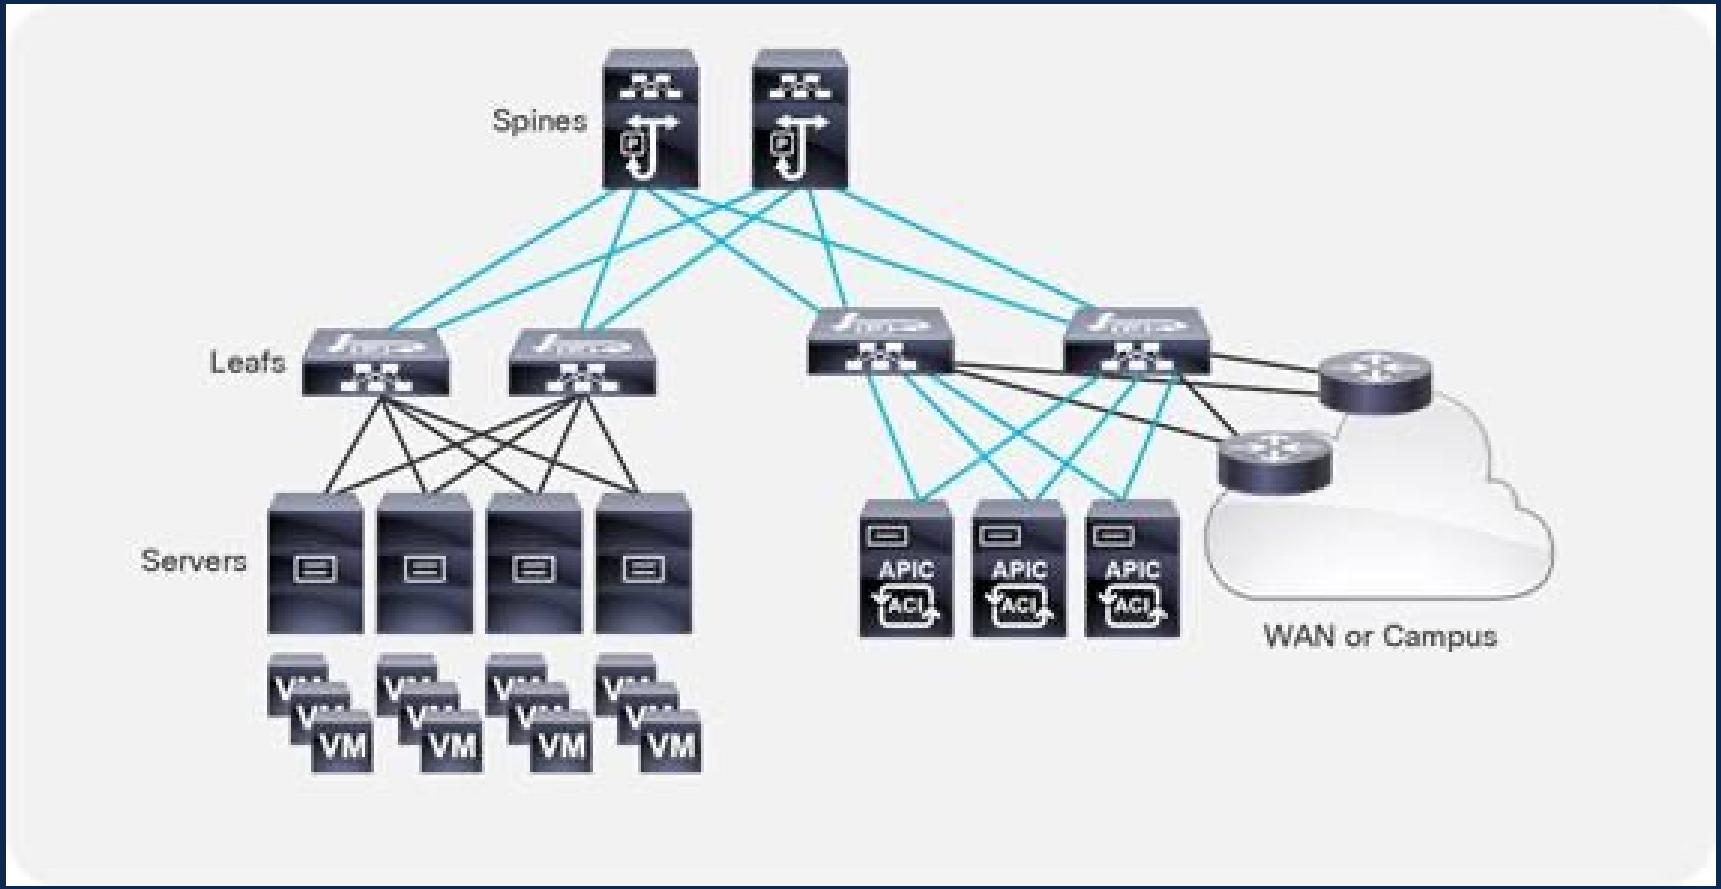

ACI Data Center Interconnect -Multipod

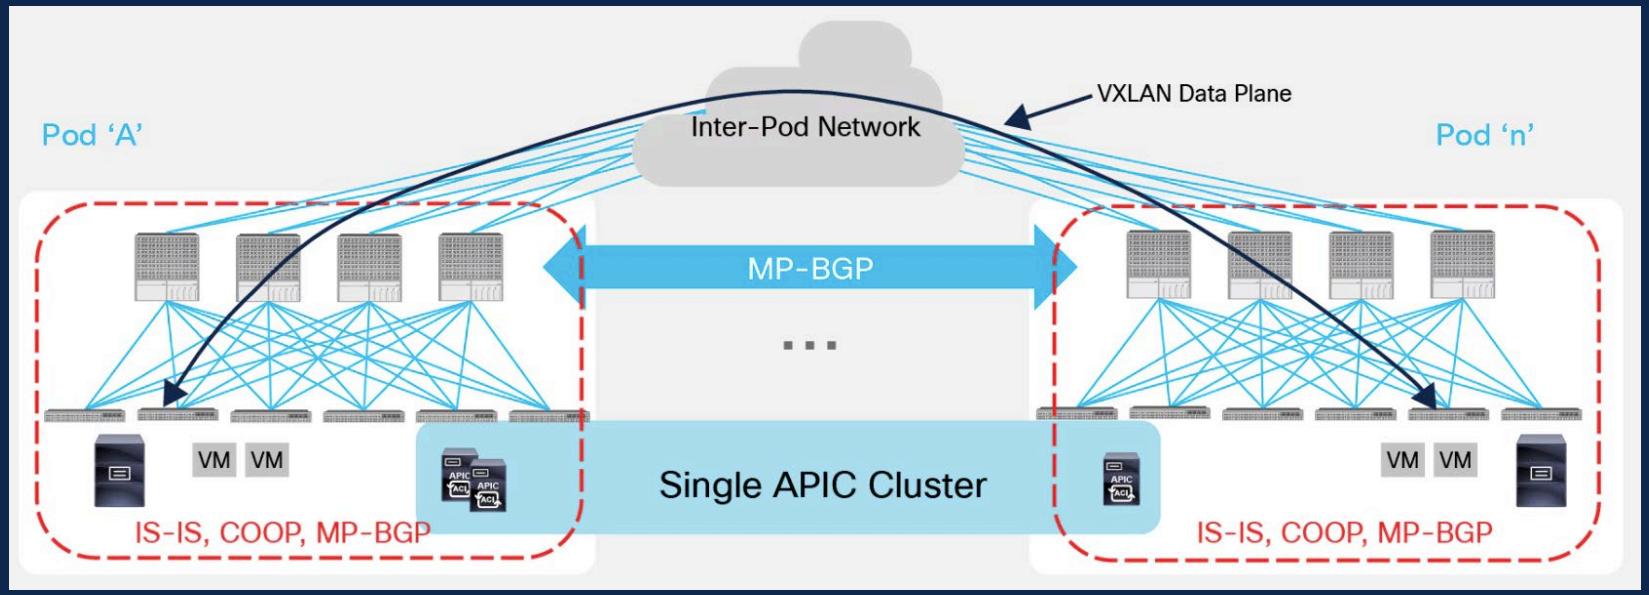

- Single APIC cluster. For a two-datacenter design, can add 4th APIC as standby
- Minimum licensing required for all nodes is Essentials
- Maximum latency between pods is 50 msec RTT. Roughly 2500 miles.
- A 3-node cluster can support up to 6 pods

