Navigating Neocloud

Security Risks and Compliance in the AI Era

© 2026 Cisco and/or its affiliates. All rights reserved.

Contents

As artificial intelligence moves from experimental labs to the core of global enterprise strategy, the infrastructure supporting it must evolve. This solution brief, Navigating Neocloud, explores the emergence of specialized, AI-first cloud providers and the unique security landscape they create. We journey from the architectural differences between Neoclouds and Hyperscalers to the specific risks inherent in high-speed Graphic Processing Unit (GPU) fabrics.

By introducing the ClusterMAX™ benchmark and detailing how Cisco's integrated security portfolio aids in achieving compliance, we provide a roadmap for organizations to achieve performance without compromising on the security, intellectual property protection, or regulatory compliance required.

This brief explores the architectural shift toward specialized GPU providers, identifies the unique risks to AI intellectual property, and demonstrates how Cisco's security portfolio enables high-performance compliance with ClusterMAX, SOC2, ISO 27001, and Cybersecurity Maturity Model Certification (CMMC).

© 2026 Cisco and/or its affiliates. All rights reserved.

What is a Neocloud?

Traditional cloud providers were built for the era of the website and the database. But the explosion of Generative AI has created a new, insatiable demand: massive, raw, parallel processing power. To meet this need, a new breed of service provider has emerged–the Neocloud.

Why Neoclouds are Emerging

The first Neoclouds didn't start in traditional data centers. They often emerged from high-performance niches like 3D rendering farms or specialized crypto-mining operations. As Large Language Models (LLMs) went mainstream, these providers pivoted their infrastructure to offer what Hyperscalers couldn't: immediate, at-scale access to dense clusters of high-throughput networks optimized specifically for the "heavy lifting" of AI training and inference.

Why Now?

The shift to Neocloud is driven by three primary market forces:

- **GPU Scarcity**: While traditional cloud struggle with supply chain constraints, Neocloud specializes in securing and deploying the latest GPU hardware (like H100s and A100s) as their primary business.
- **Architecture Bottlenecks**: Standard virtualized cloud networks are designed for general traffic, creating "bottlenecks" that slow down the massive data transfers required between GPUs and AI training.
- **The Cost of Scale**: For organizations training frontier models, the overhead and egress fees of general-purpose clouds can become prohibitively expensive compared to the streamlined, performance-focused pricing of a Neocloud.

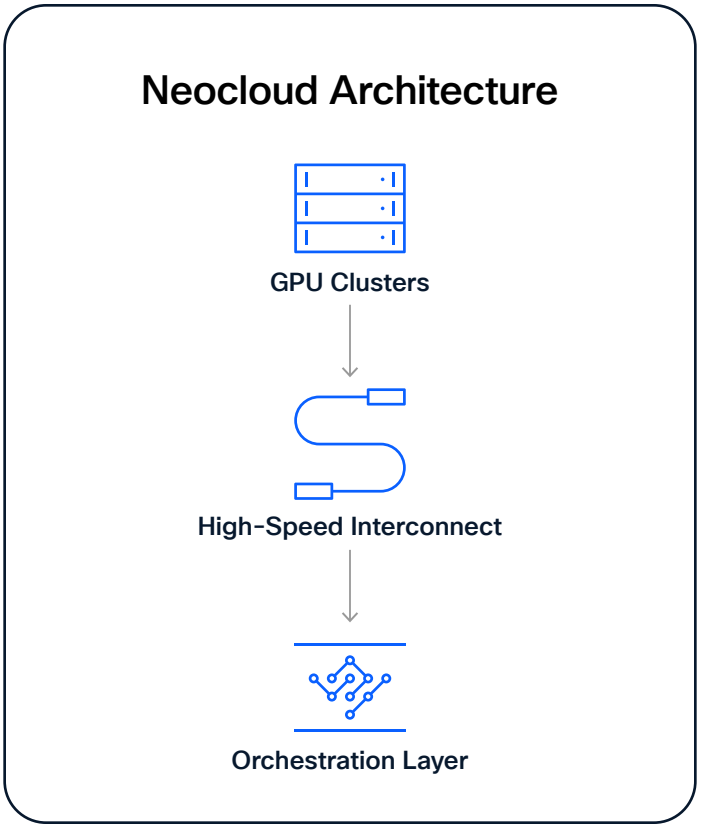

The Neocloud Advantage

Neoclouds offer a "supercomputer-as-a-service" experience that provides several distinct advantages:

