Cisco Nexus Hyperfabric

A New Data Center Experience

Martin Diviš, Solutions Engineer, mdivis@cisco.com

Cisco TechClub

Lokální edukační on-line webináře každých 14 dní

Cisco TechClub
hlavní portál

Cisco TechClub
registrační stránka

- Cisco TechClub
Webex Space

Disclaimer

Some of the features described herein remain in varying stages of development and will be offered on a when-and-if-available basis.

This roadmap is subject to change at the sole discretion of Cisco, and Cisco will have no liability for delay in the delivery or failure to deliver any of the products or features set forth in this presentation.

Customers are rapidly adopting **Cloud Managed Networking**

35%

Projected YoY growth in global NaaS market 2023-2030

4M+

Customer networks managed by the Cisco Meraki Cloud

https://www.fortunebusinessinsights.com/network-as-a-service-market-106700

Cisco has reimagined data center **design, procurement, and operations**

Introducing Nexus Hyperfabric

ORDERABLE Q2 FY25

Cisco Nexus Hyperfabric

- 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

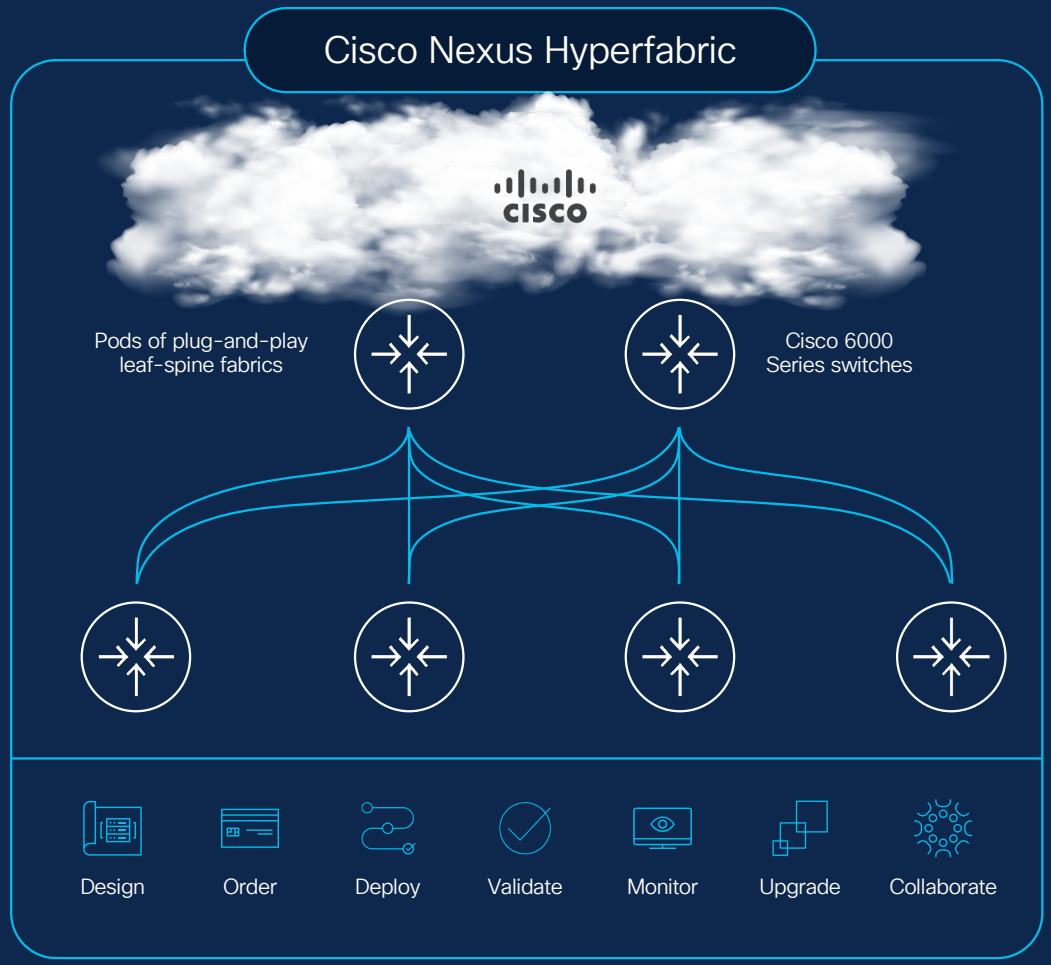

Use Cases

Single global GUI / API endpoint for all owned fabrics

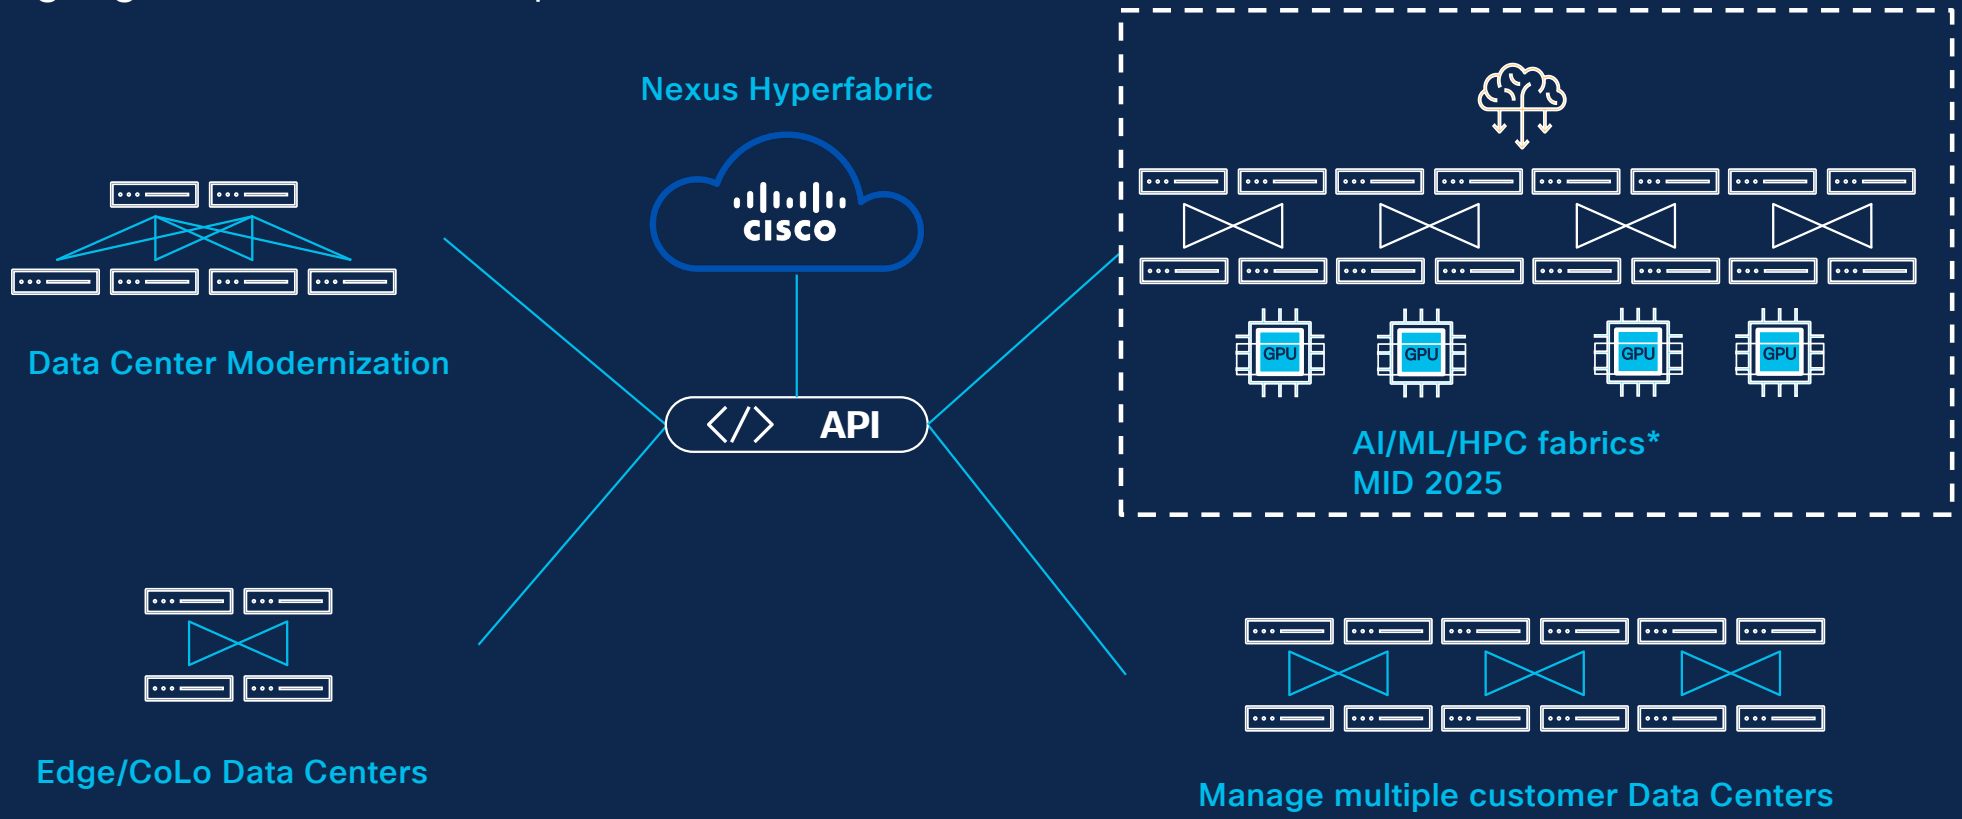

ORDERABLE 2H FY25

Cisco Nexus Hyperfabric AI

- High-performance Ethernet
- Cloud-managed operations
- Unified stack including NVAIE
- AI-native operational model
- Democratize AI infrastructure
- Visibility into full stack AI

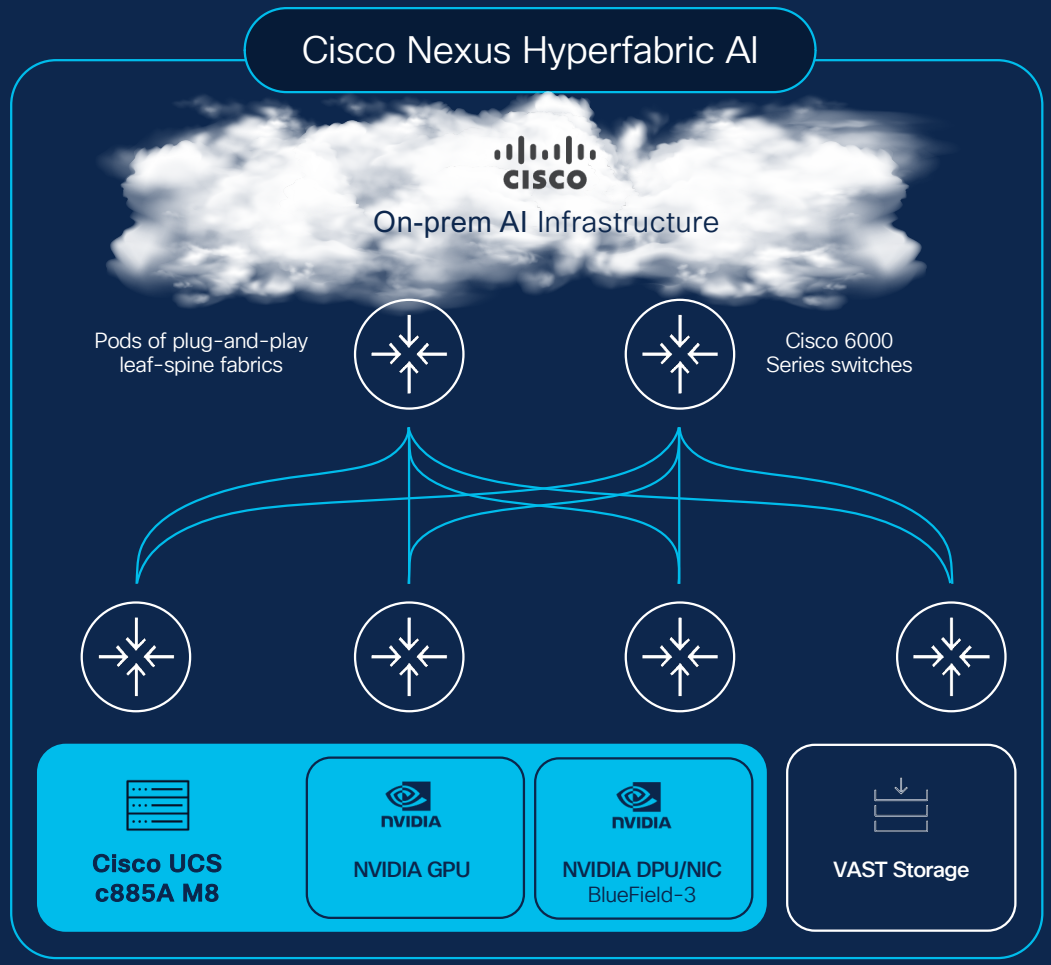

Nexus Hyperfabric	Pre-configured, cloud-managed **networking** for Cisco 6000 switches	**AVAILABLE JANUARY 2025**
Nexus Hyperfabric AI	Pre-configured, cloud-managed **networking, compute, and storage**, partnered with **NVIDIA**	**AVAILABLE MAY 2025**

