Autonomous Cyber Defense Systems Using AI/ML Intelligent Agents

Engineer Research and Development Center (ERDC)

High Performance Computing (HPC) Modernization Program
Cybersecurity for Advanced Computing and Software Modernization
Commercial Solutions Openings (CSO)

Under Solicitation Number: W912HZ25SC002

8 September 2025

Prepared for:

U.S. Army Engineer Research and Development Center (ERDC)
3909 Halls Ferry Road
Vicksburg, MS 39180-6199

Email: info@erdcwerx.org

Prepared by:

Antean Technology
5860 King Centre Drive, Suite 600
Alexandria, VA 22315
UEI: K2ZEQM1YTTV3
Contact Name: Mr. Sean Floyd, Chief Operating Officer
Contact Email: seanfloyd@anteantech.com
Contact Phone: (703) 254-0448
Website : https://anteantech.com

See details of our proposed solution at: https://anteantech.com/autonomous-cyber-defense

EXECUTIVE SUMMARY & PROBLEM ALIGNMENT

Antean Technology, LLC is pleased to offer the U.S. Army Engineer Research and Development Center (ERDC) and the High Performance Computing Modernization Program (HPCMP) the required innovative methods to protect critical computing resources from adversaries who are already leveraging artificial intelligence (AI) at machine speed. Current defenses rely on human-in-the-loop operations and centralized architectures that cannot operate effectively in contested, disconnected, or resource-constrained environments. A compromise of HPCMP resources is more than a technical failure; it undermines DoD mission advantage and mission assurance.

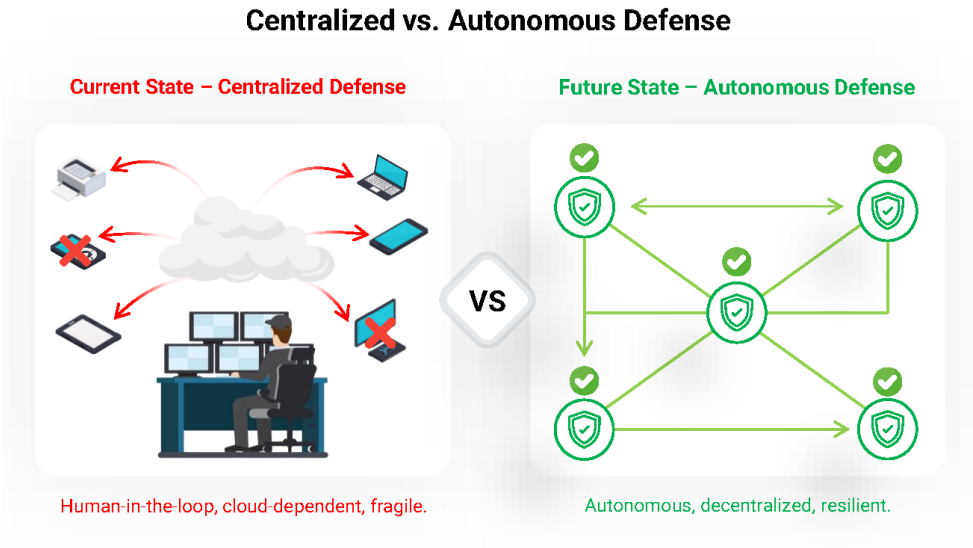

We propose a mission-aware, agentic AI framework integrated with ontology-driven Software Bill of Materials (SBOM) analytics. This system autonomously executes the full cyber OODA loop (Observe–Orient–Decide–Act) to provide real-time defense of HPC, tactical, and multi-domain systems. Agents execute trusted defensive actions under runtime policy guardrails, ensuring both machine-speed response and operator assurance. This represents an innovative application of existing agentic AI and semantic knowledge graph technologies, creating a new, mission-aware defensive paradigm that moves beyond legacy, human-in-the-loop cybersecurity systems. The design emphasizes resilience, explainability, and survivability, ensuring DoD digital infrastructure remains operational under attack.

See details of our proposed solution at: https://anteantech.com/autonomous-cyber-defense

EXPECTED OUTCOMES:

- **Decision Advantage:** Detect and respond to adversarial activity in <2s; respond in <5s. seconds (reduced from hours).
- **Mission Assurance:** Maintain ≥90% survivability of mission-critical HPC services during a cyber attack.
- **Operational Resilience:** Operate autonomously in Disconnected, Intermittent, and Limited (D-DIL) conditions without persistent cloud reliance.
- **Trusted Autonomy:** Provide explainable, auditable, and verifiable AI aligned with NIST AI RMF and DoD Responsible AI guidance.

