Operational Technology Cybersecurity Coalition

The Missing Pillar of AI Security: Critical Infrastructure Resilience

Policy Recommendations for Protecting Operational Technology Against AI-Enabled Cyberattacks

November 17, 2026

Table of Contents

Executive Summary 3

Introduction 5

Principles 6

Recommendations 8

Conclusion 17

Acknowledgments 18

Operational Technology Cybersecurity Coalition

Executive Summary

The U.S. government, AI labs, and the Operational Technology (OT) community must partner together and act now to enhance the resilience of U.S. critical infrastructure against AI-enabled cyberattacks. The past year has made clear that AI-enabled threats to OT are no longer hypothetical. Attackers are leveraging AI to discover vulnerabilities, compromise devices, and target industrial controllers at a speed and scale that traditional defenses were not designed to meet.

Yet, an essential element is missing from the national conversation about AI security: operational resilience. Much of the debate has focused on securing AI models, preventing their misuse, and developing AI-enabled cybersecurity defenses. These efforts are necessary, but alone are not sufficient. Even the most sophisticated defenses cannot guarantee that an AI-enabled cyberattack will fail. For critical infrastructure, where cyber incidents can have real-world consequences, AI security must also mean ensuring that essential services can withstand, contain, and recover from a successful attack. Operational resilience must therefore become a central objective of U.S. AI security policy. Unfortunately, according to a recent SANS Institute survey of industry security leaders, only 15% of respondents said they could fully sustain manual operations past 48 hours, and 40% did not know how.

The Operational Technology Cybersecurity Coalition (OTCC) convened four roundtables over the summer, bringing together nearly 70 representatives from at least seven critical infrastructure sectors, along with U.S. and foreign government officials and cybersecurity experts. The discussions examined how AI is changing threats to OT that can undermine critical infrastructure, where AI-enabled defenses can help, and what policymakers must do to strengthen OT resilience, thereby enhancing critical infrastructure resilience.

We concluded on a simple premise: operational resilience, grounded in fundamental cybersecurity practices and enhanced with AI-enabled cyber defenses, is the foundation of OT defense in this new age. Now, we need to ensure that AI security, cybersecurity, and critical infrastructure resilience are treated as interconnected national security objectives, supported by coordinated policies, authorities, and investments.

We believe that:

- **Prioritizing resilience is fundamental** in addressing the threat from AI in OT environments;
- **Assuming compromise is not an admission of defeat** but a recognition of the reality of the AI-enabled threat landscape; and
- **Operationalizing public-private partnerships** into OT defense with AI partners is key to the next phase of cybersecurity.

The 18 recommendations in this report build on authorities, programs, and institutions that already exist. Coupled with the further integration of the use of AI in OT cyber defense, with appropriate safeguards, Congress and the Administration can take practical steps to secure the OT systems and assets that Americans rely on for their critical infrastructure.

1. **Strengthening Baseline Cybersecurity to Counter AI-Enabled Threats**
    a. Close Regulatory Gaps in Critical Infrastructure Sectors Without Creating Duplication
    b. Establish Meaningful Secure-By-Design Requirements
    c. Create a Cyber Duty of Care Based On Secure-by-Design Requirements
    d. Provide Funding for Under-Resourced Entities
    e. Establish Whistleblower Protections
    f. Establish Compensating Controls for Legacy and Unpatchable OT Systems

2. **Enabling AI-Cyber Defenses in OT Environments**
    a. Establish an AI Governance Framework for OT Environments
    b. Expand Successful State and Regional Cybersecurity Models
    c. Ensure Long-Term Access to AI-Enabled Cybersecurity Capabilities
    d. Facilitate the Adoption of Certified COTS Cybersecurity Solutions
    e. Test Where AI Could Be Deployed Autonomously to Protect OT Environments
    f. Renew Support for the ISACs
    g. Establish a Public-Private Forum for AI Labs and OT Entities
    h. Strengthen International Norms and Coordination

3. **Enhancing Resilience in Times of Disruption**
    a. Provide Funding to the Cyber Response and Recovery Fund
    b. Reconstitute the Cyber Safety Review Board
    c. Improve the Incident Management System and Process
    d. Embrace CI Fortify Controls

The OTCC is committed to advancing these recommendations beyond this report. In the coming months, we will convene policymakers, AI developers, and the OT community to test and refine these ideas, as well as engage with the Executive Branch, Congress, state and local governments, and AI labs to support implementation.

