AI-accelerated threats don't change the cyber playbook; **they compress the clock**

Frontier AI models are driving a familiar cycle of hype and fear in cybersecurity. Some capabilities are being marketed as "expert-level" vulnerability discovery and exploitation assistance, and access to the most capable systems may be limited or uneven.

Regardless of where any single model lands, the practical implication for federal environments is straightforward: **Nothing about the core defensive mission changes, everything just moves faster and AI for defense can help shore up the existing security value streams at speed.**

That is the real level set. The same enterprise realities remain: complex hybrid estates, legacy dependencies, mission uptime requirements and determined adversaries. Now, however, AI can reduce the time and skill needed to translate weaknesses into working attacks.

It can also increase scale, allowing adversaries to run more parallel discovery, targeting and iteration than most defenders are staffed to match. The result is a higher-tempo environment where the number of capable threat actors can grow, development of exploits can accelerate and defenders face a compressed decision window.

For federal executives, the response should be cyber-focused and operational. Run the same three critical defense streams as always, but run them harder, faster and with greater discipline while using defensive AI to bolster SecOps at speed. AI, Automation, pre-authorization and repeatable playbooks engineer those security value streams for higher tempo in pace with the threat.

Three things every cyber shop can do to combat advanced AI-based threats:

01 **Quickly advance zero trust**: Shrink the attack surface, raise detectability, reduce exploitability

02 **Strengthen vulnerability management**: Assess exploitability at speed, then mitigate decisively

03 **Raise incident response velocity**: Push for higher tempo, higher credibility, better outcomes

→ Zero trust:

Shrink the attack surface, raise detectability, reduce exploitability

AI-enabled offense increases the penalty for porous boundaries, implicit trust, brittle segmentation and weak identity controls. Zero trust is not a branding exercise in this environment. It is the most direct way to reduce the number of reachable targets and limit what an attacker can do after initial access.

AI can strengthen execution by continuously analyzing identity, network, endpoint, cloud and configuration telemetry to reveal risky trust relationships, prioritize the most reachable “paths to impact” and flag drift from intended policy faster than periodic reviews.

What matters most now:

- **Reduce attack surface:** Enforce strong identity controls, least privilege access, and segmentation so fewer systems are reachable and fewer actions are permitted. AI can help identify over-privileged roles, unused access, and high-risk pathways, then recommend targeted entitlement and segmentation changes.
- **Increase detectability:** Design identity, endpoint, network, cloud, and application layers so suspicious behavior stands out sooner and can be correlated faster. AI can assist with log normalization, enrichment, and correlation, improving signal quality and speeding detection engineering.
- **Reduce exploitability:** Standardize hardened baselines, modernize where feasible, and design systems so that their compromise does not inevitably become mission impact. AI can help validate configuration against baselines, detect misconfiguration and control gaps, and accelerate hardening by turning findings into prioritized, actionable remediation steps.

Defending against AI-accelerated threats with zero trust is an ongoing task, but a simple test to mark progress is: At any given stage of the effort, do you have fewer paths for it to reach your crown jewels, more places to see the exploit in action and more ways to contain it quickly than you did before?

→ Vulnerability management:
Assess exploitability at speed, then mitigate decisively

If AI makes it easier to go from vulnerability to exploit, vulnerability management becomes a speed contest. The critical capability is rapid exploitability assessment and rapid mitigation, especially when patching is slow, risky or operationally constrained.

AI can help by accelerating the “is this exploitable here?” determination, correlating common vulnerabilities and exploits (CVEs) to real inventory and exposure, and translating findings into prioritized, verifiable actions to include proposed patches and automated pull requests to implement.

Federal leaders should push for:

- **Faster triage:** Quickly determine whether a CVE or zero-day exploit is relevant in your environment based on asset exposure, reachable paths, privilege context and existing compensating controls. AI can rapidly map advisories to your software bill of materials (SBOM)/asset inventory, configurations, and observed reachability to reduce false urgency and missed risk.
- **Exploitability measurement:** Prioritize by which CVEs can be exploited in your environment and the blast radius, not by severity score alone. AI can help model likely attack paths and highlight the handful of conditions that make exploitation practical.

