Building the AI-Ready Data Center

Gerald Richard - Director, Product Management, Cisco
Swetha Velamuri - Leader, Product Management, Cisco

PSODCN-1871

Advancing our AI infrastructure

Cisco Data Center Networking

Cisco Silicon and Optics

Cisco Unified Compute

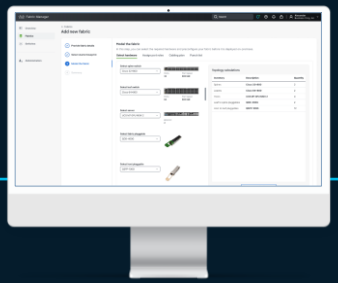

Nexus Nexus Hyperfabric

Full-stack cloud-managed solution for AI

Cisco Validated Designs for AI

Easily design, deploy, and manage AI infrastructure, including networking, optics, and compute with plug-and-play cloud-managed data center network operations.

Security

Building high-performance AI/ML Ethernet fabrics

Cisco 8000

Tier 2 web
AI aaS

Hyperscalers

Nexus 9000 with Nexus Dashboard

Enterprise
Public sector
Commercial

Service
providers

Tier 2 web
AI aaS

Nexus Hyperfabric for AI

Enterprise
Public sector
Commercial

Customizable solution
BYO management, SONiC, or BYO-NOS

Shipping

Network solution
Private cloud managed, interoperable

Shipping

SaaS solution
Public cloud managed, full stack

** H1 CY25*

Cisco's AI-ready secure data center

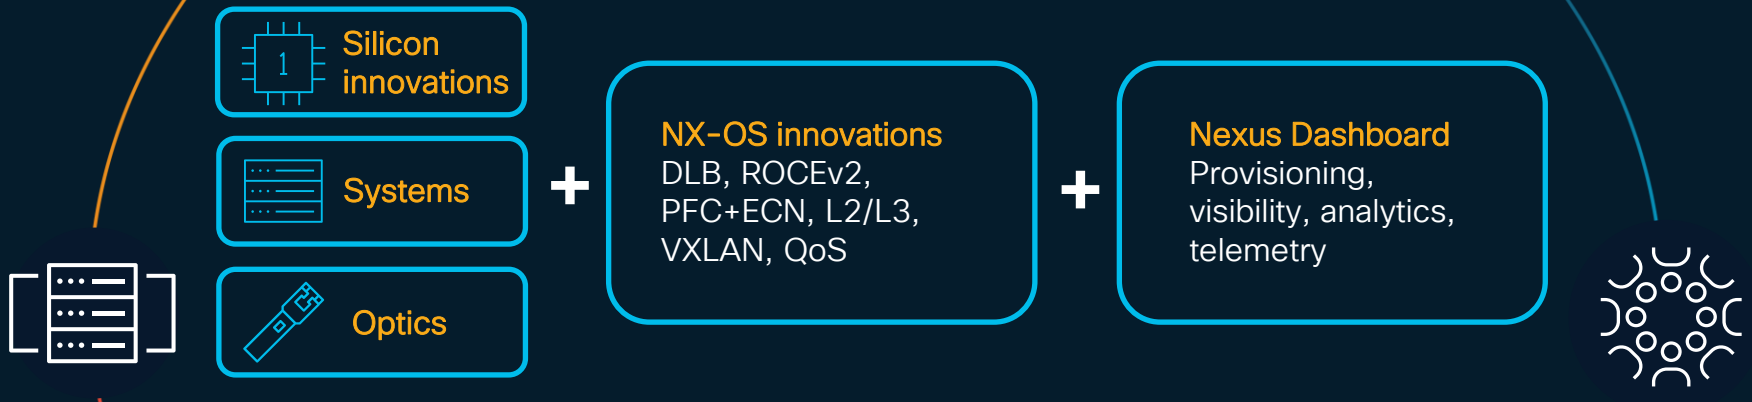

Security

Gain comprehensive threat prevention, detection, investigation, and response for organizations of any size and security maturity.

Assurance

Enable seamless end-to-end connectivity across cloud, internet, and enterprise networks to assure the delivery of applications and services.

Observability

Prevent downtime and optimize experiences with visibility and insights across end-to-end services, including owned and unowned environments.

Cisco Nexus 9300 Series 64p 800G fixed switch

Introducing the first 51.2T switch

Compact 2RU 51.2T switch

G200 ASIC (5 nm) | 100G SerDes | 256MB packet buffer

64 800G ports | Up to 128 line-rate 400G ports (2x 400G breakout)

Choice of QSFP-DD800 or OSFP ports

QSFP-DD800 ports are backward compatible with QSFP-DD, QSFP28, QSFP+

Multi-core x86 CPU | 32GB RAM | 128GB SSD

Cisco NXOS spine and AI/ML spine/leaf capable

N9364E-SG2-Q or N9364E-SG2-O

Introducing Cisco QSFP-DD800 modules

QSFP-DD 8x 100G FR

QSFP-DD 2x 400G FR4

Increased density	Double port bandwidth for single-mode fiber links up to 2 km
Investment protection	Reuse existing cabling infrastructure: dual duplex LC and dual MPO-12 SMF connectors
Backwards compatibility	Connect existing pluggable transceivers: QSFP+, QSFP28, QSFP56, QSFP112, QSFP-DD
Flexible design support	800G port to port | Breakout to 400G or 100G ports
Improved sustainability	Supports over 30W of power dissipation and riding heatsink in host platform
Standards compliant	QSFP-DD 800 MSA, IEEE 400GBASE-FR4, 100GBASE-FR1