- **Bare-Metal Performance:** By removing the "virtualization tax," Neoclouds allow AI models to run directly on the hardware, maximizing every clock cycle of the GPU.
- **Ultra-Low Latency Fabrics:** Neoclouds utilize high-bandwidth, low-latency fabrics to ensure thousands of GPUs can communicate at the speeds required for intensive model training and inference.
- **Dense Clustering:** Unlike general clouds that spread resources across a data center, Neoclouds provide high-density clusters (like ClusterMAX) where GPUs are physically and logically optimized for massive parallel workloads.

© 2026 Cisco and/or its affiliates. All rights reserved.

Architecture Matters: Neoclouds vs. General-Purpose Clouds

Not all clouds are created equal. While Hyperscalers built the modern internet, the massive compute requirements of Generative AI have exposed limitations of general-purpose architecture. Choosing the right platform depends on whether you need a "swiss army knife" or a "precision instrument".

The Hyperscaler Gap

Hyperscalers (AWS, Azure, GCP) dominate general-purpose cloud services but prioritize breadth over AI optimization.

- **Optimization:** They support broad workloads; Neoclouds are purpose-built for AI.
- **Provisioning:** GPU access can take weeks at scale; Neoclouds offer near-instant availability.
- **Networking:** Hyperscalers often rely on standard Ethernet; Neoclouds deploy high-speed, high-throughput net

The Traditional Cloud Limit

Standard cloud providers focus on virtualizing CPU and memory for web apps and storage.

- **Performance:** Traditional clouds rely on virtualization layers; Neoclouds deliver bare-metal GPU performance with specialized cooling and orchestration.
- **Use Cases:** Traditional clouds power SaaS and enterprise apps; Neoclouds enable massive-scale AI training, inference, and scientific computing.

Table 1. Neocloud vs. Hyperscaler vs. Regular Cloud Provider

Feature/Aspect	Neocloud	Hyperscaler (AWS, Azure, GCP)	Regular Cloud
Primary Focus	AI / HPC Workloads	General Purpose / SaaS	Web Apps / Storage
Compute Type	Bare-metal GPU	Virtualized GPU / CPU	Virtualized CPU
Networking	High-speed, high throughput network	Standard Ethernet	Standard Ethernet
Security	Minimal Security For advanced security, contact Cisco to tailor a solution. tailor a solution.	Standard L3/L4	Standard L3/L4
GPU Availability	Instant / High Density	Restricted / Shared	Limited to None
Performance Tax	Zero (Bare-Metal)	High (Virtualization)	High (Virtualization)

© 2026 Cisco and/or its affiliates. All rights reserved.

The Shared Responsibility Model

In the Neocloud world, the “shared responsibility model” looks different than it does in traditional cloud environments. Because Neoclouds provide high-performance, bare-metal access to GPUs, the customer inherits a much larger security footprint.

Understanding this boundary is critical: The provider secures the infrastructure, but you must secure the innovation.

The Neocloud Boundary

- **The Provider’s Responsibility:** They are responsible for “security OF the cloud.” This includes physical data center security, power, cooling, and the health of the physical GPU hardware and the underlying network fabric.
- **The Customer’s (Your) Responsibility:** You are responsible for “security IN the cloud.” Because you have direct access to the hardware, you own the security of the operating system, the GPU drivers, the data, and—most importantly—the AI Model Weights and Intellectual Property.

How Cisco Fills the Gap

Cisco’s security portfolio is designed to help you manage your side of the shared responsibility model without sacrificing the performance you came to Neocloud for.

