eclypsium /FEDERAL

TECHNICAL BLUEPRINT_

AI INFRASTRUCTURE ASSURANCE FOR MISSION CRITICAL SYSTEMS

Executive Summary_

This document details the necessary effort to implement the Eclypsium Infrastructure Assurance Platform to address critical hardware supply chain vulnerabilities, infrastructure integrity, and component-level security gaps within Department of War (DoW) information systems and high-performance computing AI environments.

Eclypsium is pleased to propose the Eclypsium Infrastructure Assurance Platform to provide proactive hardware supply chain security for critical AI hardware, firmware, and software infrastructure. The platform's key architectural features include:

- **Key Feature One:** Complete component-level visibility and automated generation of Software, Firmware, and Hardware Bills of Materials (SBOM/FBOM/HBOM) to maintain an uncompromised, verifiable inventory of all device components, including GPUs, specialized AI accelerators, and high-performance networking devices. This directly fulfills Presidential and DoW mandates for absolute supply chain transparency before hardware is cleared for active deployment in sensitive AI environments.
- **Key Feature Two:** Agentless discovery and analysis combined with a powerful REST API for flexible integration with existing SIEM, SOAR, and ITSM tools without causing operational disruption. This allows rapid security verification without altering critical mission code bases, consuming costly GPU cycles, or disrupting long running, resource intensive AI model training runs.
- **Key Feature Three:** Firmware monitoring capable of discovering hidden rootkits, bootkits, and backdoors that evade traditional security tools. The platform validates cryptographic integrity against a database of 15 million and growing known good firmware images, intercepting adversarial tampering directly at the hardware layer before threat actors can exfiltrate proprietary models, manipulate training data, or silently hijack compute resources.

Mandates & Policies Specific to AI Infrastructure

AI data center hubs, specialized clusters, and high-performance computing "AI factories" have been designated as national-level strategic assets and critical infrastructure. Due to high data volumes and the continuous leasing or swapping of raw compute nodes between different classified and unclassified customer workloads, these environments face severe threats from data poisoning, model weight leaks, and hardware-level exploitation. Eclypsium delivers the low-level visibility needed to defend these networks while satisfying current DoW AI mandates:

- **White House Executive Order (Advanced AI Innovation and Security)**
  The June 2, 2026 EO explicitly directs the DoW, the NSA, and CISA to prioritize the cyber defense of National Security Systems hosting advanced AI. It also establishes a framework for evaluating "covered frontier models" via a classified benchmarking process, requiring early secure access for trusted federal partners. Eclypsium hardens the physical infrastructure hosting these evaluations, ensuring that the underlying servers, GPUs, and network switches are free from low-level vulnerabilities.

- **National Security Presidential Memorandum 11 (NSPM-11)**
  Signed on June 5, 2026, NSPM-11 orders the national security enterprise to accelerate secure AI adoption while embedding strict assurance and accountability infrastructure. A primary directive of NSPM-11 is ensuring that no commercial entity or foreign adversary can unauthorizedly disable, degrade, or modify a fielded system that warfighters depend on. Eclypsium provides the independent hardware-level verification needed to guarantee that underlying systems cannot be covertly altered or disabled via firmware manipulation.

- **FY 2026 NDAA (Section 1513) Compliance**
  Congress has mandated a strict security procurement framework for all acquired AI/ML technologies, focusing specifically on mitigating supply chain vulnerabilities, counterfeit parts, data poisoning, and adversarial tampering. The law targets covered AI/ML technologies, including source code, software, and the underlying hardware/firmware methods used to develop or host AI. Eclypsium satisfies this mandate by continuously scanning and cryptographically verifying the foundational layers of the compute nodes running those models.

- **NIST SP 800-223 Standards**
  Eclypsium fulfills the specialized High-Performance Computing (HPC) guidelines detailed in NIST SP 800-223, specifically providing automated mechanisms to solve Compute Node Sanitization challenges between task runs on shared infrastructure.