Introduction

The U.S. government, AI labs, and the Operational Technology (OT) community must act now to enhance the resilience of U.S. critical infrastructure against AI-enabled cyberattacks. This is not a hypothetical threat. Federal agencies have warned that attackers are using AI-generated scripts to target industrial controllers across U.S. energy, water, and manufacturing sectors. Anthropic uncovered a Chinese state-sponsored campaign targeting chemical manufacturers and other organizations, with AI reportedly performing 80-90% of the work. Wiz researchers used AI to uncover exposed administrator access at a public rail operator and an unsecured hospital alert channel, leaving us to wonder what might have happened had the researchers been malicious actors. At a Senate Homeland Security and Governmental Affairs Committee hearing in September, Dragos SVP Kurt Gaudette cited an AI-assisted attack on a Mexican water utility in which an AI model autonomously redirected an attacker toward the control network and made more than 2,800 credential attack attempts in rapid succession. During his testimony, he warned that most small- and medium-sized utilities lack the basic cybersecurity fundamentals needed to defend against these evolving threats. Lastly, a SANS Institute 2026 survey of industry leaders found that 34% of respondents had an industrial control system (ICS)/OT incident in the past year, up from 22% in 2025, and detection slowed for the first time. While this rise cannot be directly attributed to AI-enabled cyber attacks, the temperature is rising on OT systems.

The Operational Technology Cybersecurity Coalition (OTCC) convened a series of four roundtables over the summer fo 2026 to identify specific recommendations policymakers can implement to enhance resilience in the age of AI. The purpose of those roundtables was to discuss AI-enabled cyber attacks against OT systems and how the OT community could use AI to defend itself. None of the participants, however, could have predicted what would transpire during those months. President Trump issued a significant executive order in June to implement voluntary cybersecurity measures for AI models. Then in July, it was revealed that OpenAI's agents broke out of a sandbox months earlier to hack Hugging Face. Just one week after that revelation, an alleged Iranian actor carried out a widespread attack on water and wastewater facilities across at least a dozen states. Although the role of AI in those attacks remains unclear, the Cybersecurity and Infrastructure Security Agency (CISA) released its first Cybersecurity Advisory two weeks later, acknowledging that AI was being used to target OT systems.

We then remembered the 25th anniversary of the 9/11 attacks. "Failure of imagination" is a foundational assessment in the 9/11 Commission report that should inform how we confront AI-enabled cyber attacks. We must imagine a scenario in which a malicious actor uses AI to target multiple U.S. critical infrastructure systems with no regard for consequences. This summer's developments make that scenario hard to dismiss.

The AI OTCC Roundtables focused on what we can do about AI-enabled threats to OT. Those sessions included nearly 70 participants from federal and foreign governments, cybersecurity companies, local utilities, large critical infrastructure owners and operators, nonprofits, and members of the OT community. Together, this diverse group identified three principles that form the basis of the 18 recommendations listed in this report. With these principles and recommendations, the OTCC will continue to educate public and private sector policymakers and work to increase the resilience of OT systems and the essential services they support in the wake of AI-enabled cyberattacks.

Principles

1 Prioritizing Resilience is Fundamental in Addressing the Threat from AI in OT Environments

While emerging AI capabilities introduce novel threat vectors to OT, the foundation of OT defense remains the same: operational resilience grounded in fundamental cybersecurity practices. AI tools can strengthen cyber defenses, yet they alone cannot compensate for a lack of baseline security measures across critical infrastructure. Securing OT environments against evolving AI threats starts with prioritizing fundamental controls such as robust access management, network segmentation, and strict least privilege protocols. The National Security Agency (NSA) doubled down on this point by issuing guidance on implementing zero trust (ZT) for Department of War OT systems to account for AI threats. The NSA concluded that "Adversaries are leveraging [super intelligence] capabilities to automate reconnaissance, accelerate exploit development, and conduct cyber campaigns at unprecedented speed and scale, making achieving ZT outcomes on OT systems more critical than ever." Ultimately, sophisticated defense strategies are only as effective as the resilient infrastructure supporting them.

This approach is particularly important for national AI security policy. Preventing the misuse of AI and strengthening cyber defenses are necessary objectives, but neither can eliminate the possibility of a successful attack. For OT environments, safety is paramount and the ultimate measure of security must also include whether essential services can continue operating during a cyber incident and how quickly those services can be restored.

A resilience approach should apply to AI security, especially to ensure OT systems and assets can withstand disruptions and support critical infrastructure, which is foundational to the nation's resilience. Resilience should be measured not only by the adoption of cybersecurity controls, but also by whether critical infrastructure operators can identify connected assets, detect anomalous activity, and sustain and restore essential operations during a cyber incident. This focus aligns with a bipartisan approach to preserving national resilience. The Biden Administration's National Security Memorandum on Critical Infrastructure Security and Resilience (NSM-22), which is still in effect, defines resilience as the "ability to prepare for threats and hazards, adapt to changing conditions, and withstand and recover rapidly from adverse conditions and disruptions." The current Administration builds off this definition through the 2026 National Resilience Strategy, which focuses on four key goals to achieve resilience: reduce risk, improve reliability of infrastructure and continuity of key services, deliver faster response and recovery, and promote investments that bring value and opportunity for Americans. The United Kingdom, too, recently declared that "AI makes cyber attacks easier, faster, and cheaper to execute…improving cyber resilience fundamentals remains the most effective way to reduce both the likelihood and impact of cyber incidents."