Cisco Nexus 9300 Series 64p 800G fixed switch

Delivering sustainable operations

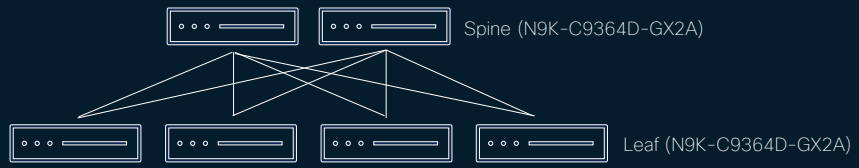

51.2 Tbps non-blocking fabric

25.6 Tbps ASICs | 50G SERDES

6x 2RU switches → 12 RU

Typical power: ~8,000W *

75% energy savings

83% rack space savings

51.2 Tbps switch

51.2 Tbps ASIC | 100G SERDES

2RU switch

Typical power: ~2,000W *

* https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-743854.html

* *https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/nexus-9364e-sg2-switch-ds.html - Typical power to be validated*

Where security meets the network
Cisco N9300 Smart Switch

Security

Cisco Hypershield

Cisco Security Cloud Control

Autonomous Segmentation | Distributed Exploit Protection | Future services

Platform

Workload and network enforcement points

Public Cloud | Private Cloud

Holistic Data Center Security

Network

Cisco N9300 Smart Switch

N9324C-SE1U

Cisco Nexus 9000 Series Switches

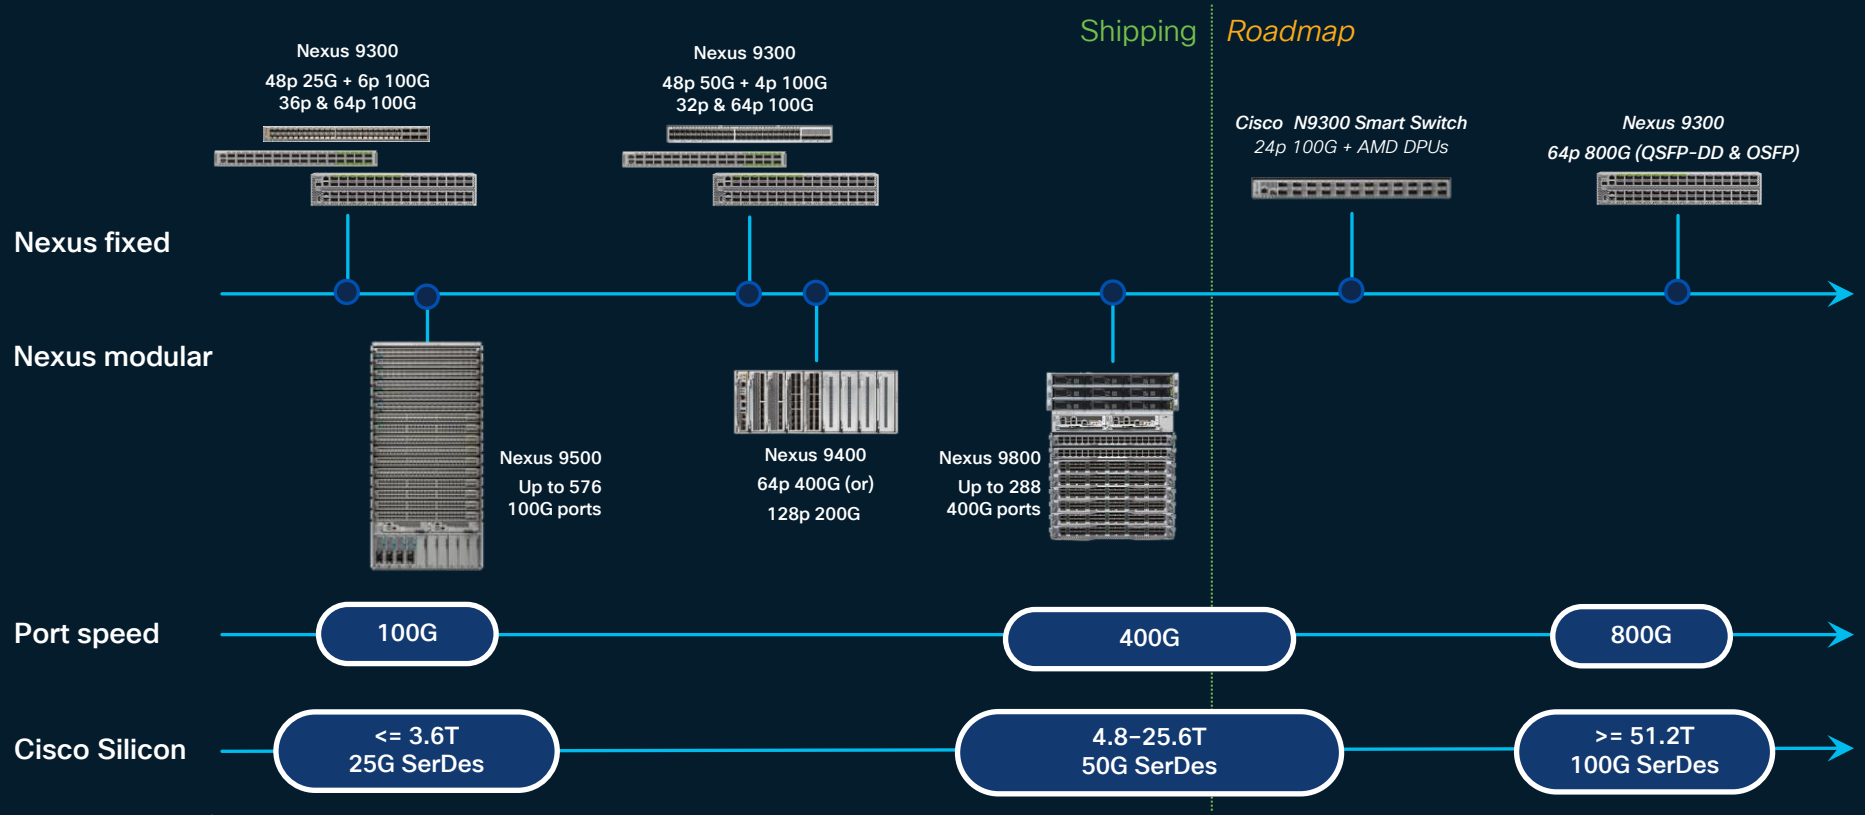

Artificial intelligence outcomes span every industry

Government

- Deliver enhanced citizen services
- Data-driven policy decisions and creation
- Modernizing and streamlining operations
- Optimizing infrastructure management
- AI-powered traffic design and public safety