AI Datacenter Architectural Domain Mapping & Attack Surface Matrix

The AI Datacenter Architectural Domain Mapping & Attack Surface Matrix provides a structured representation of modern high-performance computing (HPC) and 'AI factory' infrastructure. Securing a modern AI data center requires continuous protection across the entire technology stack—from the physical facility at the bottom to the application models at the top. Every layer depends on the integrity of the one beneath it. Because these environments dynamically reallocate high-performance compute nodes, security gaps can emerge anywhere in the stack. However, if an adversary compromises the foundational hardware or firmware layer, higher-level controls like zero-trust policies and MLOps guardrails are rendered completely ineffective. Defending this complex, highly dynamic architecture requires addressing several key technical, operational, and policy considerations:

Cross-Domain Interdependencies & Infrastructure Blind Spots

Trust in the application and model layers ultimately depends on the integrity of the underlying hardware, firmware, and management plane. If a BMC, accelerator firmware, or similar low-level component is compromised, higher-layer controls such as MLOps governance, runtime monitoring, and Zero Trust policy enforcement may be weakened or bypassed. That is why AI factories need a Domain Specialized Exposure Management (DSEM) approach that can validate low-level device integrity before workloads run.

Multi-Tenant Dynamics & Compute Node Sanitization

Unlike traditional enterprise data centers, modern "AI factories" operate under dynamic, high-throughput multi-tenant paradigms. High-performance compute nodes are continuously provisioned, partitioned, and swapped between different mission teams, contractors, and workload classifications. To maintain complete multi-tenant isolation, hardware must be automatically scanned and verified immediately following every AI session. Without post-execution hardware scanning and automated sanitization, residual data, altered microcode, or stealthy firmware implants from a previous workload can persist in non-volatile memory—exposing the next task to cross-tenant contamination and silent persistence.

Threat Actor Profiles & Adversarial Tactics in AI Infrastructure

Modern high-performance computing (HPC) environments and 'AI factories' represent high-value targets for nation-state advanced persistent threats (APTs) and sophisticated cyber adversaries. Rather than targeting visible software applications protected by traditional Endpoint Detection and Response (EDR) tools, adversaries increasingly focus on hardware and infrastructure vulnerabilities. By implanting malicious firmware into Baseboard Management Controllers (BMCs), Option ROMs, or GPU accelerator cards, threat actors establish persistent foothold positions below the operating system layer. This low-level persistence enables adversaries to bypass host security, silently eavesdrop on inter-node GPU communications, manipulate training parameters, or exfiltrate sensitive model weights without triggering high-level security alerts.

Extending Zero Trust to the Silicon & Firmware Layers

Traditional Zero Trust Architecture (ZTA) focuses heavily on user identity and network perimeters—implicitly trusting the underlying physical hardware. In dynamic AI factories where compute nodes are continuously reallocated across classified and multi-tenant workloads, this implicit trust creates a critical vulnerability.

True defense requires extending Zero Trust principles down to the silicon. By establishing cryptographic component identity, continuously monitoring for hardware and firmware configuration drift, and generating automated HBOM/FBOMs, unverified compute nodes are prevented from joining the trust fabric or processing classified workloads. This physical-layer enforcement directly satisfies core capabilities within the **Device/Endpoint** and **Data** pillars of the DoW Zero Trust Strategy.

Strategic Policy Alignment & Mandate Fulfillment

This matrix supports federal AI security objectives by providing a defense-in-depth structure for hardware provenance, anti-tampering, and operational assurance. It also creates a practical foundation for implementing DoW AI security requirements, NIST SP 800-193 Platform Firmware Resiliency Guidelines, NIST SP 800-223 HPC guidance, and related White House and DoW policy directives across the full stack.

The Vertical Pillars: Policy Enforcement & Unified Governance

As illustrated in **Figure 1**, the AI architecture consists of horizontal functional layers anchored by two cross-cutting vertical pillars:

- **Policy Enforcement:** Acts as a continuous validation boundary across the stack, inspecting and verifying data and operational transitions as they move between physical, network, and orchestration layers to block lateral threat movement.

- **Management & Unified Governance:** Consolidates administrative oversight by unifying user identity (IAM), security logging (SOC/SIEM), and compliance auditing into a single control plane for enterprise-wide visibility.

AI Workload	Frontier LLM Training	Copilots & RAG Apps	Agentic Systems	Policy Enforcement	**Management & Unified Governance**
Model Serving	Inference Endpoints	RAG & Embeddings	Runtime Model Guardrails		**Users & Credentials**
MLOPS Platform	Workbenches & Tracking	Pipeline Governance	Model Security Assurance		
Orchestration	Cluster Schedulers	GPU Partitioning	Multi-tenancy & Quotas		**Security Operations (SOC)**
Data Layer	Lakehouse & Stores	Vector DBs & Catalogs	Ingestion & Lineage		
Storage & Memory	High-Speed Memory	Parallel File Systems	NVMe & Storage Arrays		**Mandate Alignment**
Compute & Accelerators	Rack-Scale AI Platforms	Server Components	Hardware Supply Chain		
Network Fabric	Back-End CPU Fabric	Front-End Fabric	Network Acceleration		**Infrastructure Mgmt**
Edge & Interconnect	Perimeter & OOB Gateways	Network Perimeter	Cross-Domain Controls		
Facility, Physical Security & Cooling	Physical Security & Racks	Power Infrastructure	Liquid Cooling & BMS		

