

Al-Ready Data Centers

Franck BONNEAU - Cloud and Al Infrastructure Adrien LÉCHARNY - Cloud and Al Infrastructure

22 mai 2025



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.



Data Centre Strategic Priorities



Upcoming wave of growth and capacity

- Next gen switch ASICS
- Server platform updates/enhancements

Hardware design aligned with environmental goals

Core and edge data center Al use cases

- HyperFabric AI
- Al PODs

Leverage web scale investment

Partnerships, blueprints, & validated designs

Deliver common experience

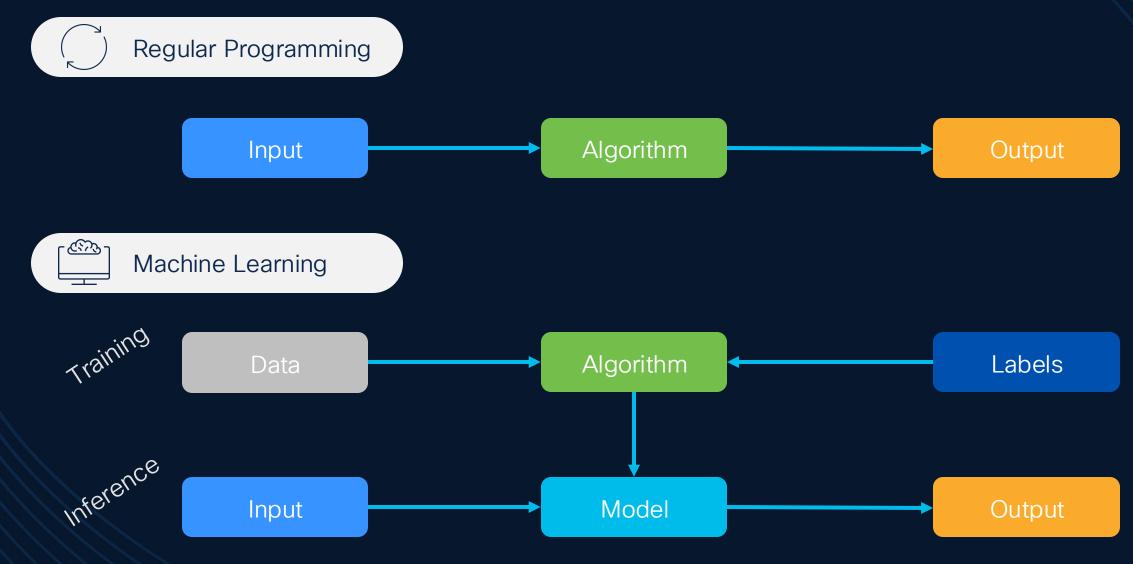
- Unified Nexus Dashboard for VXLAN EVPN and ACI fabrics
- Intersight for AlOps

Advanced use case

- RedHat and OpenShift
- NOS, controller, and API evolution for advance ACI use cases

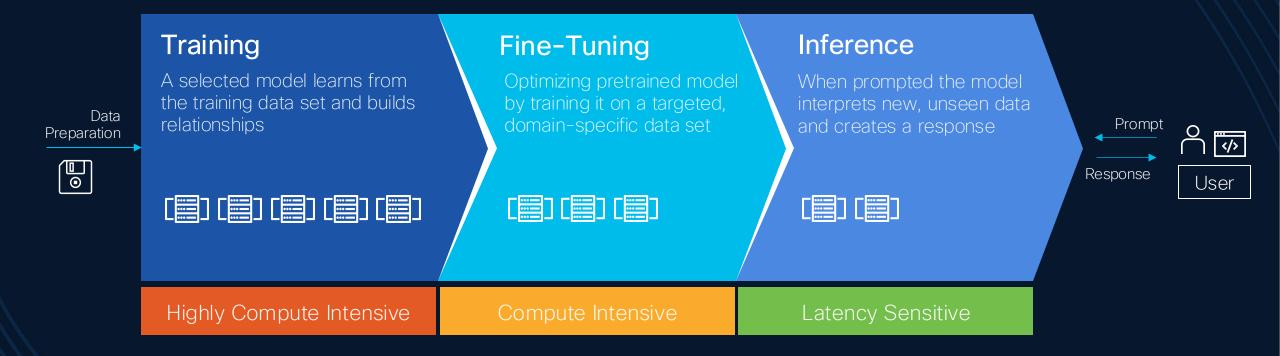


How is Al Different from Programming?