Manufacturing

- Intelligent quality control
- Proactive machine maintenance
- Digital twin creation
- Supply chain optimization and tracking
- Optimizing production processes

Finance

- Predictive trading algorithms
- Fraud detection and prevention
- Personalized financial advice
- Investment portfolio optimization
- Virtual assistants and seamless transaction experiences

Healthcare

- Medical imaging analysis
- Enhance diagnosis and treatment
- Patient management through predictive analytics
- Improved access to healthcare with remote monitoring tools
- Drug research and development

Retail

- Consumer behavior analytics
- Enhanced customer experiences
- Personalized product recommendations
- Dynamic virtual shopping experience
- Demand analysis and prediction

Build the model | Training

Optimize the model | Fine-tuning and RAG

Use the model | Inferencing

Multiple networks for AI infrastructure

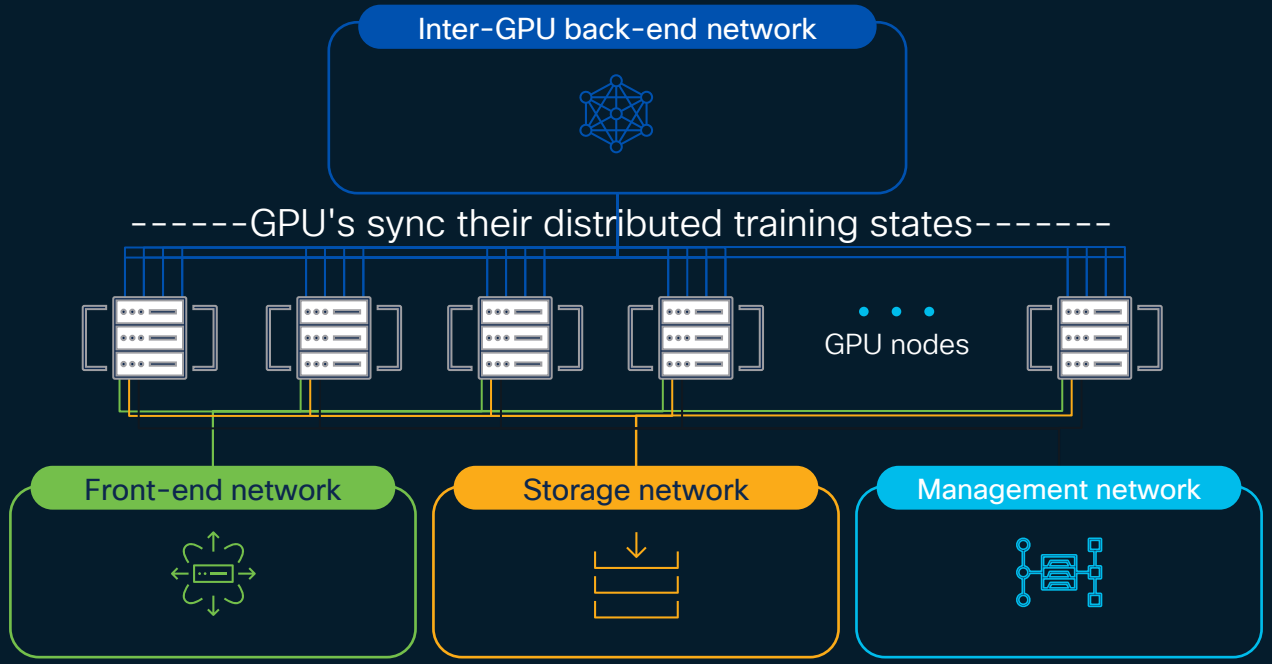

AI/ML Networks Demand More From the Network

AI workloads do not tolerate latency, so make sure your data center network delivers...

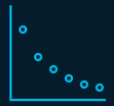

Low latency

High, lossless throughput

Greater bandwidth

Congestion management

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)
Non-blocking & Lossless network (PFC + ECN)