Nexus Hyperfabric
Experience

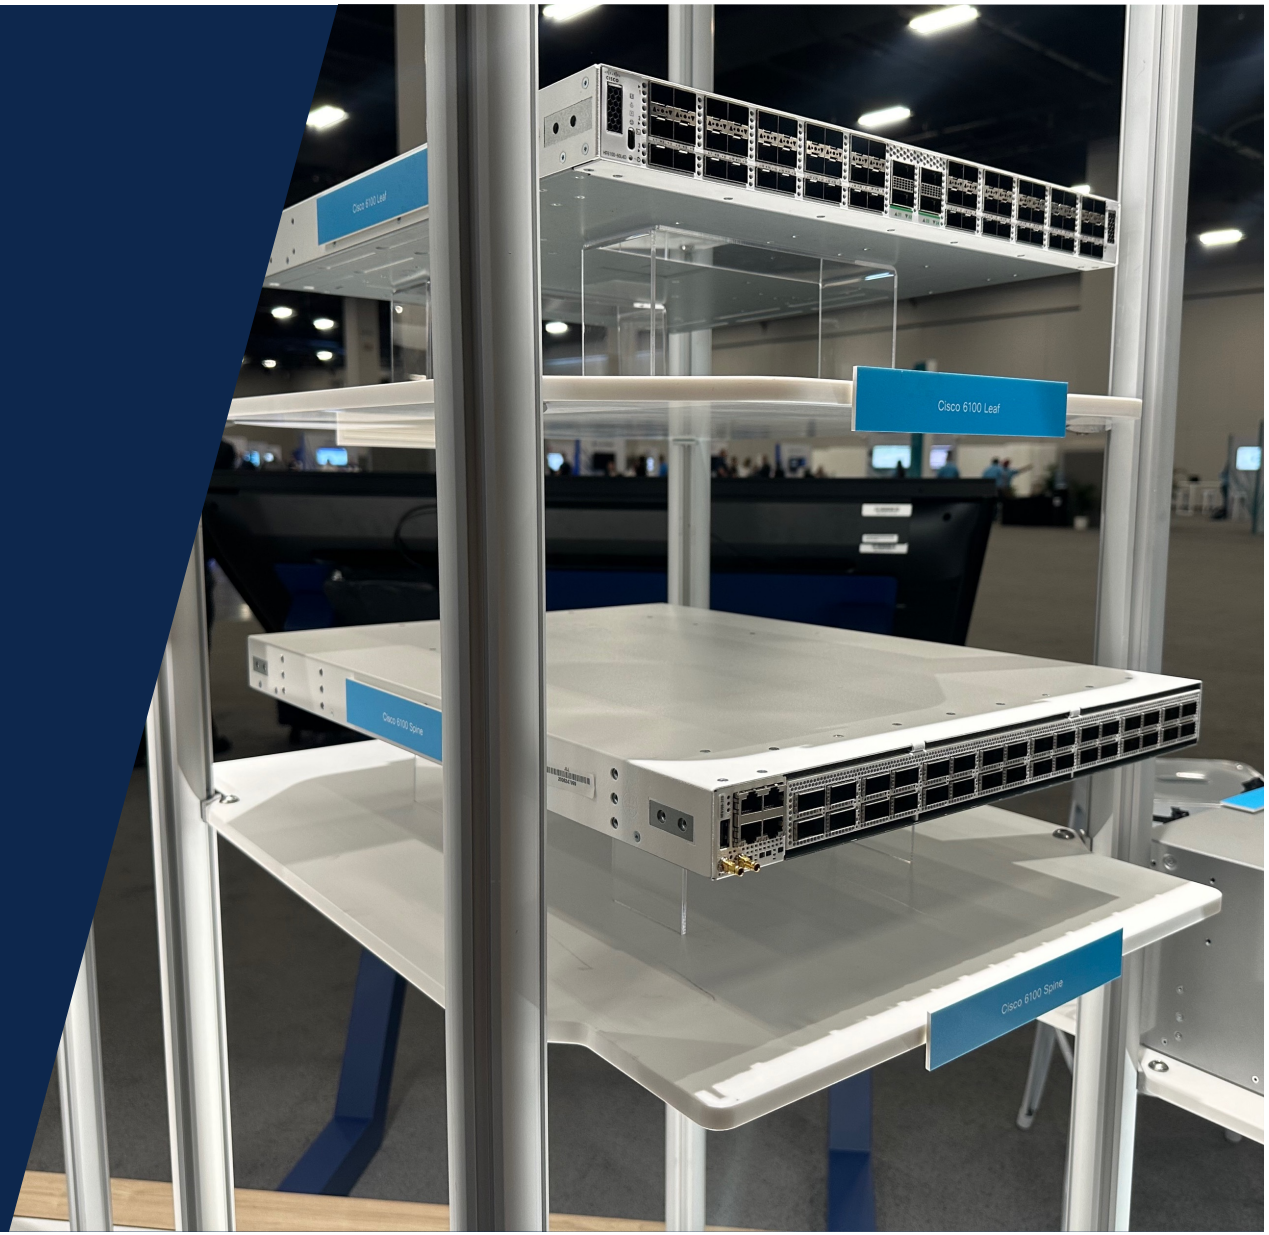

How It Works

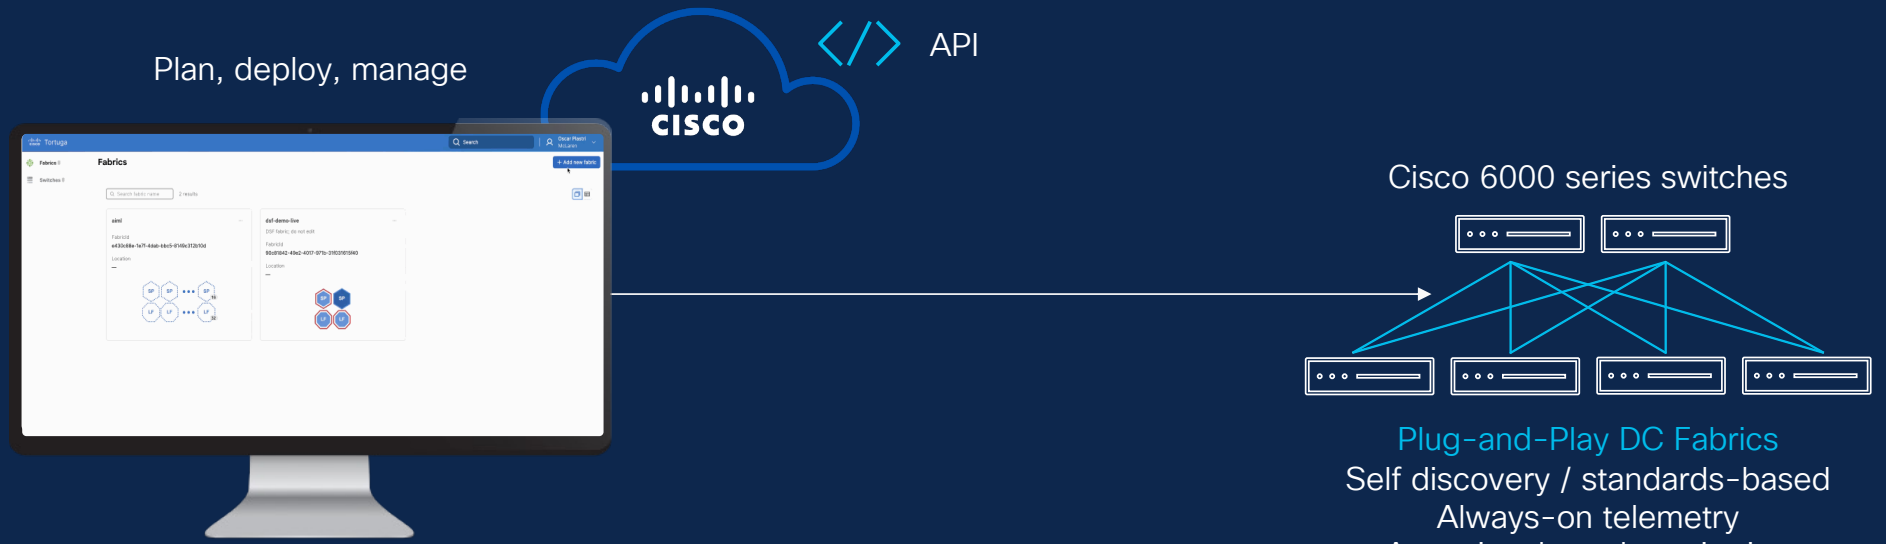

Plug-and-Play DC Fabrics
Self discovery / standards-based
Always-on telemetry
Assertion-based monitoring

Purpose-built for predictable outcomes optimized for ease of use

Complete Lifecycle Experience

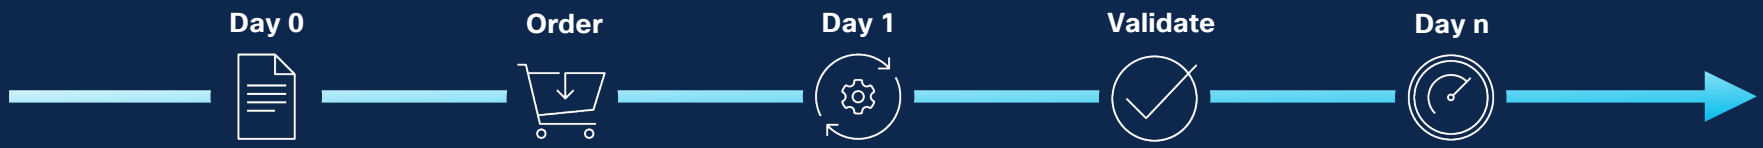

Plan with fabric designer

Configure network while equipment ships

Plug in and register to fabric

Ensure deployment matches design

Monitor, upgrade & collaborate from **anywhere**

Cloud-managed controller

- Scalable, globally distributed multi-tenant cloud service
- GUI, mobile, and API access
- **Helping hands** for smart remote hands visibility

Cloud-managed switch

- Cisco 6000 series
- Boots from cloud
- Full visibility & control from the cloud

Software Lifecycle Management

Cloud SaaS controller:

Continuous delivery model: always up-to-date

- Continuous delivery of new features and software updates to the production cloud service.
- No user testing or software maintenance required

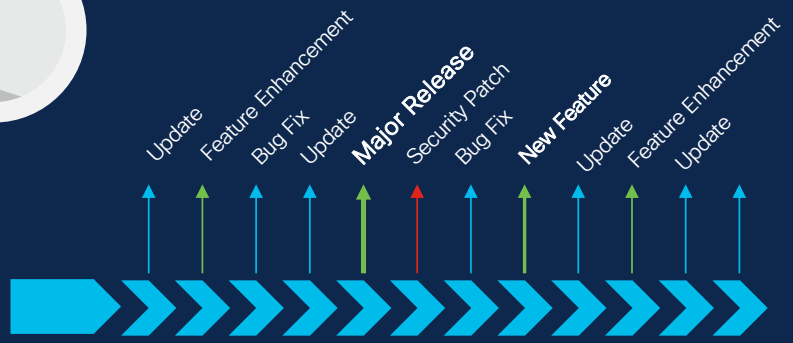

On-prem switch software

Cloud-delivered Software Upgrades: User-Driven Update Schedule

- Schedule firmware updates
- Software rollback support
- Intelligent sequencing of fabric upgrades