Al Workflow and Infrastructure Impacts



Model performance tightly coupled to infrastructure



Deployment Options for an Al Fabric





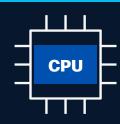


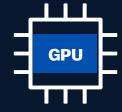
Dedicated Ethernet Fabrics











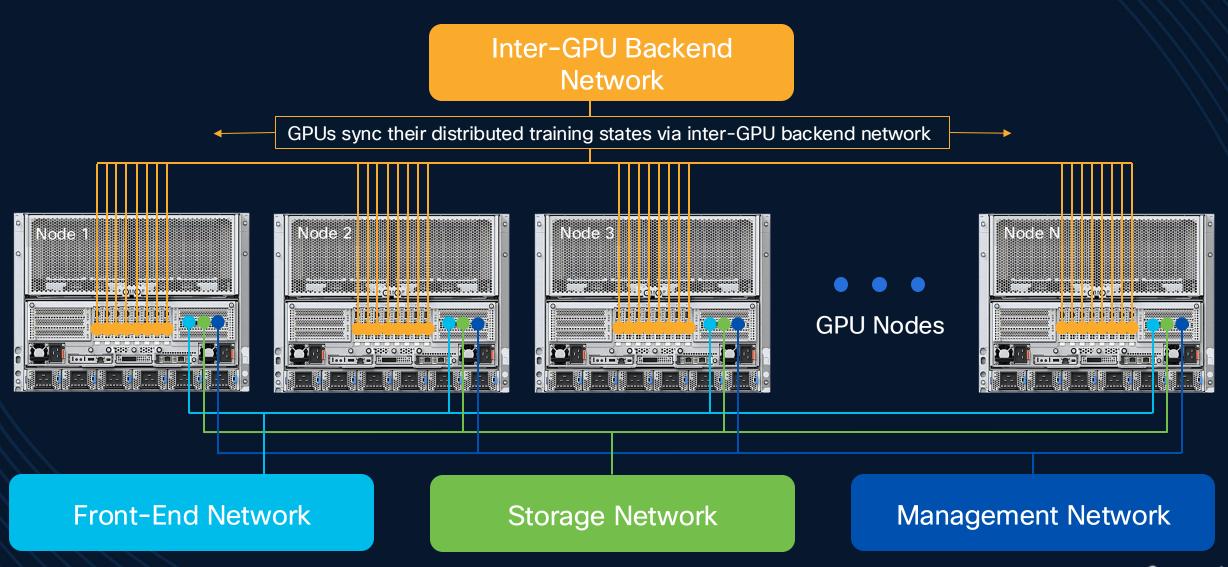
intel AMD







Al Cluster Networking Blueprint





Key Network Challenges and Requirements



Training requirements



Inference requirements



General requirements



High bandwidth:

Movement of massive datasets across the network demands high bandwidth networks



Non-Blocking Lossless:

Network inconsistencies can affect the accuracy and training time of Al models



Congestion management:

Detect potential congestion and redistribute network traffic accordingly



Low latency:

Real-time Al applications require extremely low latency



Visibility:

Comprehensive visibility tools for real-time monitoring, issue detection, and troubleshooting



Scalability:

Dynamic nature of Al workloads mean that the network needs to be agile and scalable









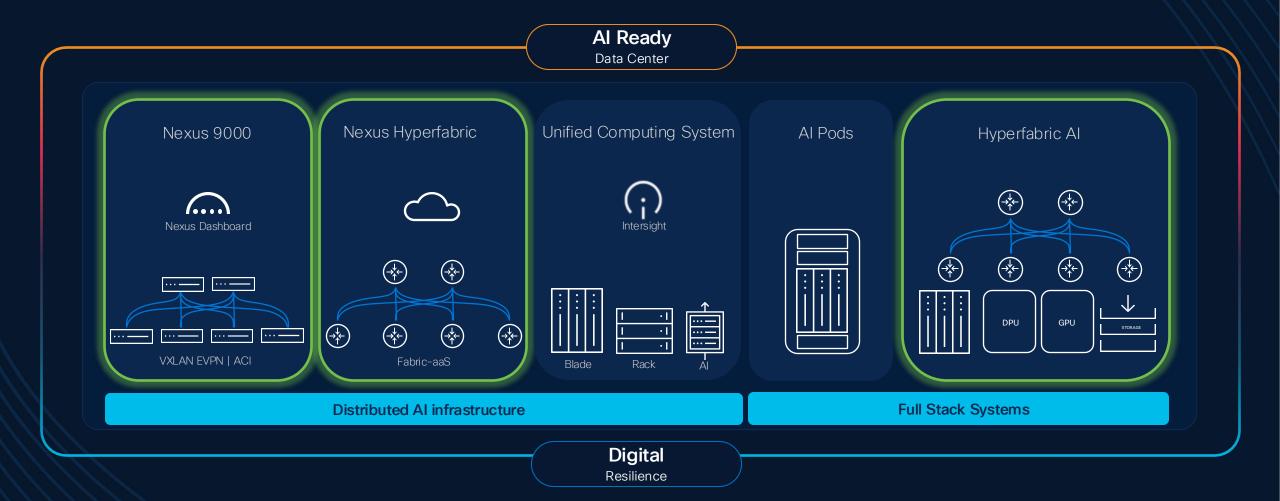








Cisco Al-Ready Data Center Building Blocks





Cisco Al/ML Approach

Building blocks

Nexus 9300 400G or 100G Switches Cisco 6000

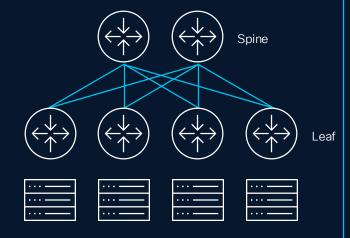




400G DR4 or 100G DR Optics



Non-blocking spine-leaf fabric



AI/ML for networking blueprint RDMA over Ethernet (RoCEv2) Lossless network (PFC + ECN)

Proven full network stack



Fabric automation and visibility made simple



Extensive interop testing with major NIC vendors



Widely deployed OS with rich feature set



25.6T and 51.2T Silicon supporting 50G/100G SerDes



Recent deployments with financial institutions, research schools, government agencies





Optimized AI/ML Fabrics with Nexus

Silicon, systems, software, operations

400/800G Ethernet transition (25.6T & 51.2T switches) High-bandwidth fabrics with reduced footprint and energy savings

RDMA over Ethernet (RoCEv2) Lossless network (PFC + ECN)

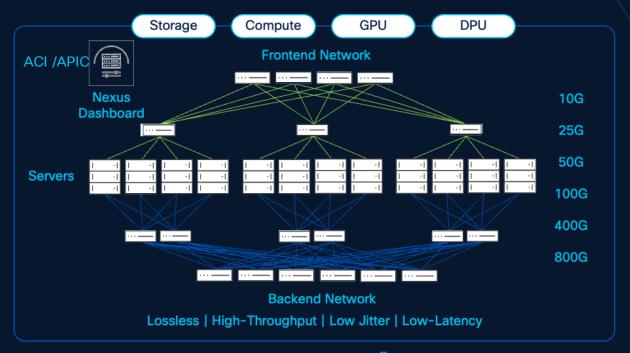
Non-blocking network Low latency, congestion management AFD.DLB, ECMP static pinning, per packet LB