Advanced Congestion Aware Load Balancing

Flow-let(DLB), Per-Packet(PLB), Flow Pinning,

WCMP based LB, Policy Driven LB

AI fabric templates, AI analytics, telemetry, congestion scores

Validated designs for networks and ecosystem partners

AI/ML blueprint, Whitepapers

PFC: priority flow control

ECN: explicit congestion notification

DLB: dynamic load balancing

PLB: per packet load balancing

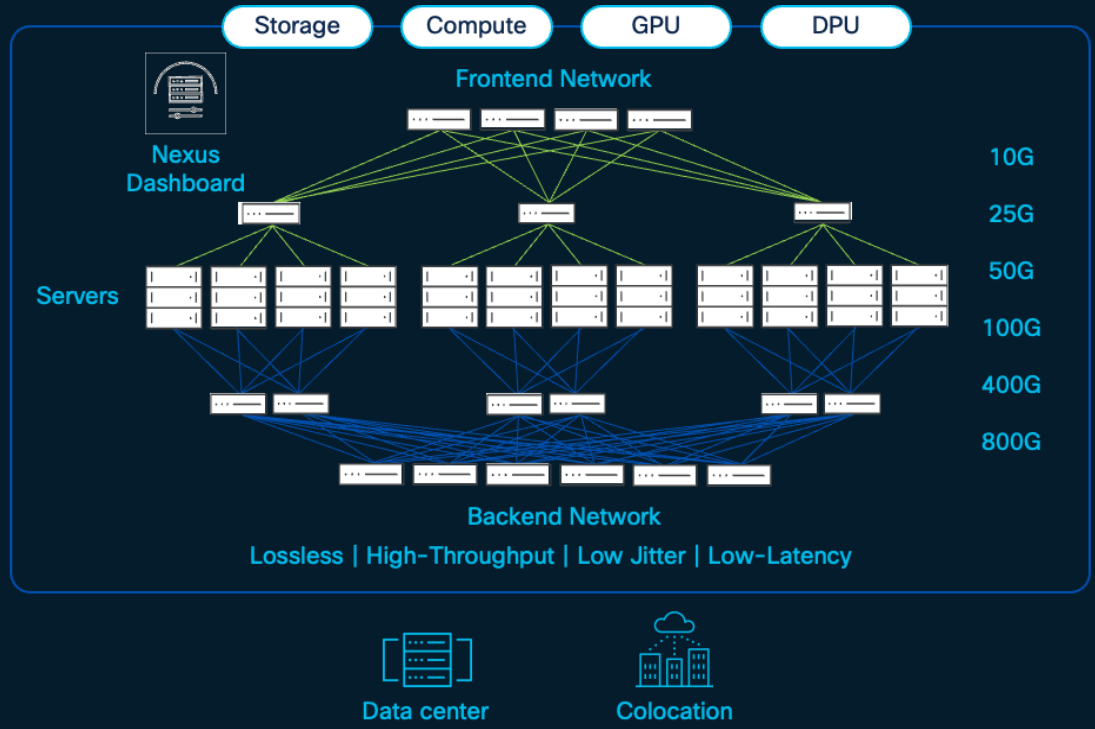

Advanced Load Balancing for AI workloads

Intelligent Load Balancing For Every Path, Every Flow

ECMP Flow Pinning

Equal Cost Multi Path Load Balancing (ECMP)

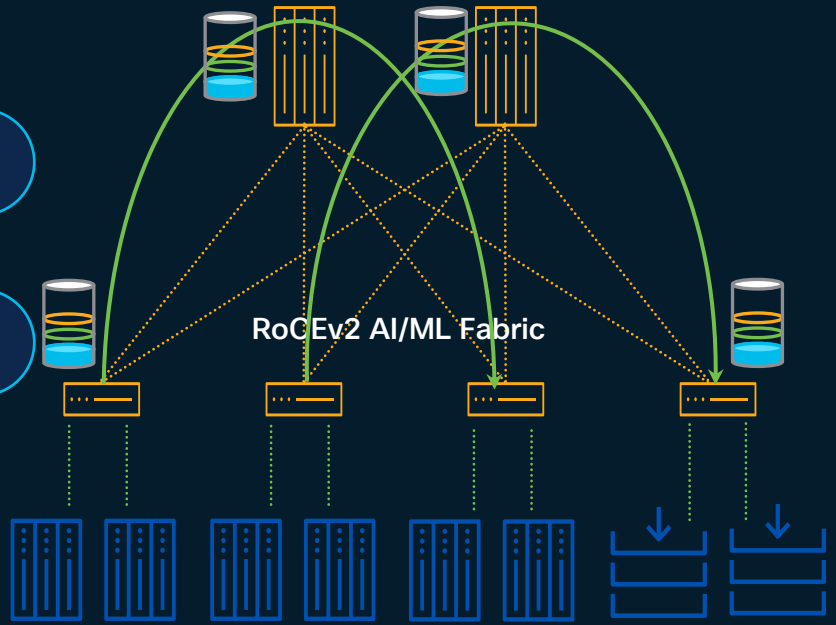

Dynamic Load Balancing (DLB)

Load Aware Congestion Aware

Per Packet Load Balancing (PLB)

Packet Spray

Weighted Cost Multi Path Load Balancing (WCMP) with DLB

Policy Based Load Balancing

*Future: UEC Features**

Load-balancing enhancements

Dynamic load balancing(DLB)

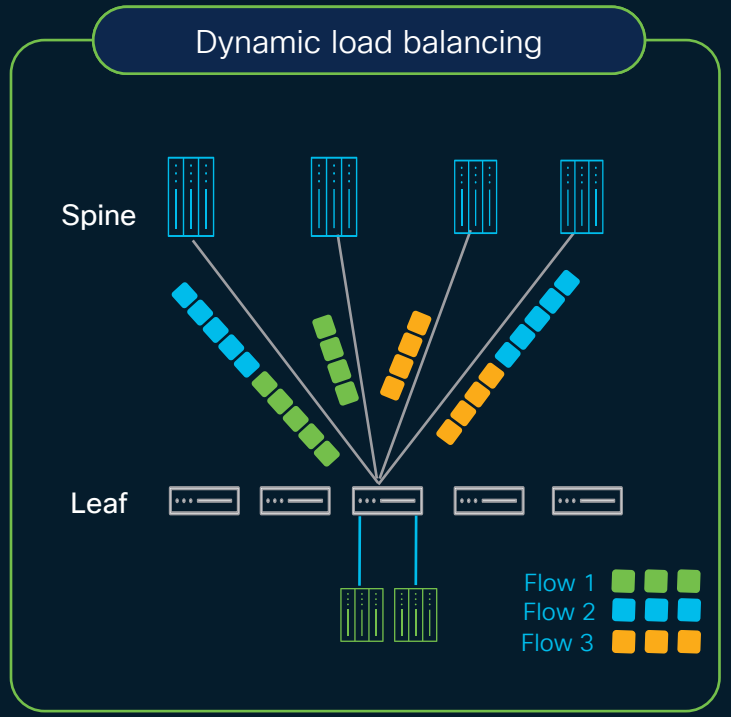

- Utilizes real-time transmission (TX) link utilization to determine the best outgoing port for network traffic, optimizing locally on the switch.
- Automatically selects the ECMP link with the lowest TX utilization for new network flows to ensure even distribution.

