Cisco Silicon One

Redefine your AI/ML Network

Yaron Agami
Senior Manager, Product Marketing

Ramesh Sivakolundu
Director, Technical Marketing

PSODCN-1005

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 June 13, 2025.

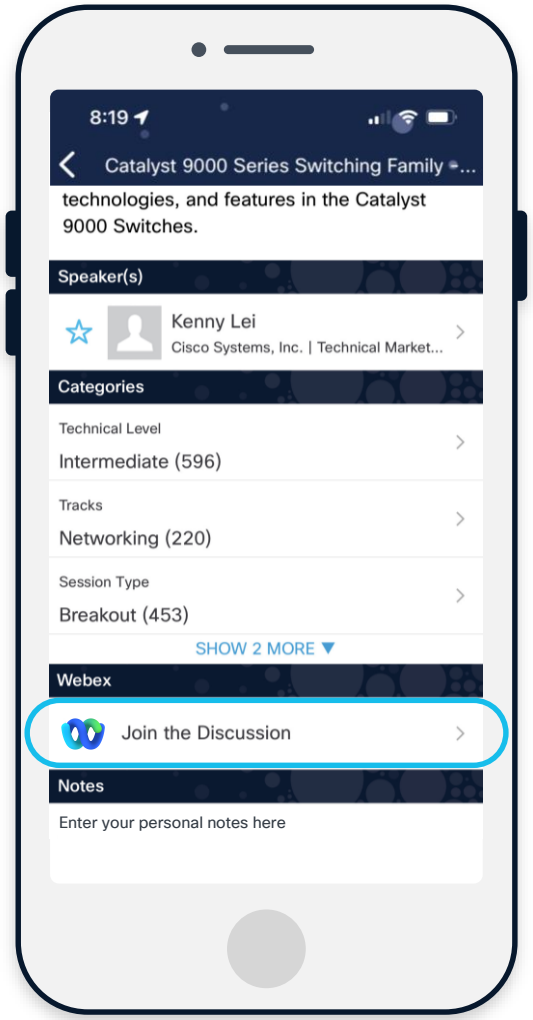

https://ciscolive.ciscoevents.com/
ciscolivebot/**#PSODCN-1005**

Agenda

01 Introduction to Cisco Silicon One

02 Networking in AI / ML clusters

03 Cisco in AI / ML clusters

04 Cisco Silicon One AI / ML differentiators

Introduction to Cisco Silicon One

Cisco Silicon One

Purpose built silicon devices
Built on a **unified** architecture

Powering **networks** that drive the future of business

To meet new **paradigm of scale**

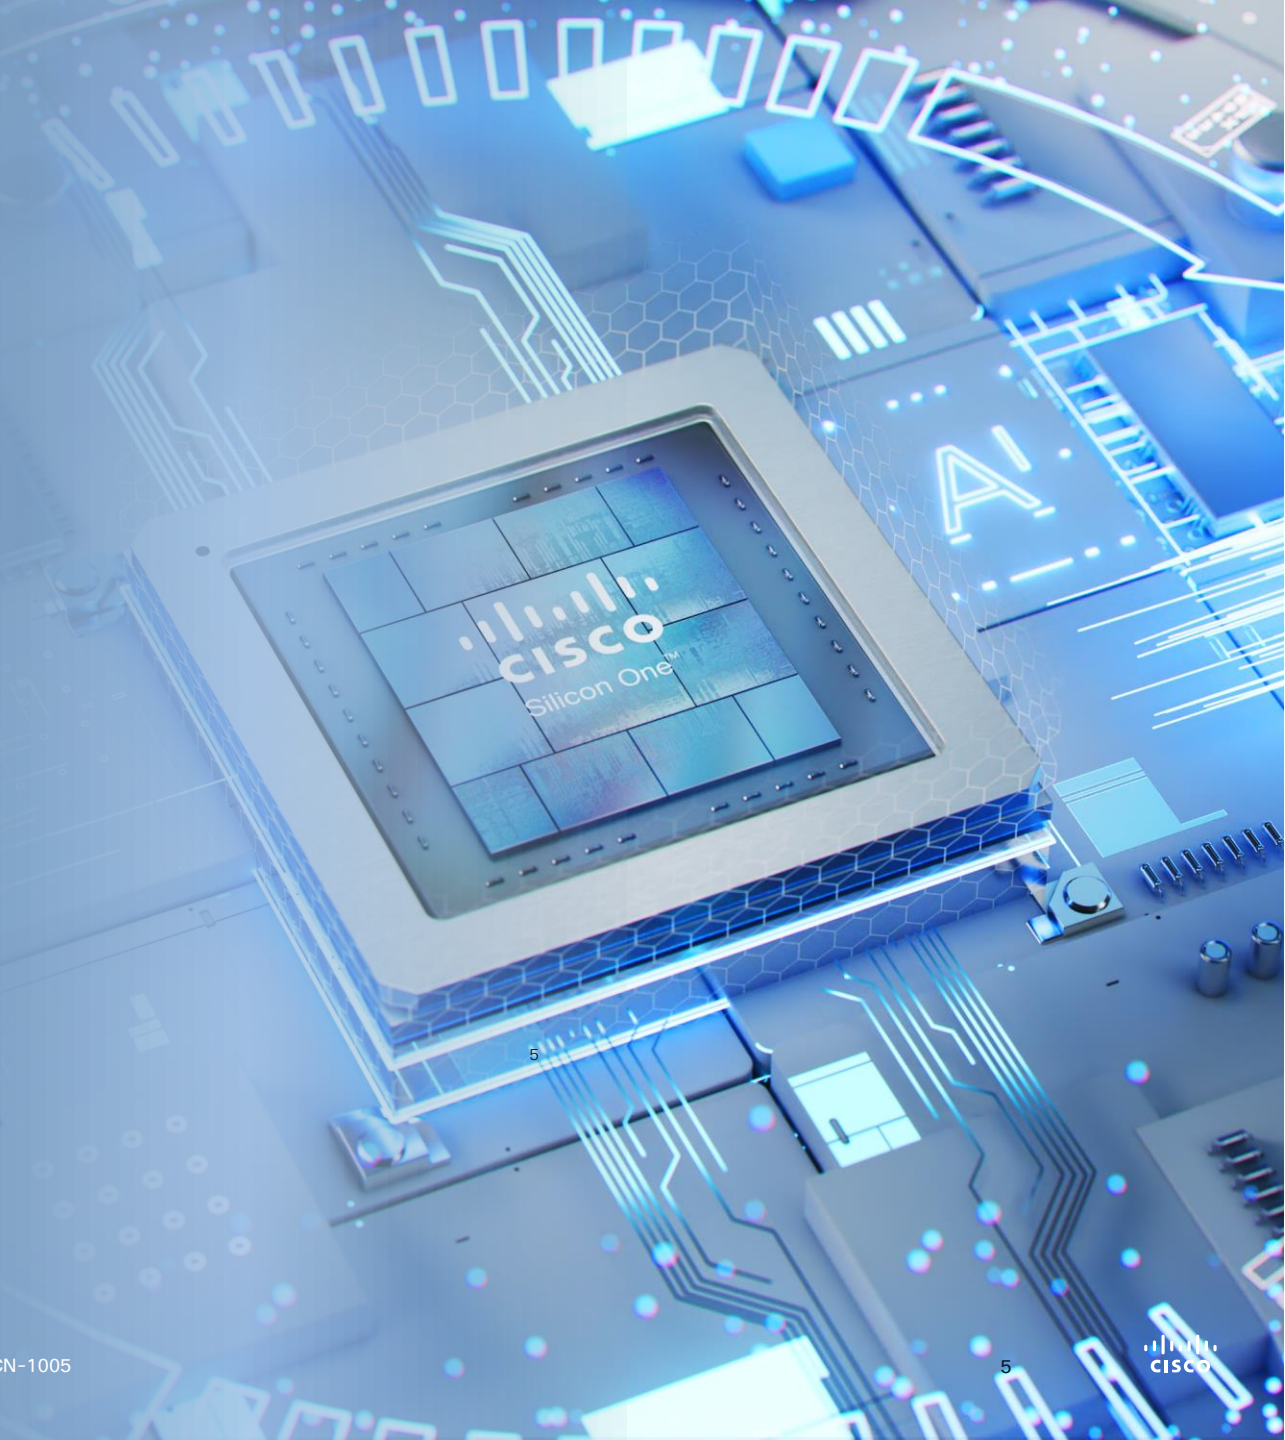

One Architecture. One Experience. Unmatched Capabilities

- 17 silicon chips were already developed
- Used in 60 Cisco platforms across all customer segments
- Powers
  - All of Cisco 8000 systems, every new Nexus and Catalyst system
  - 5 out of 6 Tier1 Web Scalers are using Cisco Silicon One, with more than 20 use cases
  - Running variety of operating systems, including SONiC OS

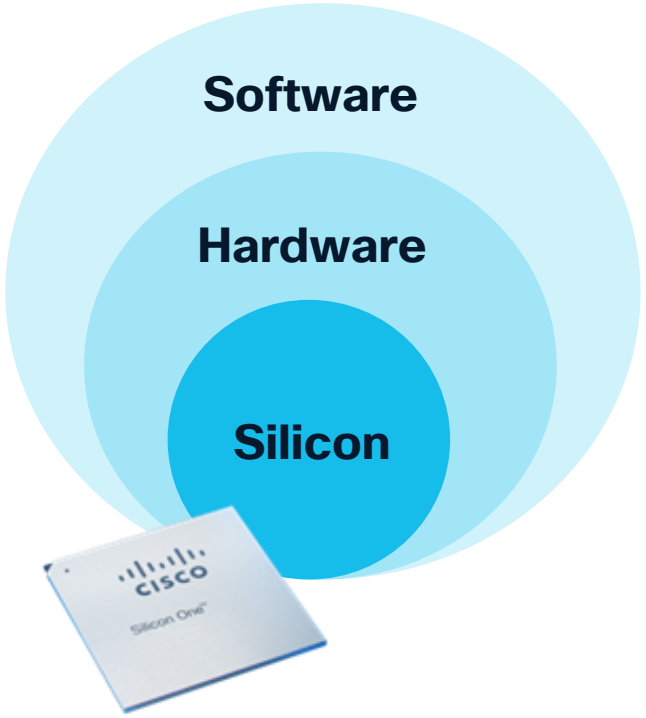

Breakthrough innovation in performance, flexibility, and sustainability

Cisco Silicon One is a unifying architecture that delivers the highest bandwidth routing and web scale switching silicon in the industry

Get performance at scale

Meet demands for scale with a choice of a high-performance, feature-rich, and low-power networking silicon—that provides a consistent experience across the entire Cisco Silicon One portfolio.

Leverage programmability

Evolve your AI/ML infrastructure with programmability while still providing low latency, cost, and power.