- **You own the Workload:** Cisco Secure Workload provides the visibility and microsegmentation the provider doesn’t.
- **You own the Traffic:** Cisco Hypershield and Secure Firewall protect the data flows that the provider simply carries.
- **You own the Access:** Cisco Duo and Identity Services Engine (ISE) ensure that only authorized researchers and admins can touch your expensive GPU clusters.

© 2026 Cisco and/or its affiliates. All rights reserved.

The New Frontier of Threat: Unique Risks in Neocloud

While Neoclouds offer unparalleled performance, their specialized architecture creates a different kind of attack surface. Traditional security tools often suffer from "blind spots" in these high-speed, bare-metal environments.

The Isolation Challenge: Lateral Movement

In dense GPU clusters, the high-speed backend fabric is designed for performance, not necessarily for security.

- **Cross-Tenant Leakage:** Without robust microsegmentation, there is a risk of traffic "leaking" between customers sharing the same physical cluster.
- **Silent Lateral Movement:** Attackers can move between GPU nodes across the high-speed fabric, often undetected by traditional OS-level monitoring.

Infrastructure Blind Spots: Fabric Persistence

Because Neoclouds provide "close to the metal" access, the underlying hardware components become primary targets.

- **Vulnerable Nodes & Drivers:** Unpatched GPU drivers or firmware can become persistent entry points for exploits that survive even if a container is deleted.
- **Orchestration Vulnerabilities:** Insecure Kubernetes configurations or unpatched images can lead to cluster-wide compromises, giving attackers access to the raw compute power.

Identity & API Vulnerability: The Management Plane

Neoclouds are heavily driven by APIs and management consoles to orchestrate massive AI jobs.

- **Over-Privileged Access:** API keys and service accounts often have excessive permissions, allowing a single compromised credential to take over an entire GPU fleet.
- **Console Abuse:** Unauthorized access to management planes can lead to "Job Abuse," where attackers hijack your expensive compute for their own purposes (like crypto-mining or unauthorized model training).

The "Crown Jewel" Threat: Model & Data Integrity

In the AI era, your most valuable asset is your model. Neoclouds introduce specific risks to this Intellectual Property.

- **Model Theft & Exfiltration:** High-speed fabrics allow for the rapid, silent exfiltration of proprietary model weights and training data.
- **Adversarial Manipulation:** If the infrastructure is tampered with, training data can be "poisoned," or model inference can be manipulated, leading to unreliable AI outputs.

The Compliance Hurdle

Many Neocloud environments lack the native logging and auditability required for regulated industries.

- **Audit Gaps:** Difficulty in hosting financial, healthcare, or public sector workloads due to a lack of controls aligned with SOC2, ISO 27001, or CMMC.

© 2026 Cisco and/or its affiliates. All rights reserved.

ClusterMAX: Benchmarking the AI Cloud

As Neocloud adoption accelerates, organizations need a standardized way to evaluate providers for security, reliability, and operational maturity. ClusterMAX, developed by SemiAnalysis, has emerged as the industry benchmark for GPU cloud environments, rating providers across five critical dimensions:

- **Security** – Isolation guarantees, GPU sanitization, pentesting rigor.
- **Orchestration** – Kubernetes support, GPU operators, SLURM, observability.
- **Reliability** – Uptime, failure resilience, hardware lifecycle management.
- **Storage** – High-performance storage integration for AI workloads.
- **GPU Availability** – Access speed and cluster scalability.