Figure 1 - AI Architecture Stack

Facility, Physical Security & Cooling:

- **Overview & Function:** The underlying physical and industrial infrastructure—comprising Operational Technology (OT), Building Management Systems (BMS), and physical perimeter controls—that supplies power, environmental liquid cooling, and site security to high-density GPU clusters.
- **Key Assets to Protect:** Liquid cooling loops & SCADA/BMS controllers, smart Power Distribution Units (PDUs), Physical Access Control Systems (PACS), IP security cameras, and local console/KVM access.
- **Primary Threat Vectors:** Cyber-attacks targeting BMS/SCADA protocols to alter cooling/power parameters, unauthorized physical access to server racks, crash-cart port exploitation, and IP camera or badge reader spoofing.
- **Potential Operational Impact:** Instant thermal shutdown or physical damage to multi-million-dollar GPU arrays, physical drive theft, and unmonitored hardware implant insertion.
- **Recommended Safeguards:** Deploy encrypted OSDP badge readers, door contact sensors with tamper-evident locks, isolated SCADA/BMS VLANs, and mandatory 802.1X port authentication on physical endpoints. Conduct continuous firmware assessments, vulnerability scanning, and integrity auditing across OT assets, PLCs, IP security cameras, smart PDUs, and connected building controllers (BMS/SCADA).
- **How Eclypsium Helps:** Assess the IP cameras for firmware vulnerabilities and supply chain integrity.

Edge & Interconnect:

- **Overview & Function:** Governs external perimeter defenses, dedicated administrative management routes, and secure cross-domain data transfers entering and exiting the AI data center.
- **Key Assets to Protect:** Perimeter Next-Gen Firewalls (NGFW), Zero Trust Network Access (ZTNA) gateways, Out-of-Band (OOB) hardware management routes (IP-KVM/serial consoles), and Cross-Domain Solutions (CDS)/ hardware data diodes.
- **Primary Threat Vectors:** Unauthenticated IP-KVM console hijacking, boundary firewall bypasses, volumetric Distributed Denial of Service (DDoS) attacks disrupting node clusters, and unauthorized data leakage during cross-domain transfers.
- **Potential Operational Impact:** Total loss of administrative control plane, unauthorized remote access to compute hardware, and exfiltration of classified AI models or training data across domain boundaries.
- **Recommended Safeguards:** Implement hardware data diodes for strict one-way cross-domain transfers, isolated Out-of-Band (OOB) administrative networks hardened by disabling legacy IPMI in favor of TLS-encrypted Redfish REST APIs, ZTNA perimeter gateways, and mandatory multi-factor authentication for console connections. Implement continuous firmware inspection and vulnerability management across perimeter firewalls, routers, switches, Cross-Domain Solutions (CDS), and Out-of-Band management controllers.
- **How Eclypsium Helps:** Provides continuous firmware inspection and vulnerability management across perimeter firewalls, routers, switches, Cross-Domain Solutions, and Out-of-Band management controllers.

Network Fabric:

- **Overview & Function:** Delivers the ultra-low-latency, high-bandwidth interconnects that link GPU and storage clusters, enabling real-time GPU-to-GPU data exchanges during distributed AI training and inference.
- **Key Assets to Protect:** High-speed Back-End GPU Fabric (InfiniBand, RoCE v2), Front-End Ethernet switches, and Data Processing Units / SmartNICs (e.g., NVIDIA BlueField DPUs).
- **Primary Threat Vectors:** Inter-node GPU side-channel traffic sniffing, unauthorized Remote Direct Memory Access (RDMA) packet injection, memory buffer exploitation over RoCE, and malicious firmware implants on DPUs/SmartNICs.
- **Potential Operational Impact:** In-flight interception of proprietary model weights, silent corruption of gradient exchanges (training data poisoning), and persistent, below-OS footholds inside network acceleration hardware.
- **Recommended Safeguards:** Enforce cryptographic line-rate encryption for RDMA/RoCE traffic where supported, alongside strict physical and logical back-end fabric isolation. Implement continuous firmware monitoring, supply chain verification, and anti-drift detection across Data Processing Units (DPUs/SmartNICs like NVIDIA BlueField), back-end fabric switches, and high-speed routers.
- **How Eclypsium Helps:** Delivers continuous firmware monitoring, supply chain verification, and anti-drift detection across Data Processing Units (DPUs/SmartNICs like NVIDIA BlueField), back-end fabric switches, and high-speed routers.