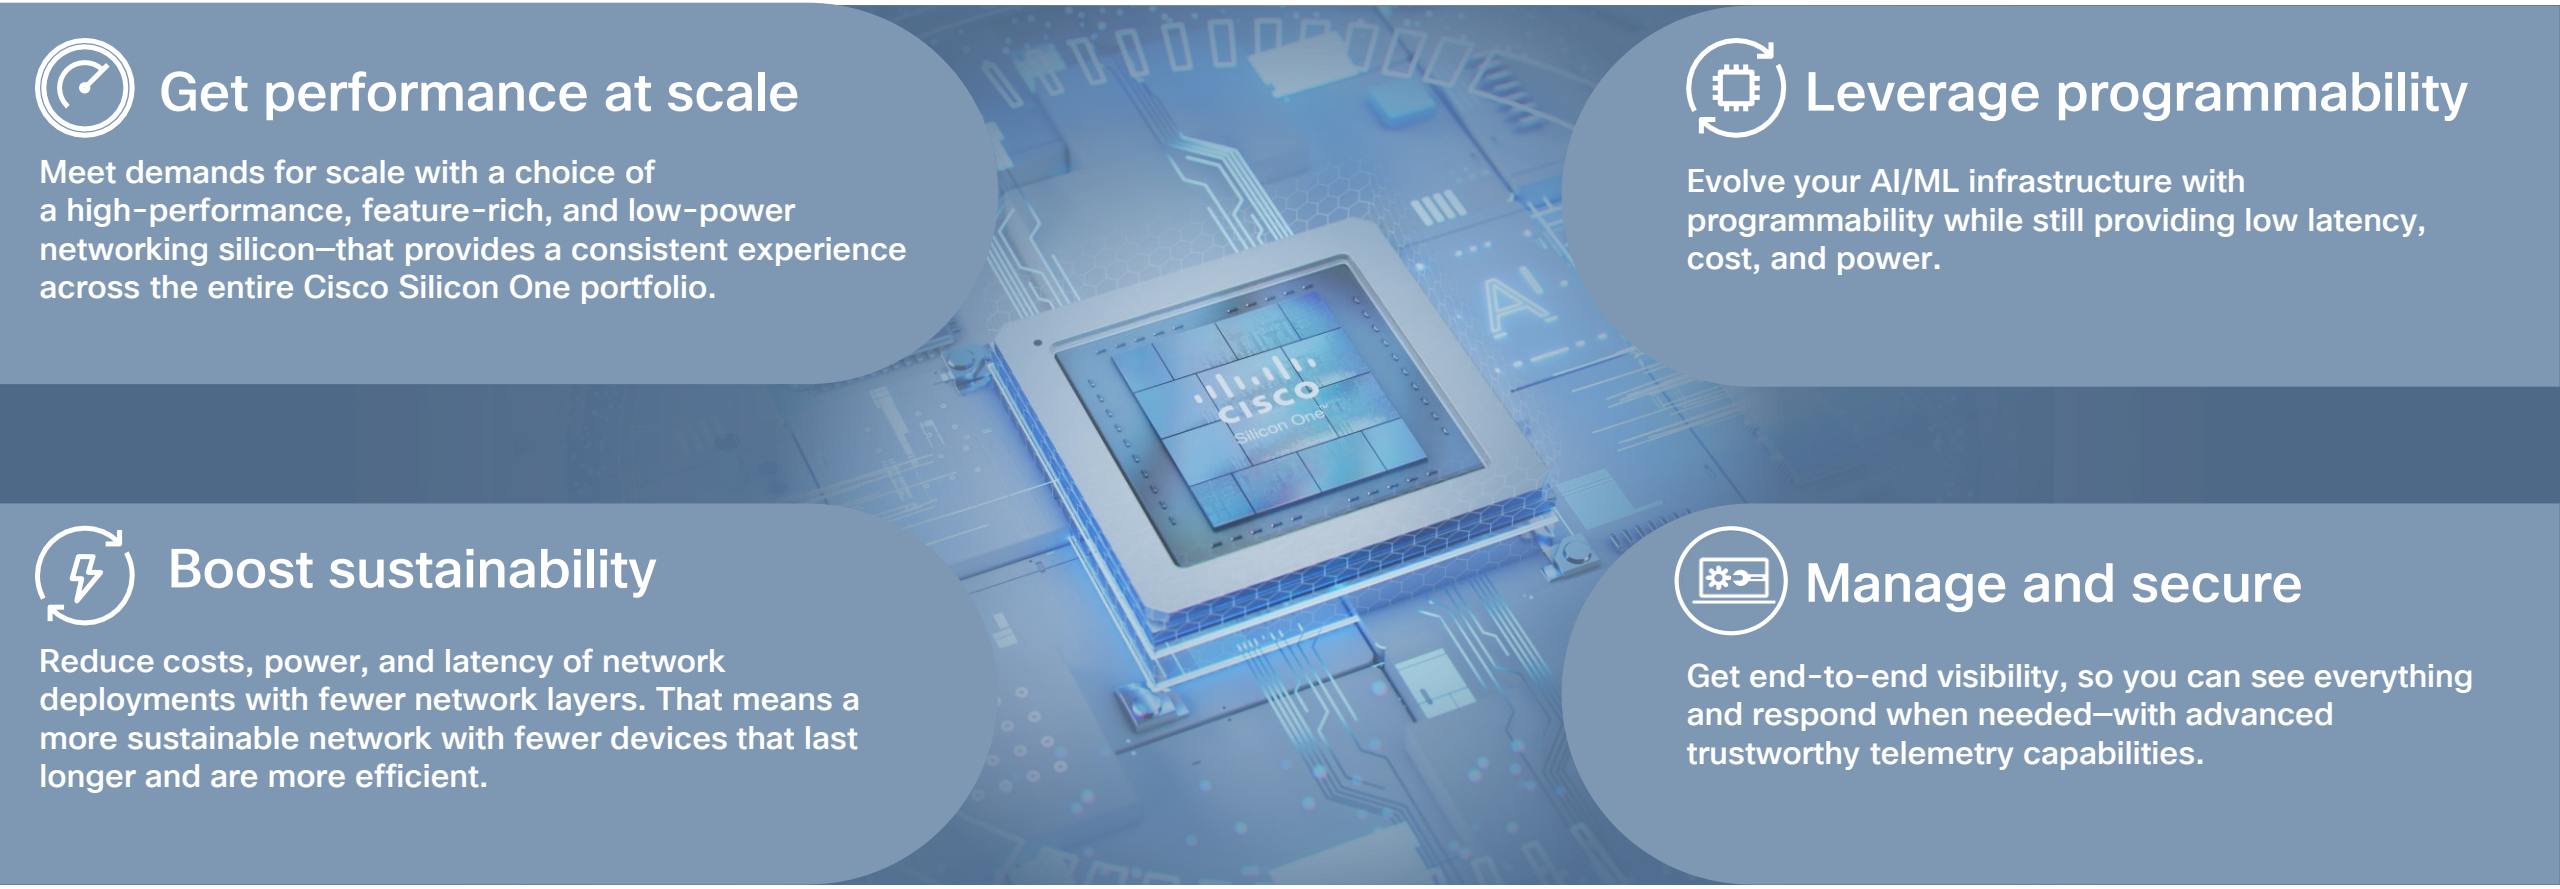

Boost sustainability

Reduce costs, power, and latency of network deployments with fewer network layers. That means a more sustainable network with fewer devices that last longer and are more efficient.

Manage and secure

Get end-to-end visibility, so you can see everything and respond when needed—with advanced trustworthy telemetry capabilities.

From top-of-rack (ToR) switches, web scale data centers and across service providers, to enterprise networks and data centers with a unified routing and switching portfolio, our devices power Cisco and third-party systems

Cisco Silicon One

One Architecture. Multiple Devices. No Compromise

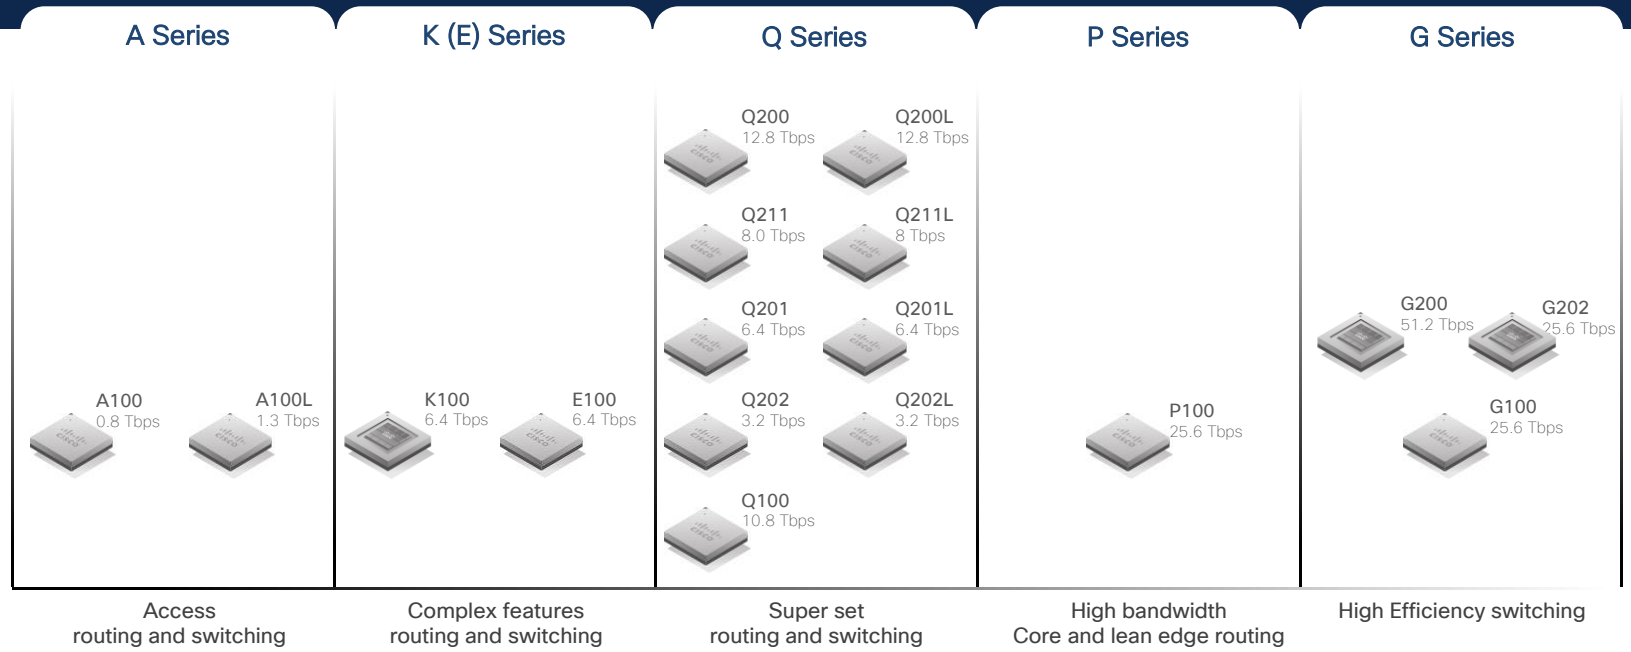

Use Cases

- Access
- Distribution
- Core
- Data Center
- Access
- Pre-aggregation
- Aggregation
- Core
- Peering
- Spine
- Leaf
- TOR
- DCI