ACI Multi-site

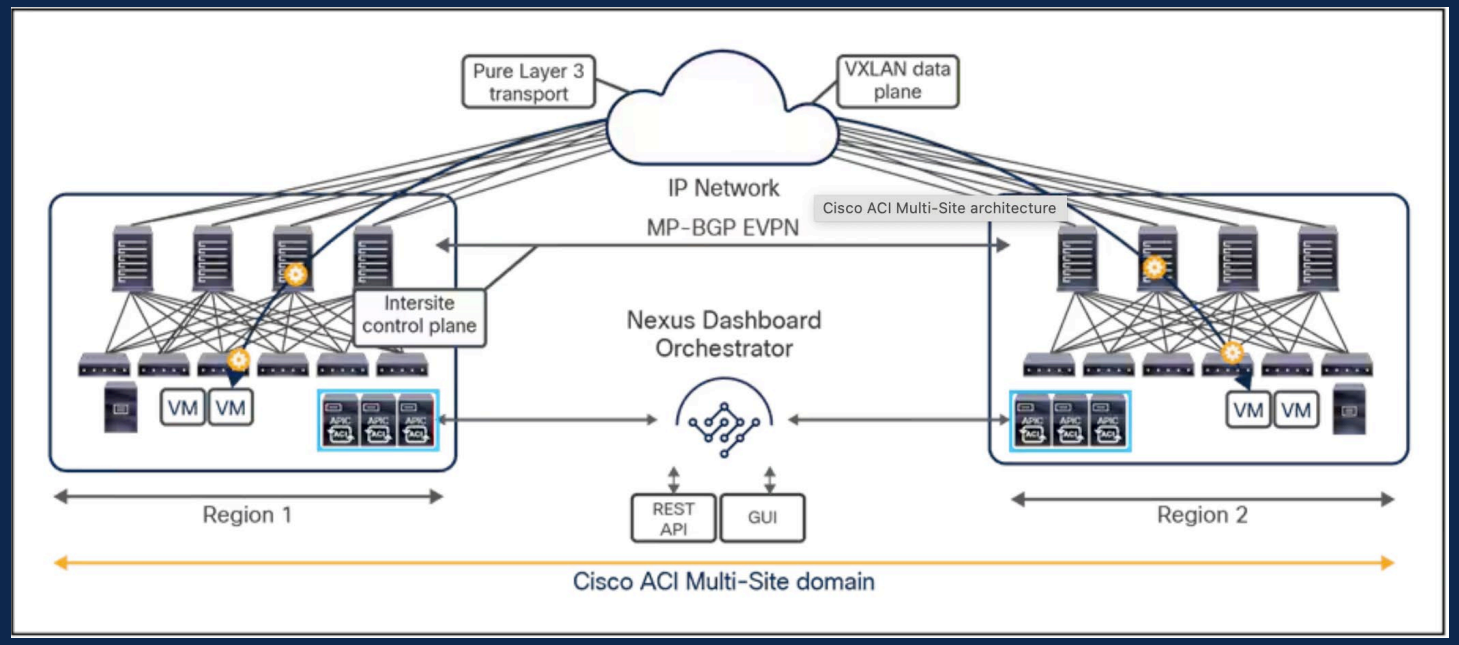

- APIC cluster per DC.
- Minimum licensing required for all nodes is Advantage.
- Maximum latency between is 500 msec RTT between ND node and APIC
- Up to 14 sites can be multi-site with EVPN sessions between them

ACI Nexus 9000 portfolio evolution

New 400G and 800G-ready products

Shipping

100G

400G

800G

CY19

Nexus 9300
16p 400G
28p 100G plus 8p 400G

CY21

Nexus 9300
32p 400G

Nexus 9500
16 port 400G line card

CY22 Q2

Nexus 9300
48p 400G
64p 400G

CY23 Q1

Nexus 9400
64p 400G
128p 200G

Future

High Density 400G and 800G Platforms

Nexus 9332D-H2R
32p 400G

Nexus 9364C-H1
64p 100G

Nexus 93400LD-H1
48p 50G + 4p 400G

N9K-X9736C-FX3
Nexus 9500
36 port 100G line card

Current APIC–APIC M4/L4 and Virtual

APIC-M4

APIC-L4

- Medium – up to 1200 endpoints
- Large – more than 1200 endpoints
- New Small virtual APIC also available
- Ability to run APIC cluster on AWS

Nexus Dashboard Fabrics

Cisco Nexus Dashboard

Simple to automate, simple to consume

Private cloud | Third-party apps | Public cloud

Cisco NDFC modes

Fabric discovery for LAN Deployments

Fabric controller for LAN and IPFM Deployments

SAN controller for MDS Fibre Channel deployments

Nexus Dashboard Fabric Controller (NDFC)

- Management controller for NX-OS based classic LAN, VXLAN, or FC SAN Fabrics
- Runs as service **hosted on Nexus Dashboard**
- Operates in one of 3 modes:
  - **Fabric Discovery** (View inventory/config only)
  - **Fabric Controller** (Manage/configure LAN Fabrics)
  - **SAN Controller** (Manage/configure SAN Fabrics)

- **Capabilities:**
  - Software image/patch/EPLD management
  - Endpoint location and tracking
  - Customizable template-based and freeform **configuration management** with rollback. Templates based on Cisco Best Practices and validated designs
  - Change control with defined user roles
  - Automation of device isolation & RMA replacement