Flexible Architectures

Deploy any fabric, anywhere

Mesh / spine-less fabrics

"A Fabric of One"

2-switch fabric
120 host ports

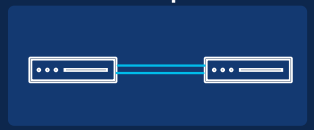

3-switch fabric

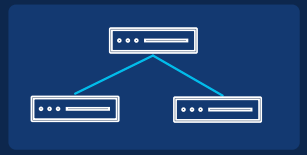

4-switch fabric
240 host ports

5-switch fabric

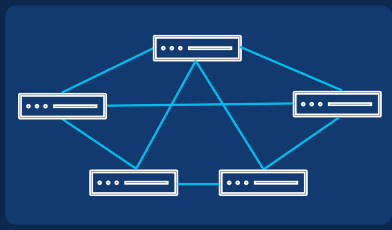

2 spine, 2 leaf

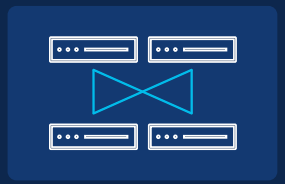

2- or 4-way spine, 2-32 leaf
Nearly 2000 host ports

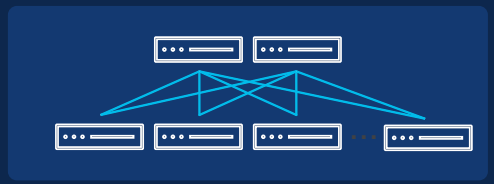

Leaf-spine DC fabrics

Cisco 6000 series switches

Leaf: HF6100-60L4D

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

Spine/Leaf: HF6100-32D

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

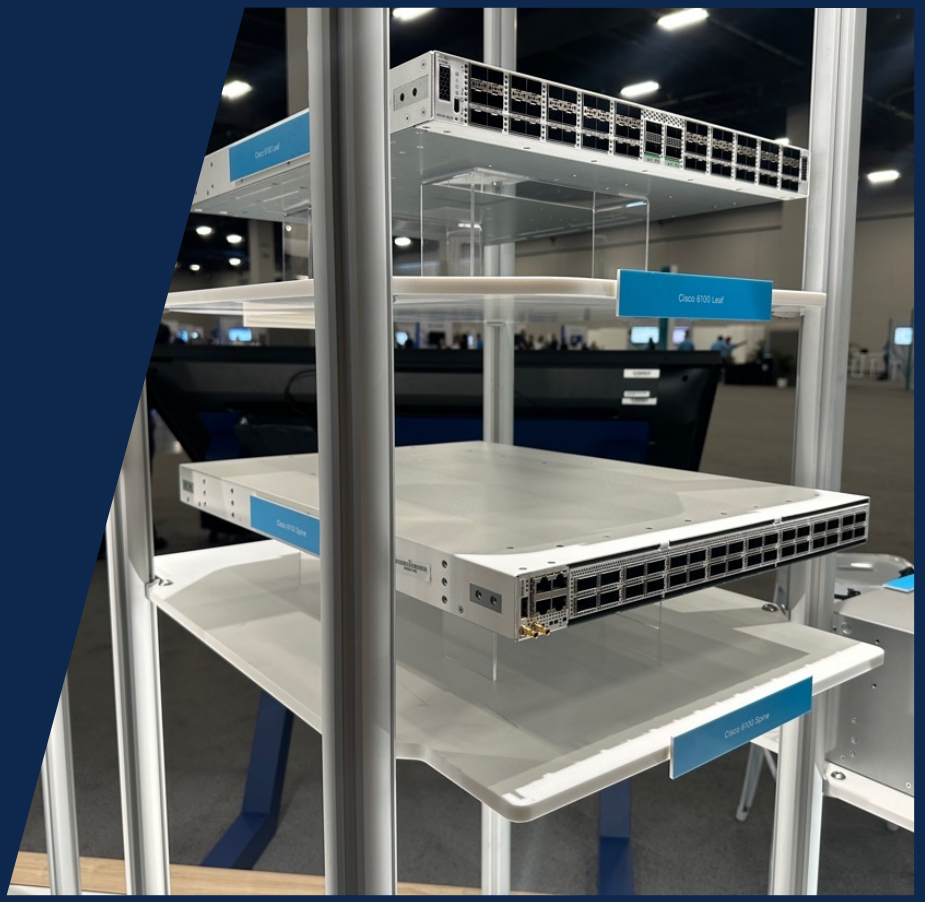

Why is Hyperfabric Designed like this?

- Hyperfabric is not a data center fabric: it is an OPERATIONAL MODEL driven by cloud
- Hyperfabric is designed to rethink network automation from a top-down model, and is designed to work like cloud IaaS
- Hyperfabric is about creating an ABSTRACTION MODEL for the network that is consistent and scalable.
- Hyperfabric is about building a supportable and repeatable fabric architectures – no snowflake fabrics
- Hyperfabric is about building and reinforcing value in the CONTROLLER/SERVICE and not the hardware/NOS.
- Hyperfabric is designed to leave legacy features behind wherever we can (Hyperfabric will never have feature parity)
- Features implemented consider the COST to deliver as a service via cloud

Hyperfabrix User Experience

Pre-release Image • Subject to change

Blueprint Creation

- Fabric Blueprints are created in an organization
- Initial Blueprint creation involves the following:
  - Fabric Name: Can include a text description, as well as Address, Location, and a Tags for automation purposes.
  - Switches/Cabling/Optics: You select the switch hardware, desired cabling, and optics based on the supported combination. You can choose to allow for multiple fabric connections between switches.
  - Topology: Choose the desired switch topology based on the deployment requirements.