Key benefits:

- ✓ Maximize the utilization of available network paths
- ✓ Increase overall network performance
- ✓ Provide faster convergence and redundancy in case of link failures
- ✓ Minimize congestion by efficiently utilizing the available bandwidth across multiple paths
- ✓ Faster JCT

Load-balancing enhancements

Per-packet load balancing

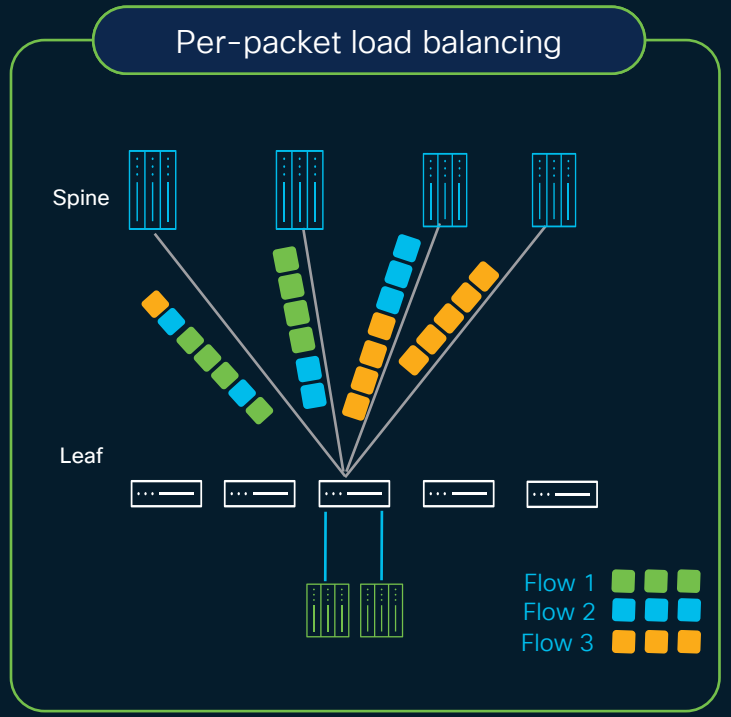

- Traffic is forwarded on a per-packet basis
- Every packet will be hashed to a different output port
- Host is expected to do packet reordering

Key benefits:

✓ Reduction in network congestion
✓ Improved resilience and fault tolerance
✓ Efficiently adapts to unpredictable traffic for optimal performance

Load-balancing enhancements

ECMP static pinning

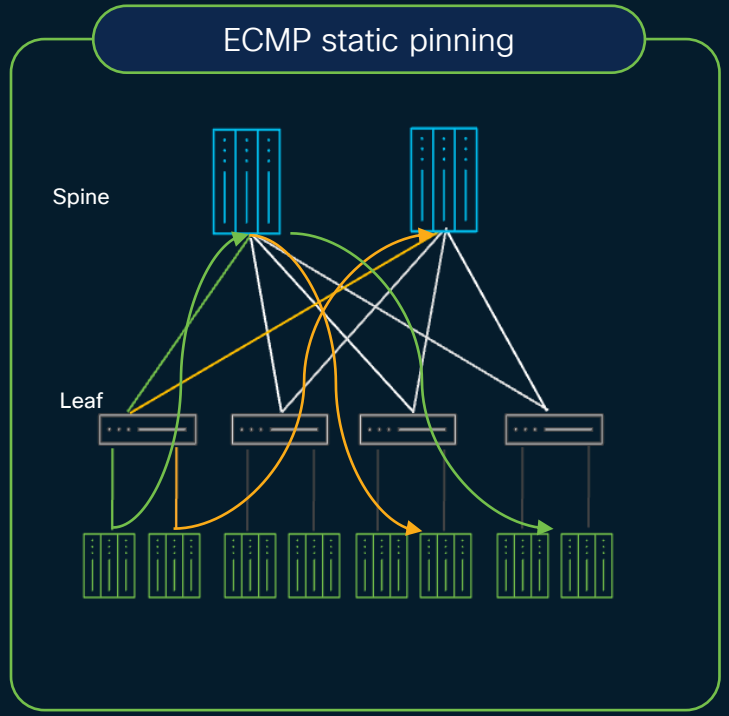

- Pins source interface to destination interface
- User controls traffic forwarding per switch
- If the pinned port is down, the next port with the least link utilization is selected

Key benefits:

✓ Predictability in traffic forwarding
✓ Elimination of over-subscription
✓ Stable network paths for latency-sensitive applications

Load-balancing enhancements

Path Based (WCMP) load balancing

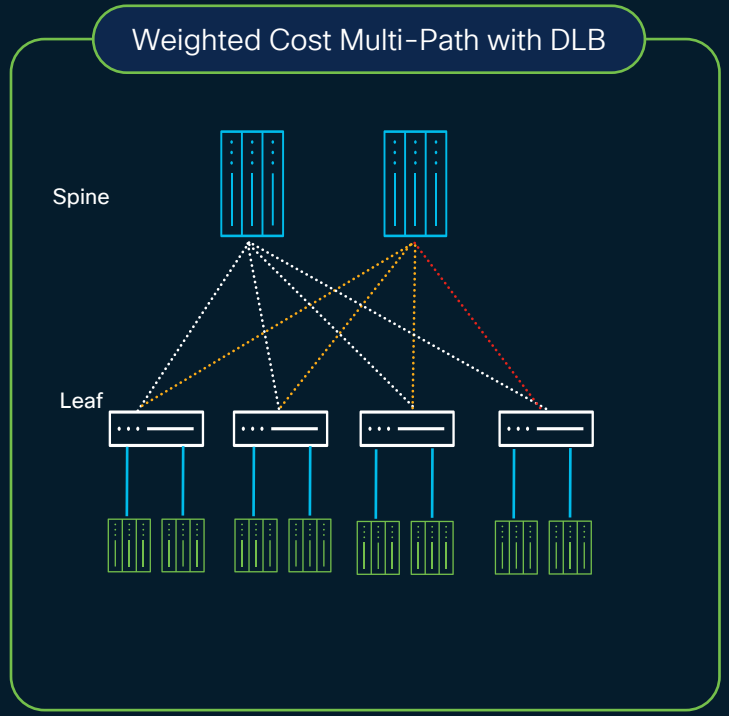

- Weighted Cost Multi-Path distributes traffic based on relative bandwidth of the links
- Uses unique weights & ranges to distribute traffic