Al fabric templates, Al analytics, telemetry, congestion scores

Validated designs for networks and ecosystem partners AI/ML blueprint

PFC: Priority Flow Control AFD: Approximate Fair Drop DLB: Dynamic Load Balancing

ECN: Explicit Congestion Notification







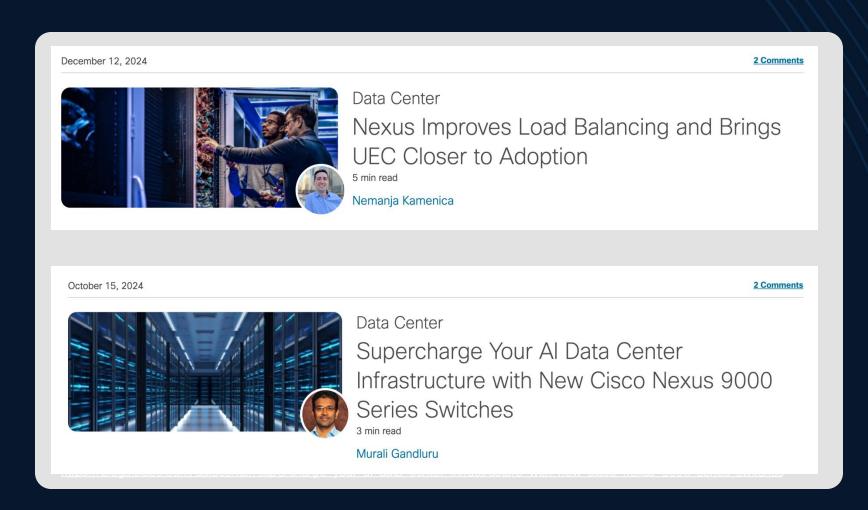




Nexus 9000 Series Switches are UEC-ready

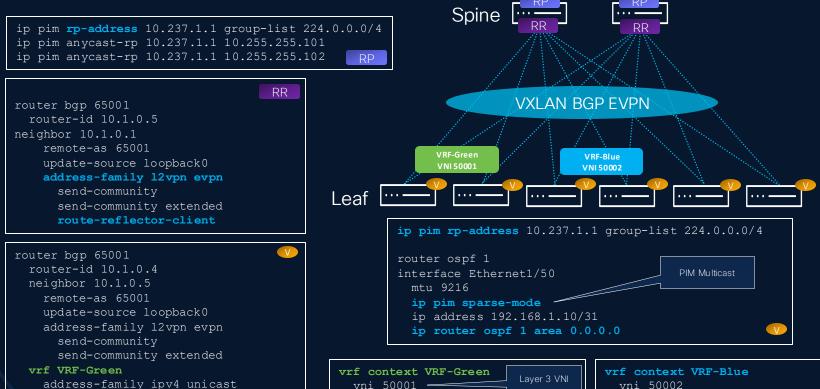
Cisco Founding and Steering Member







Network Configuration Snippet



```
vrf context VRF-Green
vni 50001
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
evpn
vni 3010 12
rd auto
route-target both auto
```

```
vrf context VRF-Blue
  vni 50002
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  evpn
  vni 3011 12
    rd auto
    route-target both auto
```

```
vn-segment 3010
                                  Layer 2 VNI
Vlan 11
vn-segment 3011
Vlan 1000
Layer 3 VNI
vn-segment 50001
Vlan 2000
Laver 3 VNI
vn-segment 50002
interface Vlan10
  no shut.down
  vrf member VRF-Green
  ip address 192.168.10.254/24 tag 12345
  ipv6 address 2001::1/64 tag 12345
  fabric forwarding mode anycast-gateway
interface Vlan11
  no shutdown
  vrf member VRF-Blue
  ip address 192.168.11.254/24 tag 12345
  ipv6 address 2002::1/64 tag 12345
  fabric forwarding mode anycast-gateway
interface nvel
  source-interface loopback0
  host-reachability protocol bgp
  member vni 3010 _____
                                 Map L2VNI to
                                    NVE
    mcast-group 239.1.1.1
  member vni 3011
                                   to NVE
    mcast-group 239.1.1.1
  member vni 50001 associate-vrf
  member vni 50002 associate-vrf
fabric forwarding anycast-gateway-mac 0002.0002.0002
```



advertise 12vpn evpn

advertise 12vpn evpn

advertise 12vpn evpn

advertise 12vpn evpn

vrf VRF-Blue

address-family ipv6 unicast

address-family ipv4 unicast

address-family ipv6 unicast

Cisco Data Center Networking Portfolio





Nexus Dashboard

On-Premises Delivered





Powered by Nexus 9000 Series





Nexus Hyperfabric

Cloud Delivered



Powered by Cisco 6000 Series

Delivering networking at speed from the cloud



Cisco Nexus HyperFabric

Al Cluster

in partnership with NVIDIA

Democratize Al Infrastructure

Visibility into full stack Al

Unified stack Including NVAIE

Al-native operational model

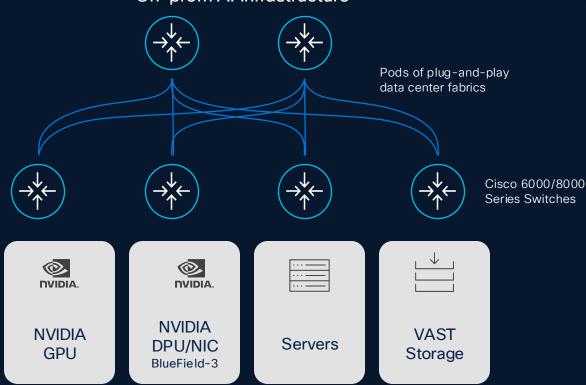
High-performance Ethernet

Cloud managed operations

A solution that will enable you to spend time on AI innovation—not on IT.

