



Cisco Nexus One

Turning Network
Capacity into
AI Infrastructure Value

Cisco Nexus One:

Turning network capacity into AI infrastructure value

Executive Summary

For production AI, raw bandwidth is not enough: the network must keep accelerators highly performant, connect every traffic domain and automate isolation, and provide the workload-to-fabric visibility required to operate securely and efficiently at scale. Cisco Nexus® One provides a standards-based, next-generation data-center networking architecture that preserves customer choice across the entire AI ecosystem.

Built on three foundational principles—high-performance networking from silicon to workload, fused security at every layer, and simplified operations powered by AgenticOps—it serves as a cornerstone for the world's most advanced data centers, supporting hyperscalers, neoclouds, sovereign clouds, enterprises, and service providers.




Figure 1. Cisco Nexus One delivers next-generation data-center networking architecture for AI.

Why Cisco Nexus One for AI networking?

Below summarizes its technical differentiators and key customer outcomes.

Differentiator 1

Cross-domain AgenticOps with human control

Cisco Nexus One advantage

[Cisco Cloud Control](#), our purpose-built unified AgenticOps platform, is designed for both humans and AI agents. It manages, monitors, defends, and remediates environments while ensuring consistent human oversight.

Cisco Cloud Control brings authorized data from these platforms into a common operational view. This data includes inventory, sites, topology, neighbor relationships, device and fabric health, alerts, lifecycle and compliance status, licensing information, and relevant telemetry.

Cisco Cloud Control provides a unified, cross-domain operational context without replacing existing controllers. From a data-center networking perspective, it seamlessly integrates with Cisco Nexus One.

Unified identity and role-based access control preserve permissions from the source platforms.

Through AI Canvas serving as an AgenticOps platform, agents correlate networking data with insights from security, compute, observability, and collaboration systems. This shared context helps teams investigate dependencies, identify likely root causes, and coordinate remediation.

Powered by [Cisco's Deep Network Model](#), it delivers networking-specific reasoning with 20% higher accuracy than general-purpose LLMs. By combining 40+ years of trusted expertise with rigorous approvals, audit trails, and human-in-the-loop controls, the system enables the deployment of autonomous actions with absolute confidence.

Customer outcomes



Faster remediation

Compress incident investigation timelines and improve collaboration between platform, network, and security teams, enabling defensible, faster remediation decisions that minimize network downtime.



Operational efficiency

Realize up to 5x fewer tool-calls for troubleshooting tasks and 3x lower latency for networking Q&A.



Rapid time-to-value

Utilize real-time safety validation to accelerate transition from deployment to value in "minutes."

Sources: [Cisco Cloud Control](#) | [Resolve Data Center Issues Faster with Agentic AI](#)

Differentiator 2

One architecture from silicon to workload

Cisco Nexus One advantage

Replace an inefficient “part-by-part” assembly with a unified, end-to-end networking stack that provides compounding value over time.


[Cisco Nexus One](#) is uniquely positioned to deliver fully integrated data-center networking by bringing together silicon, systems, optics, software, security, and automated lifecycle operations.

Customer outcomes



Increased efficiency

Reduces integration challenges, deployment risks, and lifecycle costs.



Rapid scale

Accelerates the journey from pilot to production while ensuring consistent policy, security, and visibility across multiple fabrics—fabric to workload.

This architecture is proven by [Cisco IT AI-ready data center deployment](#).

- AI fabric backend design and deployment time reduced from weeks to mere hours.
- An entire AI cluster was designed and deployed in just three months.
- 80% faster process than alternative rebuild methods.

Sources: [Cisco Nexus One Solution Overview](#) | [Cisco IT deploys AI-ready data center](#)

Differentiator 3

Investment protection through Cisco custom silicon innovation

Cisco Nexus One advantage

Cisco maintains total control over the end-to-end architecture, from the silicon to the systems and software stack.

By leveraging P4-programmable [Cisco Silicon One](#) with fully shared, deep buffers, we enable new network

functions and role shifts on existing equipment. This reduces dependency on third-party ASIC roadmaps and allows for tighter hardware and software co-engineering to meet emerging AI requirements.