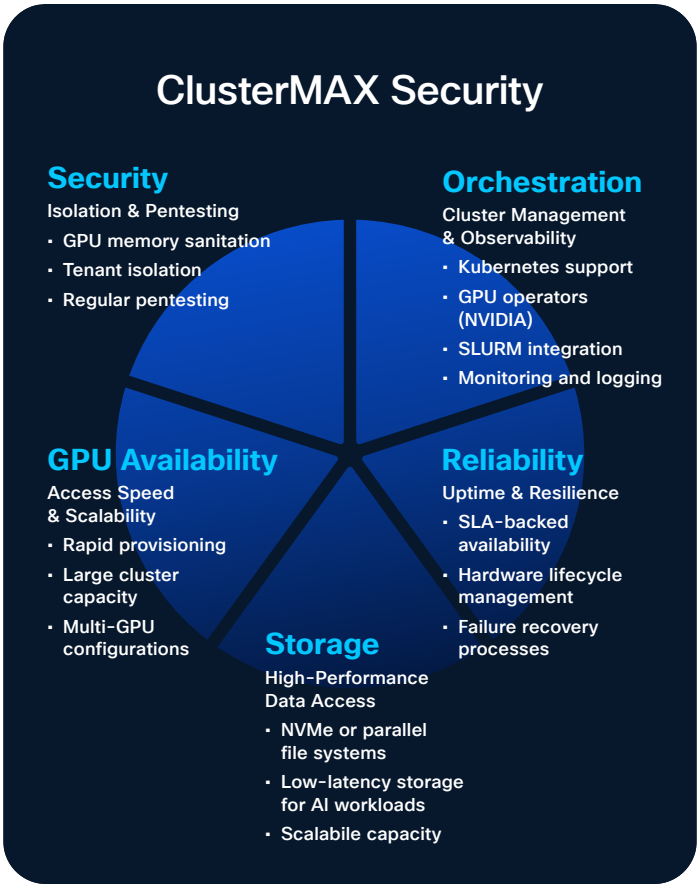

The 5 Pillars of ClusterMAX Security

To achieve a top-tier rating, a Neocloud provider must meet five rigorous security requirements that address the unique risks of AI infrastructure:

- **Tenant Isolation:** Secure segmentation and backend network separation.
- **Node Hardening:** Enforced patching for GPU nodes, drivers, and firmware.
- **Identity Control:** Per-tenant RBAC for management planes and GPUs.
- **Deep Observability:** Telemetry for East-West traffic and job behaviors.
- **Auditability:** Compliance-aligned logging (SOC2, ISO, CMMC).

Why Benchmarking Matters

ClusterMAX helps organizations move beyond “marketing promises” to objective validation:

- **Objective Comparison:** Compare providers using a standardized scoring system.
- **Risk Reduction:** Select providers with proven isolation and sanitization protocols.
- **Compliance Readiness:** Ensure the underlying infrastructure can support regulated AI customers (Financial, Healthcare, Public Sector).

© 2026 Cisco and/or its affiliates. All rights reserved.

The Cisco Unified Defense for Neocloud

As AI workloads redefine the boundaries of compute, traditional security models must evolve to protect the specialized, high-performance architecture of the Neocloud. Cisco's unified defense for Neocloud provides a performance-first, integrated security framework designed specifically for GPU-accelerated environments. Rather than treating security as a bolt-on, this architecture embeds# protection across four critical layers: the Management Plane, Perimeter Layer, Fabric Layer, and Workload Layer.

By unifying identity, infrastructure, and deep observability, Cisco ensures that security never becomes a bottleneck for AI innovation. Backed by the global intelligence of Cisco Talos, this layered approach mitigates unique risks such as cross-tenant lateral movement and model integrity attacks. Whether enforcing Zero Trust access for administrators or providing eBPF-powered isolation for massive GPU clusters, Cisco delivers the technical controls and auditable evidence required to pass stringent compliance benchmarks like ISO 27001, SOC 2, and CMMC—all without compromising the ultra-low latency that production-grade AI demands.

- **Cisco Secure Firewall**
  Cisco Secure Firewall delivers high-performance macrosegmentation and advanced threat defense. It provides "virtual patching" for vulnerable systems and inspects encrypted traffic without introducing performance-killing latency, leveraging the Encrypted Visibility Engine (EVE) to detect threats in encrypted traffic without full decryption.