Cisco Nexus HyperFabric

On-prem Al Infrastructure



Built on Cisco Silicon One and Optics innovations





AI/ML: Visibility via Nexus Dashboard

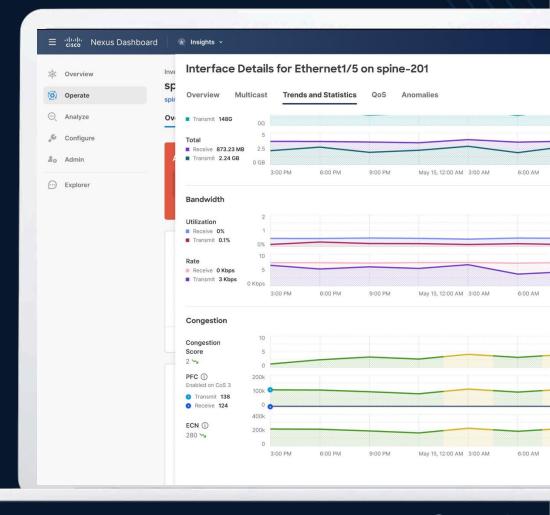
Power of telemetry and insights to optimize AI/ML fabrics

Fabric performance

Lossless Ethernet

Congestion management & scoring

Real-time troubleshooting





Cisco Expands Partnership with NVIDIA to Accelerate Al Adoption in the Enterprise



https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2025/m02/cisco-expands-partnership-with-nvidia-to-accelerate-ai-adoption-in-the-enterprise.html

Cisco Silicon (S1 and Cloudscale) & Cisco Optics fully validated as part of Spectrum-X

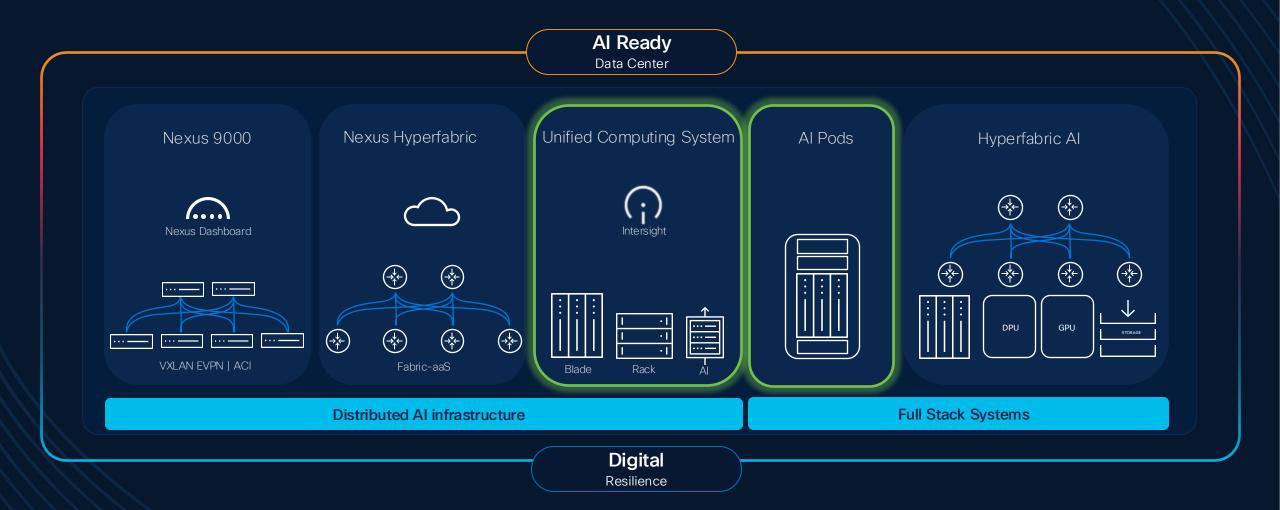
The only 3rd party switch validated as part of the reference architecture

Support packet spraying and adaptive routing capabilities in Cisco Silicon Nexus Switches