Customer outcomes



Extended longevity

Extend the usable life of qualified infrastructure and avoid premature “rip-and-replace” cycles.



Supply chain agility


Gain greater supply chain control and adopt new capabilities through software-defined programmability.



Future-ready innovation

Preserve a controlled innovation path as AI networking evolves, with built-in readiness for standards including UEC and ESUN.

Sources: [Cisco Silicon One](#) | [Cisco Silicon One Product Family](#)



Differentiator 4

The only vendor reference architecture with full-stack powerhouse

Cisco Nexus One advantage

An end-to-end infrastructure with integrated high-density compute that solves for supply chain constraints and thermal density—all while maintaining the security and observability that enterprises already trust.

Cisco expands the [Cisco Secure AI Factory with NVIDIA through a partnership with Supermicro](#), bringing its high-density, liquid- and air-cooled compute to Cisco's industry-leading full-stack, secure AI infrastructure. It unifies compute, network, power, and cooling management into a single console, eliminating operational silos and providing end-to-end visibility and

consistent policy enforcement from day 0 through day 2.

Similarly, Cisco Nexus One offers a fully validated reference architecture with AMD for AI infrastructure deployment. It combines balanced compute, high-memory capacity, and an open, Ethernet-based AI networking layer to enable predictable AI workload scaling, maximize GPU productivity, and sustain performance from pilot to large multinode clusters.


This transforms complex, multi-vendor AI environments into a cohesive, manageable factory.

Customer outcomes



Computing power for rack-scale

Deploy rack-to-fabric scale featuring liquid-cooled AI networking Cisco N9000 systems alongside Supermicro's liquid- and air-cooled servers. This unlocks trillion-parameter training and high-throughput inference use cases with platforms including [NVIDIA Vera Rubin NVL72](#), [NVIDIA Grace Blackwell NVL72](#) and [NVIDIA HGX Rubin NVL8](#).



Validated infrastructure services

Shifting focus from “building infrastructure” to “generating intelligence” by compressing deployment timelines through Cisco Validated Infrastructure Services (CVIS) to de-risk an acceleration path from design to first token, and from first token to business value.



Unmatched co-engineered ecosystem

Cisco is the only NVIDIA technology partner to utilize its own networking switches and network operating system in an NVIDIA Cloud Partner (NCP) compliant solution.

Cisco also provides a fully validated reference architecture for high-performance AI clusters with [AMD GPUs](#) and [NICs](#).



Optimized value

Customers gain a reliable, repeatable path to AI buildouts for any use cases with pre-validated designs that help ensure that the infrastructure is built exactly as tested, helping to start AI projects faster and extract more value from every watt.

Sources: [Cisco Expands Secure AI Factory with NVIDIA for the Rack-Scale Era](#) | [More Than Rack-Scale Compute: Operationalizing AI at Scale](#) | [New Cisco Secure AI Factory with NVIDIA: Built for Rack-Scale Era](#)

Differentiator 5

Network capacity converted into accelerator productivity

Cisco Nexus One advantage

Optimizes the traffic patterns that static ECMP can distribute inefficiently, leading to network bottlenecks.

[Cisco Intelligent Packet Flow](#) with [Cisco G300 Intelligent Collective Networking](#) utilize flowlet,

per-packet or policy-driven load balancing, congestion awareness, fault detection, and telemetry to help ensure that the network acts as a high-performance extension of the GPU compute cluster.

Customer outcomes



Optimized utilization

Improve path utilization and make accelerator performance more predictable, reducing the risk that expensive GPUs wait behind avoidable network bottlenecks.




Industry leading hardware

The [Cisco Silicon One G300 Series](#) delivers 102.4 Tbps capacity, fully shared memory architecture that provides [2.5x burst absorption, and fault detection 100,000x faster](#) than traditional mechanisms.



Maximize capacity

Minimizes network-induced GPU stall time, reclaiming up to 28% of total cluster capacity that would otherwise be lost to idle time.