2 In the AI Era, We Must Assume Compromise

AI is accelerating vulnerability discovery and exploitation, placing additional pressure on critical infrastructure operators that already struggle to secure aging OT systems. Unlike traditional IT environments, industrial controllers and other OT assets often support continuous physical operations and cannot be patched, rebooted, or replaced without extensive testing or operational disruption. Roughly 70% of OT vulnerabilities reside in legacy or deeply embedded systems that cannot be safely patched without disrupting operations. While patching remains an essential cybersecurity practice, it cannot serve as the primary defense for every OT environment. CISA applied the same logic to federal civilian agencies in its Binding Operational Directive on Prioritizing Security Updates Based on Risk (BOD 26-04), recognizing that it is impossible to patch everything, everywhere, all at once. Further, CISA has tracked over 12,000 ICS vulnerabilities in its advisories, yet the pool of engineering talent qualified to develop and test industrial patches has not grown to match this exponential pace, creating severe engineering bottlenecks.

Assuming that an OT system will eventually be compromised is not admitting defeat. Rather, it requires us to reframe our approach to enhancing OT resilience and the critical infrastructure they support.

3 Operationalizing Public-Private Partnerships into OT Defense with AI Partners

To build meaningful resilience in the AI era within the OT community, all parties must continue to engage proactively. The federal government has the intelligence on how bad actors will misuse AI. The labs know how best their tools can be misused for bad and used for good. Cybersecurity vendors have commercial off-the-shelf (COTS) technologies that protect organizations against AI-powered attacks. OT owners and operators also understand how best to build resilience into critical infrastructure. We must bring these groups together with state and local partners to secure the nation's critical infrastructure.

AI companies should also engage OT and federal partners to identify mechanisms to safely train AI models on sensitive OT data while maintaining appropriate data safeguards. For electric utilities, for example, data that falls under Bulk Electric System Cyber System Information under Critical Infrastructure Protection (CIP) Standard CIP-011 and Critical Energy/Electric Infrastructure Information under Federal Energy Regulatory Commission (FERC) rules should remain in an owner's environment or in a controlled test bed. OT partners must be included in AI-industry-led initiatives and forums to help inform key policy decisions and AI defensive capabilities. Ultimately, the greatest national security risk posed by AI is one where an OT system that supports critical infrastructure is compromised and degraded. These two communities should therefore encourage active defensive experimentation with AI to bridge the gap between traditional engineering practices and emerging threat realities.

Recommendations

The following 18 recommendations translate these principles into three areas for action: (1) establishing baseline cybersecurity requirements, incentives, and accountability to reduce the likelihood of compromise; (2) building partnerships between the AI and OT communities to accelerate the safe adoption of AI-enabled cyber defenses; and (3) strengthening operational resilience to ensure critical infrastructure can withstand, contain, and recover from successful cyberattacks. Together, these recommendations recognize that while stronger cybersecurity and AI-enabled defenses can prevent many attacks, no defensive strategy can eliminate all risk. National AI security policy must therefore address not only how to prevent a cyberattack, but also how to minimize its consequences when prevention fails.

1 Strengthening Baseline Cybersecurity to Counter AI-Enabled Threats

For years, malicious cyber actors have compromised entities and OT systems through well-known security gaps. They exploited weak passwords, the absence of multi-factor authentication, over-privileged access, and inadequate network segmentation.

AI will only exacerbate this challenge. Recent attacks show that bad actors, whether using AI-enabled tools or not, compromise systems through these same gaps. While cyber defenders could use AI to identify these flaws and address them using established practices, there is no guarantee they will. In other words, too few incentives or requirements exist to compel organizations to act, despite increased AI-enabled cyber attacks.

Meanwhile, organizations are simultaneously deploying their own AI agents and other non-human identities (NHIs) to fight back against AI-enabled attacks. Yet, they are adopting them at a rate that far exceeds their ability to govern them. In this new landscape, where NHIs can outnumber human employees by more than 100 to 1, the very definition of identity has changed. This creates a security concern in its own right by potentially having an exponential increase in insider threats within an organization. Therefore, any organizational cybersecurity framework must account for securing AI agents by extending common security controls such as zero trust and identity access management protocols.

Congress and the Executive Branch must establish a national framework that aligns enforceable cybersecurity requirements, incentives, accountability, and financial assistance to increase the adoption of baseline cybersecurity and identity security controls across critical infrastructure. This includes mandatory, harmonized cybersecurity requirements across critical infrastructure; secure-by-design requirements for hardware and software producers to increase resilience at scale; liability protections for those that fall outside critical infrastructure sectors or hardware and software providers; and financial assistance to under-resourced entities. The framework should establish a common cybersecurity floor based on CISA's Cross-Sector Cybersecurity Performance Goals (CPGs), while allowing sector-specific regulators to impose additional, risk-based requirements. These requirements must also account for AI agents and other NHIs operating with privileged access to sensitive systems. This framework is not meant to add another layer of cybersecurity compliance, but to ensure every critical infrastructure sector meets a minimum level of security commensurate with the risks it poses to national security, public safety, and essential services.

a. **Close Regulatory Gaps in Critical Infrastructure Sectors Without Creating Duplication:** Creating new cybersecurity mandates is possible, and it is not wholly a new idea. NERC CIP has been mandatory since FERC approved version 1 in 2008, the longest-running mandatory OT cyber regime in the country. After the Colonial Pipeline ransomware incident, TSA established cybersecurity requirements for certain critical pipeline and transportation operators. The Coast Guard followed suit by establishing minimum cybersecurity requirements for covered entities in the Marine Transportation System. A participant in the roundtable series representing an entity regulated by these agencies said the regulations were common sense, long overdue, and helped the participant's company get its house in order.