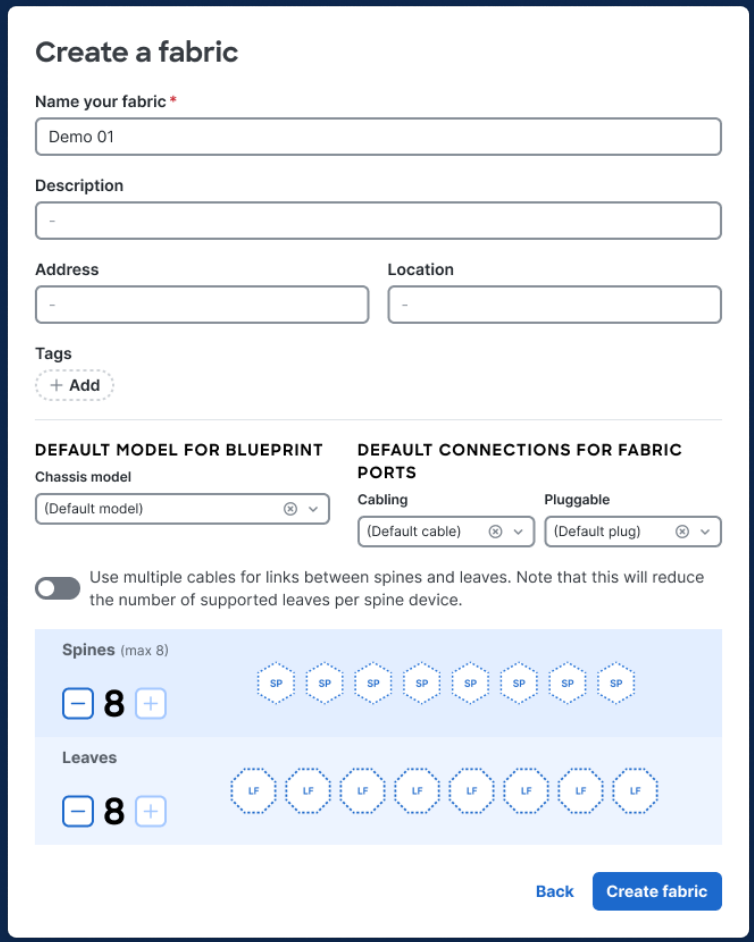

Pre-release Image • Subject to change

Per switch inventory information

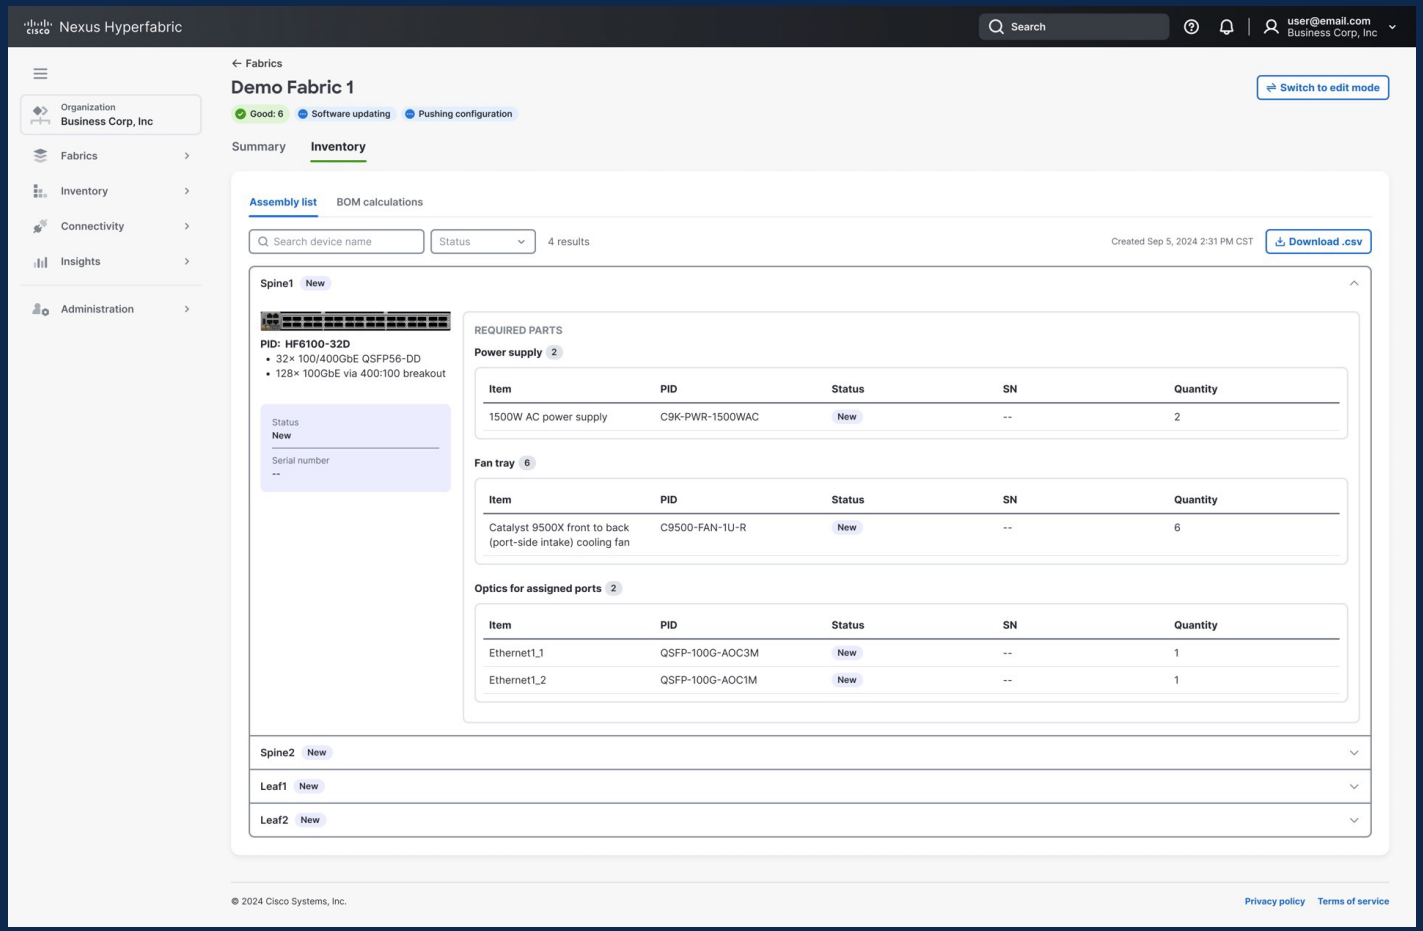

Pre-release Image • Subject to change

BOM (total number of switches and optics)