Predictable scalability

Provides stable, predictable performance for massive AI clusters (up to 1M+ GPUs), ensuring that expensive compute resources are prioritized for actual training rather than waiting for data.



Superior speed

Achieves Collective Completion Times (CCT) within 2% to 4% of the theoretical ideal, maintaining consistency even during complex link failures. Delivers an [82% reduction in Job Completion Time \(JCT\)](#) compared to standard Ethernet (ECMP) and a 28% speedup over random packet spraying.

Sources: [AI network performance with Cisco Intelligent Packet Flow](#) | [Cisco G300 Intelligent Collective Networking transforming AI Infrastructure TCO](#) | [Cisco Silicon One G300: The Next Wave of AI Innovation](#)

Differentiator 6

Choice across silicon, systems, and software

Cisco Nexus One advantage

Provide multiple architectural paths' flexibility instead of forcing one vertically integrated stack.

[Cisco Nexus One](#) allows organizations to tailor their fabric to unique workload, scale, and ecosystem requirements, helping to ensure that the network aligns with operational preferences.

The [Cisco N9000 Series switches](#) span Cisco Silicon One, Cisco Cloud Scale, and [NVIDIA Spectrum-X](#)

[Ethernet switch silicon](#) options, with Cisco NX-OS, Cisco ACI, or enterprise-grade SONiC operating models.

Cisco Nexus One supports the distinct requirements of backend, frontend, storage, and management networks across scale-out and [scale-across](#) environments.

Customer outcomes



Future-ready infrastructure

Built on an open, modular foundation to help ensure that infrastructure remains adaptable to changing technology while avoiding unnecessary dependency on a single stack.

Maximize your commercial leverage and avoid the risks of proprietary lock-in.



Superior architecture

Deliver predictable training and inference services without leaving application access, data movement, management, or recovery as architectural gaps.



Consistent scale

Achieve consistent, end-to-end performance at any scale without complexity.

The [Cisco Silicon One G-Series](#) powered switches deliver a massive, high-density AI fabric supporting 1.6Tbps-ready connectivity, while the [Cisco Silicon One P-Series](#) powered switches provide deep-buffers with 800G performance for seamless Data Center Interconnect (DCI) and [GPU cluster-to-cluster communication across campus, metro, and regional environments](#).

Sources: [Cisco Nexus One Solution Overview](#) | [Cisco N9000 Series switches](#) | [Cisco N9100 Series switches](#) | [Networking for the Agentic Era: Cisco Unveils New Innovations in Scale and Simplicity](#) | [Unlock the power of scale-across with Cisco](#) | [One open NOS, any workload: SONiC on Cisco](#)

Differentiator 7

Open Ethernet and multi-vendor AI ecosystem

Cisco Nexus One advantage

Preserve accelerator, NIC, server, and storage choice through an open, standards-based Ethernet and end-to-end validated design.

Cisco Nexus One, a fully validated [Ethernet-based architecture](#), integrates seamlessly into a broad

ecosystem including NVIDIA, AMD, Intel®, Supermicro, VAST, WEKA, DDN, Pure Storage, and Lenovo—enabling AI factories with best-in-class components.



Customer outcomes



Unified operations

[Ethernet-based Cisco Nexus One](#) enables a unified operational model across frontend, storage, and backend networks. This means one talent pool, one set of tools, and one management process, reduced supply-chain lead times, and eliminating the “proprietary tax” of maintaining two distinct network fabrics.



Flexibility

Select technology of choice according to workload, availability, and economics while retaining a consistent experience with network, heterogeneous technology environments, and the operating model.

Sources: [Uncompromised Ethernet: Performance and Benchmarking for AI/ML Fabric](#) | [Scaling the future: Why Ethernet is the backbone of AI Supercomputing](#)

Differentiator 8

Kubernetes-to-fabric networking, observability, and security

Cisco Nexus One advantage

A unified, distributed, and workload-aware networking model that seamlessly connects application identity to the physical network path.