   Congress has recognized the need for a consistent national approach to cybersecurity through the Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA). It must now apply a similar principle to baseline cybersecurity requirements. Congress will need to navigate different committee jurisdictions and existing sector requirements, but it must agree on and establish an enforceable baseline for all 16 critical infrastructure sectors. Where mature cybersecurity regulatory frameworks already exist, including in the nuclear, bulk electric (NERC CIP), financial, aviation, rail, and pipeline sectors, Congress should preserve existing regulators and recognize equivalent or stronger requirements. For sectors that lack enforceable cybersecurity standards, Congress must identify the appropriate federal authority responsible for implementation and oversight. The Office of the National Cyber Director (ONCD), in coordination with CISA, OMB, and Sector Risk Management Agencies (SRMAs), should identify regulatory gaps, harmonize overlapping requirements, and establish mechanisms for measuring compliance and effectiveness across sectors. Congress should require regular reporting on implementation and remaining gaps to ensure that these requirements produce measurable security improvements rather than additional compliance burdens. Congress should also establish a process for determining when compliance with an existing sector-specific requirement satisfies an equivalent cross-sector obligation.

b. **Establish Meaningful Secure-By-Design Requirements:** Both the Biden and Trump administrations have embraced the concept of building cybersecurity into hardware and software at their point of origin rather than as an added feature. Just as all cars come with seatbelts rather than as an added cost after purchase, so too must OT hardware and software include built-in cybersecurity features. Congress should direct ONCD to coordinate with the National Institute of Standards and Technology (NIST), CISA, and SRMAs, and work with industry, to identify where minimum secure-by-design requirements are warranted for OT hardware, software, and AI-enabled products used in critical infrastructure. These requirements should prioritize secure default configurations, strong authentication, vulnerability disclosure and remediation processes, logging capabilities, and mechanisms for securely updating products throughout their supported lifecycle. Requirements should also extend to supply chain systems, assets, hardware, and software, and complement existing executive orders, such as EO 14420-Declaring a National Emergency to Secure the United States Bulk-Power System.

   Congress could also require a mix of internal and independent third-party assessments against recognized national and international cybersecurity standards, including applicable ISA/IEC 62443 requirements. This approach would shift more responsibility to the manufacturers and providers best positioned to improve security at scale, rather than placing the entire burden on critical infrastructure operators that may lack dedicated cybersecurity personnel.

AI labs providing tools and services to critical infrastructure should also incorporate safeguards that protect sensitive operational data, enforce appropriate access controls, and prevent unauthorized or unsafe autonomous actions.

c. **Create a Cyber Duty of Care Based on Secure-by-Design Requirements:** Building off recommendation 1. b., ONCD should work with Congress to establish a duty of care for OEMs, integrators, and OT owners and operators that requires entities to implement reasonable, risk-based cybersecurity safeguards appropriate to their role in the OT ecosystem. The framework should align with CISA's CPGs while recognizing that manufacturers, integrators, and operators have different responsibilities and capabilities to manage cyber risk. For OT owners and operators, these obligations should include fundamental protections such as identity and access management, enforceable network segmentation, secure remote and third-party vendor access, vulnerability management, and incident response planning. For manufacturers and integrators, the duty of care should emphasize secure-by-design and secure-by-default practices, timely vulnerability and incident notification, responsible vulnerability management, and the secure configuration and integration of OT products.

   The framework should also account for the use of AI in operational environments, including appropriate governance and protections for deploying NHIs that adhere to established security principles and least-privilege access policies. The duty of care should complement, rather than replace, existing and future sector-specific cybersecurity regulations. Liability protections should therefore incentivize entities that meet or exceed applicable requirements, rather than serve as an alternative to regulatory compliance.

d. **Provide Funding for Under-Resourced Entities:** Smaller critical infrastructure owners and operators may lack the resources to meet the requirements set forth in government-mandated regulations. Congress must reauthorize and fund the State and Local Cybersecurity Grant Program and make CPG implementation an eligible expense. Federal and state leaders must also explicitly expand allowable uses for existing funding streams such as the DHS Homeland Security Grant Program and State Revolving Funds so they can cover core OT cybersecurity costs like asset discovery, risk assessments, network segmentation design and implementation, continuous monitoring, managed detection and response services, workforce training, compensating controls, legacy OT modernization, and incident response and recovery planning. Congress and administering agencies should assess whether existing funding rules adequately accommodate both recurring service costs and the substantial capital investments required to upgrade or replace aging OT equipment. Funding should prioritize high-consequence assets where unaddressed vulnerabilities pose the greatest risk to essential services.