Co-develop future Nexus based switches based on Spectrum-X ASIC

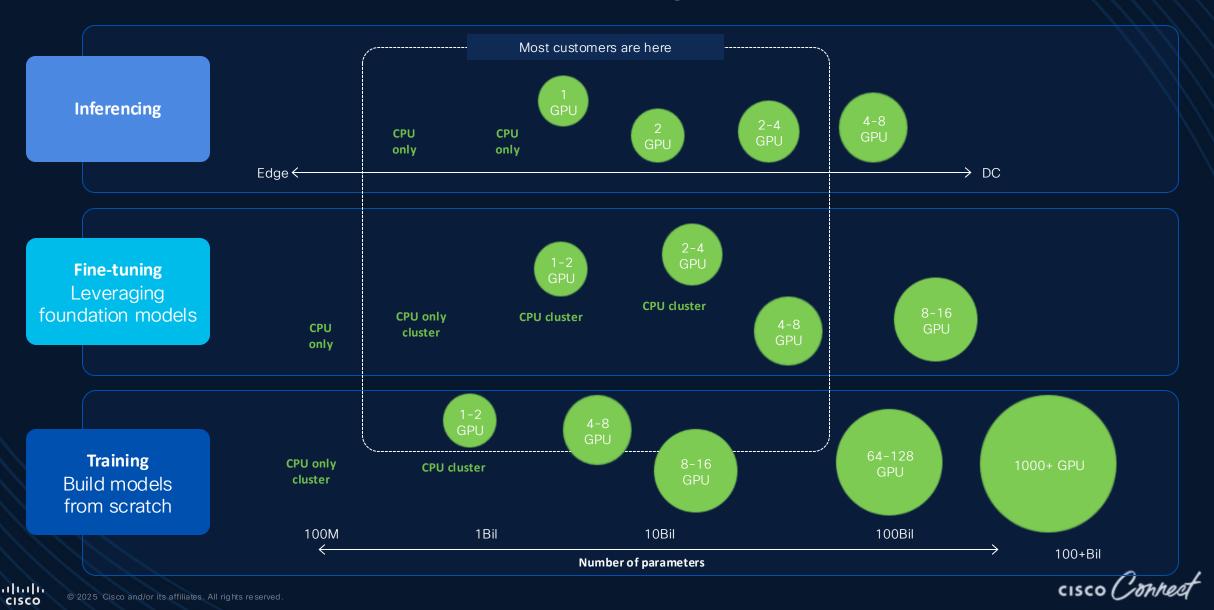


Cisco Al-Ready Data Center Building Blocks





Enterprise Al Infrastructure Mapping



Cisco Unified Computing System

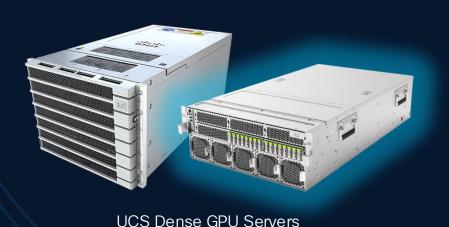
Address AI workloads with visibility, consistency, and control

Validated solutions for Al with compute, network, storage, and software

BUILD THE MODEL
Training

OPTIMIZE THE MODELFine-tuning and RAG

USE THE MODEL
Inferencing





To Come!

UCS Modular (w/GPU Expansion) and Rack

Enterprise Al Edge

Dense compute for demanding Al

Full stack AI with compute and networking

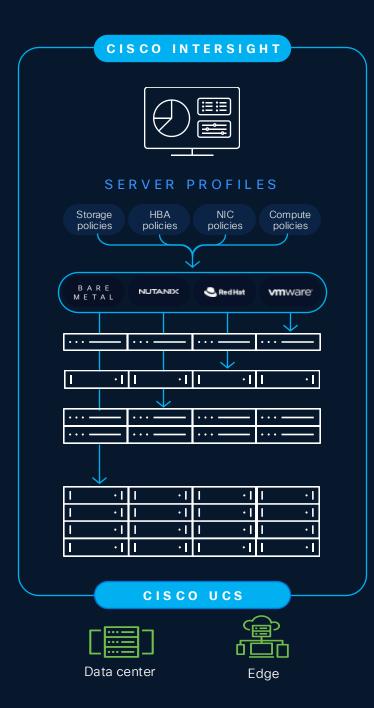




Cisco UCS

Programmable, Flexible & Unified Computing

Delivering industry differentiation with Cisco's SaaS-managed, fabric-based data center



WHY CUSTOMERS CARE

Programmable

Shape infrastructure to any workload, depending on business needs

Flexible

Adapt to the future with unparalleled simplicity, performance and resiliency

Unified

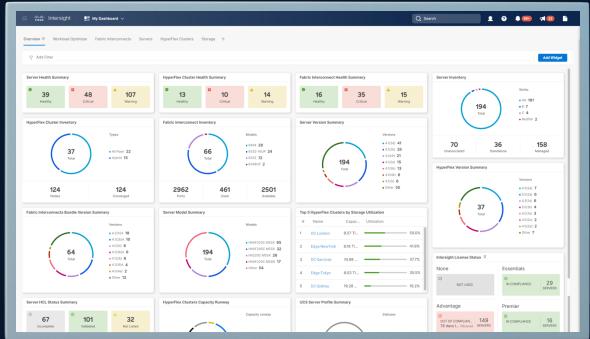
Fabric-based architecture with unified fabric connectivity





Cisco Intersight

Unified Experience for Infrastructure Operations





See your global on-premises, cloud, and edge environments.

Connect your infrastructure operations across compute, storage and networking.

Secure operations with built-in advisories and continuous risk mitigation.

Automate workflows, APIs and SDKs to automate day 0 to day N operations





Cisco GPU-Accelerated UCS Platforms



C-Series Rack Servers

Up to 8x L4 GPUs or 3x H-Series GPUs per 2U Rack Server



X-Series Modular Servers

Up to 24x L4 GPUs or 8x H-Series GPUs per Chassis



Dense GPU Servers

Up to 8x H-Series GPUs per Server



intel

AMDA

L4

L40S

H100

H200

FLEX140

FLEX170

MI210

MI300X

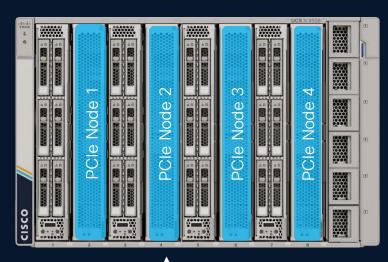
X-Series Modular Architecture with X-Fabric Technology

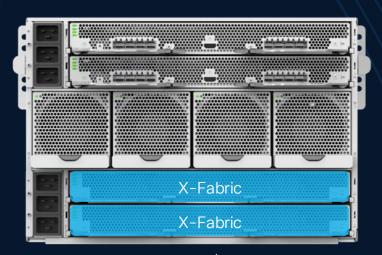
Modular and easily upgradeable solution (no backplane or cables)

Increase operational efficiency and reliability through shared power, cooling, and management

Standardize on a single system to support a wider range of applications, reducing support costs