- **Rapid compensating controls:** When patches lag, deploy mitigations such as segmentation changes, temporary feature disablement, web application firewall (WAF) rules, endpoint controls, privilege reduction and targeted monitoring. Track from closure to completion. AI can recommend the most effective compensating controls for your architecture and verify whether mitigations are in place.
- **Surge capacity:** Plan for periods where multiple critical issues land at once, with predefined owners, emergency change lanes and verification steps that keep mission uptime in view. Use AI to auto-generate tasking, validate proposed fixes against policy and accelerate evidence collection for closure.

This also means staying honest about where intrusions often begin. Many real-world incidents still start with identity compromise, misconfiguration or third-party access. High-tempo vulnerability management must sit alongside identity and supplier hardening because adversaries will take the fastest path in your environment.

This is where faster is tangible and defensive AI is key: Time-to-assess, time-to-mitigate and time-to-verify become executive metrics, not technical footnotes.

→ Incident response:
Higher tempo, higher credibility, better outcomes

As threats accelerate, incident response must become more practiced, more automated where appropriate, and more decisive. Speed matters, but so does accuracy. Federal organizations cannot afford to chase noise, nor can they afford delays that let an intrusion mature.

AI can help incident response teams move faster by correlating signals across tools, summarizing what is known with supporting evidence, preparing responsible upward reporting and accelerating scoping and decision support. Meanwhile, humans remain in control of approvals and high-impact actions.

The incident response capability should emphasize:

- **Early containment and scoping:** Limit spread quickly, identify affected identities and systems and preserve evidence. AI can speed scoping by linking related alerts, identifying likely patient zero paths and proposing the highest-value containment steps based on observed behavior.
- **Repeatable playbooks and pre-authorization:** Define in advance which containment actions can be executed immediately, by whom, and under what conditions, so teams can move without delay in the moments that matter. Leverage AI to recommend the right playbook for the scenario, check prerequisites and generate consistent tasking and documentation.
- **Stronger investigation discipline:** Establish facts quickly and credibly so leadership can act with confidence and communicate without speculation. AI can assist analysts by building event timelines, surfacing relevant telemetry, and highlighting gaps in collection that need to be filled to confirm or refute hypotheses.
- **Recovery readiness:** Rebuild and restore at pace, verify integrity and close the gaps that enabled entry. AI can support recovery by helping validate configurations against hardened baselines, identifying residual persistence risks, and accelerating evidence collection for closure.

In an AI-accelerated landscape, the best response is a machine-ready operational cadence. Diligence, persistence and proactive strategy are more important than in-the-moment heroics.

Fighting AI with AI

While rapidly evolving AI capabilities open new threats and risk vectors, they can also provide a strong defense against those same threats. AI can aid cyber defense by:

Continuously analyzing identity, network, endpoint, cloud, and configuration telemetry

Enforcing strong identity controls, least privilege access, and segmentation

Validating configuration against baselines, detecting misconfiguration and control gaps, and accelerating hardening by turning findings into prioritized, actionable remediation steps.

Accelerating determination of CVE applicability by rapidly mapping advisories to your SBOM/asset inventory, configurations, and observed reachability

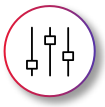

Recommending the best compensating controls for your architecture

Auto-generating tasking and validating proposed fixes against policy, for surge readiness

Speeding response by correlating signals across tools and summarizing what is known with supporting evidence

Linking related alerts and identifying likely patient zero paths, and proposing containment steps

Assist analysts by building event timelines, surfacing relevant telemetry, and highlighting gaps

How CGI Federal can help

Our perspective on this moment is shaped by decades of delivering cybersecurity outcomes in complex federal environments, and by **operating at national scale where speed and repeatability are crucial**. Across the full cybersecurity lifecycle, CGI Federal supports agencies from risk assessment through detection, response and modernization.