   ONCD should also review all grants to identify where Notice of Funding Opportunities could be amended to explicitly clarify that cybersecurity and AI token use are allowable expenses. This framework should also permit operators to reserve dedicated cyber-hardening portions from broader infrastructure grants, simplify administrative procurement by reducing cost-share burdens, give states flexibility in distributing funds, and tie all funding directly to CPGs. Additionally, ONCD and the Department of Energy should work with the National Association of Regulatory Utility Commissioners on cost-recovery guidance for OT security and resilience.

e. **Establish Whistleblower Protections:** Congress should establish comprehensive federal whistleblower protections for employees and contractors across the AI and OT ecosystems who identify and report cybersecurity vulnerabilities, unsafe AI capabilities, or security practices that they reasonably believe could threaten national security, public safety, or critical infrastructure resilience. California's Senate Bill 53, the Transparency in Frontier Artificial Intelligence Act, provides an important starting point by protecting certain employees of frontier AI developers who report significant AI-related safety risks or violations of the law. However, these protections do not extend broadly to contractors or the wider cybersecurity and OT ecosystem, where vulnerabilities in industrial control systems, insecure products, and the deployment of AI-enabled technologies could have significant real-world consequences.

   Congress should build on this approach by establishing a consistent federal framework that protects employees and contractors at AI labs, cybersecurity companies, OT manufacturers, system integrators, and critical infrastructure owners and operators who reasonably believe they are reporting serious security risks, including the concealment of known vulnerabilities, inadequate safeguards, or unsafe deployment of autonomous AI systems in critical infrastructure environments. Such a framework should prohibit retaliation, establish confidential and secure reporting channels to appropriate federal authorities, including CISA and relevant sector regulators, and provide meaningful enforcement mechanisms and remedies for affected individuals. By empowering those closest to these technologies to raise security concerns without fear of retaliation, Congress can strengthen accountability throughout the technology supply chain and help prevent cyber incidents that could disrupt essential services or threaten public safety.

f. **Establish Compensating Controls for Legacy and Unpatchable OT Systems:** Many critical infrastructure operators rely on aging OT systems that cannot be safely patched, upgraded, or replaced without significant costs or operational disruption. CISA, in consultation with NIST, SRMAs, manufacturers, and owners and operators, should establish guidance recognizing compensating controls as a risk-based alternative when patching is not feasible (NERC CIP provides a good precedent through its Technical Feasibility Exceptions). These measures should include network segmentation, industrial firewalls, secure remote access, application allow-listing, and passive monitoring, aligned with CISA's CPGs and ISA/IEC 62443 standards.

   Congress and sector regulators should recognize documented, risk-based compensating controls as an acceptable means of meeting applicable cybersecurity requirements under Recommendations 1.a and 1.c, provided they adequately mitigate the underlying risk. Federal funding mechanisms outlined in Recommendation 1.d should also prioritize compensating controls and the modernization or replacement of high-consequence legacy OT assets. Manufacturers should provide clear end-of-support timelines and security guidance for products that can no longer be updated.

2 Enabling AI-Cyber Defenses in OT Environments

Forging connective tissue between AI labs, the federal government, state and locals, and OT partners will be critical to strengthening resilience against AI-enabled cyber attacks. AI developers possess unique knowledge of emerging model capabilities and potential misuse, while OT manufacturers, cybersecurity providers, and critical infrastructure operators understand the operational realities of securing industrial systems. Yet, AI adoption rates within OT have been slow. Nozomi Networks found that fewer than 8% of organizations they surveyed have deployed AI across multiple OT cybersecurity functions.

Congress, the Executive Branch, and private-sector entities should establish sustainable partnerships that bring these communities together to identify high-value applications of AI, test defensive capabilities in realistic OT environments, and accelerate the adoption of proven technologies across critical infrastructure.

These partnerships should include OT owners and operators, industrial automation providers, and cybersecurity service providers alongside AI developers, so that pilots address real-world deployment barriers, including legacy infrastructure, workforce constraints, interoperability, and the need to maintain operational continuity. Federal agencies should use pilot findings to develop practical, risk-based guidance that helps operators evaluate and procure effective AI-enabled defensive capabilities.

a. **Establish an AI Governance Framework for OT Environments:** NIST, in consultation with CISA and SRMAs, and working with AI developers, OT manufacturers, integrators, and critical infrastructure owners and operators, should develop an AI governance framework specifically tailored to OT environments. Building on NIST's AI Risk Management Framework, Cybersecurity Framework 2.0, and existing OT security guidance, the framework should establish best practices for securely deploying, managing, and monitoring AI systems and NHIs in operational environments. NHIs should be assigned unique, verifiable identities, subject to least-privilege access controls, and continuously monitored and logged to ensure accountability for their actions. NHIs should also be subject to network segmentation policies that restrict their ability to access or move between sensitive OT environments. The framework should address human oversight, restrictions on autonomous actions affecting physical processes, AI defenses deployed in OT that should keep working when isolated, and safeguards to ensure AI systems do not compromise operational safety or resilience.