TECHNICAL APPROACH & ARCHITECTURE

To achieve these outcomes, the proposed solution employs a **layered cognitive architecture** aligned with the OODA loop and engineered for resilience and operational adaptability.

Core Components

- **Autonomous Intelligent Agents:** Distributed agents deployed across HPC nodes, tactical devices, and forward-deployed assets. Each agent senses, reasons, and acts independently to ensure local defense even when networks are disrupted.
- **Ontology-Driven SBOM Analytics:** A semantic knowledge graph links software vulnerabilities to specific mission functions, enabling agents to prioritize defense actions based on operational impact rather than generic severity scores. This includes SBOM transparency across containerized and microservice environments, reducing supply chain risk while embedding mission context into agent decisions.
- **Federated Learning:** Agents improve their models by sharing updates instead of raw data, conserving bandwidth and protecting sensitive information.
- **Multi-Agent Swarm Defense:** Agents collaborate peer-to-peer, instantly propagating new threat intelligence to one another and enabling self-healing resilience across the network.
- **Explainable AI (XAI):** Each autonomous action provides a human-readable rationale and a verifiable audit trail, building trust and satisfying governance requirements. All actions are simultaneously validated against runtime Zero Trust guardrails (LAM Gate/Engine), producing auditable, policy-validated autonomy
- **Adversarial Robustness:** Agents are designed to operate efficiently without requiring GPUs, enabling deployment on edge devices with constrained compute. When available, agents can also leverage remote or hybrid models, ensuring flexibility across HPC, tactical, and disconnected environments. Lightweight Edge Nodes enforce policies locally, sustaining autonomous defense in disconnected or resource-constrained conditions without cloud reach-back.

Operational Relevance & Use Cases

Our autonomous cyber defense is most effective in ERDC's operational contexts, where high latency, intermittent connectivity, and contested environments are the norm.

HPC Security: During a weapons simulation on an HPCMP cluster, embedded agents baseline the workflow, detect data corruption attempts, and identify the targeted process as mission-critical. The agent quarantines the threat within seconds and generates an explainable report, preserving simulation integrity.

Tactical Edge: In a comms-denied environment, malware introduced via USB is autonomously neutralized by the local agent. A peer-to-peer hash exchange inoculates nearby kits—achieving response without cloud reach-back.

Multi-Domain Operations (MDO): Autonomous agents across joint nodes detect anomalous targeting traffic, triangulate the attack source, and reroute communications—preserving fast, resilient kill chain coordination.

Alignment to DoD Priorities:

- **Zero Trust:** Continuous verification and autonomous segmentation
- **DoD AI Adoption Strategy:** Secure, trustworthy AI data pipelines
- **NIST AI RMF / OMB Guidance:** Transparent, auditable autonomous behavior

Together, these use cases demonstrate measurable enhancements to ERDC and HPCMP mission assurance by ensuring availability, integrity, and survivability of critical compute assets under contested conditions.

MATURITY, OUTCOMES, AND KPIS

Technology Readiness Level (TRL)

- **TRL 5:** Agentic AI framework validated in cyber range environments and as an approved capability for the U.S Navy. The proposed framework builds on capabilities originally validated in operationally realistic training environments, where autonomous agents were stress-tested against adversary TTPs. This heritage provides empirical grounding for our current TRL 5-6 assessment, while our roadmap advances the capability to TRL 6 in operational HPC/tactical demonstrations.
- **TRL 4–5:** Ontology-driven SBOM analytics operational in pilot deployments.
- **Path to TRL 6:** A 12–24-month plan including lab validation, HPC/tactical demonstrations, and an ERDC pilot deployment.

Expected Outcomes

- **Speed of Detection and Response:** Detect adversarial activity in <2 seconds; respond in <5 seconds. ≥80% reduction in adversary dwell time compared to baseline SIEM/SOAR.
- **Trusted and Explainable Autonomy:** ≥99% of autonomous actions include a rationale and immutable log entry; ≥95% validated against policy rules prior to execution; ≥95% operator validation rate for autonomous decisions.
- **Mission Assurance and Survivability:** ≥90% accuracy in prioritizing mission-critical HPC workloads; ≥50% fewer mission-disruptive false positives versus anomaly-only detection; ≥95% uptime of mission-critical HPC workloads under simulated adversarial attack.
- **Operational Resilience in Contested Environments:** ≥95% of defensive actions completed without cloud reach-back; ≤10% performance degradation during 24-hour blackout scenarios.
- **Cross-Cutting Performance Metrics:** TRL 3–5 maturity validated in relevant environments; interoperability with ≥3 data formats (JSON, XML, OSCAL, SBOM) and ≥3

heterogeneous systems (e.g., SIEM, HPC scheduler); ≤1 hour adaptation to new adversary patterns; ≤20% CPU and ≤25% memory overhead per agent under peak load.