Key Benefits

✓ Adaptive Traffic Distribution
✓ Optimized Performance & link utilization
✓ Faster Convergence & Redundancy in case of link failures

Load-balancing enhancements

Policy Based load balancing

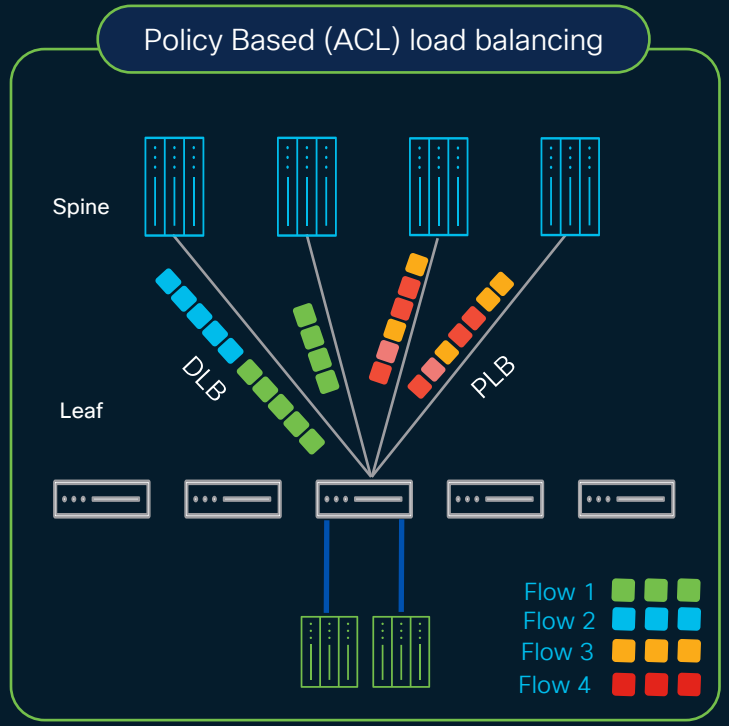

- Provides the ability to support either Flow-let based (DLB) or Packet Spray Mode (PLB) or both simultaneously for a subset of traffic based on policy.
- DSCP or ACL based mode selection

Key Benefits

- ✓ Flexible Traffic Steering
- ✓ Optimized Performance & link utilization
- ✓ Application Aware Load Balancing