Our capabilities include 24x7 security operations, advanced threat intelligence and hunting, incident response, governance, risk and compliance (GRC) aligned to federal requirements, DevSecOps and secure continuous integration/continuous development (CI/CD), identity and access management with zero-trust-aligned architectures, and high-tempo vulnerability management using automation.

That “faster fundamentals” approach is also reflected in how CGI Federal helps agencies drive enterprise-wide visibility and prioritization. Work supporting federal continuous monitoring and risk visibility programs has reinforced an executive truth: Real resilience comes from shortening the time from signal to decision to action, and doing it consistently across diverse systems and stakeholders.

As AI increases attacker speed and scale, agencies are increasingly pairing strong cyber process with AI-enabled assistance on the defensive side as well. CGI is investing in and applying agentic AI capabilities to compress the time between “a change occurs” and “we know whether it is safe.” In the software development

lifecycle and CI/CD pipeline, this means using automated analysis to review code, infrastructure-as-code, and making configuration changes at a granular level.

Flagging likely exploit paths and recommending specific control improvements before deployment ensures tighter security. In operations, the same approach can be applied to continuously evaluate zero trust posture and security configurations against intended policy, highlight drift and misconfigurations that expand exposure, and accelerate hardening activities by translating findings into actionable tickets, updated baselines and prioritized remediation sequences.

Investigating security breaches depends on understanding the story of the event across identity, endpoint, network and cloud. Agentic AI can assist in telling that story through log onboarding and normalization, and by improving signal quality through smarter correlation and enrichment.

In the security operations center (SOC) and incident response context, these capabilities can reduce manual bottlenecks in alert triage, hypothesis generation, scoping and containment planning.

These advances provide a much more thorough and more quickly compiled data set to enable leaders to make informed mission risk decisions much faster.

Whether the mission need is integrating security into the SDLC, strengthening SOC operations, improving vulnerability response, securing enterprise and back-office systems, or standing up cyber managed services, CGI Federal works alongside both public and commercial sector organizations to execute these fundamentals at speed in an evolving threat environment.

If you would like to discuss how to apply this high-tempo defensive approach to your environment, reach out to a CGI Federal cybersecurity subject matter expert.

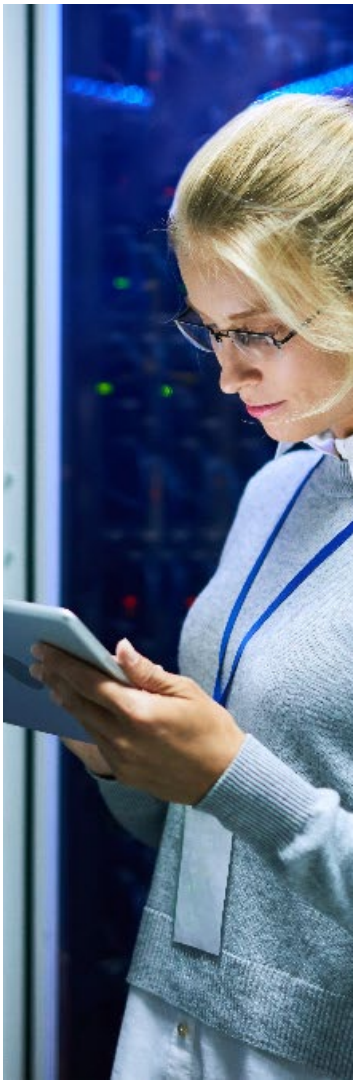

About CGI Federal

CGI Federal Inc. (CGI) is a leading U.S.-based technology and professional services company that serves federal agencies across defense, civilian, healthcare, justice, intelligence and international affairs. With nearly 8,000 professionals, CGI works with its clients to modernize government through innovative technology solutions, flexible delivery models and a commitment to achieve mission outcomes. CGI Federal is a wholly owned subsidiary of CGI Inc.

Visit: www.cgifederal.com
Contact: info@cgifederal.com