The Layered Neocloud Stack

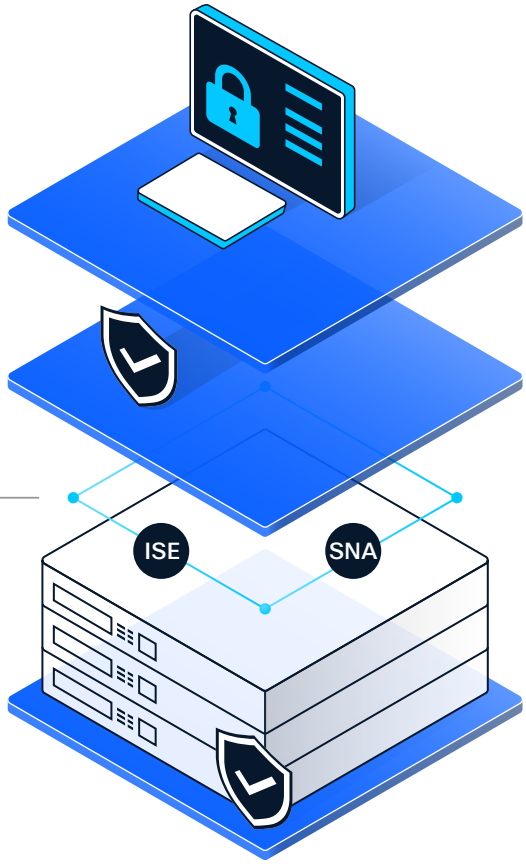

Management Plane
Unified policy management and verified identity access.
- Security Cloud Control
- Secure Access
- Duo

Fabric Layer
Identity-based segmentation and telemetry across the high-speed fabric.
- Identity Services Engine (ISE)
- Secure Network Analytics (SNA)

Perimeter Layer
Foundational macrosegementation and threat defense at the network.
- Secure Firewall

Workload Layer
Kernel-level protection and microsegmentation.
- Secure Workload
- Isovalent Enterprise Platform

© 2026 Cisco and/or its affiliates. All rights reserved.

- **Cisco Secure Workload**
  Cisco Secure Workload automates microsegmentation and least-privilege access across multicloud environments. It provides deep visibility into application behavior to secure East-West traffic and prevent lateral movement within dense GPU clusters.
- **Isovalent Enterprise Platform**
  Isovalent Enterprise Platform offers scalable, high-performance networking and security tailored for dynamic cloud-native infrastructure. It enhances visibility and threat defense across hybrid and multicloud deployments, providing identity-based security, zero-trust networking, and efficient observability for Kubernetes and cloud environments without compromising performance
- **Cisco Security Cloud Control**
  Cisco Security Cloud Control is a unified management platform that centralizes visibility and policy enforcement. It unifies threat response and ensures a consistent security posture across all Neocloud infrastructure.
- **Secure Access**
  Cisco Secure Access is a comprehensive Zero Trust Network Access (ZTNA) solution that secures access across applications and environments for any user, device, and location. It consolidates multiple security functions such as Zero Trust Network Access, Secure Web Gateway, Cloud Access Security Broker, cloud firewall, DNS-layer security, and more into a single cloud-managed service. This solution simplifies user experience, enhances security with granular access controls, and streamlines IT operations with a unified management console. It is designed to protect hybrid workforces and supports secure access to SaaS, private apps, and the internet regardless of user location or device.
- **Cisco Secure DDoS Protection**
  Cisco Secure DDoS Protection provides automated, real-time mitigation of volumetric and application-layer attacks at the network edge. It ensures the continuous availability of GPU infrastructure and management APIs, preventing costly service disruptions during intensive AI training and inference cycles.
- **Cisco Identity Services Engine (ISE)**
  Cisco ISE delivers centralized, identity-based access control and policy enforcement across the physical and virtual infrastructure. It ensures that only authorized administrators and compliant devices can access critical hardware, providing the foundational visibility and control needed to meet stringent enterprise security and audit requirements.
- **Cisco Duo**
  Cisco Duo enforces strong identity verification and device health checks across all administrative and developer access points. It secures the Neocloud management plane against credential-based attacks, ensuring that only trusted users and compliant devices can manage high-value GPU resources and AI APIs.
- **Cisco Secure Network Analytics (SNA)**
  Cisco SNA uses machine learning and telemetry to detect “silent” threats like data exfiltration and lateral movement. It identifies hidden threats in encrypted traffic without requiring decryption, maintaining both security and privacy.

