HORIZON3

CyberCom 2.0 and the Revolution in AI-Enabled Offensive Cyber Operations

The Next Evolution of American Cyber Power in the AI Revolution of Military Affairs

WHITEPAPER

The Secretary of Defense's announcement of CyberCom 2.0 should be viewed as more than an organizational restructuring or personnel reform initiative. It represents a recognition that cyberspace has fundamentally changed as a warfighting domain and that the institutions, operational concepts, and technologies developed over the past two decades must evolve accordingly.

Throughout the announcement, the Secretary returned to several central themes: operational readiness, technical excellence, meritocracy, speed, and the elimination of bureaucracy. While discussed primarily through the lens of force generation, these principles ultimately point toward a broader transformation in how offensive cyber power is generated, sustained, and employed.

The strategic challenge facing the United States is no longer simply building capable cyber forces. It is ensuring those forces can consistently outpace adversaries who are already leveraging artificial intelligence to compress operational timelines, automate technical tradecraft, and scale campaigns beyond what human operators alone can achieve. Future cyber superiority will depend less on the size of the force and more on the ability of that force to integrate AI into every stage of offensive cyber operations.

The AI era of Offensive Cyber Operations has already begun

Much of the public discussion surrounding AI in cybersecurity has focused on defensive applications such as threat detection, security operations centers, and autonomous pentesting which allows more complete CTEM capabilities for defense. These capabilities are important, but understate the strategic reality of the AI use-cases for offensive cyber operations.

Over the past eighteen months, multiple incidents have demonstrated that sophisticated threat actors are successfully integrating large language models and autonomous agents into real-world operations. These are not isolated laboratory demonstrations. They represent the emergence of a new operational model in which AI dramatically increases the speed, scale, and adaptability of offensive cyber campaigns.

Anthropic's disclosure of the first publicly reported AI-orchestrated cyber espionage campaign provides critical evidence. A Chinese state-sponsored threat group employed Claude Code, Model Context Protocol (MCP)-connected tools, automated exploit generation, credential harvesting, lateral movement, and data exfiltration targeting approximately thirty organizations. Anthropic concluded that AI performed the vast majority of operational tasks, with successful compromises occurring at multiple victim organizations.

The campaign used multiple coordinated AI instances through a custom automation framework, completing 80 to 90 percent of the operational workflow. While human operators remained in command, AI became the primary executor of offensive tradecraft.

The cyberattacks targeting municipal water and wastewater systems across multiple U.S. states, including Minnesota, during July and August 2026, illustrate the profound asymmetry of modern cyber conflict. While these campaigns targeted small, localized utilities, the impacts were amplified by the strategic objectives of the attackers, believed to be linked to Iran and its cyber proxies.

These incidents reveal a critical national security vulnerability: state-sponsored actors are weaponizing the digital exposure of thousands of under-resourced local organizations to generate disproportionate physical and emotional disruption to United States critical infrastructure and its citizens. This demonstrates that the modern cyber battlespace is no longer confined to major government warfighting institutions; instead, the frontline now encompasses any internet-connected system that can be leveraged to generate geopolitical effects, effectively drawing municipal infrastructure into the center of international state-on-state conflict.

The economics of Offensive Cyber Operations are changing

Offensive cyber operations traditionally required scarce technical expertise across reconnaissance, vulnerability research, and execution. AI fundamentally changes this equation by automating technical tradecraft and scaling campaigns beyond human-only capabilities.

Highly specialized tasks like reconnaissance, code generation, attack-path analysis, reverse engineering and malware engineering are increasingly accelerated through AI, reducing the total cost in time and expertise for generating offensive capability.

Research into the HexagonalRodent campaign demonstrated how North Korean operators used commercial AI platforms including ChatGPT, Cursor, and other coding assistants to develop malware, construct phishing infrastructure, and build fraudulent software companies supporting cryptocurrency theft. Investigators linked the campaign to the compromise of more than 2,000 developer systems and the theft of approximately $12 million in digital assets.

Similarly, researchers documented campaigns in which AI systems reportedly generated and executed the majority of remote commands during intrusions targeting multiple government agencies. While those findings remain researcher-reported rather than independently confirmed by victims, they illustrate how rapidly offensive workflows are becoming AI-assisted.

Cyberspace has become a machine-speed battlespace

In the cyber domain, artificial intelligence alters decision-making timelines. Reconnaissance and target analysis that once required days can now be completed in minutes, with operational plans evolving dynamically at machine speed.