Automate VXLAN-EVPN deployments

Cisco best practice templates for VXLAN EVPN templates

Fabric builder

Support for both brownfield and greenfield deployments

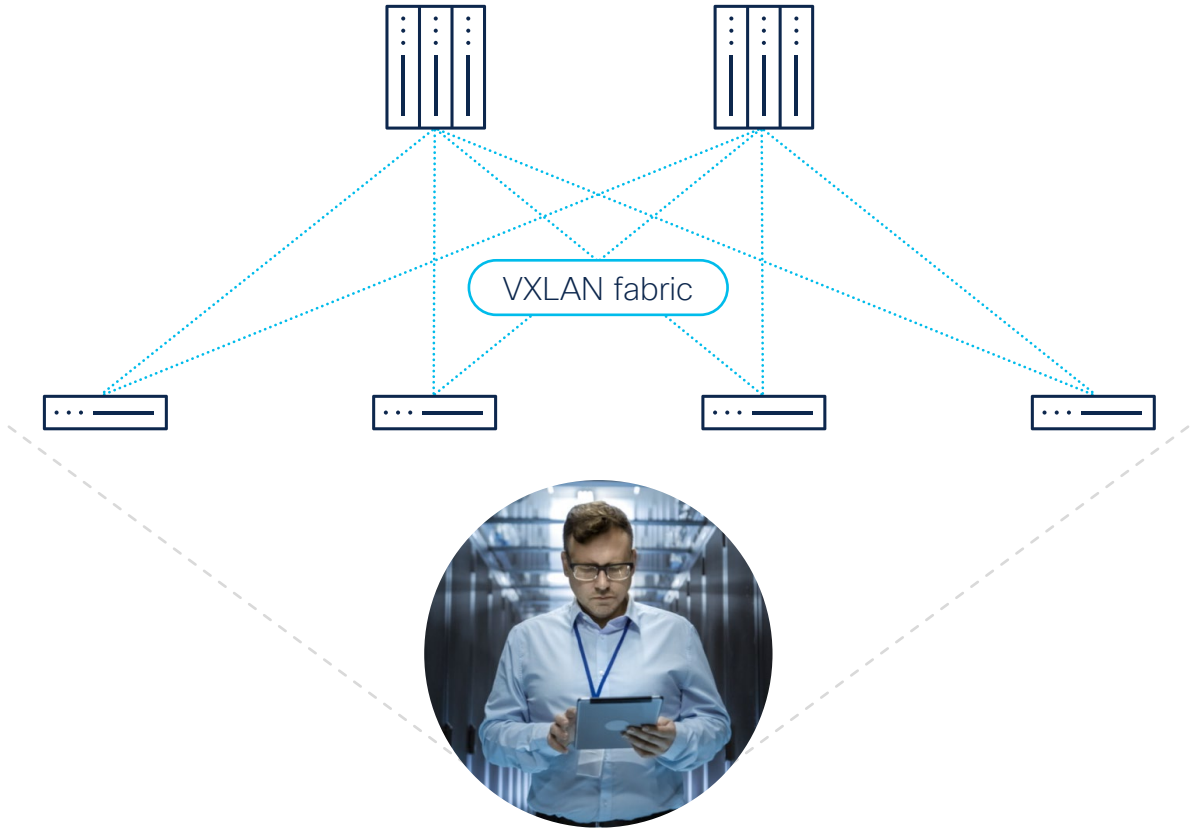

Benefit

Simplify deployment time, reduce chances of errors

Enhanced Classic LAN

Q3CY23

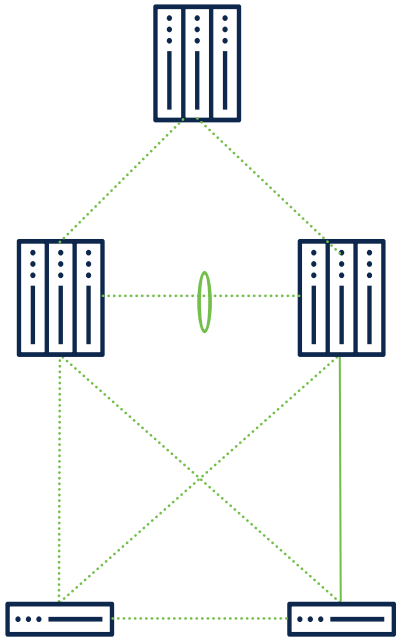

Classic LAN Fabric

- Fully automated fabric – Enhanced Classic LAN
- Support for greenfield and brownfield deployments
- Provisioning of 3tier architecture/ L2/L3 Networks and VRFs
- VRF-Lite Between Agg and Core

Benefits

Best Practice Templates | Simplified workflows | Flexibility based on customer needs

Nexus Dashboard Fabric Components

- Nexus 9000 switches
- Nexus Dashboard
  - Runs on 3-node Nexus Dashboard Cluster in physical or virtual options. 1-node cluster options also available.