Market segments

Enterprise | Service Provider | AI-ML Web Scale

One Architecture, SDK and P4 Forwarding Code

Networking in AI / ML clusters

Traditional Data Center traffic patterns

- **Numerous server applications** generate diverse traffic patterns and timing
- Flows are **random, asynchronous, and low-bandwidth**
- These small flows result in a **consistent average network load** over time

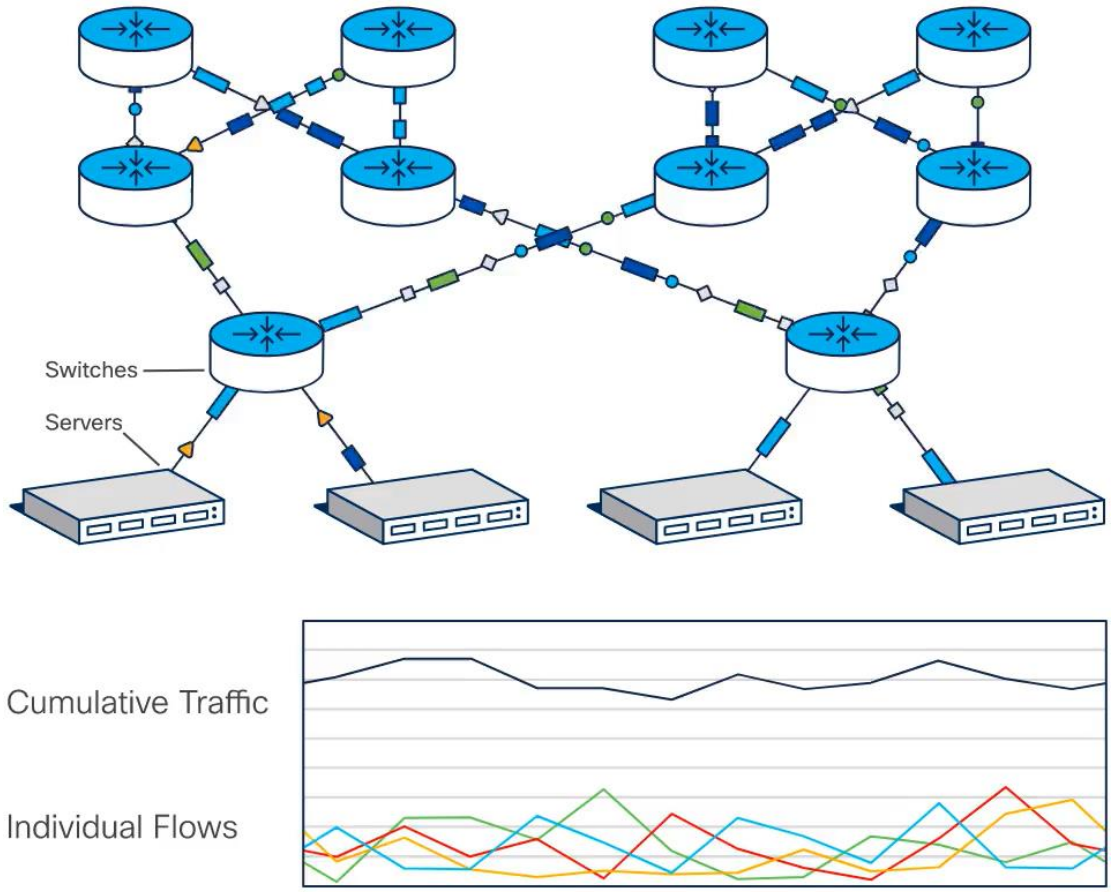

AI Data Center traffic patterns – optimized network

- AI/ML workloads depend on fewer; high-bandwidth flows with sharp fluctuations in network load
- Three execution phases drive traffic patterns:
  - Computation: GPUs process data
  - Notify: Results are exchanged between GPUs
  - Synchronization: GPUs wait for the slowest path to complete
- Optimized networks reduce delays, improving Job Completion Time (JCT) and maximizing GPU utilization.

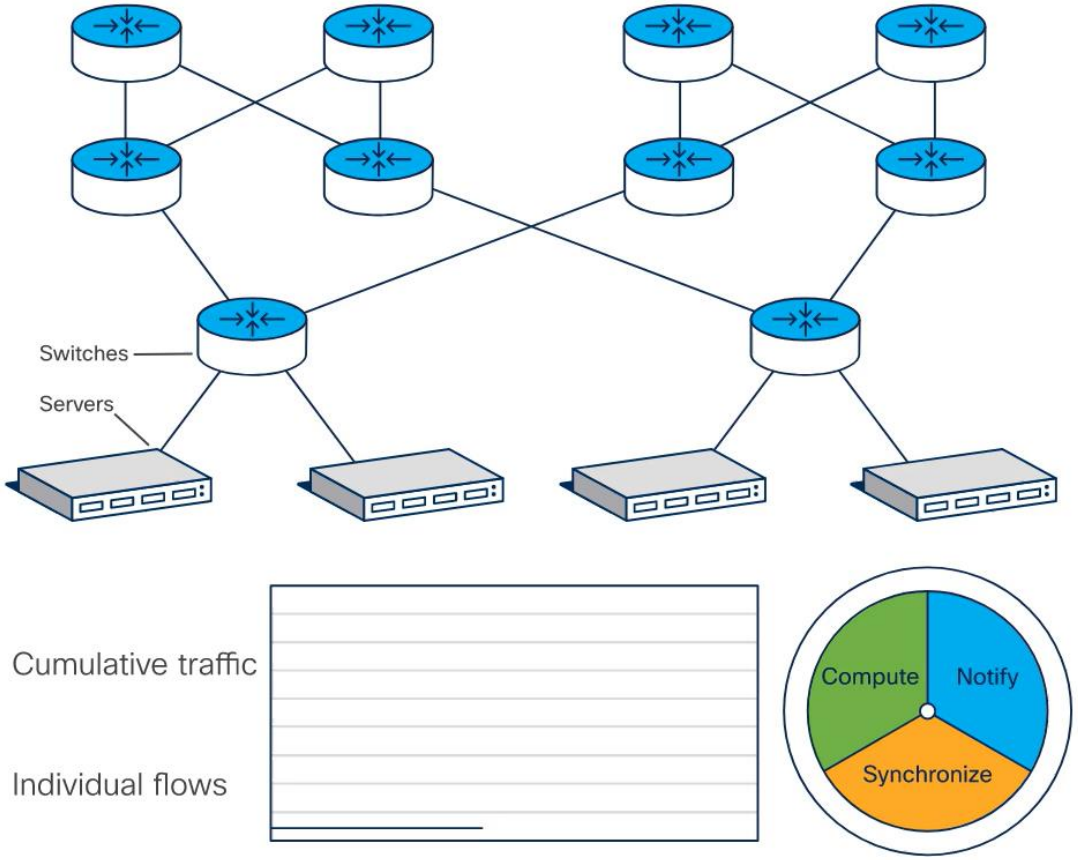

AI Data Center traffic patterns – network issues

- Network issues **disrupt synchronized flows**, causing delays in AI/ML workloads.
  - Congestion
  - Link failures
  - Uneven load balancing.
- Extended synchronization times lead to **increased Job Completion Time (JCT)**
- Inefficiencies can add **20–45% to AI execution time**

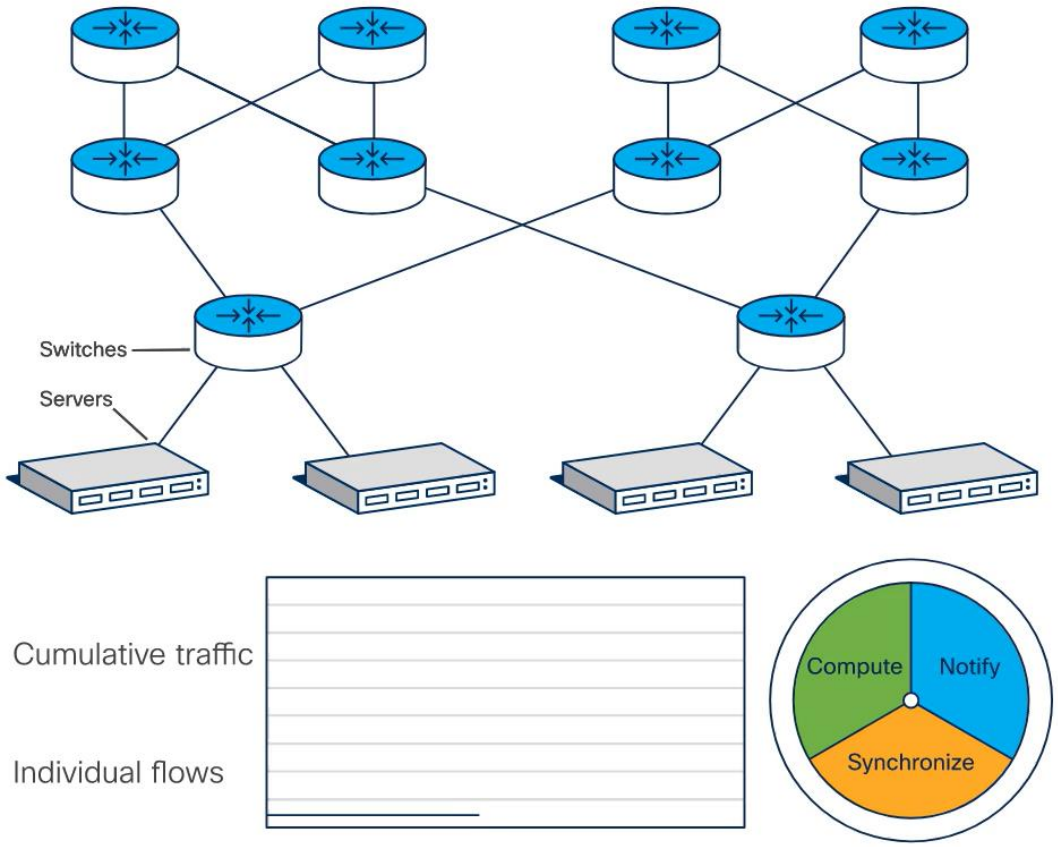

AI Networking Ethernet Fabric Choices

Lower latency, high bandwidth and superior multi-job efficiency, for optimal AI workload performance

Standard Ethernet	Enhanced Ethernet	Ultra Ethernet consortium	Scheduled Ethernet
Foundational networking, for general-purpose data communication	Upgraded Ethernet for better reliability and efficiency for AI networking	AI-focused Ethernet initiative optimizing Ethernet for large-scale AI workloads	Ethernet with deterministic scheduling, providing short job completion time
Scalability – support high speed connections, cost-effectiveness, broad industry adoption	Lossless transport, improved Quality of Service (QoS), better congestion control	AI-aware congestion control, low latency, scalability, industry compatibility, improved RDMA performance	Predictable latency, time-synchronized packet delivery
High bandwidth but lacks low-latency and congestion management for AI workloads	Reduces packet loss, improves stability for AI training clusters	Designed for multi-job AI training, dynamic resource allocation, scalability, large AI training clusters	Useful for AI real-time and time-sensitive applications

Silicon One programable architecture gives customers flexibility to program networks into various modes

Where does Ethernet fit across AI Landscape?