AI amplifies operator effectiveness, allowing a focus on mission objectives and judgment while machines perform repetitive analytical work. This enables the command of AI-enabled operations with the discipline applied to other advanced weapon systems.

AI changes the scale of cyber warfare

Modern AI systems can coordinate parallel activities such as intelligence collection, exploit generation, and infrastructure adaptation which shifts cyber warfare from individual missions to coordinated operational ecosystems. The organizations capable of orchestrating these fleets of AI-enabled agents will enjoy significant operational advantage.

The battlespace now includes AI systems themselves

As AI becomes embedded within military, government, and commercial systems, those systems themselves become part of the operational environment.

Researchers have already demonstrated vulnerabilities unique to AI-enabled infrastructure. Microsoft's EchoLeak research showed that indirect prompt injection could manipulate an enterprise AI assistant into exfiltrating sensitive organizational information without direct user interaction. While Microsoft mitigated the vulnerability before widespread exploitation, the research demonstrated that AI systems introduce entirely new attack surfaces that traditional cybersecurity architectures were never designed to defend.

Researchers have also demonstrated supply-chain attacks specifically designed to manipulate AI agents.

The FakeGit campaign created thousands of fraudulent GitHub repositories, hundreds of which masqueraded as AI Skills or MCP servers. Multiple leading AI coding assistants were shown to recommend these malicious repositories during development workflows, effectively creating a new software supply-chain attack vector in which attackers target AI assistants rather than human developers.

Similarly, the emerging technique of AgentBaiting demonstrates that adversaries can intentionally optimize malicious repositories and documentation to influence AI recommendations, poisoning software development workflows without directly compromising the underlying AI models.

The human operator becomes the operational commander

CyberCom 2.0's emphasis on meritocracy, technical excellence, and elite talent reflects this future operating model. The operator of tomorrow will command AI-enabled operational capabilities rather than manually execute every technical task.

- → Human operators establish mission objectives.
- → They determine acceptable operational risk.
- → They understand geopolitical consequences.
- → They integrate cyber effects with broader military campaigns.
- → They exercise authority that machines cannot.

CyberCom 2.0 as the foundation for AI-enabled cyber power in the current revolution of military affairs

CyberCom 2.0 signals an understanding that future cyber superiority cannot be achieved through legacy organizational models, industrial-age personnel systems, or operational concepts designed for a slower digital environment.

The United States has historically maintained military advantage by integrating emerging technologies into doctrine before competitors could fully exploit them. Precision strike, stealth, space-based capabilities, and network-centric warfare all reflected this pattern. Artificial intelligence and its rapid adoption represents the next revolution in military affairs.

CyberCom 2.0 establishes the organizational conditions necessary to capitalize on AI. The next challenge is operational: developing doctrine, training, authorities, and command structures that enable AI-augmented offensive cyber forces to operate with the speed, precision, adaptability, and discipline required for modern conflict.

The future of cyber superiority will not belong solely to the nation with the largest cyber force or the most advanced algorithms. It will belong to the nation that most effectively integrates exceptional human operators with artificial intelligence to generate operational advantage faster than any adversary can respond. This is the strategic opportunity presented by CyberCom 2.0, and it is one that should shape the next generation of American offensive cyber doctrine.

About the author

Corey Brunkow is Director of Federal Operations at Horizon3, where he brings decades of leadership, engineering, and operational experience to helping federal organizations strengthen cybersecurity.

A U.S. Army Ranger and Special Operations veteran, Corey held senior engineering, operations, and command positions throughout his military career, including service with the 75th Ranger Regiment. He is a graduate of the United States Military Academy at West Point and holds master's degrees from National Defense University and the Naval Postgraduate School.

References

Secretary of Defense Pete Hegseth – CyberCom 2.0 Announcement

Anthropic – Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign (Original Announcement)

Anthropic – Full Technical Report (PDF)

MITRE ATT&CK – Campaign C0062: Anthropic AI-Orchestrated Campaign

WIRED – AI Tools Are Helping Mediocre North Korean Hackers Steal Millions

TechRadar – Hackers Use Claude and ChatGPT to Breach Government Agencies

TechCrunch – OpenAI Says Hugging Face Was Breached by Its Pre-Release Models

SecurityWeek – EchoLeak: AI Attack Enabled Theft of Sensitive Data via Microsoft 365 Copilot

The Hacker News – FakeGit Campaign Uses 7,600 GitHub Repositories to Spread SmartLoader Malware

SecurityWeek – Zero-Click AI Browser Hacking: Claude and ChatGPT Atlas Hijacked