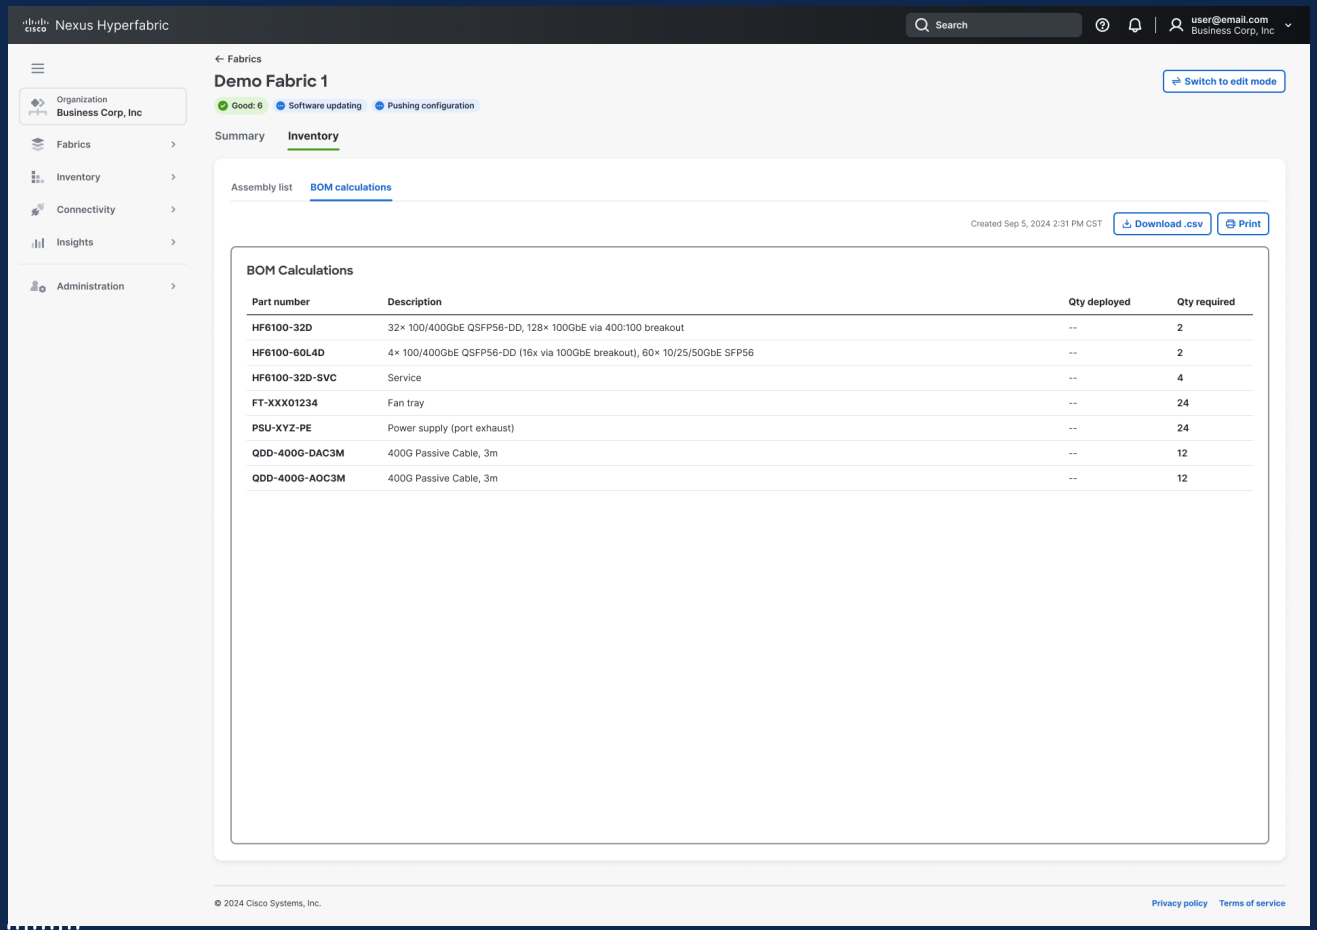

Pre-release Image • Subject to change

Cable Map

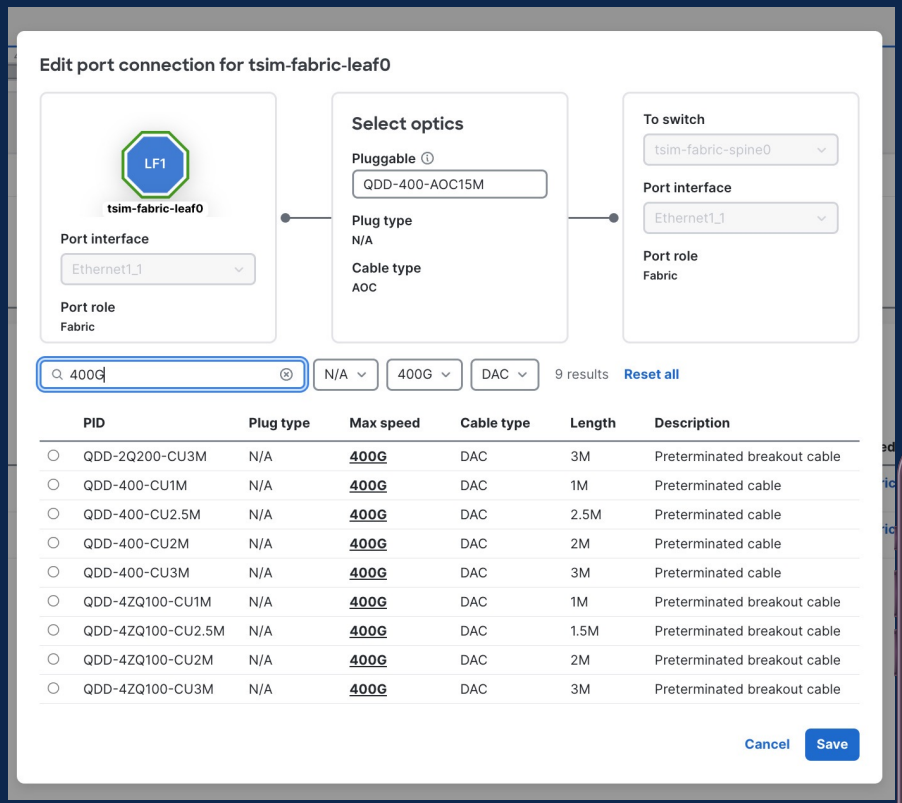

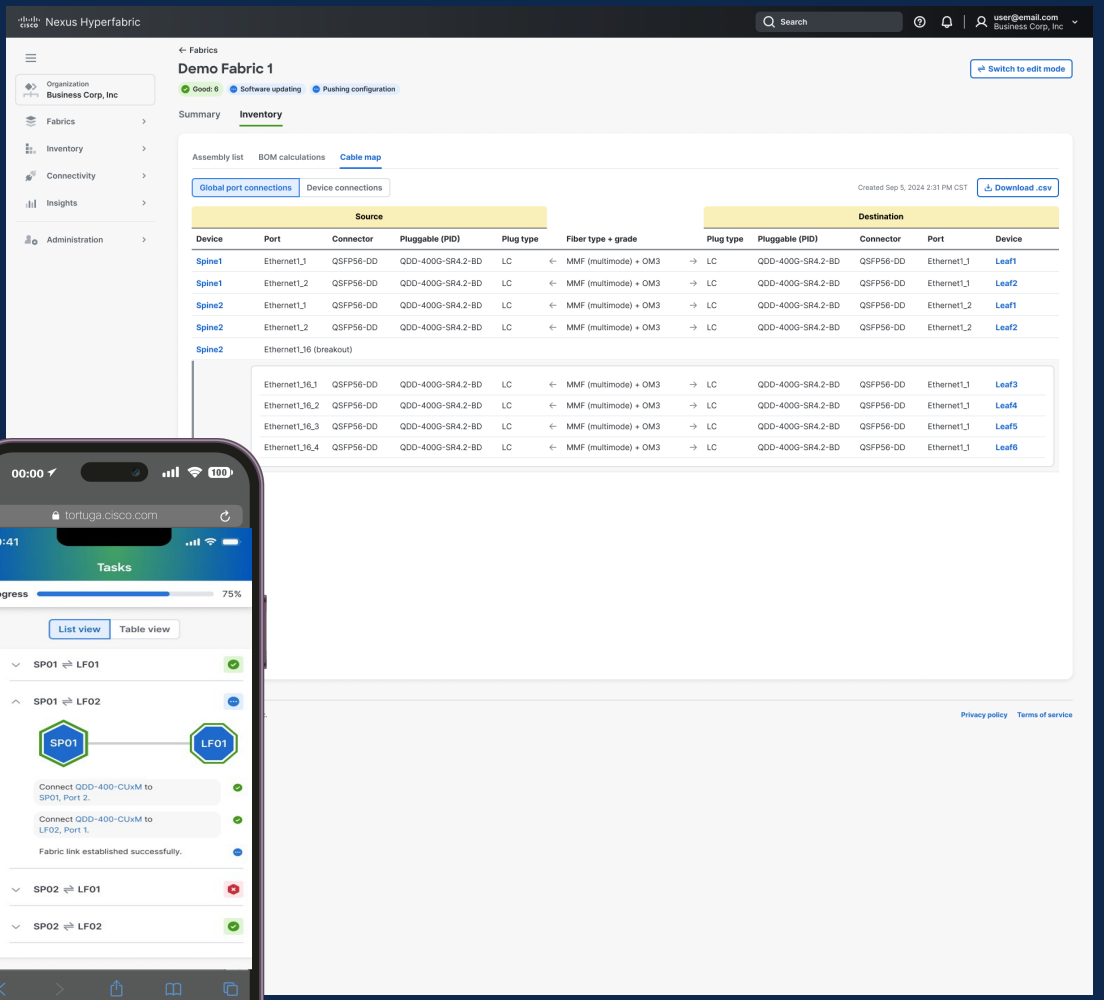

Bootstrapping the device

- Assign the switch with
  - management IP address/mask (IPv4/IPv6)
  - Gateway
  - DNS
  - and if needed, proxy
- Methods of initial configuration
  - Via management port / console -> dialog-based configuration
  - Via USB stick with configuration attached on first boot
    - USB stick can contain configuraition for multiple devices
  - DHCP

Connecting to the cloud controller

- Connectivity via
  - Management port
  - In-band
  - Proxied via another switch in the same fabric
  - One is enough, more is better, montitored via assertions
- Forward proxy support with HTTPS decryption/inspection
  - Approve proxy certs on the switch
  - HTTPS within HTTPS

The fabric

- VXLAN BGP EVPN
- IPv6 IP unnumbered uplinks
- All of the underlay is managed by Hyperfabric, no CLI, no direct visibility
- VRFs supported
- VNI assignement – auto or manual
- SVI for VNI’s, anycast GW

Assertions

- Tests on configured entities/features
- Prevents alert storms
- Latched assertions

Interfaces and VLANs

- VLAN to VNI mapping
- VLANs are port-significant
- MLAG – EVPN ESI
- Loop prevention via STP
  - all leaf switches are participating in the same L2GSTP

Demo

Demo Link

External L3 connectivity

- Static routing
- BGP
  - Routed interfaces and subinterfaces supported
  - VRF aware
  - Multiple sessions in multiple VLANs supported
- No dynamic routing on SVIs, static only
- No L3 multicast within fabric