An Oversimplified Clustering of AI Technology Requirements

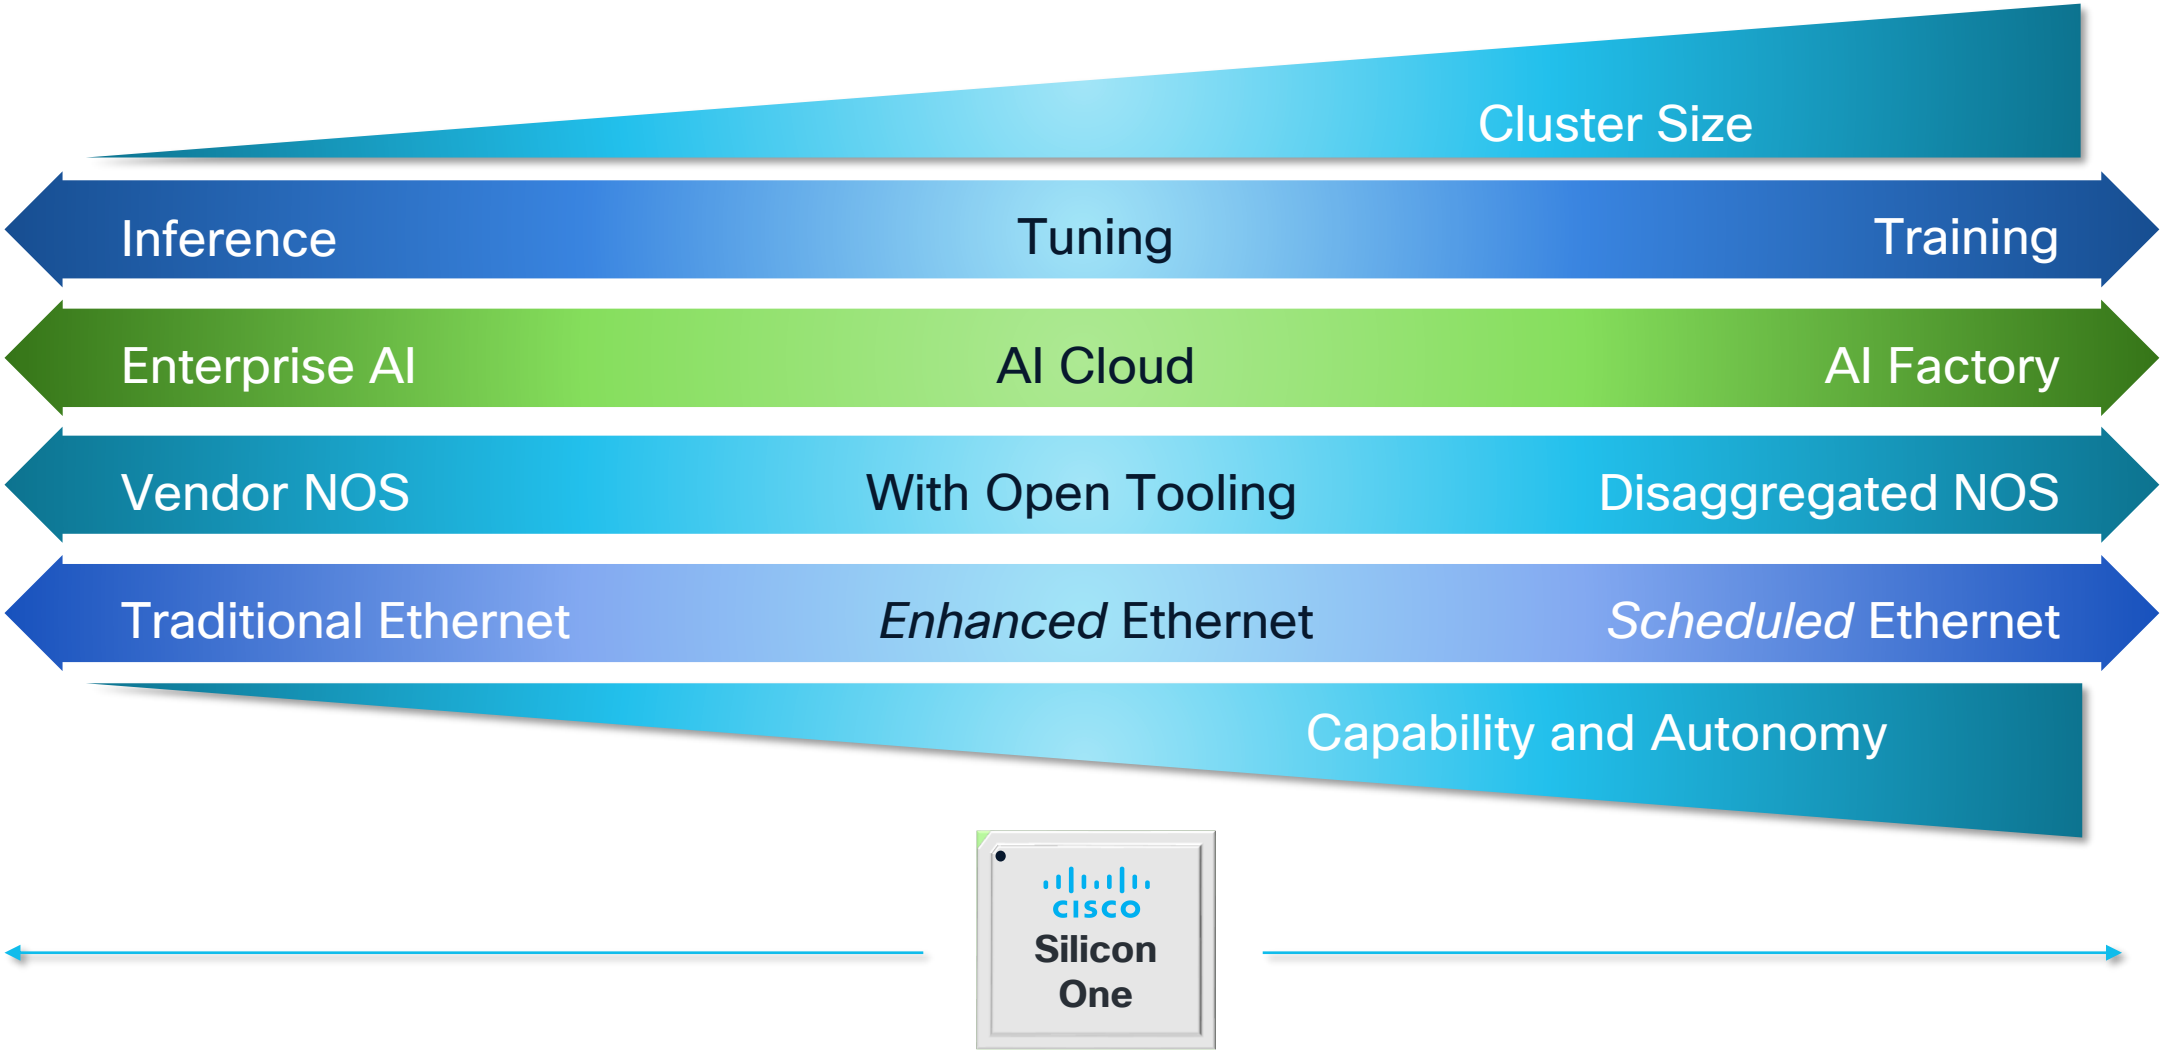

AI back-end network switch market share

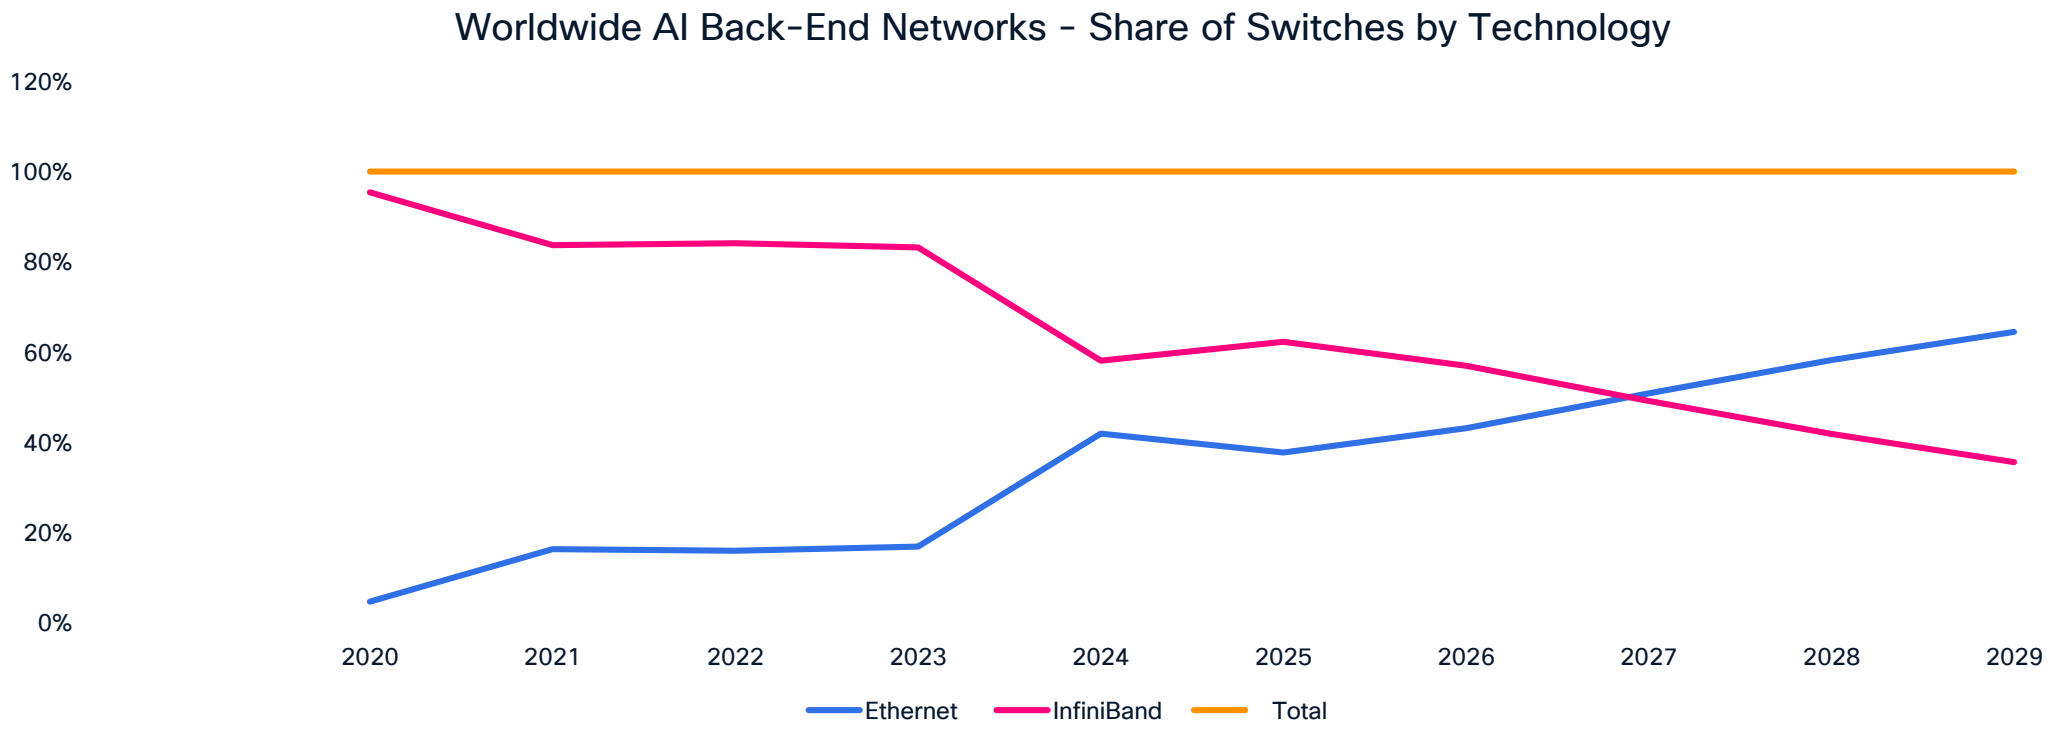

Source: Dell'Oro Group ADVANCED RESEARCH REPORT AI NETWORKS FOR AI WORKLOADS MARKET FORECAST (January 2025)

Cisco in AI / ML clusters

Breaking Down the AI Stack

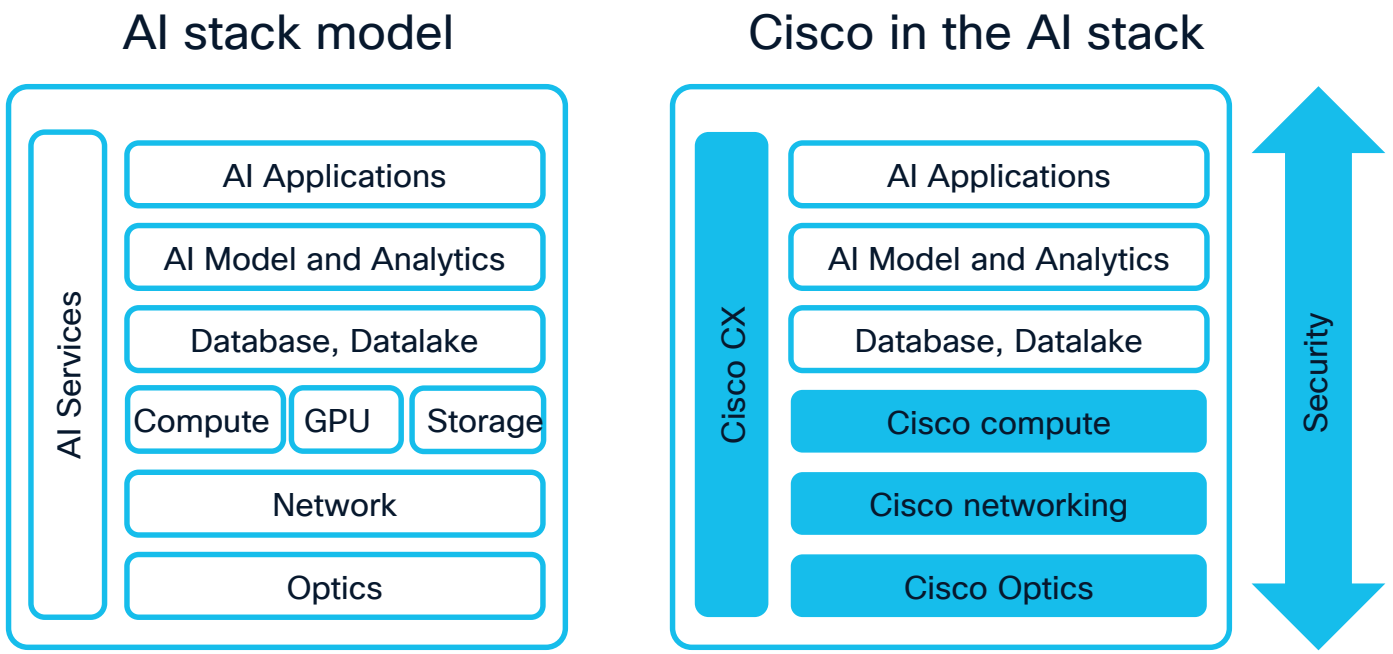

AI Native Infrastructure

The demands of AI/ML workloads are surging, and the solution is clear: modern Ethernet-based AI-native infrastructure. You can now deliver open, high-performance, scalable infrastructure that supports AI/ML workloads.

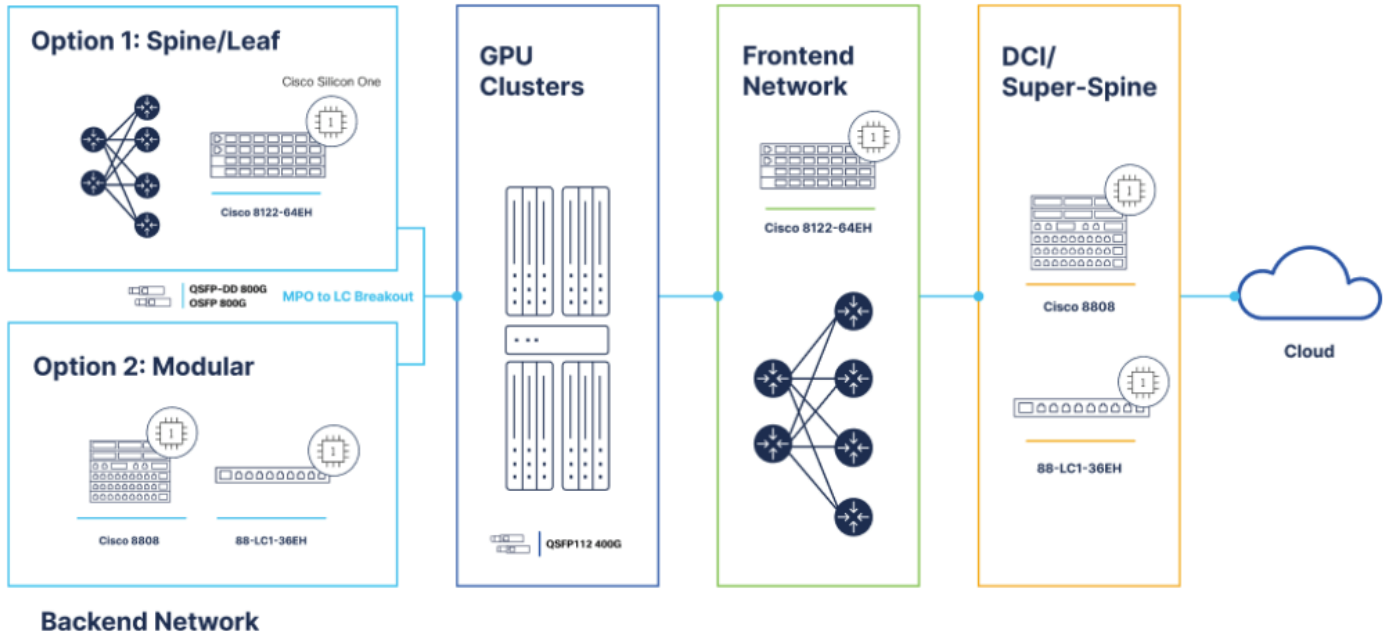

Cisco Nexus Hyperfabric AI

- 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

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

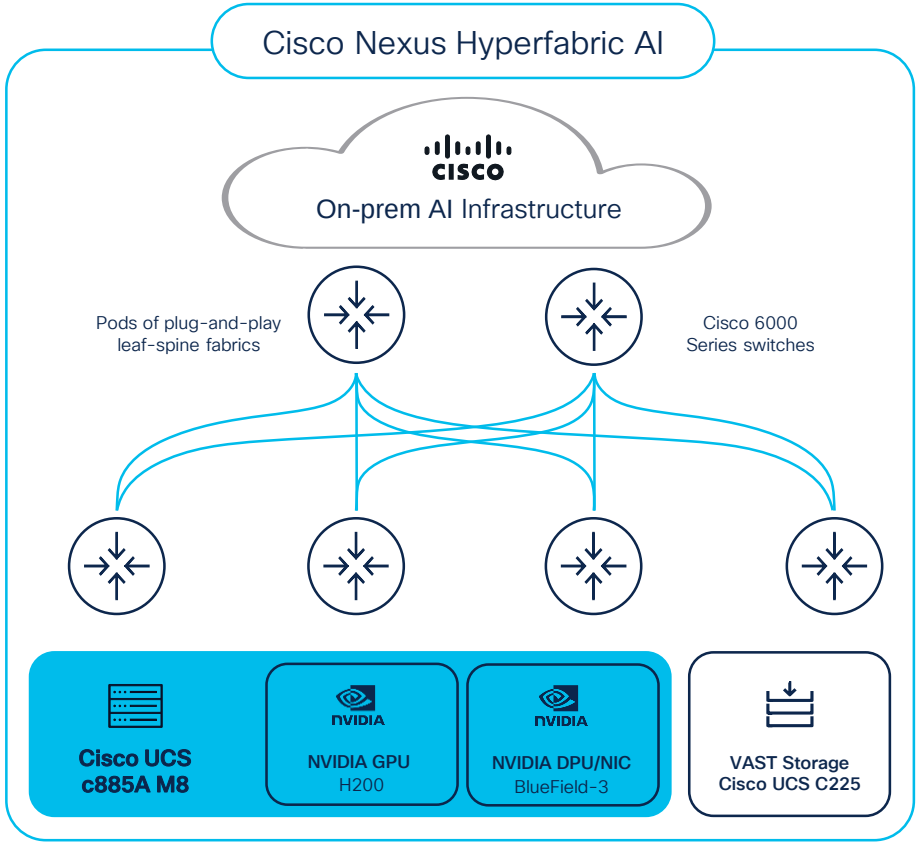

Cisco Silicon One AI/ML differentiators

Cisco Silicon One Differentiators For AI networking

Efficient Workload Processing

Optimized for AI/ML and modern networks

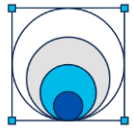

Scalable Clusters

Flexible support for any scale of infrastructure

Programmable Infrastructure

Adaptable to evolving workloads and use cases

Sustainable Data Centers

Energy-efficient and cost-effective

Management & Control

Telemetry for seamless operations

Security & Reliability

Built-in encryption and secure networks

Efficient workload processing

Load balancing alternatives

- **Traditional – ECMP**
  - Distributes data evenly across multiple network paths for smooth flow
  - Uses stateless Equal-Cost Multi-Path (ECMP) routing
  - Relies on a unique code from data source and destination
  - Ensures data travels the best path, maintaining order