Through [Cisco Nexus One fabric integration with Isovalent®](#), experience deep correlation across containers, services, nodes, fabric paths, policies, and external dependencies.

As Kubernetes becomes the standard for GenAI inference ([used by 66% of organizations](#)) and [production workloads \(82% of container users\)](#), our solution delivers a standards-based network fabric that


extends directly into the container environment. This ensures native connectivity, identity-aware policies, and unified visibility across both cloud-native and traditional workloads.

AI and agentic workloads are reshaping data centers through machine-scale, east-west traffic and rapidly evolving applications. The [Cisco N9300 Series Smart switches](#) combines trusted Nexus® networking with programmable DPUs, enforcing security inline to protect data exactly where workloads run.

Customer outcomes

**Accelerated troubleshooting**

Move rapidly from an inference symptom to the specific affected workload and network path.

**Streamlined operations**

Drastically reduce cross-team handoffs and significantly improve Mean Time to Resolution (MTTR).

**Unified Security**

Achieve a faster, more resilient network. Consolidate high-performance networking, stateful segmentation, and policy enforcement with [Cisco N9300 Series Smart switches](#). Modernize infrastructure for AI and strengthen security—without compromising performance or increasing complexity.

Sources: [Isovalent Networking for Virtualization \(INV\) Integration with Cisco Nexus One](#) | [End-to-end AI networking: Cisco's answer to the inferencing era](#) | [Fuse Security with Data Center Networks](#)

Differentiator 9

Vulnerability response built into network operations

Cisco Nexus One advantage

Add a mitigation path between vulnerability disclosure and permanent patching by embedding runtime protection directly into the operating systems of modern switches.

Leveraging eBPF and Tetragon technology, [Cisco Live Protect](#) applies compensating controls instantly. This bridges the gap between the initial vulnerability disclosure and the final patch, enabling security responses without requiring switch reboot or waiting for scheduled maintenance windows.

Customer outcomes



Minimize vulnerability window

Reduce vulnerability window from [months to minutes](#), effectively neutralizing AI-speed threats.



Preserved uptime

Protect infrastructure while avoiding emergency outages or unscheduled reboots.



Disciplined remediation

Secure systems immediately while maintaining the agility to complete permanent fixes on your own schedule.

Sources: [Cisco Live Protect](#) | [Bringing eBPF-Powered Protection into Network Infrastructure](#)





Differentiator 10

A proactive, phased path toward post-quantum readiness

Cisco Nexus One advantage

Extend cryptographic planning across links, administrative access, and hardware identity.

[Cisco Nexus One integration of security standards](#) across the entire stack—from hardware identity (SUDI)

to transport encryption (MACsec/IPsec), ensures that quantum-safe security updates are manageable at the infrastructure level.

Customer outcomes

The evolving threat landscape—especially threats posed by AI and quantum computing—demands a proactive response from data center operators.



Minimize data exposure

Eliminates the vulnerability window for data encrypted today that must remain secure for years with long-lived cryptographic risk.



Proactive cryptographic planning

Provides a clear, low-friction migration path, avoiding sudden infrastructure overhaul or operational risks when quantum computing advances.

Sources: [Protecting against rising cybersecurity risks in data centers](#) | [Post-Quantum Cryptography, Quantum-Safe Computing](#)

Differentiator 11

Unified operational consumption models

Cisco Nexus One advantage

Support both on-premises control and cloud-managed lifecycle operations.

[Cisco Nexus Dashboard](#) and [Cisco Nexus Hyperfabric](#) allows organizations to align AI infrastructure management with their governance, skills, and consumption preferences.

Native Splunk in Cisco Nexus One features a resilient multi-node architecture that integrates with Cisco Data Fabric to bring analytics directly to your network telemetry. By analyzing data where it resides, you eliminate data gravity, bridge silos across NetOps

and SecOps, and maintain full data sovereignty while gaining unified, high-availability visibility.