NDFC DCI – Multisite Domain

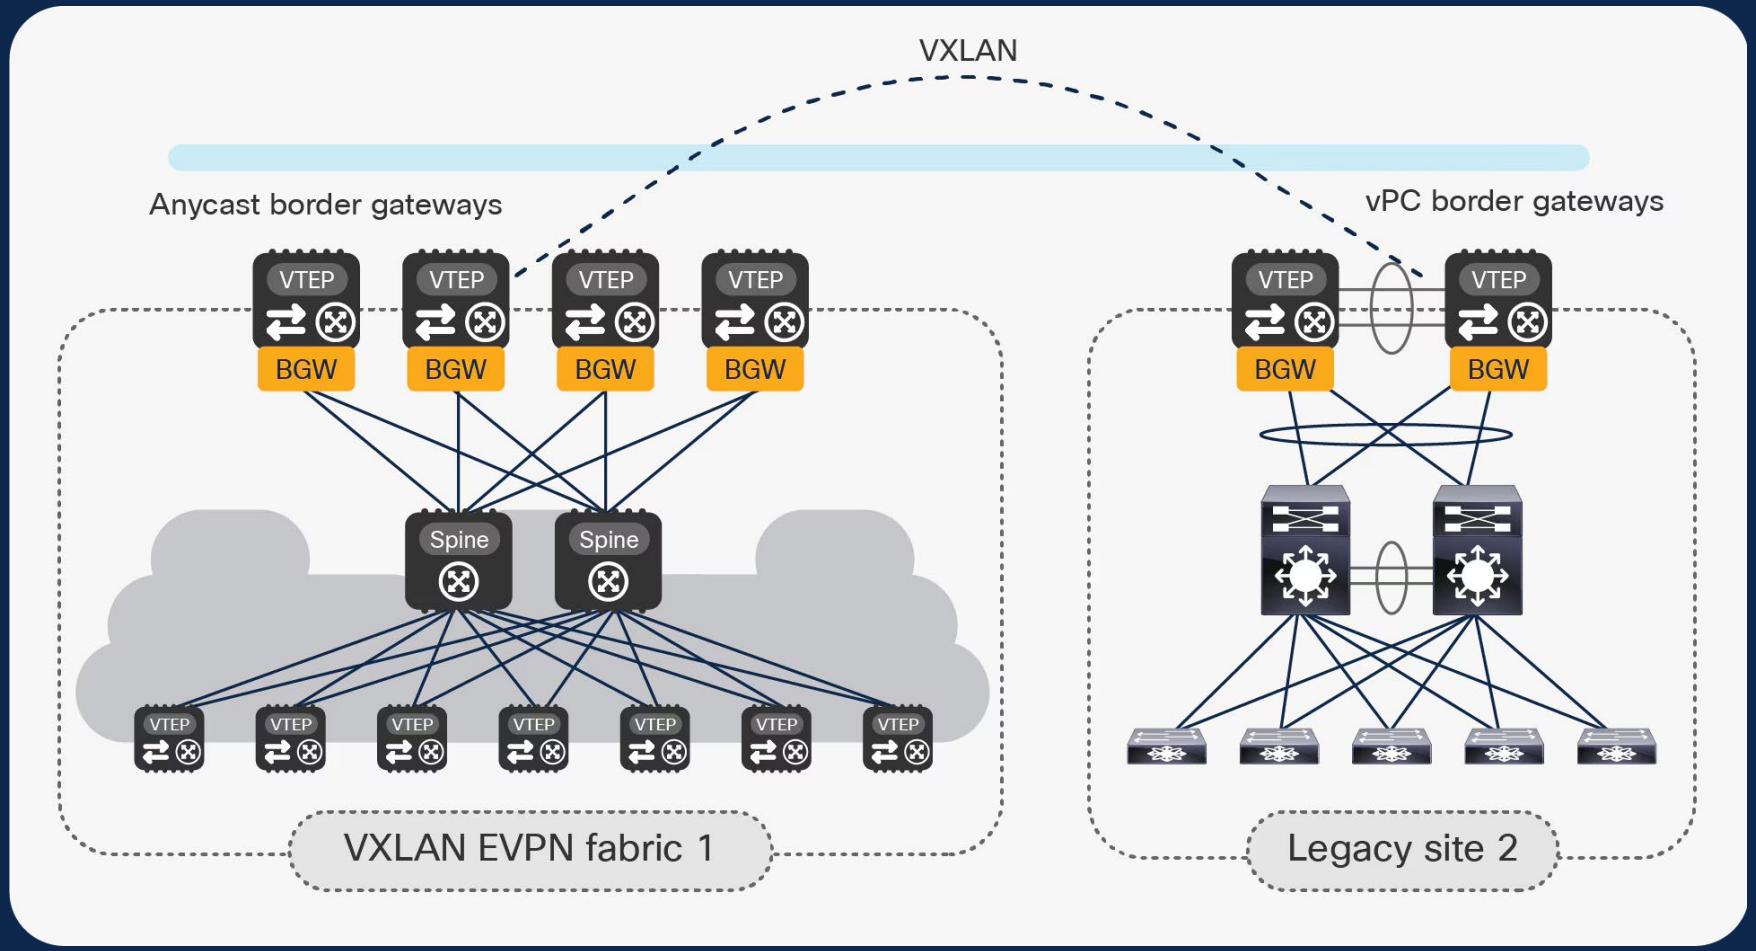

Nexus Hyperfabric

Cisco Nexus HyperFabric

Data center network fabric-as-a-service

Design, deploy and operate on-premises fabrics located anywhere

Easy enough for IT generalists, application and DevOps teams

Outcome driven by a purpose-built vertical stack