Compute & Accelerators:

- **Overview & Function:** Represents the core hardware processing engine of the AI factory, housing bare-metal server silicon, specialized AI accelerators, and low-level system management controllers.

- **Key Assets to Protect:** High-density GPU/TPU accelerator clusters (e.g., NVIDIA Blackwell/HGX), server host CPUs, motherboard BIOS/Option ROMs, and Baseboard Management Controllers (Dell iDRAC, HPE iLO, Supermicro/ATEN).
- **Primary Threat Vectors:** Firmware implants and rootkits in BMCs or GPUs, persistent bootkits in Option ROMs, unpatched GPU host drivers, and counterfeit supply chain hardware.
- **Potential Operational Impact:** Silent exfiltration of proprietary model weights directly from GPU memory, covert training data corruption, and total EDR/OS security bypass via persistent low-level footholds.
- **Recommended Safeguards:** Enforce GPU Confidential Computing and hardware-backed Trusted Execution Environment (TEE) attestation to keep model weights and prompt contexts encrypted in memory during active execution. Implement automated hardware scanning and cryptographic integrity verification of GPUs, host CPUs, BMCs (Dell iDRAC, HPE iLO, Supermicro), and server components immediately following every AI compute session or task run—validating binary hashes against 15M+ known-good firmware repositories, generating real-time xBOMs (SBOM/FBOM/HBOM), and providing continuous low-level anti-drift monitoring.
- **How Eclypsium Helps:** Automates hardware scanning and integrity verification of GPUs, host CPUs, BMCs and server components, while providing continuous low-level anti-drift monitoring.

Storage & Memory:

- **Overview & Function:** Delivers high-throughput data staging and persistent storage for massive training datasets, live compute memory, and saved model checkpoints across high-performance computing clusters.
- **Key Assets to Protect:** Volatile GPU High-Bandwidth Memory (HBM3e/DDR5), Parallel File Systems (Lustre, GPFS), and high-speed NVMe storage arrays containing model checkpoints.
- **Primary Threat Vectors:** Exposure of unencrypted checkpoint files, multi-tenant memory bleed during node swapping, residual memory crypto-erase failures, and physical NVMe drive theft or diversion.
- **Potential Operational Impact:** Direct theft of proprietary model checkpoints, cross-tenant data contamination between mission runs, and loss of training dataset integrity.
- **Recommended Safeguards:** FDeploy FIPS 140-3 validated hardware-backed encryption for storage media, fine-grained access controls for scale-out parallel file systems, and automated cryptographic erasure of residual memory artifacts. Audit storage controller and NVMe drive firmware posture, validating post-execution node sanitization workflows to ensure hardware and firmware layers are clean and compliant with NIST SP 800-223 HPC guidelines.
- **How Eclypsium Helps:** Audits storage controller and NVMe drive firmware posture, while validating post-execution node sanitization workflows.

Data Layer:

- **Overview & Function:** Aggregates, indexes, and streams structured and unstructured data assets into AI training pipelines and real-time inference engines.
- **Key Assets to Protect:** Data lakehouses/raw data stores, Vector Databases (Pinecone, Milvus) and catalogs, and automated data ingestion and lineage pipelines.
- **Primary Threat Vectors:** Adversarial training data poisoning, vector store embedding injection/tampering,

unverified data pipeline sources, and data catalog access control bypasses.

- **Potential Operational Impact:** Corruption or biased skewing of model outputs, compromise of Retrieval-Augmented Generation (RAG) context, and ingestion of malicious payloads during training.
- **Recommended Safeguards:** Implement automated dataset classification tagging, data poisoning scanners, cryptographic pipeline lineage tracking, and fine-grained RBAC/ABAC access controls for vector databases and data catalogs. Validate the hardware and firmware integrity of all underlying storage nodes and data ingestion servers to ensure datasets are streamed securely into the cluster.
- **How Eclypsium Helps:** Provides hardware and firmware integrity assurance across underlying storage node hardware and data ingestion servers streaming datasets into the cluster.