Future-ready for new X-Fabric devices as they are made available





UCS X-Fabric Technology

Up to 24x L4 GPUs or 8x H-Series GPUs per Chassis





Cisco UCS C885A High Density Rack Server

Bringing high-density GPU servers to the Cisco UCS family and to Cisco's Al solution portfolio

Purpose-built for compute and data-intensive use cases

Training—large/small models
Fine-tuning
Large model inferencing
Retrieval-augmented Generation (RAG)

- Service Providers
- Financial Services
- Manufacturing

- Healthcare and Life Sciences
- Automotive





Cisco UCS® C885A M8 Rack Server

NVIDIA HGX with 8 NVIDIA H100/H200 GPUs or

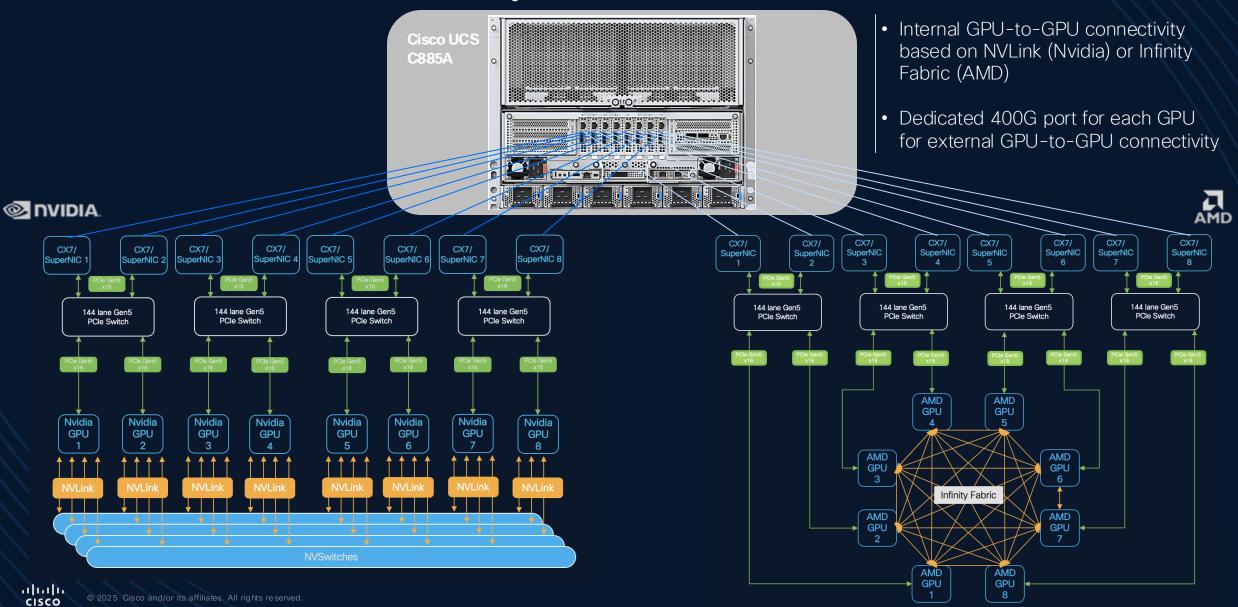
8 AMD MI300X GPUs

2 AMD 4th or 5th Gen EPYC[™] Processors

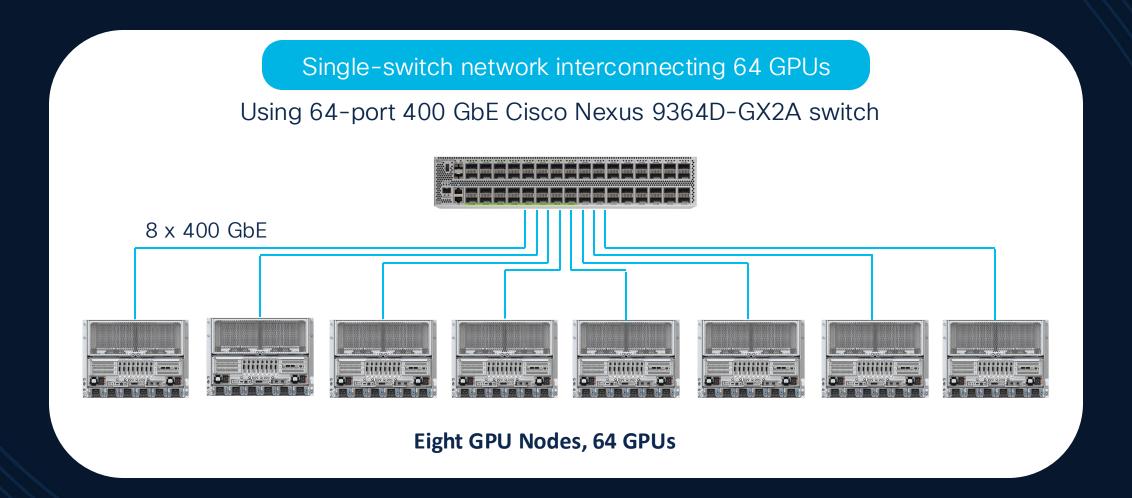




GPU-to-GPU Connectivity



Designing a Smaller Inter-GPU Backend Network



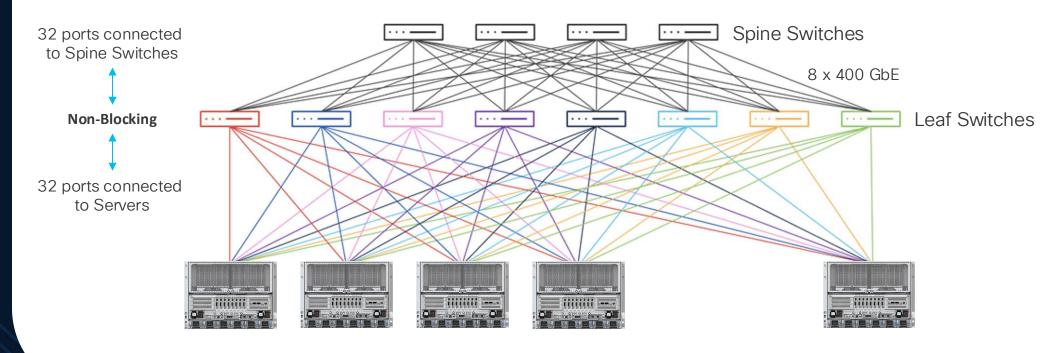




Designing a Larger Inter-GPU Backend Network

Rail-optimized network interconnecting 256 GPUs

Using 64-port 400 GbE Cisco Nexus 9364D-GX2A switch



Port 1 on all nodes connect to Leaf 1, Port 2 on all nodes connect to Leaf 2, and so on





Bringing flexibility and scalability to your Al workloads

Versatile, accelerated, 4RU rack server for broad set of Al use cases