Cisco’s unified framework helps meet ClusterMAX standards by integrating protection across the network fabric, workloads, and management plane. These ten solutions provide a foundational, layered defense that secures AI infrastructure at scale, ensuring your Neocloud remains high-performing, secure, and compliant.

© 2026 Cisco and/or its affiliates. All rights reserved.

Mapping Security to Compliance

For organizations operating in regulated industries, performance is only half the equation; compliance is the other. As AI moves from the lab to production, Neocloud environments must prove they can protect sensitive data under the world's most rigorous regulatory frameworks.

Cisco's security architecture is designed not just to defend, but to provide the granular evidence and enforcement required to meet these global standards. The following sections map our seven core security solutions to the specific requirements of SOC 2, ISO 27001, and CMMC, providing a blueprint for building a compliant and trusted AI infrastructure.

SOC2: Establishing Trust in the AI Cloud

SOC2 is the industry benchmark for demonstrating that a service organization has the necessary controls to protect customer data. In a Neocloud environment, the ClusterMAX™ requirements for Multi-Tenant Isolation and Deep Observability are the technical bedrocks of the SOC 2 "Security" and "Confidentiality" criteria.

By implementing Cisco's layered defense, organizations can prove to auditors that East-West traffic is strictly isolated and that every GPU job is monitored for abnormal behavior. This mapping illustrates how Cisco provides the granular auditability and protection required to host sensitive AI workloads while meeting all five Trust Services Criteria.

Table 2. SOC2 Compliance for Neocloud

AICPA SOC 2	Secure Firewall	Secure Workload	Isovalent Enterprise Platform	Security Cloud Control	Secure Access	Secure DDoS Protection	Secure Network Analytics (SNA)	Duo	Identity Services Engine (ISE)
Security	Enforces network policies, prevents attacks	Enforces policies isolates compromised workloads	Enforces microsegmentation, runtime security	Enforces secure posture, detects vulnerabilities	Enforces unified zero-trust security	Blocks sophisticated volumetric attacks	Detects threats, anomalies, and forensics	Enforces MFA, device trust, access	Enforces micro/macro segmentation, controls access
Availability	Ensures uptime via threat prevention	Ensures authorized application execution	Ensures network resilience, observability	Prevents exploits, ensures system stability	Ensures resilient, security connectivity	Ensures maximum service uptime	Detects availability threats, enables response	Prevents unauthorized access, disruptions	Limits incident impact, ensures uptime
Processing Integrity	Validates traffic, prevents data corruption	Ensures authorized application execution	Protects data, process integrity	Secures nodes, prevents data corruption	Validates secure, authorized access	Maintains legitimate traffic flow	Detects process anomalies, unauthorized changes	Ensures authorized access, valid actions	Secures processing paths, prevents interference
Confidentiality	Prevents data leaks, secures traffic	Isolates sensitive data, prevents leakage	Enforces encryption, egress controls	Protects nodes, prevents data access	Prevents unauthorized data disclosure	Protects data-hosting infrastructure	Detects data exfiltration, suspicious movement	Controls access to sensitive data	Controls confidential data access
Privacy	Protects PII from network breaches	Isolates PII, prevents unauthorized access	Restricts access to personal data	Secures nodes, protects personal data	Protects user, data privacy	Safeguards data access points	Detects PII exfiltration, suspicious access	Protects PII via strict access	Secures PII access, isolates data

Green = Key solution core to the requirement

© 2026 Cisco and/or its affiliates. All rights reserved.