b. **Expand Successful State and Regional Cybersecurity Models:** Congress and the Executive Branch should identify, evaluate, and expand successful state and regional initiatives that strengthen cybersecurity for critical infrastructure, particularly for small and under-resourced utilities. CISA, in coordination with SRMAs, and working with state governments, should evaluate these initiatives to identify approaches that can be replicated across different critical infrastructure sectors and geographic regions. Congress should support pilot programs that bring together state governments, AI developers, OT cybersecurity providers, military installations, and critical infrastructure operators to deploy and evaluate defensive capabilities. These pilots should measure improvements in vulnerability remediation, threat detection, incident response, and continuity of operations, with successful approaches made available for broader adoption. Federal support should prioritize scalable models that can serve both large infrastructure systems and smaller rural operators that lack dedicated cybersecurity resources.

c. **Ensure Long-Term Access to AI-Enabled Cybersecurity Capabilities:** OpenAI and Anthropic have announced initiatives to help critical infrastructure entities use AI-enabled cyber tools to defend against cyber attacks. These initiatives present an important opportunity to expand access to advanced defensive capabilities, particularly for organizations that could not otherwise afford them. However, short-term commitments of free or subsidized model access will not provide a sustainable foundation for critical infrastructure cybersecurity.

   Public-private initiatives should prioritize sustainable adoption rather than short-term access alone. Congress and relevant agencies should examine ways to support ongoing evaluation, integration, training, monitoring, and maintenance of AI-enabled defensive tools, particularly for smaller operators that lack dedicated OT cybersecurity personnel. In terms of the latter, SANS 2026 survey found that 68% of respondents said losing one or two key people would delay or severely impair their OT response and recovery. Congress and AI developers should establish mechanisms to ensure that under-resourced critical infrastructure operators can continue accessing effective AI-enabled cybersecurity tools after initial private-sector commitments expire. These mechanisms could include eligible expenses under existing federal cybersecurity grant programs, shared-service arrangements, public-private partnerships, and additional in-kind contributions from AI developers.

d. **Facilitate the Adoption of Certified COTS Cybersecurity Solutions:** To accelerate the hardening of critical infrastructure against AI-driven threats, the federal government should streamline the procurement and deployment of proven COTS cybersecurity solutions for OT owners and operators. Many entities, particularly those that are under-resourced, struggle to identify and afford effective security tools. Congress and CISA should promote the use of cybersecurity solutions independently certified against recognized OT product security standards, such as ISA/IEC 62443-4-1 and 62443-4-2.

e. **Test Where AI Could Be Deployed Autonomously to Protect OT Environments:** While malicious actors will increasingly use AI to augment their cyber attacks, OT owners and operators may be wary of deploying AI within their operational environments for a variety of reasons, one of which is the inability to control the data that an AI tool operates on and the need to restrict training and use of sensitive data. The federal government should therefore establish a coordinated research and testing program to determine where AI-enabled cyber defenses can safely operate autonomously in OT environments and where human authorization must remain required.

   DOE, Federally Funded Research and Development Centers, National Laboratories, AI developers, OT cybersecurity vendors, OEMs, and critical infrastructure owners and operators should use secure test environments and digital twins to evaluate AI-enabled defensive capabilities against realistic cyber threats. This should build upon existing capabilities rather than create a new testing infrastructure from scratch. INL, for example, is one of many National Labs that maintain extensive OT cybersecurity test environments and is developing AI-enabled capabilities for anomaly detection, autonomous cyber defense, digital twins, and resilient control systems. These capabilities could provide the foundation for a broader public-private testing program across multiple critical infrastructure sectors. Testing should also evaluate whether AI-enabled monitoring and detection can strengthen compensating controls for legacy OT systems that cannot be safely patched. The government should ensure these efforts are coordinated and complementary.

Congress or the Executive Branch must establish appropriate safeguards for sensitive OT data used in testing, including restrictions on unauthorized model training, retention, and disclosure. These efforts should provide a foundation for the safe adoption of AI-enabled defensive technologies across multiple critical infrastructure sectors, allowing defenders to benefit from increasingly capable AI systems without introducing unacceptable risks to essential services.

f. **Renew Funding for the Multi-State Information Sharing and Analysis Center (MS-ISAC) and Continue Support for Relevant Sector ISACs:** Congress should restore federal support to the MS-ISAC to continue providing services and tools to publicly owned critical infrastructure owners and operators. Federal funding should support core shared capabilities such as cyber threat intelligence, 24/7 security operations, incident response, vulnerability management, information sharing, and technical assistance. This will also help under-resourced public critical infrastructure implement CISA's CPGs as detailed in Recommendation 1. Further, this will enable the MS-ISAC to serve as a trusted intermediary between the AI labs and local utilities. OpenAI's Daybreak initiative, for example, includes a public sector and water-focused pilot with the MS-ISAC to train and support State, Local, Tribal, and Territorial (SLTT) cyber defenders. The pilot pairs Daybreak access with guided training and hands-on assistance for an initial group of public sector and water system defenders, helping them validate and prioritize findings, coordinate remediation, and develop a repeatable approach that can be expanded over time. Additionally, sector ISACs, such as the electricity ISAC (E-ISAC), water ISAC (WaterISAC), and others, should continue to receive support and be reviewed to determine how to best continue serving their customers.