AVAILABLE NOW





Cisco UCS® C845A M8 Rack Server

NVIDIA MGX platform with 2/4/6/8 NVIDIA H100 NVL/H200 NVL/L40S GPUs 2 AMD 5th Gen EPYC processors

Scalable AI performance

- Choose the number of GPUs that match your use case
- Increase the number of GPUs when your workloads demand it
- Address a variety of Al use cases

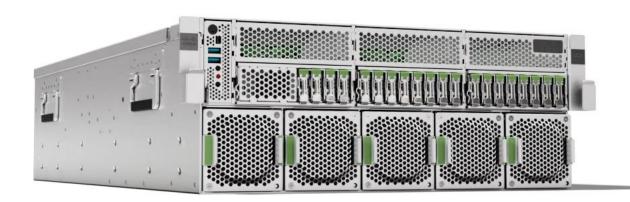
Consistent management

- Have unified management of all your Cisco UCS® servers through Cisco Intersight®
- Manage your Al servers with the same tool as your traditional servers

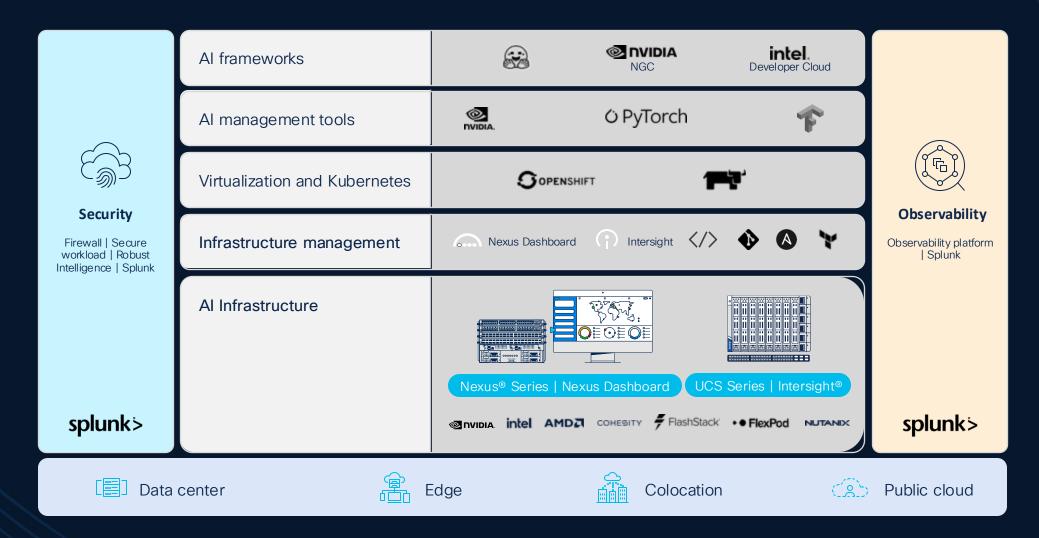




Introducing Cisco UCS C845A M8



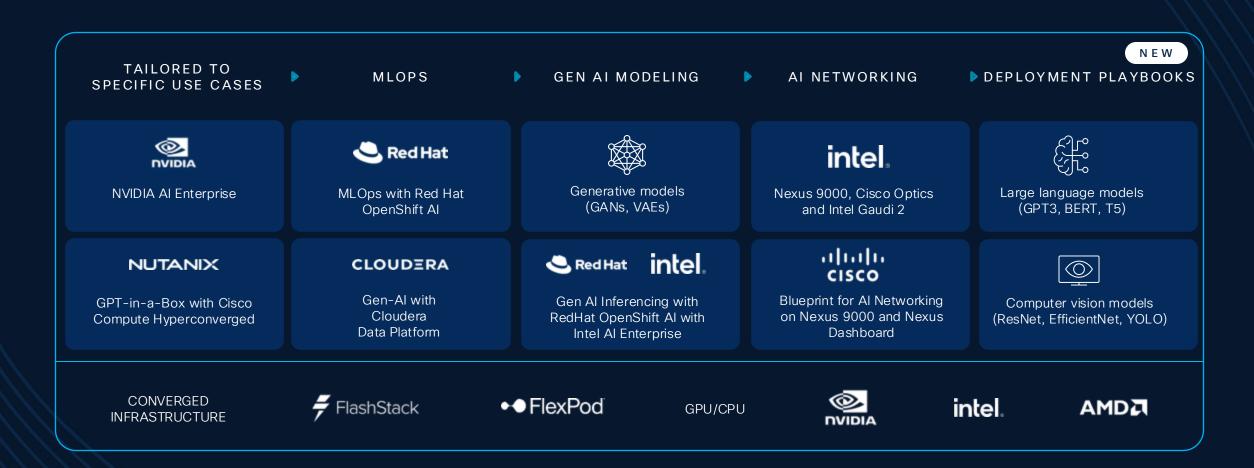
Cisco Al Ecosystem Stack





Solutions to Simplify and Automate Al Infrastructure

Cisco Validated Designs for the Data Center





Al PODs

Deploying AI with confidence

- Confidently deploy Al-ready infrastructure with pre-designed full stack architecture bundles for targeted Al use cases.
- Leverage automation frameworks for rapid deployment and adoption of infrastructure.
- Operate with best-in-class single-support model for your Al deployment architecture, include enterprise support for select Operations Support System (OSS) tools and libraries.