ISO 27001: Global Standards for Information Security

ISO 27001 provides a globally recognized, risk-based framework for managing information security. The ClusterMAX pillars of Infrastructure Hardening and Node Lifecycle Management align directly with ISO 27001's focus on operational security and system integrity. Cisco solutions support these requirements by automating "virtual patching" and enforcing identity-based access across the GPU fabric. This table demonstrates how Cisco's architecture maps to key Annex A controls, ensuring that high-performance GPU clusters are operated within a mature, internationally standardized security management system that prioritizes both performance and risk mitigation.

Table 3. ISO 27001 Compliance for Neocloud

	Secure Firewall	Secure Workload	Isovalent Enterprise Platform	Security Cloud Control	Secure Access	Secure DDoS Protection	Secure Network Analytics (SNA)	Duo	Identity Services Engine (ISE)
A.9 Access Control	Enforces network-level access policies	Enforces workload access policies	Controls granular workload access	Enforces secure access policies	Enforces zero-trust identity access	n/a	Detects unauthorized access attempts	Enforces verified, granular access	Enforces strong identity-based access
A.10 Cryptography	Enforces encrypted tunnels and TLS inspection	Monitors workload encryption	Enforces high-speed encryption	Centralizes encryption management	Secures encrypted sessions	Protects encrypted traffic	Analyzes encrypted patterns	n/a	n/a
A.12 Operations Security	Provides logging and threat protection	Secures workload operations, policies	Enforces runtime security, posture	Enforces secure operational baselines	Secures cloud-based operational traffic	Ensures operational service resilience	Monitors operational network activity	Secures operational system access	Controls operational access, devices
A.13 Communications Security	Secures network boundaries and segmentation	Segments workload communications	Secures, segments network communications	Secures communication endpoints	Protects all data communications	Safeguards network communication availability	Detects system anomalies, misuse	Protects communication system access	Secures communications via segmentation
A.14 System Acquisition, Development & Maintenance	Blocks exploits during system lifecycle	Protects workload integrity, policies	Enforces policy-as-code for clusters	Centralizes infrastructure policy management	Secures development environment access	n/a	Monitors new system behavior	Secures build pipeline access	Enforces new node posture
A.16 Information Security Incident Management	Detects and blocks network threats	Provides deep workload forensic data	Provides real-time process-level forensics	Centralizes incident reporting and logs	Logs and detects access incidents	Migrates volumetric security incidents	Detects and analyzes network anomalies	Logs and reports authentication incidents	Logs and manages access violations
A.17 Business Continuity	Ensures resilient network security	Maintains workload security continuity	Ensures resilient, distributed networking	Centralizes recovery policy management	Ensures resilient remote access	Ensures maximum service availability	Detects availability-impacting anomalies	Ensures resilient identity services	Ensures resilient network authorization
A.18 Compliance	Logs network policy compliance	Validates host-level compliance	Provides auditable compliance evidence	Centgralizes unified compliance reporting	Logs zero trust compliance	Provides availability SLA compliance	Provides network hehavioral evidence	Provides strong authentication, audit	Enforces network policy compliance

Green = Key solution core to the requirement

© 2026 Cisco and/or its affiliates. All rights reserved.

CMMC: Securing the Defense Industrial Base

For organizations handling Controlled Unclassified Information (CUI) within the federal supply chain, CMMC compliance is a non-negotiable requirement. In a Neocloud, the GPU fabric itself must act as a hardened enclave. The ClusterMAX requirements for Secure Backend Networks and Per-Tenant RBAC are essential for meeting CMMC domains like Access Control and System & Communication Protection.

Cisco's security portfolio ensures that federal AI initiatives are protected by the same level of rigor required for traditional defense systems. This table highlights how our solutions provide the continuous monitoring and incident response capabilities necessary to maintain CMMC eligibility and secure the nation's most sensitive AI innovations.