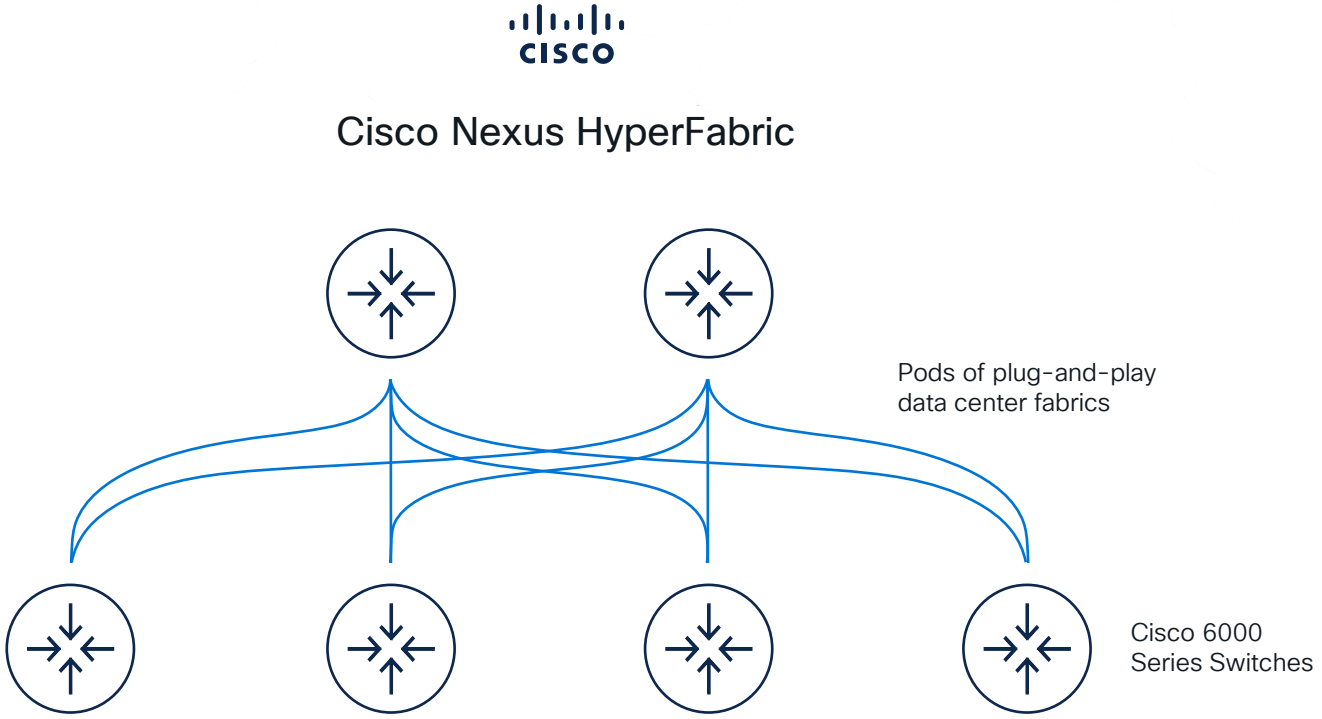

Design › Order › Deploy › Validate › Monitor › Upgrade › Collaborate

Cisco Nexus HyperFabric Components

Cloud Controller

- Scalable, globally distributed multi-tenant cloud service
- Design, plan, control, upgrade, and monitor your fabrics
- Browser, API, and mobile access