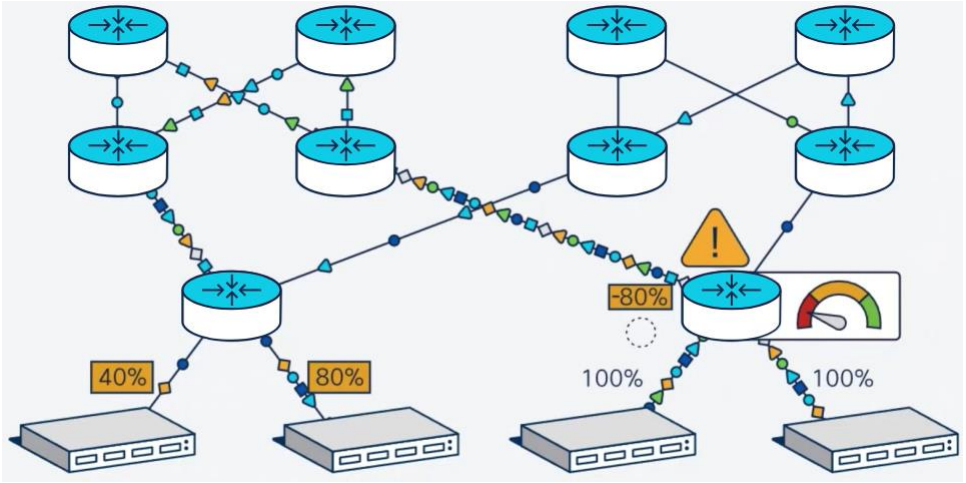

Efficient workload processing

Load balancing alternatives

- Traditional – ECMP
  - Distributes data evenly across multiple network paths for smooth flow
  - Uses stateless Equal-Cost Multi-Path (ECMP) routing
  - Relies on a unique code from data source and destination
  - Ensures data travels the best path, maintaining order
- **Dynamic stateful – Flowlet**
  - Unlike stateless ECMP, maintains state per flow and re-balances paths during idle times
  - Adapts to conditions to optimize link utilization and performance

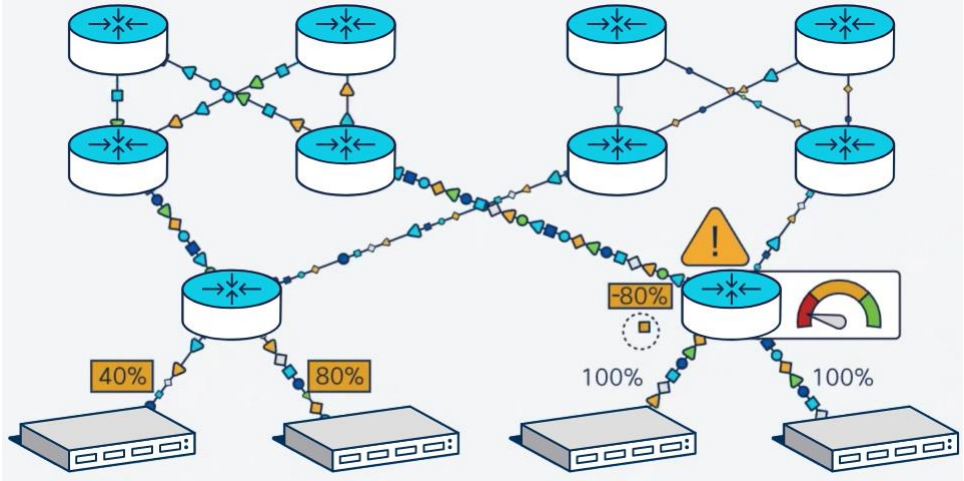

Efficient workload processing

Load balancing alternatives

- Traditional – ECMP
  - Distributes data evenly across multiple network paths for smooth flow
  - Uses stateless Equal-Cost Multi-Path (ECMP) routing
  - Relies on a unique code from data source and destination
  - Ensures data travels the best path, maintaining order
- Dynamic stateful – Flowlet
  - Unlike stateless ECMP, maintains state per flow and re-balances paths during idle times
  - Adapts to conditions to optimize link utilization and performance
- **Dynamic stateless - packet spray**
  - The switch sprays data packets randomly in a round-robin fashion across all available network paths based upon port utilization
  - Optimizes resources and maintains packet order for reliable transmission
  - Enhances performance by adapting to network conditions and maximizing link use

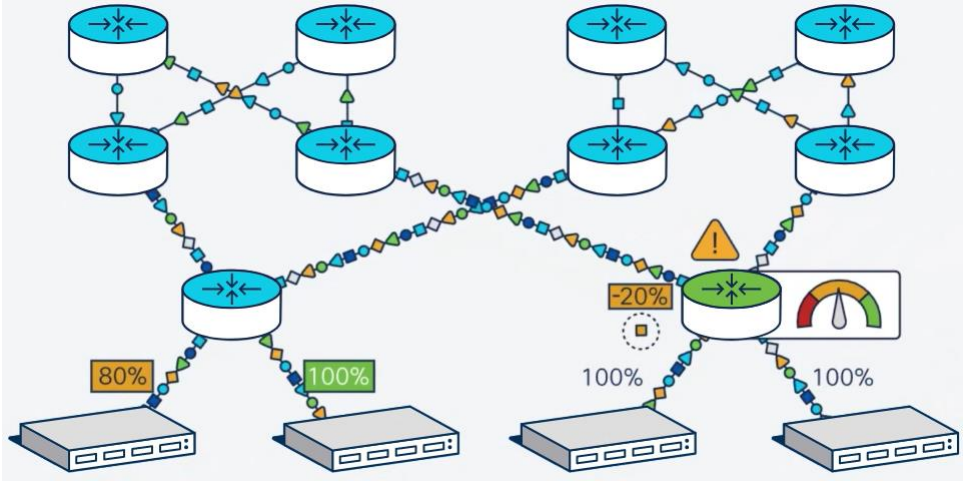

Efficient workload processing

Packet trimming manages congestion

- **Cisco Silicon One's Packet Trimming** ensures critical feedback is delivered during severe congestion
- **Payloads are removed** from dropped packets, but **headers are retained** and prioritized
- Retained headers: **Forwarded to the destination** or **returned to the sender**
- Reduces **retransmission delays** and provides explicit drop notifications
- **Endpoints react quickly** to adjust traffic flow and prevent further congestion

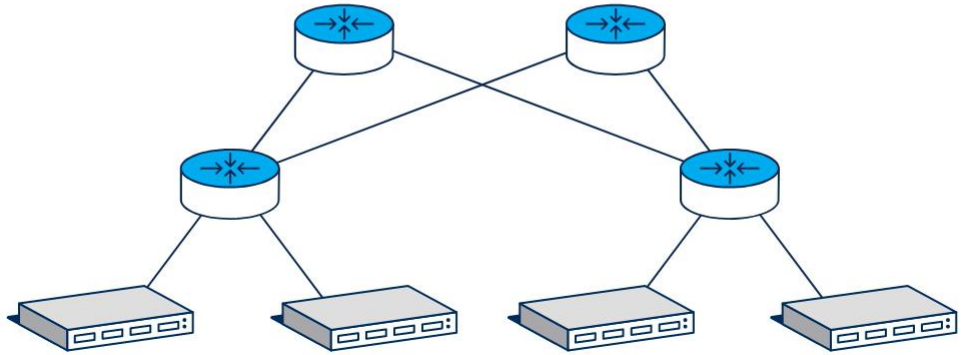

Efficient workload processing

Packet trimming - Forwarding enables faster recovery

- **Forward to Destination mode** sends trimmed packet headers to the destination via a **high-priority queue**
- The **destination notifies the source** of dropped packets for immediate retransmission
- Minimizes **retransmission timeouts** and accelerates congestion recovery

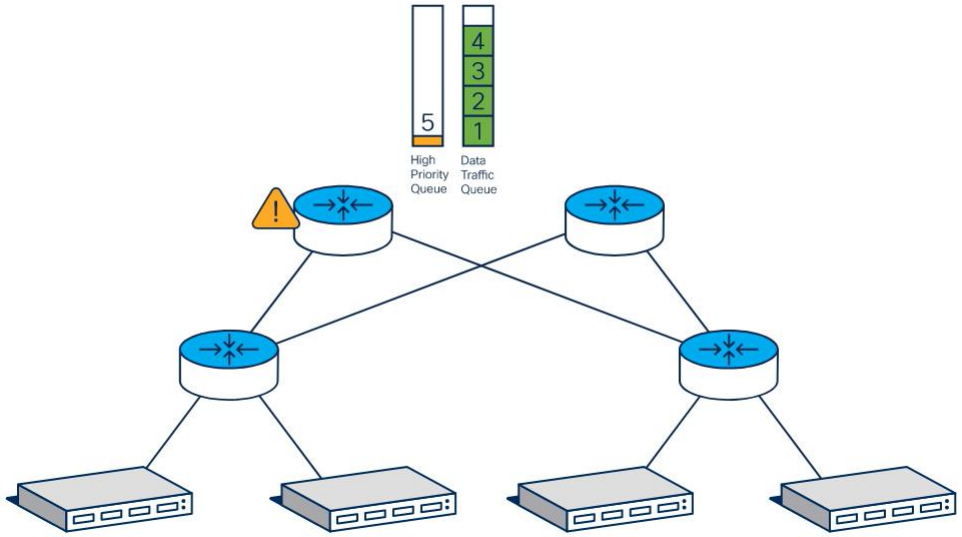

Scalable clusters

Silicon One architectural superiority

Traditional Memory Struggles with AI Traffic

- **Fixed buffer allocations** in traditional memory architectures are inefficient for bursty AI traffic
- **Congestion causes buffers to fill quickly**, leading to packet drops and interrupted data flows
- AI clusters amplify the issue due to **uneven bursts** from GPUs completing tasks at different times
- Traditional architectures cannot scale to meet the **high-bandwidth demands** of AI workloads or multi-terabit Ethernet-based AI fabrics