Cisco also provides [Network as Code](#) (NaC), which treats network infrastructure and services as code, enabling automation, version control, and continuous integration, deployment, and testing (CI/CD/CT) practices to manage and provision networks programmatically. This methodology applies DevOps principles to replace manual network configuration with code-based automation, improving reliability, scalability, and efficiency

Customer outcomes



Flexible operating model


Deploy AI infrastructure on your terms. Adopt the operating model that fits the organization without redesigning the underlying AI networking strategy.



Operational efficiency

Shorten deployment and troubleshooting cycles while enabling existing network teams to operate AI infrastructure with established controls.

Sources: [Cisco Nexus Dashboard](#) | [Cisco Nexus Hyperfabric](#) | [Automate Network Infrastructure using Network as Code](#)





Differentiator 12

Security granularity extending from tenant to AI job

Cisco Nexus One advantage

Advance job-aware observability and segmentation beyond infrastructure boundaries.

Nexus One [AI Job-ID-aware segmentation](#) associates job identity with fabric policy using a VXLAN Endpoint Security Group (ESG) and securely hosts multiple AI Jobs on a single shared fabric.

Customer outcomes



XPU monetization

Enable shared AI infrastructure to support jobs with different data-access, regulatory, and security requirements without creating dedicated infrastructure islands.



Granular segmentation and observability

Provides AI job-level segmentation with deeper visibility and faster troubleshooting for AI fabrics.

Sources: [From tenant-aware to job-aware: scaling shared AI clusters with Cisco Nexus One](#) | [Full Stack Security and AI Job-ID-aware Segmentation: Cisco Raises the Bar for AI Fabric Efficiency](#)

Differentiator 13

High-speed optics included in the architecture

Cisco Nexus One advantage

AI networks present unique challenges, requiring infrastructure that can handle high data throughput, low latency, and continuous reliability under extreme workloads. Cisco's approach stands out by combining advanced technologies, a comprehensive portfolio, and rigorous testing methodologies to address these demands. Unlike many vendors, Cisco provides seamless interoperability across multivendor environments, delivering [Cisco Optics](#) tested for both Cisco and third-party platforms.

Cisco continues to expand its portfolio of high-speed optics designed for AI applications in front-end and

back-end networks, including OSFP 800G optics for AI switch connectivity and QSFP-DD 800G optics for AI server connectivity to the next generation of 1.6T+ optics. Cisco co-designs silicon, systems, and optics, utilizing Cisco Linear Drive Pluggable Optics (LPO) to shift power-intensive DSP functions directly into the switch ASIC, significantly reducing power consumption.

Cisco is also the coherent market-share leader and pioneer in coherent pluggables, providing scale-across capabilities for connecting AI data centers across hundreds of kilometers.

Customer outcomes**Efficiency and Flexibility**

Unlock superior power efficiency and deployment flexibility across both liquid-cooled and air-cooled environments. By reducing optical module power consumption by 50% compared to traditional retimed modules, LPO enables more reliable and sustainable data center operations.

**High-Bandwidth AI Scaling**

Achieve massive scale with up to 1.6T switch-to-NIC connectivity. Our high-density fabrics are designed to support the largest AI cluster builds, delivering the robust reliability and consistent performance necessary for next-generation workloads.

**Proactive stability**

Simplify optics fleet management using advanced telemetry, diagnostics, and firmware updates supported by a diversified supply chain. Faster identification of marginal links prevents costly downtime, keeping your mission-critical AI workloads stable, predictable, and protected from physical-layer instability.

Sources: [Cisco Optics](#) | [Cisco Introduces Optical Innovations to Power the Backbone for AI Networking](#)

Differentiator 14

Power efficiency and management

Cisco Nexus One advantage

Power demand is rising enormously and, with 800G/1.6T links that multiply optics and thermal load, AI networking must deliver significantly more bandwidth per watt.

Cisco co-designs silicon, systems, and optics and utilizes [Cisco Linear Drive Pluggable Optics \(LPO\)](#) to shift power-intensive DSP functions directly into the switch ASIC.

Drive physical layer efficiency by unifying visibility and lifecycle control across air- and liquid-cooled 102.4T systems with the [Cisco Nexus One](#) operating model.