Orchestration:

- **Overview & Function:** Handles workload scheduling, virtual sandboxing, resource allocation, and multi-tenant isolation across shared HPC compute clusters and GPU fabrics.
- **Key Assets to Protect:** Cluster schedulers (Kubernetes, Slurm), GPU partitioning engines (NVIDIA MIG), and multi-tenancy/resource quota managers.
- **Primary Threat Vectors:** GPU partition escape / sandbox bypass, cross-tenant container contamination, scheduler privilege escalation, and compute quota exhaustion attacks.
- **Potential Operational Impact:** Unsanitized node reallocation across classification boundaries, tenant-to-tenant model hijacking, and denial of compute for mission-critical training runs.
- **Recommended Safeguards:** Enforce hardware-based GPU partitioning (NVIDIA MIG), multi-tenant namespace/container isolation, and GitOps infrastructure policy enforcement. Perform pre- and post-execution hardware integrity checks at every workload boundary, using automated API integration with cluster schedulers (Kubernetes/Slurm) to automatically cordon and quarantine compute nodes immediately upon detecting firmware drift or a failed sanitization check.
- **How Eclypsium Helps:** Executes hardware integrity checks at every workload boundary.

MLOps Platform:

- **Overview & Function:** Manages the end-to-end data science lifecycle, experiment tracking, software dependency pipelines, and model artifact deployments.
- **Key Assets to Protect:** Data science workbenches (Jupyter), experiment tracking platforms (MLflow, W&B), automated CI/CD build pipelines, and signed model registries.
- **Primary Threat Vectors:** Open-source software supply chain exploits (PyPI/Hugging Face package poisoning), malicious third-party code execution, signed model registry tampering, and experiment tracking compromise.
- **Potential Operational Impact:** Deployment of backdoored or compromised models to production, execution of untrusted code on high-value GPU clusters, and pipeline execution hijacking.
- **Recommended Safeguards:** Implement NIST Secure Software Development Framework (SSDF) build assurance, static application security testing (SAST), automated software dependency scanning, and cryptographic model provenance attestation. Automated FBOM and HBOM telemetry into MLOps build pipelines, verifying that model development environments execute on validated hardware.

- **How Eclypsium Helps:** Supplies xBOM telemetry and verifies that model development environments execute on validated hardware.

Model Serving:

- **Overview & Function:** Governs active inference execution, real-time query handling, and context enrichment for deployed AI models.
- **Key Assets to Protect:** High-throughput inference endpoints & APIs (vLLM, Triton), RAG & vector retrieval services, and runtime model guardrail gateways.
- **Primary Threat Vectors:** Indirect prompt injection, guardrail bypasses, model inversion and weight extraction attacks, unauthorized API access, and adversarial input manipulation.
- **Potential Operational Impact:** Exfiltration of proprietary model weights, unauthorized execution of system commands via agentic tools, and generation of corrupted or classified outputs.
- **Recommended Safeguards:** Deploy runtime model guardrail gateways, real-time input/output prompt sanitization, cryptographic model weight verification before execution, and Zero Trust API rate-limiting. Validate underlying host and accelerator firmware integrity, guaranteeing that inference endpoints (vLLM, Triton) execute within uncompromised physical environments.
- **How Eclypsium Helps:** Validates underlying host and accelerator firmware integrity.

AI Workload:

- **Overview & Function:** Represents the uppermost operational layer, hosting end-user applications, large-scale frontier training runs, and autonomous agentic systems.
- **Key Assets to Protect:** Frontier LLM training runs, enterprise copilots, Retrieval-Augmented Generation (RAG) applications, and autonomous agentic AI systems.
- **Primary Threat Vectors:** Agentic goal hijacking & logic manipulation, malicious prompt output execution, compromised copilot plugin execution, and workload privilege escalation.
- **Potential Operational Impact:** Unintended execution of destructive system commands by autonomous agents, exfiltration of sensitive enterprise data via copilots, and subversion of frontier LLM decision logic.
- **Recommended Safeguards:** Enforce autonomous agent sandboxing, operational boundary limits, real-time prompt output filtering, strict workload isolation policies, and least-privilege API scope enforcement.
- **How Eclypsium Helps:** Ensures that all workloads run on verified hardware free from persistent rootkits or backdoors.

Although distinct from the core AI computing stack, this operational layer represents a foundational element that is inherently present in the broader deployment environment.

Management & Unified Governance:

- **Overview & Function:** While distinct from the core AI compute stack, Management & Unified Governance is an integral part of the broader AI Data Center environment. It acts as the cross-cutting administrative pillar delivering centralized identity management, unified security telemetry, and automated mandate enforcement across all architectural layers.
- **Key Assets to Protect:** User Identity & Access systems (IAM/PAM), Security Operations telemetry engines (SOC/

SIEM/SOAR), infrastructure management consoles, and automated compliance auditing systems.