g. **Establish a Public-Private Forum for AI Labs and OT Entities:** AI labs have launched various initiatives to convene trusted partners to access their highly capable models (e.g., Daybreak, Glasswing). Yet these do not necessarily include every group needed to secure critical infrastructure, such as CISA, other AI labs, local utilities, OEMs, integrators, MSSPs, and others. CISA, through ANCHOR-CI, should create a council that convenes these stakeholders to identify high-value applications of AI for OT cybersecurity, address adoption barriers, and determine where federal policy, resources, or partnerships can accelerate the safe deployment of AI-enabled cyber defenses across critical infrastructure.

   CISA should use the council as a forum to develop recommendations for federal research priorities, procurement guidance, technical assistance, and public-private pilot programs. NIST should complement this effort by expanding relevant AI consortium activities to include greater participation from OT owners and operators. Through these efforts, government and private-sector partners can produce actionable recommendations and scalable approaches that accelerate the safe adoption of AI-enabled cybersecurity capabilities.

h. **Strengthen International Norms and Coordination:** The State Department, in coordination with ONCD, should continue to engage in international collaboration, leveraging bilateral and multilateral cyber diplomacy to align on international norms and standards for building cyber resilience and crisis coordination. The U.S. should lead efforts to improve information sharing on threats and incident reporting and to clarify how existing international law and voluntary norms apply to AI-enabled cyber operations.

The EU's Cyber Resilience Act (CRA) shows why this matters. The CRA requires manufacturers of products with digital elements to report actively exploited vulnerabilities and severe incidents through a single reporting platform. As a result, the EU is building a large and growing body of vulnerability and incident data to which the U.S. may have limited access, while the EU sets the pace on global product security regulation. ONCD and the State Department, in coordination with CISA, should pursue a reciprocal information-sharing arrangement with the European Union Agency for Cybersecurity (ENISA) so that U.S. defenders can benefit from this data, and should work with allies to establish clear handling and protection requirements for it, since aggregated vulnerability information is itself a high-value target. They should also seek to harmonize CRA and CIRCIA reporting requirements where possible so that companies operating in both markets are not managing duplicative or inconsistent obligations. Finally, the United States should coordinate with allies on secure-by-design expectations for OT products and AI-enabled technologies, so that it helps shape global product security practices rather than adopt standards set elsewhere.

3 Enhancing Resilience in Times of Disruption

Malicious actors will increasingly use AI to identify vulnerabilities, compromise systems, and execute attacks with greater speed and scale. Even organizations that implement strong cybersecurity controls cannot eliminate every vulnerability. While implementing basic cybersecurity controls will reduce the overall risk landscape, it will not eliminate risk.

Federal cybersecurity policy should emphasize the ability to withstand, contain, and recover from a successful attack. OT owners and operators and the broader U.S. government should adopt ZT principles for AI models and agents, embrace CISA's CI Fortify approach, and assume compromise in either their system or in a system they rely upon. Congress and the Executive Branch should strengthen national cyber response and recovery capabilities, improve coordination between cybersecurity professionals and emergency responders, and support investments in operational resilience.

Federal resilience initiatives should encourage critical infrastructure owners and operators to regularly assess recovery readiness, including whether their security teams can detect incidents across operational assets, coordinate OT-specific response, and restore priority functions with minimal disruption. Policymakers should consider technical assistance and voluntary exercises that evaluate these capabilities in realistic operating conditions.

a. **Provide Funding to the Cyber Response and Recovery Fund:** The Cyber Response and Recovery Fund allowed CISA to contract with private sector companies to provide vulnerability assessments and mitigation, incident mitigation, malware analysis, and other services when a significant cyber incident was likely or had occurred. As AI enables malicious actors to conduct attacks with greater speed and scale, the federal government must be prepared to rapidly deploy specialized expertise and capabilities to assist affected organizations, particularly small and under-resourced critical infrastructure operators. Congress should ensure that the fund can support the procurement and deployment of appropriate AI-enabled defensive tools to analyze malicious activity, identify compromised assets, and accelerate containment and recovery.

b. **Reconstitute the Cyber Safety Review Board (CSRB):** Several AI labs and policymakers have called for a National Transportation Safety Board body to investigate AI hacks and call for improvements. The CSRB was initially intended to do this type of investigation for all forms of cyber attacks, which could include AI incidents. Yet, it lacked enforcement mechanisms, authority, and resources. CISA should re-establish the CSRB, include members from the OT and AI communities, and perform a broad investigation into AI-enabled cyber attacks conducted by AI agents. Congress should then codify the CSRB into law to give the board the institutional stability, funding, and statutory mandate needed to analyze complex security incidents. Congress should also establish a mechanism requiring relevant federal agencies to formally respond to CSRB recommendations and report on their implementation.