Assertions on intra-fabric connectivity

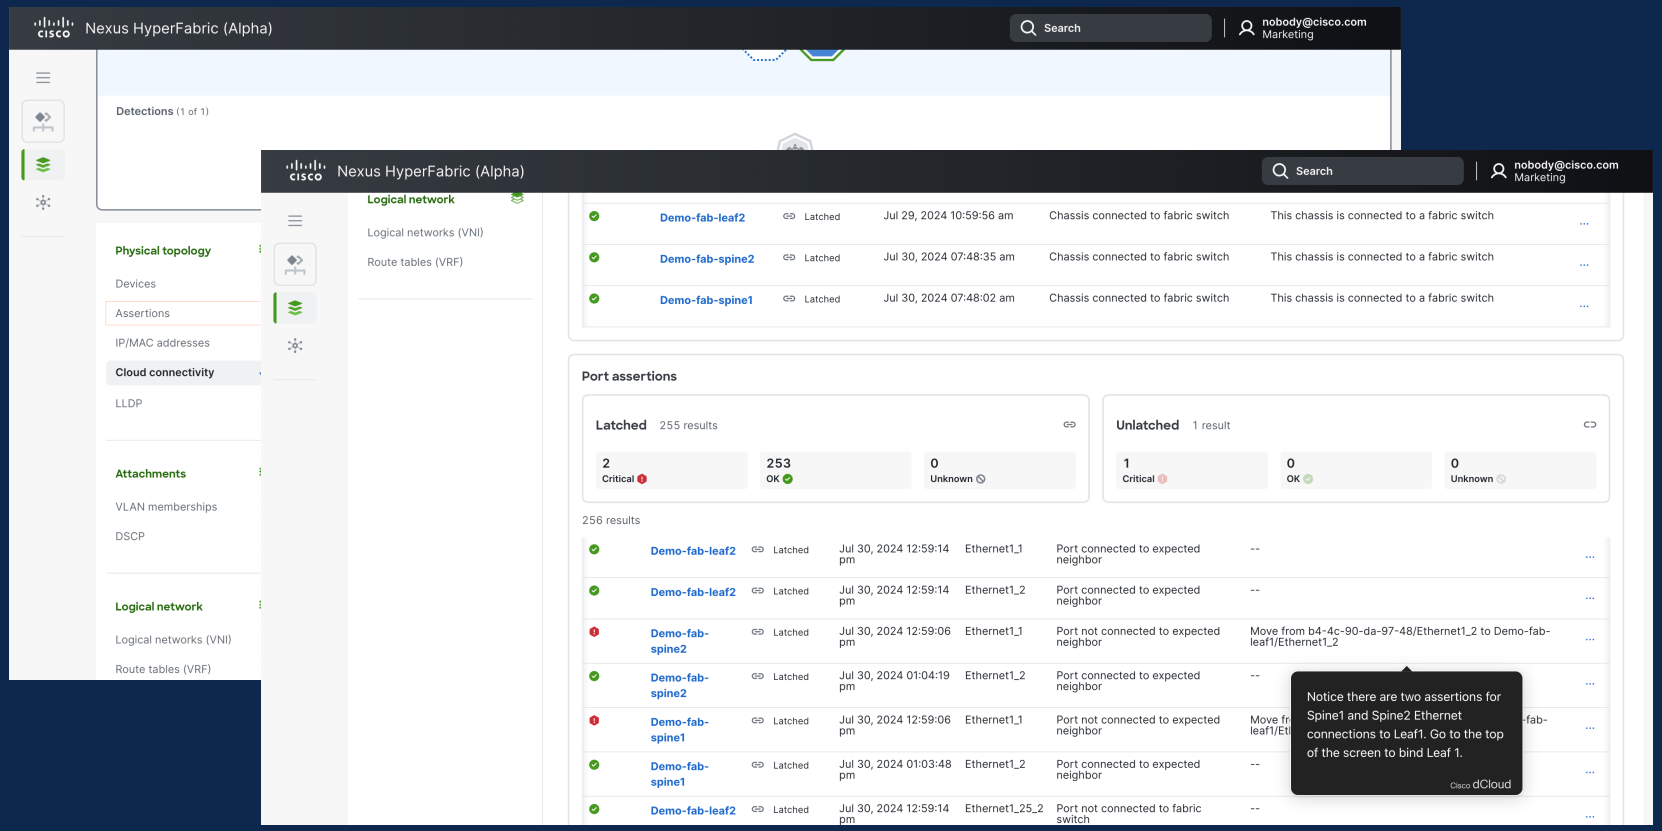

Monitoring and troubleshooting

- Metric collection
- Ad-hoc packet capture and analysis via cloud controller
- No Netflow to the cloud controller – BYO-NfC planned

Hyperfabric AI

ORDERABLE 2H FY25

Cisco Nexus Hyperfabric AI

- High-performance Ethernet
- Cloud-managed operations
- Unified stack including NVAIE
- AI-native operational model
- Democratize AI infrastructure
- Visibility into full stack AI

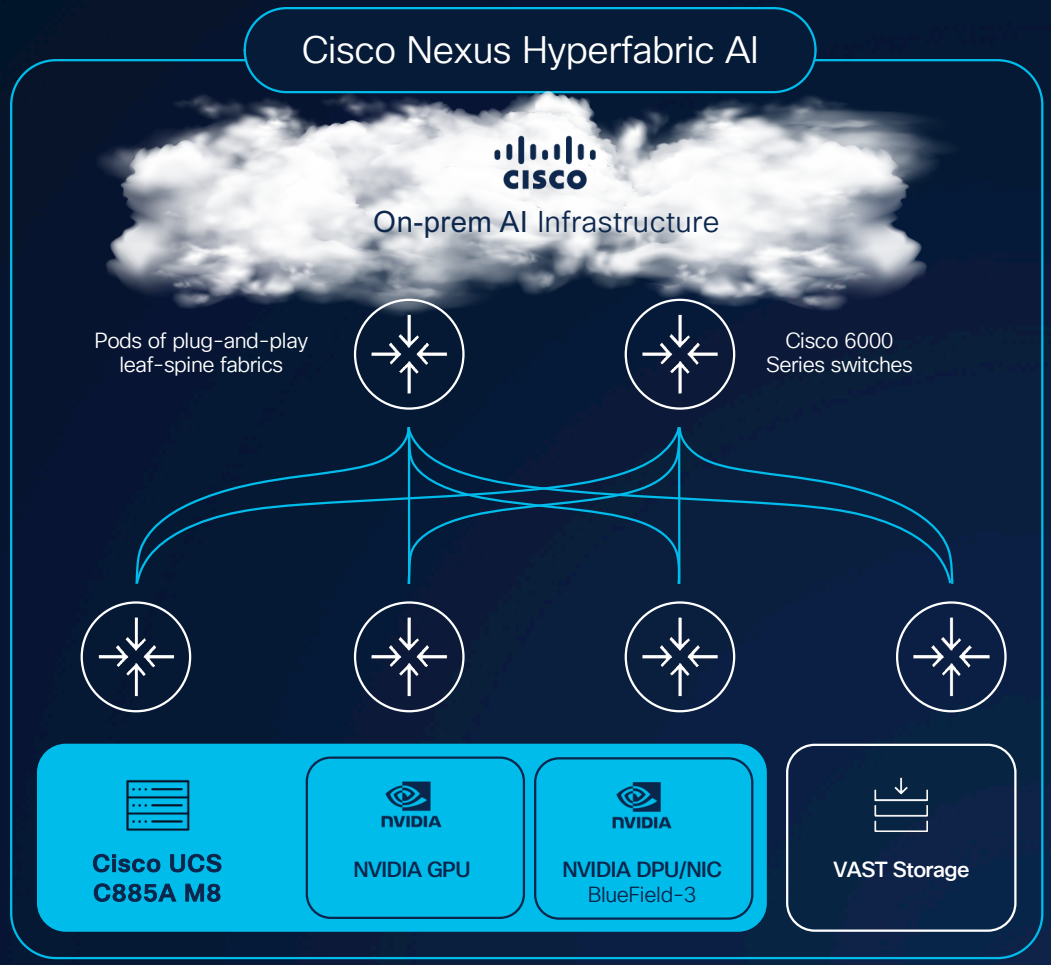

Building high-density GPU servers to the Cisco UCS family and to Cisco's AI solution portfolio