Table 4. CMMC Compliance for Neocloud

CERTIFICATION	Secure Firewall	Secure Workload	Isovalent Enterprise Platform	Security Cloud Control	Secure Access	Secure DDoS Protection	Secure Network Analytics (SNA)	Duo	Identity Services Engine (ISE)
AC (Access Control)	Enforces network-level access policies	Enforces workload access policies	Controls granular workload access	Enforces secure posture, limits entry	Enforces zero-trust application access	Protects availability of access services	Detects unauthorized access attempts	Enforces granular, verified access	Enforces identity-based access policies
AU (Audit & Accountability)	Logs network traffic and events	Monitors workload behavior, provides logs	Logs network activity, violations	Monitors posture, logs security state	Logs all user/application activity	Logs attack and mitigation events	Provides detailed network telemetry, anomalies	Logs authentication, device posture	Logs access, policy enforcement decisions
CM (Configuration Management)	Manages secure network configuration policies	Ensures authorized workload configurations	Enforces secure workload configurations	Enforces baselines, detects drift	Enforces secure access policies	Maintains availability of management systems	Detects configuration drift, anomalies	Verifies device security posture	Provides access based on device posture
IA (Identification & Authentication)	Integrates identity-based network access	Protects workload identities, access	Enforces identity-based workload security	Secures nodes, protects credentials	Validates user/device identity strictly	Safeguards authentication service availability	Detects compromised identities, misuse	Provides strong MFA, identity verification	Provides strong identity, authentication
IR (Incident Response)	Detects threats, enables incident response	Isolates threats, contains breaches	Aids incident detection, forensics	Detects vulnerabilities, prevents exploits	Provides threat visibility and context	Automates volumetric incident mitigation	Detects incidents, provides forensics	Blocks compromised users, devices	Contains incidents, quarantines devices
RA (Risk Assessment)	Blocks exploits, segments networks	Scans vulnerabilities, enforces posture	Provides high-performance kernel isolation	Centralizes policy and reporting	Secures remote CUI access	Ensures critical service availability	Detects anomalous network behavior	Enforces multi-factor identity verification	Orchestrates identity-based network policy
SC (System & Communications Protections	Secures boundaries and network segments	Segments workload communications	Secures, segments network communications	Secures nodes, protects communications	Secures remote/cloud communication channels	Protects boundary communication availability	Monitors network for threats	Protects access to systems	Secures communications via segmentation
SI (System & Information Integrity)	Prevents malware and malicious traffic	Prevents unauthorized workload changes	Maintains system, information integrity	Detects vulnerabilities, enforces patching	Protects data from malicious threats	Prevents resource-exhaustion integrity failures	Detects threats, data exfiltration	Prevents untrusted access	Prevents compromised access

Green = Key solution core to the requirement

© 2026 Cisco and/or its affiliates. All rights reserved.

Building a Trusted AI Infrastructure

The shift to Neocloud represents more than just a change in hardware; it is a fundamental shift in how we build and scale intelligence. As AI models become the most valuable intellectual property on earth, the infrastructure that hosts them must be as resilient as it is powerful.

By aligning with the ClusterMAX standard and leveraging Cisco's security portfolio, organizations can move past the "security tax" and embrace a future where performance and protection are inextricably linked. Whether you are training frontier models or deploying real-time inference, the goal remains the same: to innovate with the confidence that your data, your models, and your reputation are secure.

Ready to Secure Your AI Future?

Navigating the complexities of Neocloud security and global compliance doesn't have to be a solo journey. Cisco is ready to help you architect a secure, high-performance foundation for your AI initiatives.

Take the Next Step:

- Work with our experts to identify gaps in your current GPU cloud architecture.
- Get a deep dive into the configurations for the Cisco AI-security suite.
- Contact your Cisco Account Team to discuss how we can help you meet CyberMAX standards and achieve your compliance goals.

Resources

For more information, please refer to the following:

- Explore fundamentals of AI Security
- See why Cisco is the critical infrastructure for the AI era
- Secure your AI transformation and power your cyber defenses

© 2026 Cisco and/or its affiliates. All rights reserved. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (P0405031518) 04/26

© 2026 Cisco and/or its affiliates. All rights reserved.