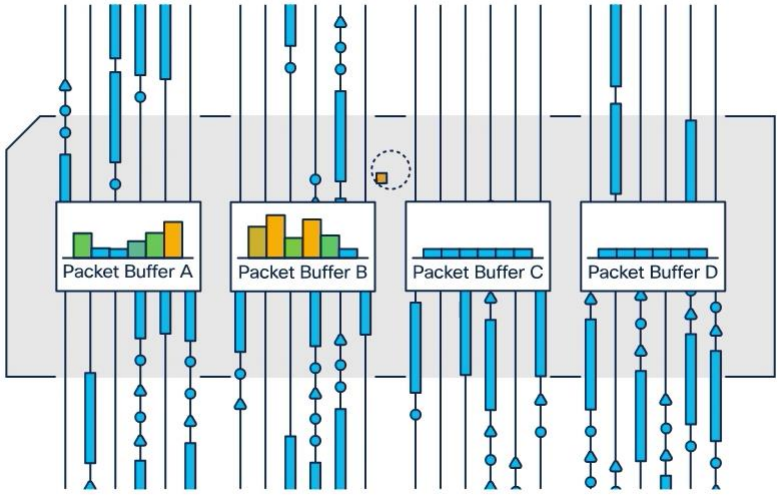

Scalable clusters

Silicon One architectural superiority

Shared Memory Transforms Scalability

- **Dynamic shared memory** optimizes resource utilization across all ports, eliminating packet drops
- Absorbs **traffic bursts** and ensures **fair memory distribution** in real time for smooth data flow
- **Sophisticated queuing mechanisms** prioritize critical traffic, reduce latency, and handle multi-terabit speeds
- Minimizes **power consumption** and **silicon redundancy**, enabling cost-effective, high-performance AI networking

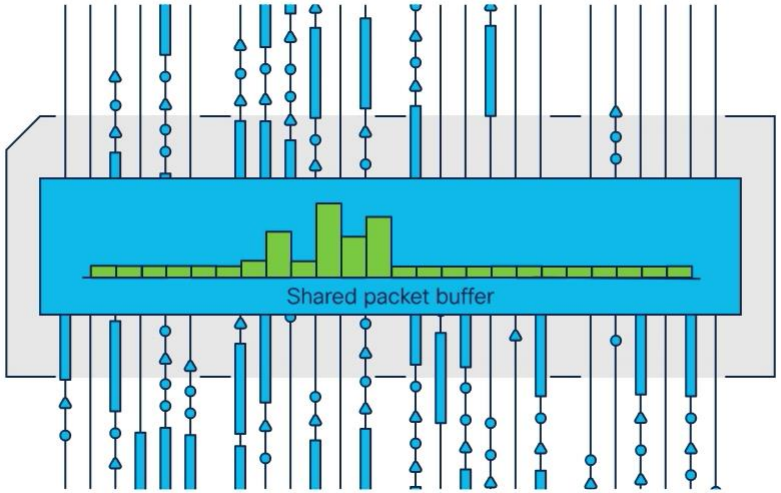

Programable network

Network programmability

What is Programmable Networking?

- **Redefines flexibility** by enabling dynamic updates to ASIC features and hardware behavior without hardware changes
- Seamlessly integrates new capabilities, ensuring networks evolve with emerging needs.

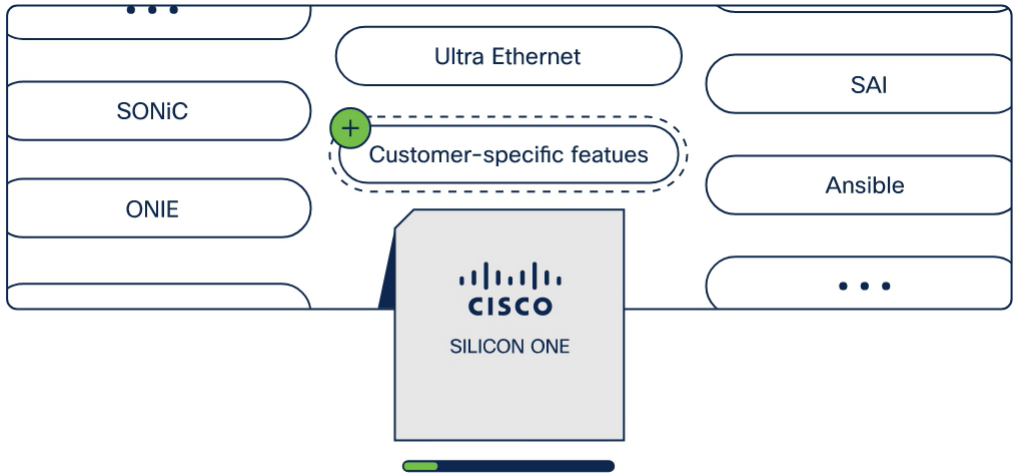

Programable network

Network programmability

How It Works?

- **Powered by P4**, Cisco Silicon One adapts hardware behavior to evolving requirements
- Key capabilities include:
  - **Precise packet inspection** to determine actions
  - **Real-time packet manipulation** to resize packets or add metadata for new protocols.
  - **Advanced routing techniques,** ensure optimal path selection
  - **Telemetry updates** provide actionable insights for traffic optimization
- Delivers unmatched **flexibility** to adapt to changing standards and workloads

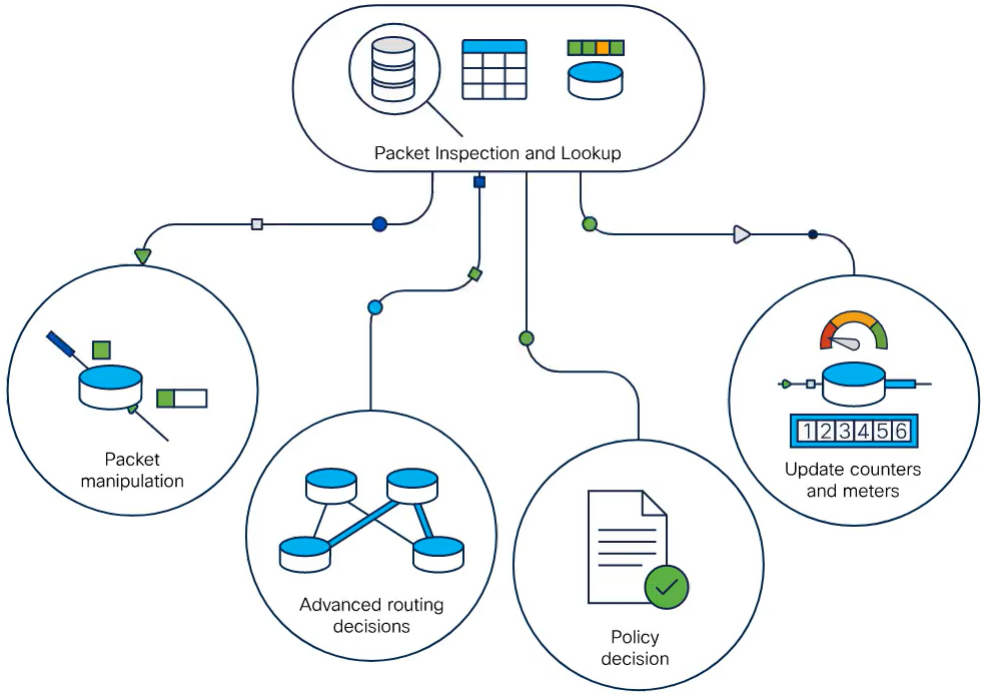

Programable network

Cisco Silicon One SDK Empowers Programmability

- Provides a **unified, consistent development experience** across all Silicon One systems
- Simplifies network management, accelerates feature deployment, and enhances serviceability
- Key features include:
  - **Consistent debug shell, shared APIs, and a common abstraction layer** for streamlined troubleshooting and seamless migrations
  - **Granular control** over packet handling, telemetry, and load balancing
- **Hitless software upgrades** ensure uninterrupted operation during updates, critical for environments requiring continuous uptime
- Enables **reliable, scalable, and future-ready networks**

Advance Data Center Sustainability

Advanced SerDes innovations

- AI-optimized networks offer 100 to 1000 times more bandwidth than traditional ones
- Increased bandwidth leads to substantial power demands
- Focus on reducing system power, especially optics
- Cisco's Silicon One uses SerDes for efficient data conversion

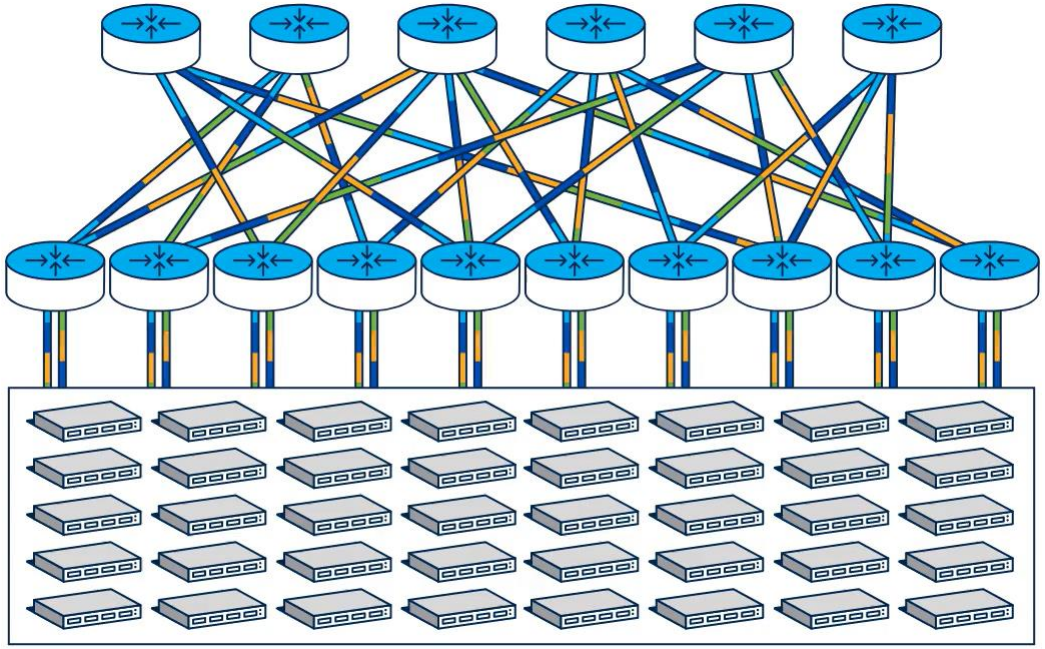

Optics account for approximately 50% of the total power used in AI networking.