- **Primary Threat Vectors:** Administrative credential theft or privilege abuse, SIEM logging blind spots at the hardware layer, silent firmware/hardware configuration drift, and regulatory non-compliance.
- **Potential Operational Impact:** Undetected low-level persistence across GPU clusters, complete administrative control plane takeover, and loss of federal Continuous Authority to Operate (cATO).
- **Recommended Safeguards:** Centralize User Identity & Access Management (IAM/PAM), enterprise security logging, and administrative governance under a unified control plane. Integrate streaming below-OS hardware and firmware telemetry directly into enterprise SIEM/SOAR tools, real-time anti-drift alerting, alignment with DoW Zero Trust Device/Data pillars, and automated xBOM compliance auditing for cATO packages.
- **How Eclypsium Helps:** Delivers centralized platform integration that streams below-OS hardware and firmware telemetry directly into enterprise SIEM/SOAR/ITSM tools, triggers real-time anti-drift alerts, maps directly to DoW Zero Trust Device/Data pillars, and generates automated xBOM compliance packages for Continuous Authority to Operate.

Bridging Domain Risks to Hardware and Component Layer Defense

Securing this full-stack architecture demands a unified security control plane that anchors trust directly at the silicon layer. The next section outlines the core technical capabilities engineered into the Eclypsium platform to proactively detect, isolate, and mitigate hardware-level vulnerabilities across the AI compute ecosystem.

Core Technical Capabilities for AI Infrastructure Security

To defend against sophisticated state-sponsored and AI-driven threats, the Eclypsium platform operationalizes Preemptive Exposure Management (PEM) at the foundational hardware layer. Recognized by industry analysts as a DSEM platform for infrastructure and hardware supply chains, Eclypsium provides the deep, context-rich visibility and validation needed where generalist enterprise security tools are blind.

The key technical pillars built into the Eclypsium architecture are tailored specifically for high-capacity AI environments:

AI Server & GPU Risk Analysis

- **Advanced Accelerator Mapping:** Discovers, inventories, and maps critical components deep within AI clusters, offering native support for advanced infrastructure such as NVIDIA GPUs.
- **Multi-Tenant Validation:** Continuously monitors the physical and cryptographic integrity of GPUs before and after they are leased or allocated to different missions, teams, or external defense contractors, preventing cross-tenant data contamination or persistent firmware modifications.
- **Automated Risk Reporting:** Generates comprehensive risk profiles detailing hardware vulnerabilities, outdated firmware, and component risks for individual accelerators or scaled across data center fleets.

AI Infrastructure Integrity Assurance & Supply Chain Security

- **Firmware Verification:** Scans and verifies production GPUs, accelerators, and underlying silicon to guarantee that only authorized, “known-good” configurations are executing in the environment.
- **Anti-Drift Monitoring:** Instantly alerts security operations centers the moment an unauthorized modification, unexpected configuration drift, or unauthorized patch distribution is detected.

- **Automated xBOM Generation:** Continuously generates and updates SBOMs, FBOMs, and HBOMs for AI infrastructure assets to fulfill NDAA Section 1513 supply chain provenance requirements.

Deep Vulnerability Management & Threat Detection

- **Component-Level Vulnerability Finding:** Uncovers deep seated structural vulnerabilities that traditional network or operating system-level scanners overlook, including outdated firmware and vulnerable GPU drivers.
- **Low-Level Threat Interception:** Proactively catches backdoors, malicious hardware manipulation, and stealthy implants attempting to bypass EDR barriers to subvert high-performance compute nodes.

Broad Lifecycle & Ecosystem Support

- **Technology Agnostic:** Delivers uniform security coverage across an array of technology stacks, from x86 and ARM server architectures to firewalls, core switches, and data center networking routers.
- **Full Lifecycle Protection:** Can be deployed fluidly at any infrastructure phase: before installing the gear, between heavy AI model training runs, or during final asset disposition.
- **Fully supports air-gapped, isolated, and classified networks:** The platform features offline database update channels, allowing security teams to import updated cryptographic firmware signature feeds (containing 15M+ known-good binaries) into disconnected environments without requiring outbound internet connectivity.