Nexus Dashboard: automation

Simplifying network operations

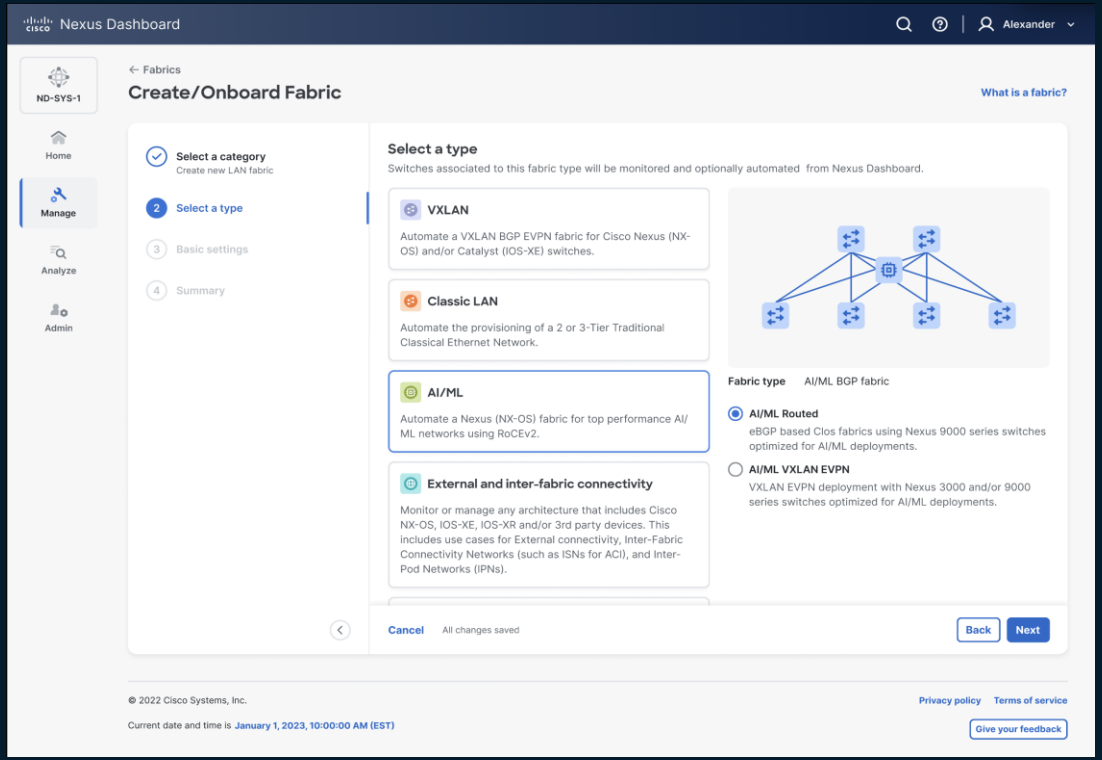

Nexus Dashboard: automation

Simplifying network operations

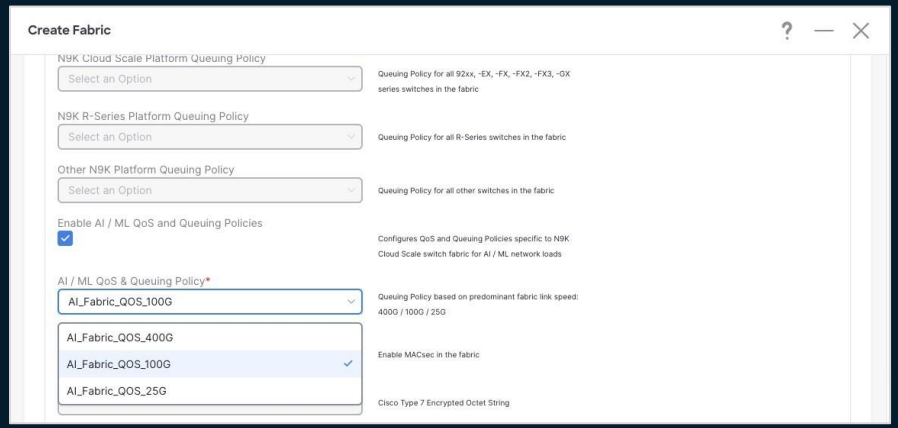

QoS and queuing policy templates for 25G,100G, and 400G fabric speeds

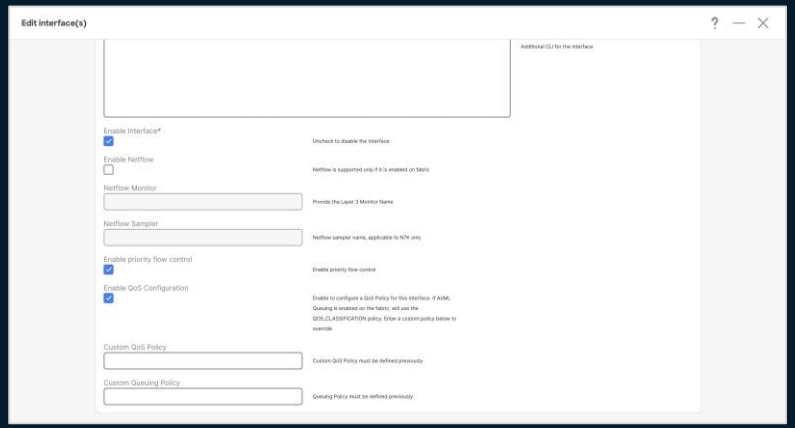

Custom QoS classification marking for RoCEv2 traffic

Nexus Dashboard: visibility & analytics

Simplifying network operations

- AI network visibility
- UX/UI dashboard
- Visibility–lossless Ethernet
- Monitoring (ECN, PFC)
- Congestion score
- Application to network performance correlation
- Telemetry and NetOps

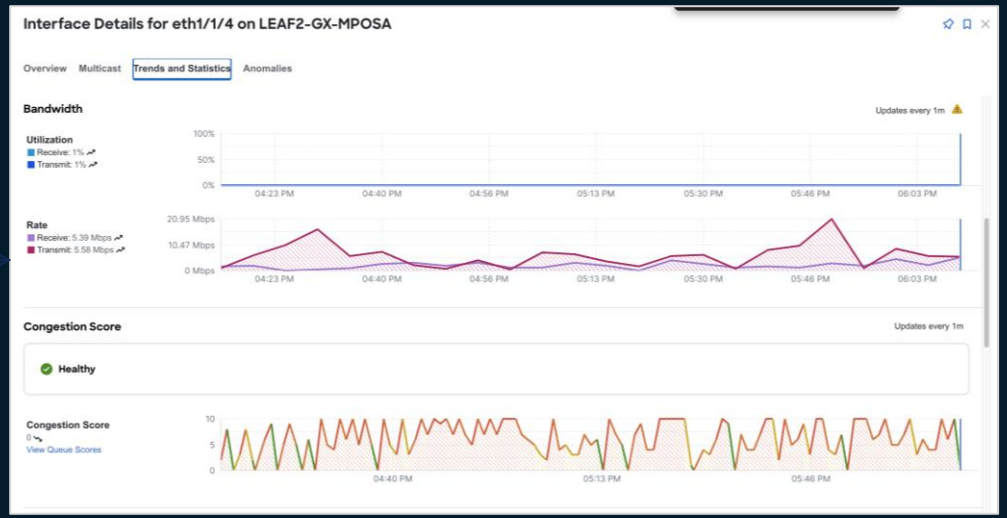

With the granular visibility provided by Cisco Nexus Dashboard Insights, the network administrator can observe drops.

Tune thresholds until congestion hot-spots clear and packet drops stop in normal traffic conditions.

This is the first and most important step to ensure that the AI/ML network will cope with regular traffic congestion occurrences effectively.

Ultra Ethernet vision

Cisco founding and steering member

Deliver an Ethernet-based open, interoperable, high-performance, full-communications stack architecture to meet the growing network demands of AI and HPC at scale.

The new era needs a new network

- UE from a switch perspective does not require Layer 1 or Layer 2 requirements
- Cisco Silicon complies with UE fabric baseline requirements (PFC, ECN, multiple traffic classes etc.)
- Cisco Silicon One is future proof because of its programmable architecture

Resources

Cisco AI/ML Blueprint

Intelligent Buffer Management on Cisco Nexus 9000 Series Switches

Cisco Data Center Networking

RoCE Storage Implementation over NX-OS VXLAN Fabrics

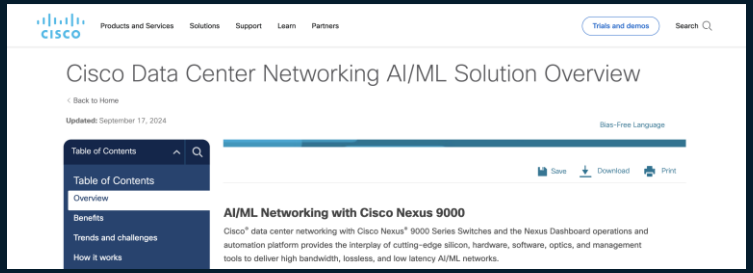

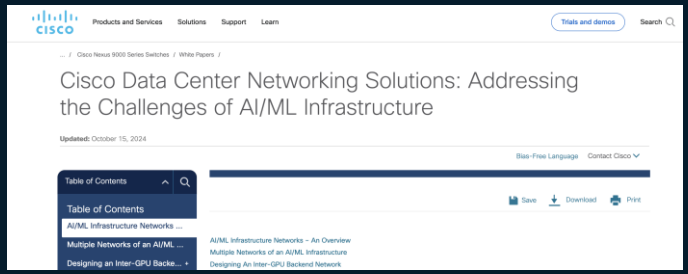

Webex App

Questions?