c. **Bridge the Gap Between Emergency Responders and Cybersecurity Professionals by Expanding the National Incident Management System (NIMS) and Incident Command System and Clarifying PPD 41 and 44:** NIMS and the Incident Command System are established FEMA frameworks that standardize how the government and private sector prepare for, respond to, and recover from incidents. Emergency managers and workers across state and local governments use this process to ensure seamless workflows during an incident. As cyber incidents increasingly target OT systems and assets that could cause physical disruptions akin to a natural disaster, cybersecurity professionals need training in NIMS and the Incident Command System. FEMA, CISA, and Sector Risk Management Agencies should expand NIMS and Incident Command System training opportunities for SLTT and private-sector cybersecurity personnel and incorporate cyber-to-physical scenarios into exercises involving emergency managers, public-safety agencies, OT operators, and federal cyber responders.

   Additionally, the White House must issue a national policy that clarifies how PPD-41 "United States Cyber Incident Coordination" transitions into PPD-44 "Enhancing Domestic Incident Response." Currently, no national policy or plan articulates how the U.S. government will manage a national crisis that starts in the cyber domain and transitions into physical consequences. This need was highlighted during a Senate hearing on the threat of AI agents, where Dragos's Kurt Gaudette called for a unified federal OT incident response plan. Clarifying the lead federal agency, resources, and coordination mechanisms will ensure improved response and recovery operations for an AI-enabled cyber attack that causes widespread physical disruptions.

d. **Embrace CI Fortify Controls:** In addition to the security controls mentioned in Recommendation 1, OT owners and operators should embrace CI Fortify controls. This includes identifying critical operational dependencies, implementing enforceable IT/OT network boundaries, isolating safety-critical systems, restricting lateral movement between network segments, and maintaining the ability to isolate compromised equipment while preserving essential operations. Operators should also establish manual operating procedures where feasible and regularly test restoration and recovery capabilities. In the SANS 2026 survey, only 15% of respondents said they could fully sustain manual operations past 48 hours, and 40% did not know how. Only 16% have ICS/OT cyber resilience fully integrated into business continuity and disaster recovery, and 29% test their resilience plans at least annually. Organizations doing both expected same-day recovery 46% of the time, against 27% overall.

   These measures are particularly important as AI enables malicious actors to identify vulnerabilities and move through compromised environments faster. Congress should direct CISA and Sector Risk Management Agencies to develop sector-specific implementation guidance and support exercises

that demonstrate whether operators can maintain essential services during a significant cyber incident. AI-enabled defensive tools could further strengthen these capabilities by identifying anomalous activity, mapping operational dependencies, prioritizing critical assets, and recommending containment and recovery actions.

Conclusion

The past several months have made clear that AI-enabled threats to operational technology are no longer hypothetical. Attackers are leveraging AI to discover vulnerabilities, compromise devices, and target industrial controllers at a speed and scale that traditional defenses were not designed to meet. The recommendations presented in this report rest on a simple premise: operational resilience, grounded in fundamental cybersecurity practices and enhanced with AI-enabled cyber defenses, is the foundation of OT defense in this new age. Many of these actions build on authorities, programs, and institutions that already exist. What is needed now is prioritization, sustained funding, and both the will and imagination to act before the next crisis arises. We are cognizant that CISA is responsible for several of these recommendations, yet is dealing with low staffing levels and insufficient resources. Rebuilding and supporting CISA and the SRMAs is imperative for the nation's cybersecurity. As AI continues to lower the barrier to sophisticated attacks, the cost of delay will only grow in size and scale.

The Operational Technology Cybersecurity Coalition is committed to advancing these recommendations beyond this report. In the coming months, the OTCC will publish follow-up papers that examine the most pressing of these issues in greater depth, host events that bring together policymakers, AI developers, and the OT community to test and refine these ideas, and engage directly with the Executive Branch, Congress, state and local governments, and AI labs to support implementation. Although the window to build resilience before the next major AI-enabled cyber attack is narrowing, it is not yet closed. The OTCC looks forward to working with partners across the public and private sectors to make the most of it.

Acknowledgments

The OTCC would like to thank Michael Garcia, Thomas Buckley, Tatyana Bolton, Katy Blake Burch-Hudson, and Faith Grant in the development of this report. We would also like to thank the participants who engaged in our roundtables for their insights and contributions to this report.

Operational Technology Cybersecurity Coalition

About the OT Cyber Coalition

The Operational Technology Cybersecurity Coalition (OTCC) is dedicated to advancing cohesive, risk-based, and globally interoperable cybersecurity policy for industrial and operational technology environments. OTCC brings together leading experts to advocate for consensus-based standards that reduce regulatory fragmentation and strengthen critical infrastructure resilience worldwide.

Info@otcybercoalition.org | www.otcybercoalition.org