KPI Framework

Metrics will be validated at three levels:

- **Technical:** Speed (MTTD/MTTR), accuracy (false positives/negatives), efficiency.
- **Operational:** Service survivability, dwell-time reduction, resilience in D-DIL environments.
- **Mission:** Continuity of mission-critical workloads, operator trust/adoption, mission success in contested scenarios.

Commercialization: Our proposed solution builds on commercially deployed components from Team member Animate's Hive platform and Digital Workforce capability, which are already in use within open-market cyber training environments and enterprise security deployments. These systems will provide the autonomous agent foundation. Additional capabilities—ontology-driven SBOM analytics and federated learning—have been exercised in environments against realistic adversarial behaviors. Together, these integrated elements form a clear commercialization pathway that supports ERDC mission alignment while maintaining a high level of transition readiness.

TEAM DIFFERENTIATORS AND STRENGTHS

Our open, API-driven design ensures rapid transition into HPCMP environments and seamless compatibility with existing DoD cyber infrastructure, enabling ERDC to scale beyond prototype. Unlike cloud-dependent approaches, our autonomy persists in disconnected conditions, ensuring resilience under ERDC's most stressed environments.

- **Mission-Aware Autonomy:** Agents prioritize actions based on mission impact rather than generic threat severity.
- **Decentralized, Cloud-Independent:** Fully operational in D-DIL environments without reliance on cloud services.
- **Explainable, Trusted AI:** Every action is transparent, auditable, and policy-bound, reinforcing operator trust.
- **Scalable Swarm Defense:** Self-healing, distributed collaboration eliminates single points of failure.
- **Hardware-Agnostic Autonomy:** Operates efficiently on CPU-only nodes and scales with local, remote, or hybrid models, removing dependence on GPUs or persistent cloud connectivity.
- **Beyond Static Orchestration:** Unlike conventional frameworks (e.g., Kubernetes/Istio), our solution enforces runtime Zero Trust and policy compliance across distributed systems.

Team Antean members has a proven expertise delivering DoD cyber training environments, high-performance computing security solutions, and autonomous agent frameworks, with experience supporting DoD and civilian agencies in AI governance, Zero Trust, SBOM analytics, and governance models. More strengths are as follows:

- **Animate Cyber Training Heritage:** AI-orchestrated, pattern-of-life training and Digital Workforce technology proven in DoD cyber ranges, enabling advanced user emulation, realistic threat scenarios, and scalable readiness exercises. These capabilities directly support Zero Trust initiatives and resilience in contested environments.
- **Innovation Assets:** Team Antean provides an ontology-driven SBOM analytics platform integrated with the multi-agent AI defense system.
- **Proven Federal Performance**

ROUGH ORDER OF MAGNITUDE (ROM)

Total Estimated Cost: $2.4M

Phase 1 – Lab Validation (12 Months, achieving TRL 5): $1.2M
• R&D labor: $0.9M
• Testing & demonstration: $0.2M
• Program management, travel, training: $0.1M

Phase 2 – HPC/Tactical Demonstrations and Pilot Deployment (12 Months, advancing to TRL 6): $1.2M
• R&D labor: $0.9M
• Testing & demonstration: $0.2M
• Program management, travel, training: $0.1M

ACRONYMS

- **AI:** Artificial Intelligence
- **CybORG:** Cyber Operations Research Gym
- **D-DIL:** Disconnected, Intermittent, and Limited
- **DARPA:** Defense Advanced Research Projects Agency
- **DoD:** Department of Defense
- **ERDC:** Engineer Research and Development Center
- **HPC:** High Performance Computing
- **HPCMP:** High Performance Computing Modernization Program
- **MDO:** Multi-Domain Operations
- **MTTD:** Mean Time to Detect
- **MTTR:** Mean Time to Respond
- **NIST:** National Institute of Standards and Technology
- **OMB:** Office of Management and Budget
- **OODA:** Observe–Orient–Decide–Act
- **R&D:** Research and Development
- **RMF:** Risk Management Framework
- **ROM:** Rough Order of Magnitude
- **SBOM:** Software Bill of Materials
- **TRL:** Technology Readiness Level
- **XAI:** Explainable Artificial Intelligence

See details of our proposed solution at: https://anteantech.com/autonomous-cyber-defense