Al Model

Al Tooling

Containers

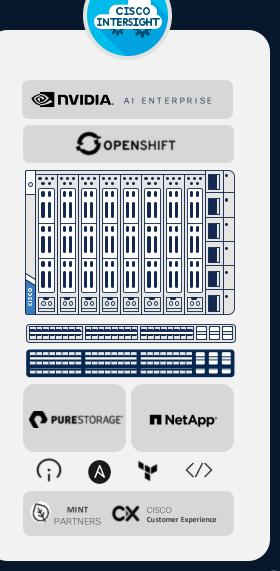
Accelerated Compute

Networking

Converged Infrastructure

Management & Automation

Adoption & Support Services





Cisco Al PODs for Inferencing

Edge Inferencing RAG Augmented Inferencing Large-Scale TYPICAL **RAG Augmented Inferencing** (7B-13B Parameter) (13B-40B+ Parameter) USE CASE Medium Small Large 1x X210C compute node 2x X210C compute nodes 2x X210C compute nodes 2x Intel 5th Gen 6548Y+ 4x Intel 5th Gen 6548Y+ 4x Intel 5th Gen 6548Y+ Hardware 1 TB System Memory 1 TB System Memory 512 GB System Memory **Specification** 5x 1.6 TB NVMe drives 10x 1.6 TB NVMe drives 10x 1.6 TB NVMe drives 1x X440p PCle 2x X440p PCle 2x X440p PCIe 4x X440p PCIe 1x NVIDIA L40S 4x NVIDIA L40S 4x NVIDIA H100 NVL X-Series FI9108 100G 2x Fabric Interconnect 2x Fabric Interconnect 6536 100G 6536 100G 6536 100G Software Cisco Intersight Cisco Intersight Cisco Intersight Cisco Intersight **Essentials** Essentials Essentials Essentials specification **Nvidia Al Enterprise** Essentials **Nvidia Al Enterprise** Essentials **Nvidia AI Enterprise** Essentials OpenShift OpenShift OpenShift OpenShift **Default** OpenShift Container Platform OpenShift Container Platform OpenShift Container Platform **Components** X210c Control Plane Cluster Single-Node Controller X210c Control Plane Cluster

Add-On









Scale-Out Inferencing Cluster

(Inferencing Multiple Models)

Scale-Out

4x X210C compute nodes

8x Intel 5th Gen 6548Y+

4 TB System Memory

20x 1.6 TB NVMe drives

8x NVIDIA L40S

2x Fabric Interconnect

Nvidia AI Enterprise Essentials

- OpenShift Container Platform
- X210c Control Plane Cluster



CI Storage





CI Storage





CI Storage





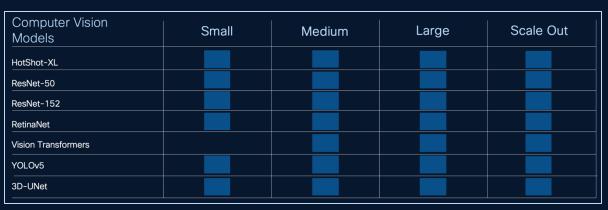


Al PODs Compatibility Matrix

NLP Models	Small	Medium	Large	Scale Out
BERT Base				
BERT Medium				
BERT Large				
BLOOM				
DLRM				
Falcon-7B				
Falcon-40B				
GPT-2B				
GPT-J				
GPT-NeoX-20B				
Galactica-Large				

Text-to-Image Models	Small	Medium	Large	Scale Out
Dreamlike Diffusion 1.0				
Openjourney				
Stable Diffusion				
Multi-Task Models	Small	Medium	Large	Scale Out
FLAN T5 XL				
FLAN T5 XXL				
Miscellaneous Models	Small	Medium	Large	Scale Out
Code Llama - 34B				
Defog SQLCoder				

NLP Models	Small	Medium	Large	Scale Out
Llama 2 - 7B				
Llama 2 - 13B				
Llama 3.1 - 8B				
Llama 3.1 - 70B				
Llama 3.1 - 405B				
Mistral-7B				
MPT-30B				
Nemotron-3-8B-chat-4k				
OPT-2.7B				
RoBERTa Large				







CX and One Cisco

Applying the Core CX Offers to Al-Ready Data Centers

★····

Software & Solution Adoption

Foundational

Success Tracks

Premium Support

Drive adoption and value realization for Data Center hardware and software with reactive or proactive support to reduce downtime, ensuring assets run at the speed and capacity needed for Al innovation

Better

Solution Attached Services + Splunk

Set Up Assist and Adoption

Day 0-1 implementation with targeted, ongoing Day 2 activities to help decrease migration risk and ensure the data center is prepared to process Al workloads

Best

LCS

Optimization and Transformation

Meet desired Al outcomes with consultative services and analytics that measure the environments' readiness, simplify fabric operations with automation, and ensure data center security from ground to cloud

Delivering Enhanced Value







Silicon One





Unified Compute





Unified Compute

cisco Conrect



Data Centre Networking



Data Centre Networking



Data Centre Networking



Best-in-class





Unified Compute