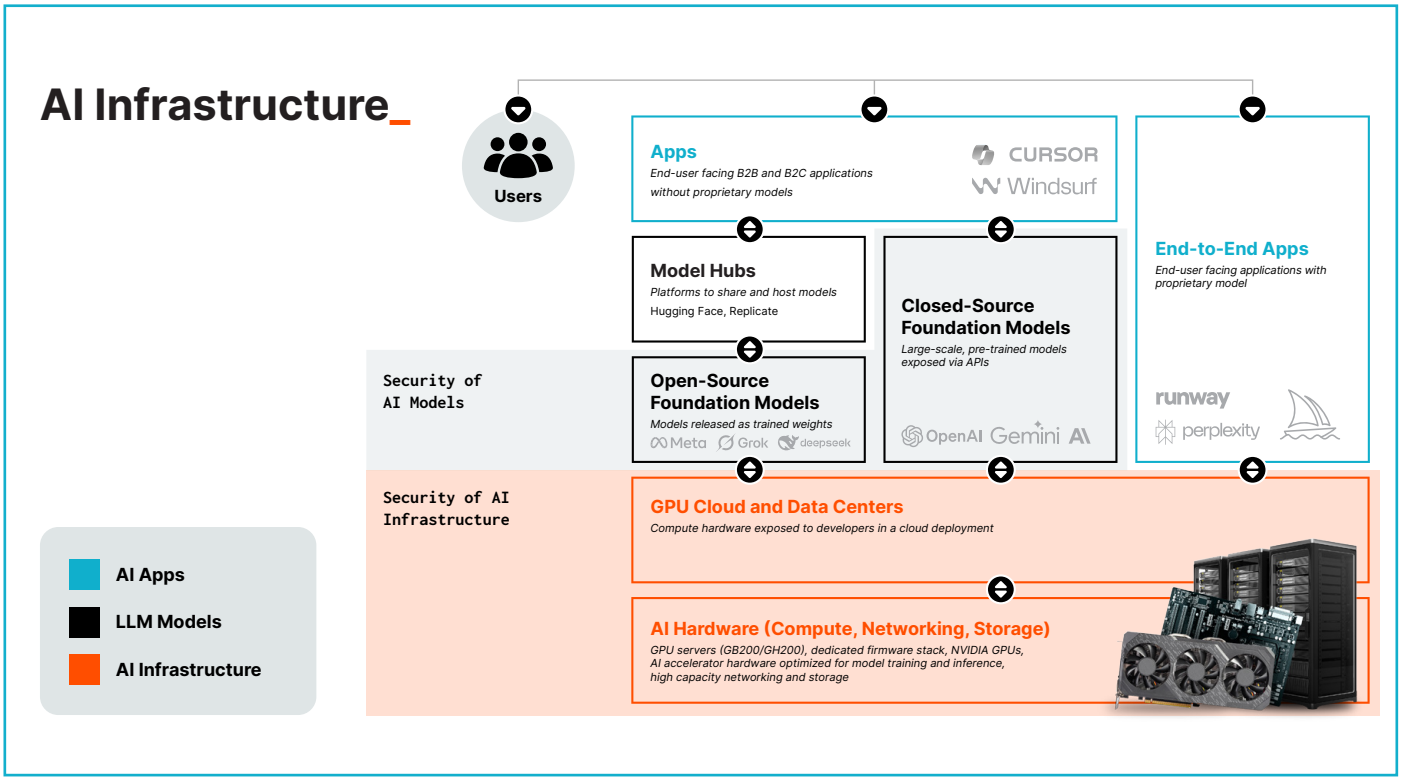

info@eclypsium.com
eclypsium.com

Objectives & Key Problems Solved

The Eclypsium platform serves as an enterprise-grade Infrastructure Assurance solution engineered to anchor the physical trust layer of critical AI infrastructure. By securing the underlying hardware, microcode, and firmware, Eclypsium directly protects the integrity of advanced machine learning models and heavy compute resources. This operational implementation explicitly addresses the following strategic objectives and operational gaps within the defense enterprise:

- **Benefit 1 (Hardware-Enforced Model Integrity and Supply Chain Provenance):** Eradicates operational blind spots at the hardware layer to prevent data loss, unauthorized access, and low-level tampering targeting proprietary AI model weights, hyper-parameters, and sensitive training datasets. Because traditional security tools cannot see down to the GPU or baseboard management controller (BMC) firmware, adversaries can persist in these components, leveraging them to silently exfiltrate models or poison data directly in memory. Eclypsium stops this by cryptographically verifying the infrastructure, ensuring the physical hardware and component firmware has not been manipulated to leak sensitive information or alter model behavior. This guarantees continuous compliance with federal supply chain risk management standards, specifically enforcing the strict hardware/software provenance and anti-tampering directives mandated under the June 2026 White House Executive Order on AI, NIST SP 800-53, NIST SP 800-223, and the FY 2026 NDAA (Section 1513).
- **Benefit 2 (Optimized Multi-Tenant Compute & Accelerated AI Deployment Timelines):** Eliminates manual vendor compliance tracking by providing automated, continuous scanning and vulnerability analysis of complex AI hardware clusters, including GPUs, TPUs, and high-performance interconnects without requiring specialized firmware engineers. This shortens implementation timelines, enabling cyber defense and data science teams to rapidly validate high-performance compute infrastructure. This accelerates safe multi-tenant compute environments, allowing infrastructure to be quickly cleared, sanitized, and redeployed for different classified workloads without introducing cross-contamination risks.
- **Benefit 3 (Streamlined SCRM Evidence & Accelerated RMF Assessment for Hardware Controls):** Eclypsium automates hardware, firmware, and supply chain evidence collection that can support RMF assessment and authorization packages for relevant controls. By continuously generating xBOMs (SBOM/FBOM/HBOM) and monitoring below-OS hardware and firmware posture aligned with NIST SP 800-53 controls, 800-161 C-SCRM practices, and 800-223 HPC security guidance, Eclypsium reduces manual hardware auditing effort, supports SCRM reviews, and provides telemetry that can contribute to the continuous monitoring component of a cATO program for AI compute clusters.

Vendor Qualifications

Eclypsium has a proven track record of providing Infrastructure Security solutions to various U.S. Government agencies and organizations, addressing critical needs related to operational continuity, risk reduction, and compliance with stringent cybersecurity mandates. As an established market leader included on CISA's Continuous Diagnostics and Mitigation Approved Products List, Eclypsium serves as a trusted partner across global defense networks, intelligence agencies, and major federal cloud providers to protect critical infrastructure and high-performance AI compute clusters. The platform delivers targeted security for modern AI factories by continuously inventorying, managing, and scanning AI infrastructure and NVIDIA GPU hardware down to the component level. By delivering low-level visibility and automated vulnerability analysis below the operating system layer, the Eclypsium platform hardens defense networks and data science pipelines against advanced persistent threats, proactively identifying outdated firmware, unpatched GPU drivers, and hardware-level risks that traditional scanners miss. Backed by an extensive and growing database of over 15 million cryptographic firmware hashes, Eclypsium provides the real-time threat intelligence needed to counter sophisticated state-sponsored supply chain

exploits. This ensures foundational hardware remains resilient against unauthorized modifications, adversarial tampering, or low-level exploits targeting proprietary AI models.

Procurement Guidance

The Eclypsium Infrastructure Assurance Platform is widely accessible through Carahsoft and our authorized reseller network and their multiple Government-wide Acquisition Contracts (GWACs), including GSA Multiple Award Schedule (MAS), NASA SEWP V, and ITES-SW2 vehicles. Additionally, Eclypsium is on the CISA CDM APL.

There is established precedent for a brand name procurement of Eclypsium by U.S Government agencies. While other commercial products offer limited hardware security, none possess the comprehensive hardware security capabilities required to secure high-performance AI infrastructure and fulfill specific mission requirements.

Summary_

Eclypsium appreciates the opportunity to offer this modernized AI Infrastructure security capability blueprint. We are entirely confident that the Eclypsium Infrastructure Assurance Platform will allow the U.S Government to achieve unprecedented speed, cryptographic validation, and process optimization across its ongoing operations while answering the nation's call to secure critical AI infrastructure.

Eclypsium provides the essential technical ecosystem required to fulfill this operational requirement: a unified infrastructure assurance platform delivering on supply chain security requirements that automatically scans, maps, and analyzes firmware binaries across servers, endpoints, networking assets, and AI accelerators, such as GPUs.

By continuously inspecting hardware against its proprietary vulnerability database, the platform detects configuration drift, flags malicious implants, and automatically manages firmware patches directly through a centralized console. These capabilities allow the DoW to sanitize compute nodes and continuously audit its high-value AI hardware assets throughout their active operational lifecycles.

How to Buy Eclypsium_

Federal Agencies

Eclypsium is available through Federal VARs and Carahsoft on the following contracts:

GSA Multiple Award Schedule (MAS) and SCRIPTS BPA

GSA Schedule | NASA SEWP V | CHESS ITES-SW2

Eclypsium is also on the CDM APL.

State and Local Governments
Carahsoft for Eclypsium SLED Contracts

Canadian Federal Government
Carahsoft Canadian Cyber Security Procurement Vehicle

Vendor Information
Eclypsium, Inc.
Address: 919 SW Taylor Street | Portland, OR 07205
DUNS/CAGE: 081023218| 8GXX1