Use the Webex app to chat with the speaker after the session

How

1. Find this session in the Cisco Events 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 February 28, 2025.

Fill Out Your Session Surveys

Participants who fill out a minimum of 4 session surveys and the overall event survey will get a unique Cisco Live t-shirt.

(from 11:30 on Thursday, while supplies last)

All surveys can be taken in the Cisco Events mobile app or by logging in to the Session Catalog and clicking the ‘Participant Dashboard’

Content Catalog

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at ciscolive.com/on-demand. Sessions from this event will be available from March 3.


cisco
Thank you
cisco Live!

cisco Live!

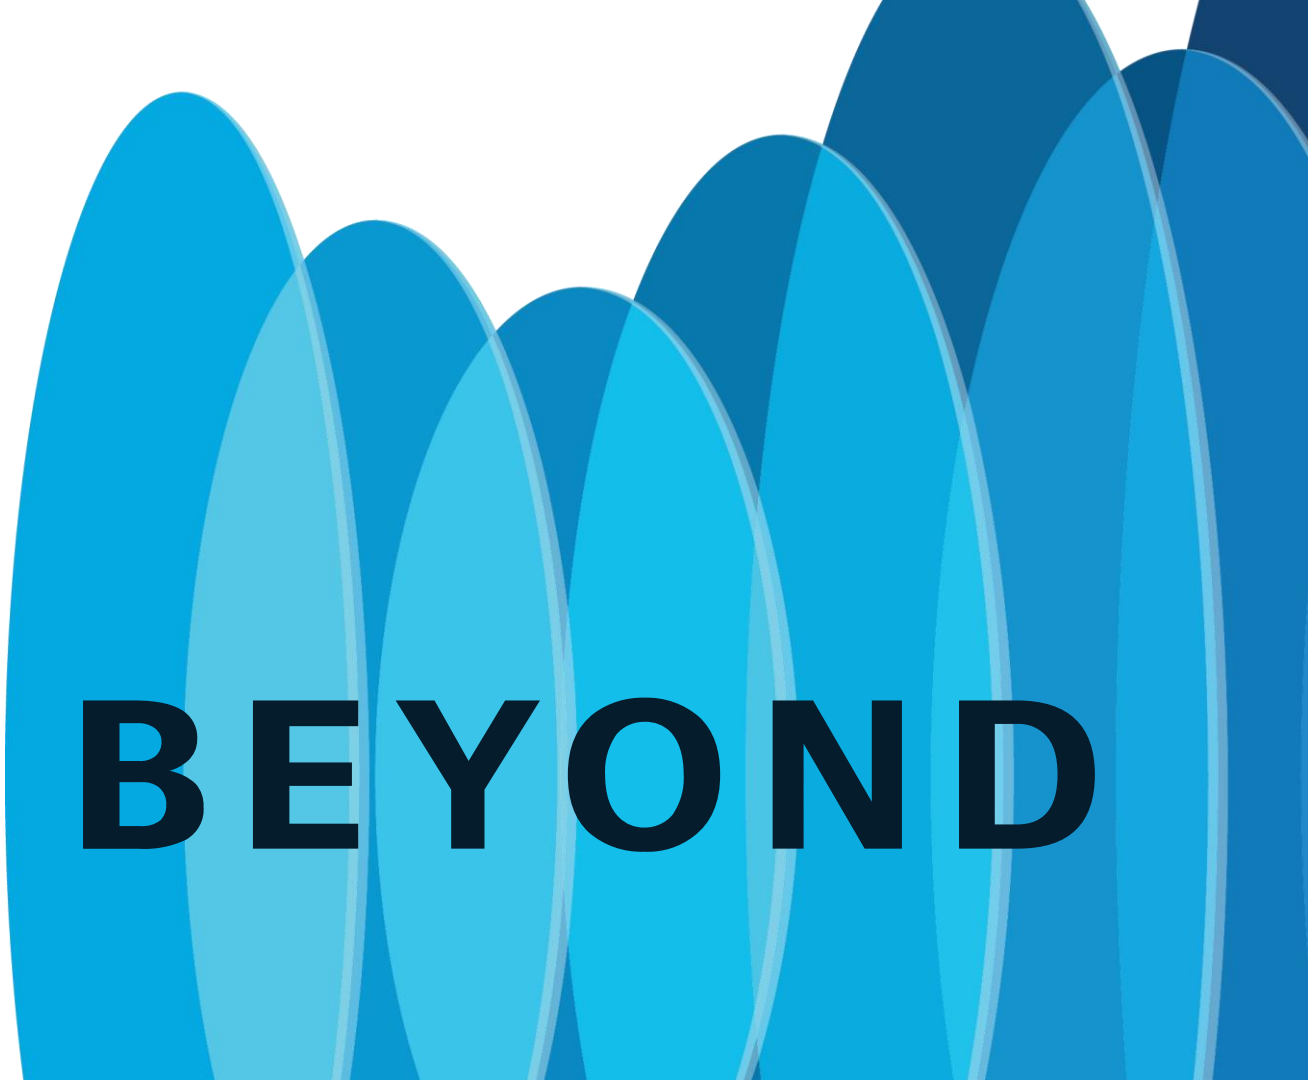

GO BEYOND