Cisco Nexus One integrates seamlessly with Panduit Intelligent Power Distribution Units (iPDUs) to bridge the gap between IT operations and facility management. This unified solution delivers granular, real-time power monitoring and emissions tracking across your entire data-center environment, from fabric-wide trends down to the individual device outlet.

Customer outcomes



Operational longevity

Lower your cooling costs, increase rack density, accelerate your progress toward net-zero initiatives, and extend the lifespan of your data center facilities.




Mission-critical reliability

Beyond power savings, our rigorous qualification—comprising over 1600 test sequences and 10,000+ Bit Error Rate (BER) hours—helps ensure the reliability that is needed to protect long-running AI jobs from physical-layer failure.



Superior energy efficiency

Cisco 800G LPO delivers a [30% reduction in system power](#) compared to traditional retimed optics. Additionally, it replaces six 51.2T systems with a single high-radix; the 102.4T liquid-cooled G300 system drives an estimated 70% improvement in energy efficiency.



Bridged IT and DC facilities

With granular visibility at the site, infrastructure, or PDU-outlet level, you can proactively identify power anomalies, forecast future energy requirements, and make data-driven decisions to optimize rack density and hardware lifecycles.

Sources: [Build data-center fabrics—quickly, easily, anywhere](#) | [Energy Optimized and Scalable Breakout Architecture for your Data Centers](#) | [Cisco Silicon One G300: The Next Wave of AI Innovation](#) | [Networking for the Agentic Era: Cisco Unveils New Innovations in Scale and Simplicity](#) | [More Scale, More Intelligence, and More Control: New Cisco Solutions for Accelerating AI Networking](#)

Differentiator 15

World-class support with Cisco Professional Services

Cisco Nexus One advantage

Unmatched global expertise, AI-driven support, and a vast logistics network deliver rapid, predictive, and comprehensive technical assistance.

[Cisco Support and Professional Services](#) teams act as trusted partners, providing proactive advisory services, lifecycle management, and strategic guidance to help align infrastructure with specific business goals.

[Cisco IQ](#) simplifies environment management by centralizing asset visibility and surfacing smarter, actionable insights. Through the Cisco IQ AI Assistant, teams shift from reactive to proactive, streamlining case management and decision-making to focus on driving better business outcomes.

Customer outcomes**Proven reliability**

Cisco puts the customer first. Cisco Support powerhouse manages 1.4M+ service requests annually. As a result, customers benefit from faster issue resolution, minimized downtime, and enhanced operational efficiency, ensuring business continuity and satisfaction.

**Accelerated success**

[Cisco Professional Services](#) enables up to 98% automated change success, 5x faster implementation, and access to a global team of more than 9500 certified engineers dedicated to your unique goals.

**Security resilience**

Achieve data sovereignty in air-gapped environments while accelerating crypto-agility to mitigate quantum risks with AI-powered Cisco IQ.

By leveraging secure peer benchmarking, organizations prioritize security investments and strengthen defenses against industry-specific threats.

Sources: [Cisco Services](#) | [Cisco Professional Services - Turn IT strategy into business value](#) | [Resilient infrastructure starts with Cisco IQ](#)

Conclusion

Success in the agentic era is measured by economic efficiency. Cisco Nexus One transforms your network into a strategic competitive moat by unifying fabric performance, inference operations, multitenancy, security, and lifecycle management into a single, secure, and co-designed full-stack architecture.

By owning the entire innovation stack—from silicon, systems, optics, software to security and observability at every layer—we have built a cohesive platform advantage where every component is engineered to enhance the whole. This eliminates the friction that drives up costs and operational complexity.

Furthermore, Cisco preserves your freedom through a standards-based, open ecosystem, helping to ensure that your interests remain protected as you scale. This approach optimizes your tokenomics and improves efficiencies, delivering a predictable total cost of ownership (TCO).

With Cisco Nexus One, you gain the production deployment velocity and resilient, scalable foundation required to turn AI potential into sustained, profitable innovation.