Discover data-intensive use cases like model training and deep learning

Orderable in October 2024

UCS Accelerated

UCS C885A M8

Nvidia HGX with
8 Nvidia H100 GPUs

AMD Mi300X

2 AMD 4th Gen
EPYC™ Processors

Generative AI Spectrum

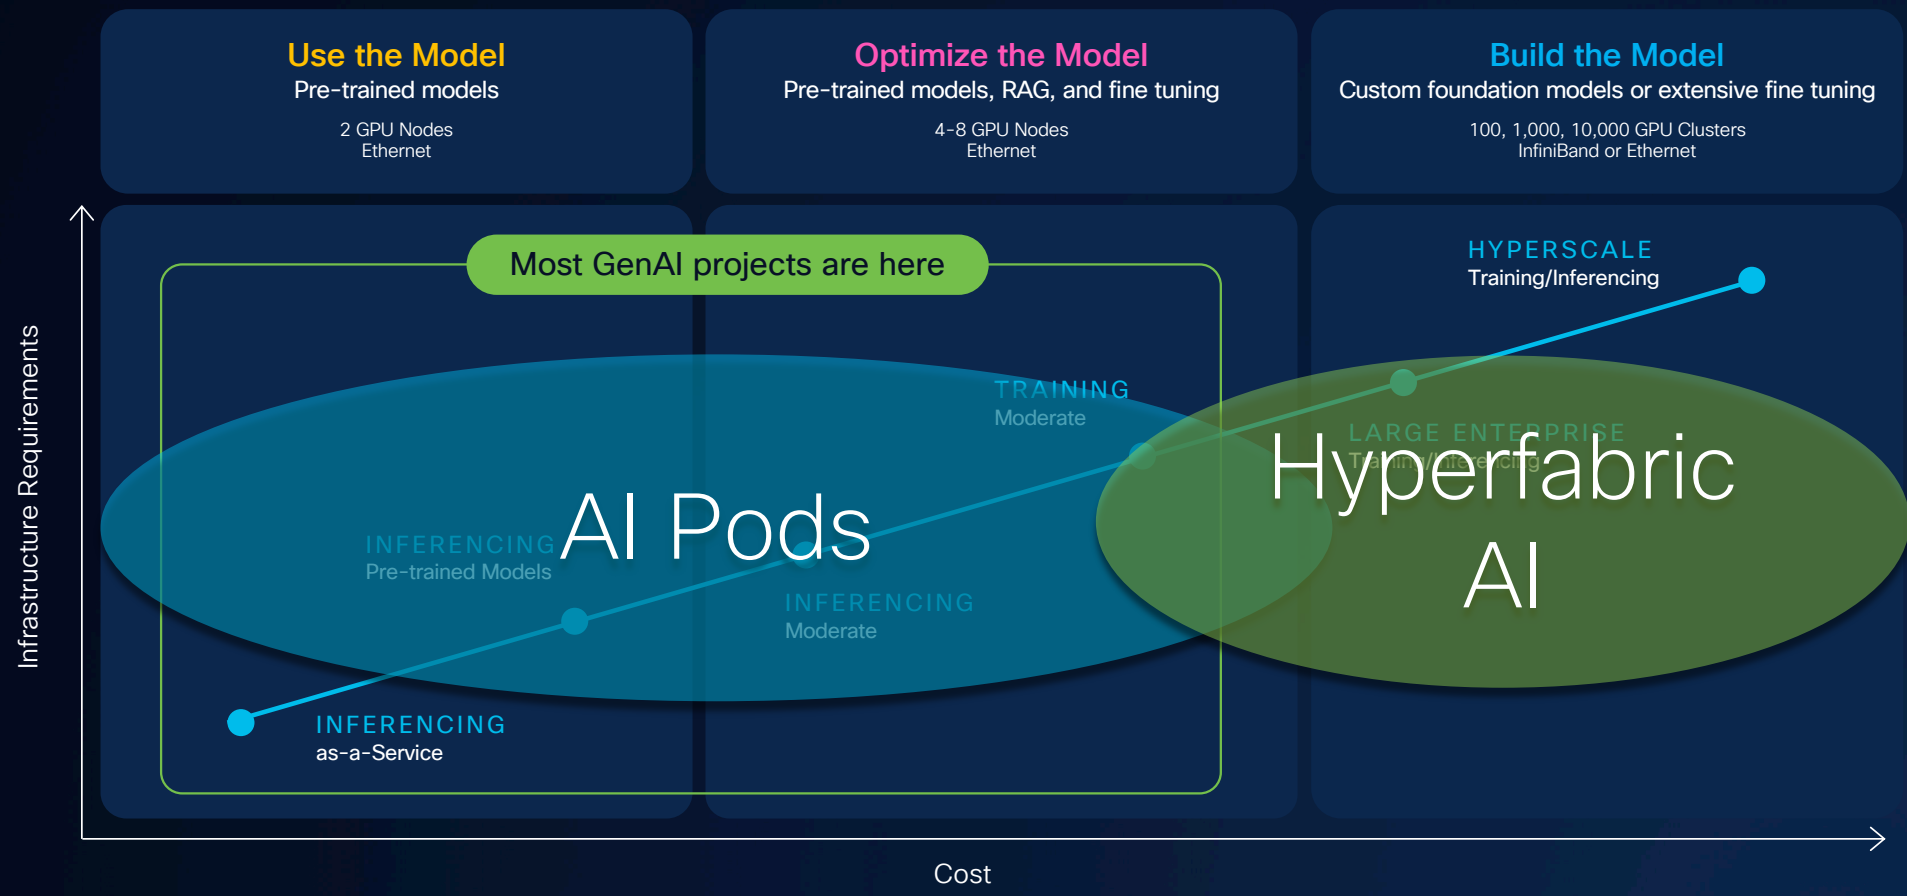

Looking Ahead

Product Direction

	FCS	Future
Designs	Mesh & spine/leaf	DCI (Border GW) Super-spine AI fabric
Fabric Scale	<50 switches <2K host ports	Increase in scale
Hosting	US	EMEA, then APJC
Features	EVPN, VXLAN underlay Static & BGP routing MLAG, IPv4/IPv6	Hypershield Common policy (SGT/SXP)
License Tier	Essentials	Advantage (advanced features) Premier (for Hyperfabric AI)
Models	HF6100-32D (32-port 100/400G) HF6100-60L4D (60-port 10/25/50G, 4x400G)	800G

Subject to change

New Pricing Model

Additional Resources

[Cisco.com - Cisco Nexus Hyperfabric]

Other DC Switching News

Secure & Futureproof Infrastructure

Programmability Performance Flexibility Efficiency

Cisco Nexus 9000 Services Accelerated Switch

Silicon One

AMD PENSANDO

- Rich NX-OS Features and Services
- High-speed connectivity and scalable performance
- Optimized for latency and power efficiency

Routing/ Switching | EVPN/MPLS/ VXLAN/SR | Rich Telemetry | Line-rate Encryption | Power Efficiency

- Software-defined Stateful Services
- Programmable at all layers: add new services without HW change
- Scale-out services with wire-rate performance

Large-Scale NAT | IPSEC Encryption | Distributed Firewall | Event-Based Telemetry | DoS Protection

800GE DC Switching

Cisco Nexus 9364E-SG2

- 64x 800GE, capable of 100GE and 400GE breakouts
- high-performance compute (HPC), webscaler, media provider, and Telco SP segments
- NX-OS, and Nexus Dashboard

Cisco 8122-64EH/EHF

- 64x 800GE , capable of 100GE and 400GE hyperscalers and webscalers adopting the hyperscaler model
- open customizable solutions
- SONiC network operating system or other 3rd party NOS
- API support for open tooling

Cisco TechClub

Lokální edukační on-line webináře každých 14 dní

Cisco TechClub
hlavní portál

Cisco TechClub
registrační stránka

- Cisco TechClub
Webex Space

Cisco TechClub 26.11.

Nové use-cases a zařízení v SP portfóliu

"Novinky v HW, managemente a monitoringu siete, prefixov a sluzieb"


CISCO