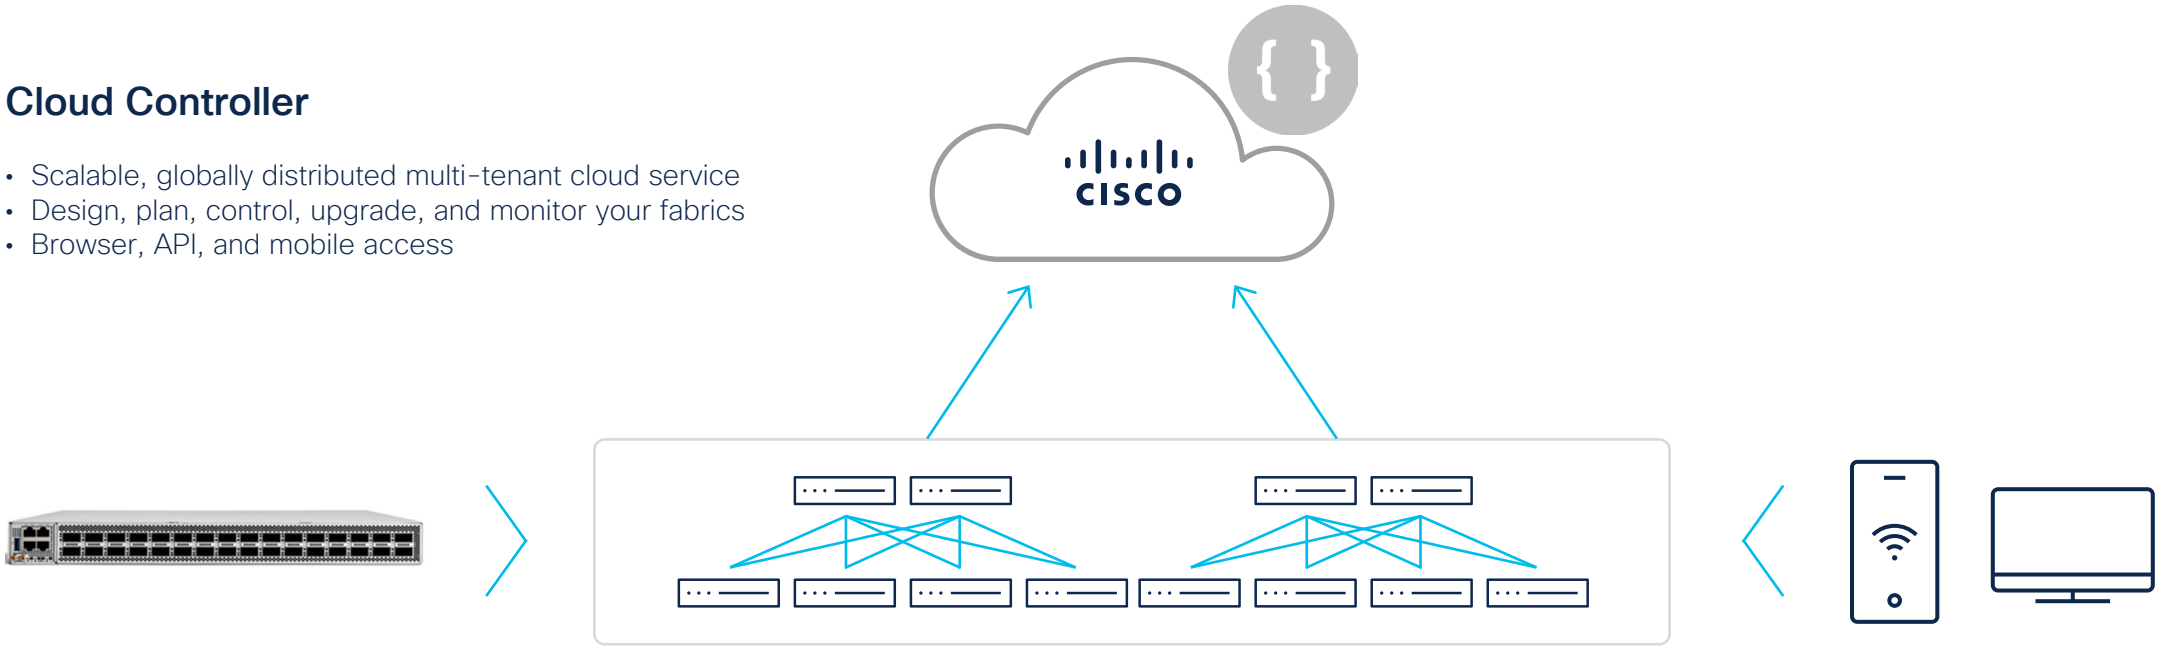

Cloud-managed 6000 switch

- Boot-strapped from cloud
- Full visibility and control from the cloud

High-performance fabrics

- Initially thousands of 10/25/100/400 GbE host ports
- EVPN/VXLAN, layer 2 VLANs, IPv4/IPv6 routing
- Mesh and spine leaf fabrics

Helping hands app

- Step-by-step deployment tasks
- Registration and cabling
- Real-time validation

Small/Remote Data Center Fabrics

Full Mesh, Spine-less

Cisco HF6100-60L4D (Q200)

- 4x 100/400GbE QSFP56-DD (16x via 100GbE breakout)
- 60x 10/25/50GbE SFP56

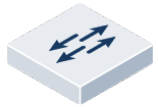

"A Fabric of One"
- Host Ports: 60x 10/25/50GbE
- No dedicated fabric links
- Gateway Ports: 4x 100/400GbE

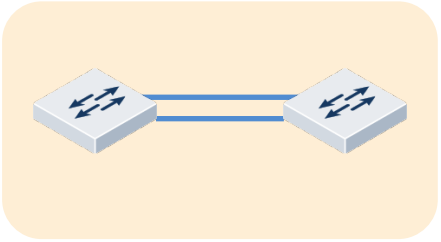

2-Switch Fabric (Remote Fabric)
- Host Ports: 120x 10/25/50GbE
- Fabric Links: 2x400GbE
- Gateway Ports: 4x 100/400GbE