Advance Data Center Sustainability

Advanced SerDes innovations

Unique capabilities of Silicon One enable the use of lower power linear pluggable optics (LPO) or passive direct attach copper cables (DAC)

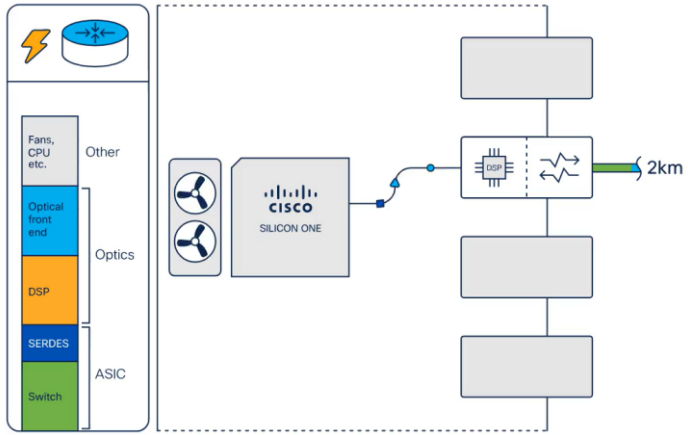

Traditional optics

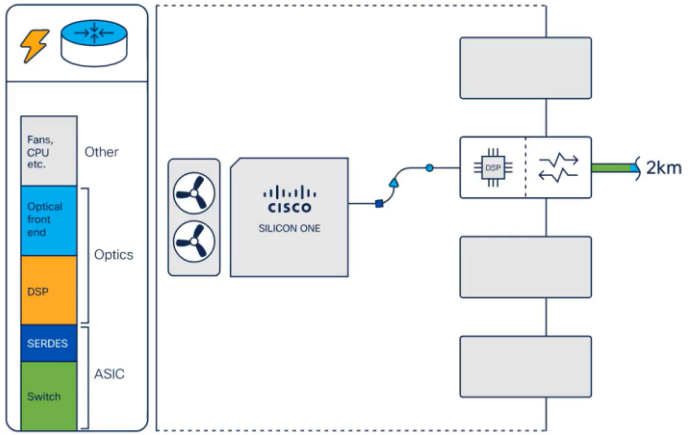

Linear pluggable optics

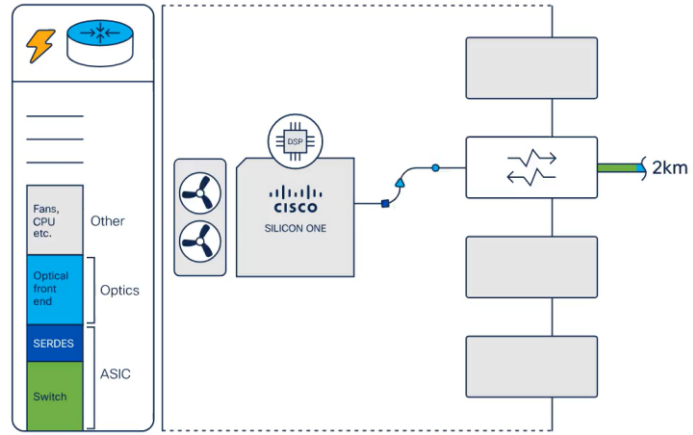

Direct attached cable

Manage and Control

Streaming telemetry – proactive network optimization

- Proactive network maintenance optimizes performance by addressing weak points and risks
- Telemetry data collection from the control system, particularly from the Cisco Silicon One chip, is vital for this maintenance approach
- Cisco Silicon One collects data on temperature, voltage, link status, congestion, traffic type, active flows, and microburst detection
- The control system updates routing paths using this data to enhance job completion time and reduce latency.

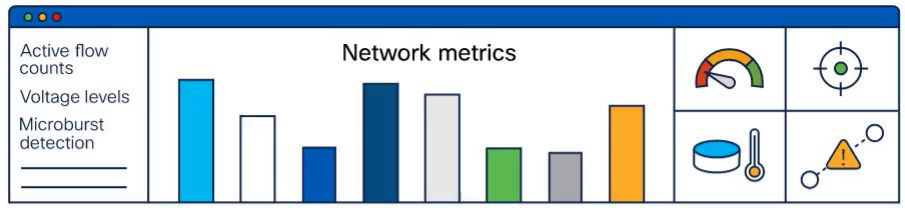

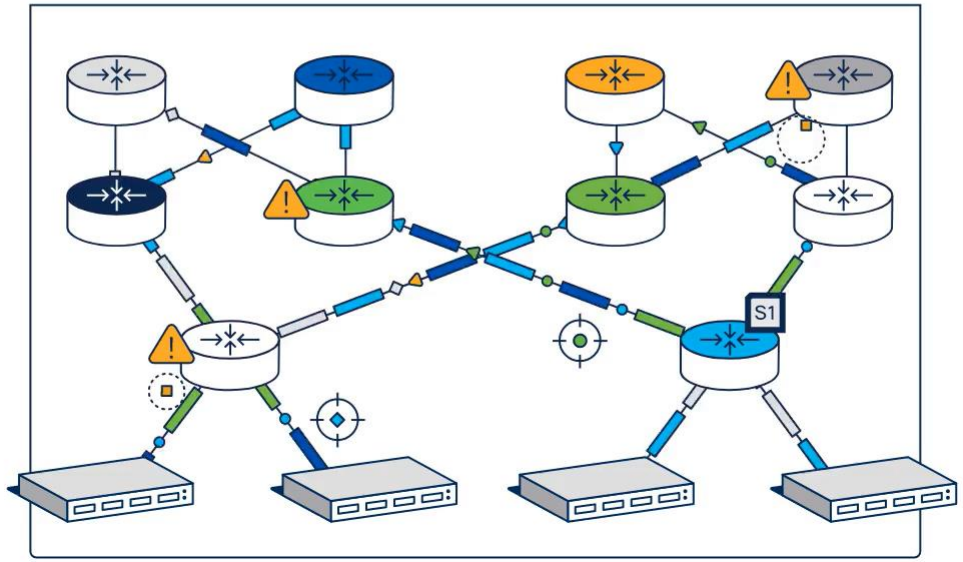

Security

Establishing Network Trust

- Cisco Silicon One ensures device security with built-in mechanisms to verify chip authenticity.
- Trust Anchor Module checks CPU software, and the Silicon One device identifies itself to the software.
- Confirms correct installation in intended systems, ensuring trusted systems and networks.

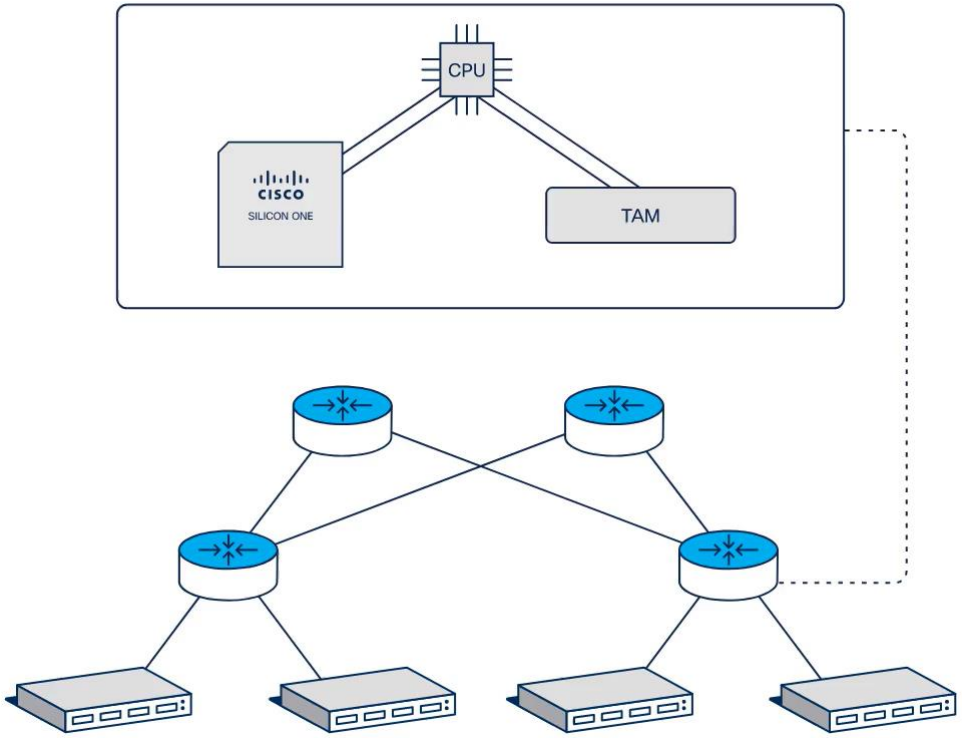

Security

Optimized Encryption and Performance

- Cisco Silicon One supports IPSec and MACSec for secure AI workloads in scale-out networks
- Features a unique format combining MACSec and IPSec benefits, tailored for RDMA workloads
- Provides native RDMA support and encapsulates all packet formats beyond IPSec limitations
- Ensures low latency with cut-through functionality and improved load balancing with visible entropy fields
- Enhances job completion time (JCT) compared to IPSec and MACSec

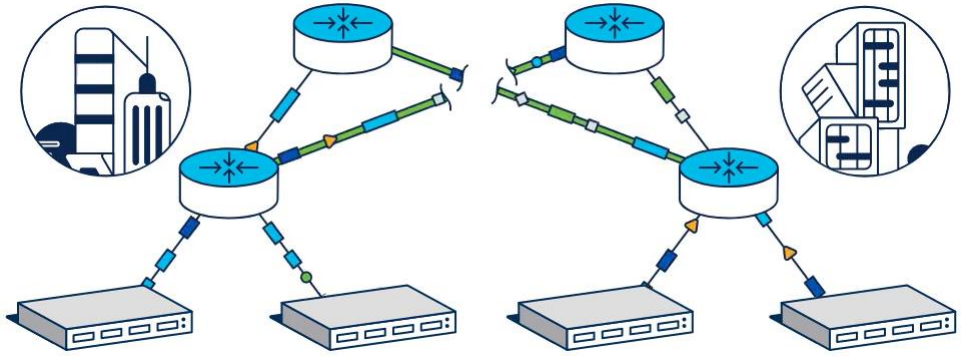

Complete your session evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to win 1 of 5 full conference passes to Cisco Live 2026.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn exclusive prizes!

Complete your surveys in the Cisco Live mobile app.

Continue your education

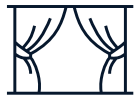

Visit the Cisco Showcase for related demos

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 www.CiscoLive.com/on-demand

Contact me at: yaagami@cisco.com

Thank you

cisco Live !


CISCO