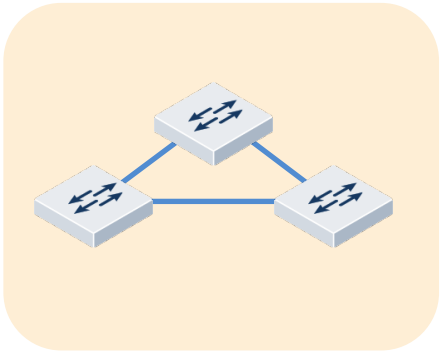

3-Switch Fabric
- Host Ports: 180x 10/25/50GbE
- Fabric Links: 3x400GbE
- Gateway Ports: 6x 100/400GbE

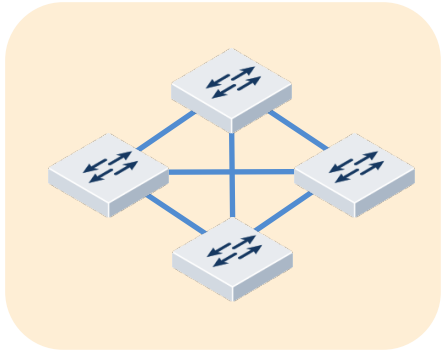

4-Switch Fabric
- Host Ports: 240x 10/25/50GbE
- Fabric Links: 6x400GbE
- Gateway Ports: 4x 100/400GbE

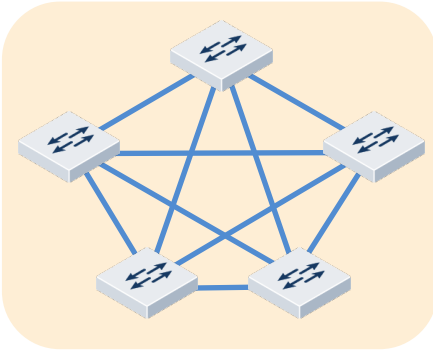

5-Switch Fabric
- Host Ports: 300x 10/25/50GbE
- Fabric Links: 10x400GbE
- Gateway: Host ports only (10/25/50GbE)

Data Center Leaf/Spine Fabrics

10/25/50G Hosts

Spine: Cisco HF6100-32FH (Q200)

- 32x 100/400GbE QSFP56-DD
- 128x 100GbE via 400:100 breakout

Leaf: Cisco HF6100-60L4D (Q200)

- 4x 100/400GbE QSFP56-DD (16x via 100GbE breakout)
- 60x 10/25/50GbE SFP56

Spine switches support host and gateway port modes

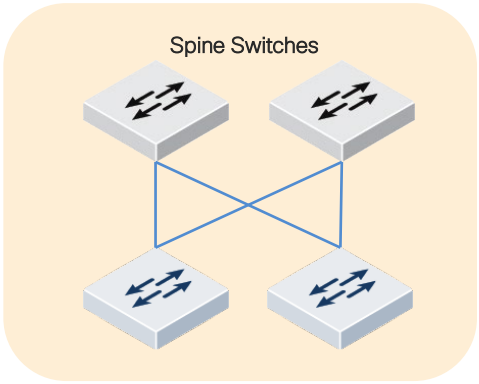

2x Spine · 2x Leaf (min)
- 400GbE Fabric Links
- 120x 10/25/50GbE Ports
- 3.75:1 to 1.87:1 Max Oversubscription

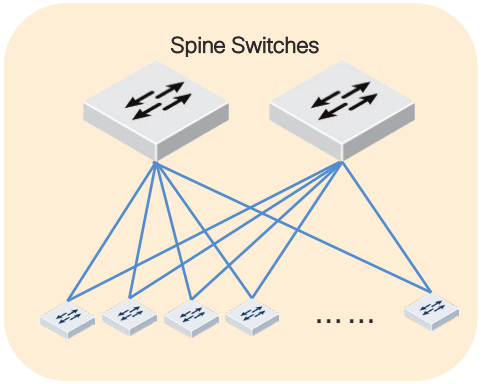

2x Spine · 2-32x Leaf (typical)
- 400GbE Fabric Links
- 120-1920x 10/25/50GbE Ports
- 3.75:1 to 1.87:1 Max Oversubscription

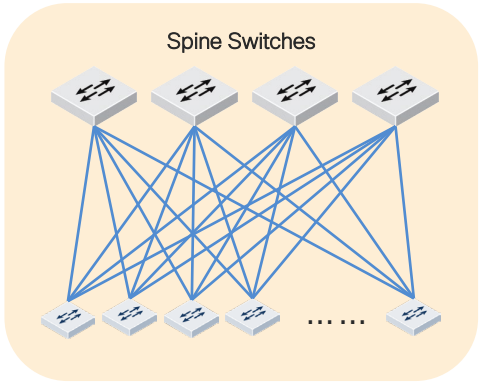

4x Spine · 8-32x Leaf (max)
- 400GbE Fabric Links
- 480-1920x 10/25/50GbE Ports
- 1.87:1 Max Oversubscription


CISCO