Velocity

V5 2026

Booz Allen

INSIGHTS FOR INNOVATORS

AI vs AI

REIMAGINING CYBER FOR A FASTER FIGHT

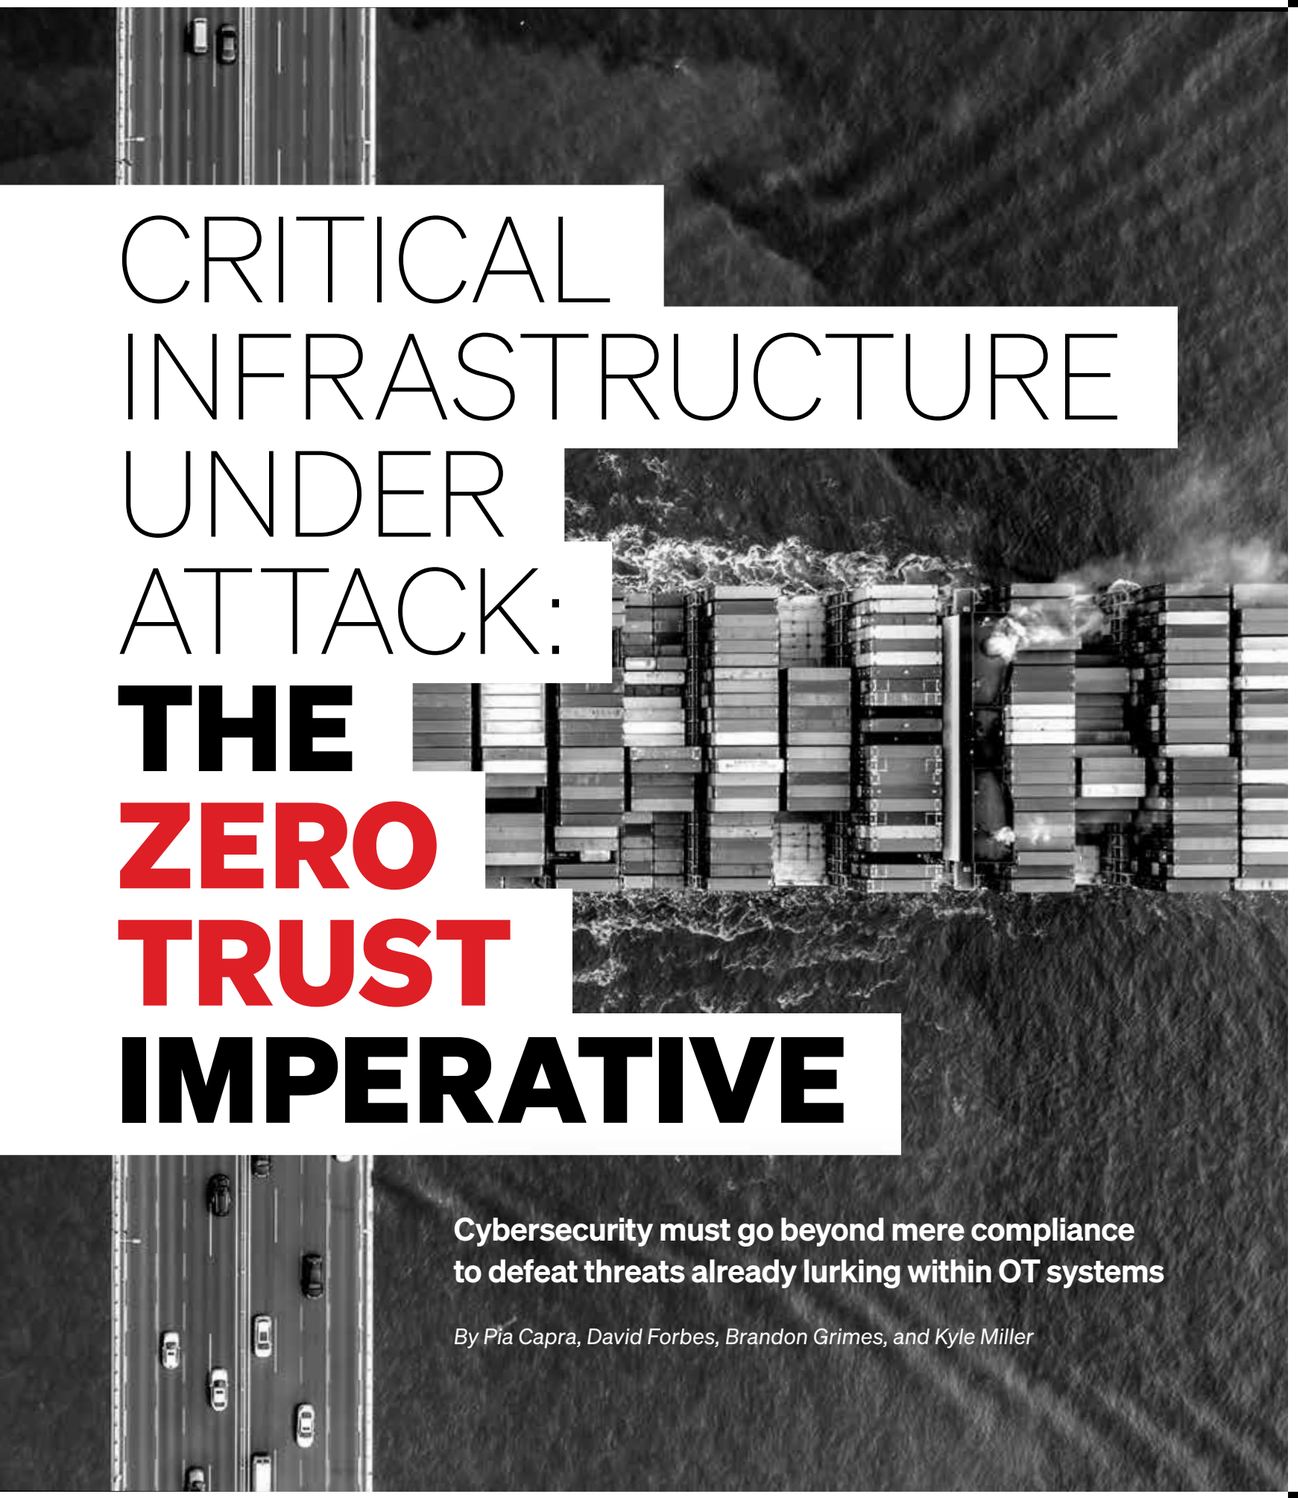

CRITICAL INFRASTRUCTURE UNDER ATTACK: THE ZERO TRUST IMPERATIVE

Cybersecurity must go beyond mere compliance to defeat threats already lurking within OT systems

By Pia Capra, David Forbes, Brandon Grimes, and Kyle Miller

The Department of War (DoW) deserves credit for moving zero trust from concept to implementation in enterprise IT. A 2025 Pentagon memorandum institutionalized governance, set department-wide expectations, and pushed zero trust adoption as an operational priority, not just a cybersecurity goal.

But that progress also reveals a sharper truth: mission assurance does not stop at traditional IT boundaries. The ability to support deployment, sustainment, and combat increasingly depends on operational technology (OT) and on privately owned critical infrastructure, such as power, water, transportation, logistics, and critical manufacturing systems. Adversaries are already targeting these systems. For some civilian networks, at a minimum, they have successfully infiltrated and targeted assets with pre-positioned threats that could be used for disruptive attacks or coercive leverage in a crisis.

That is a critical reason why DoW's recent *Zero Trust for Operational Technology Activities and Outcomes* matters so much now: it formalizes the application of zero trust principles to industrial control systems and other OT environments. The question is no longer whether zero trust principles belong in OT environments but how to apply them in systems where safety, availability, and mission continuity are non-negotiable. The importance of these systems was reinforced by the recently released *President Trump's Cyber Strategy for America*, which elevates protection of critical infrastructure and the defense industrial base as strategic imperatives.

In this article, we assess how viable the new zero trust for OT mandates are in practice, where agencies can streamline implementation across IT/OT boundaries, and what it will take to translate compliance activity into the outcome that matters most: resilient operations under real-world conditions in a contested environment.

The Erosion of Long-Term Security Expectations and Emergence of New Paradigms

For decades, a primary security model for OT was straightforward: *keep it isolated*. Organizations engineered industrial control systems (ICSs), supervisory control and data acquisition (SCADA) platforms, building-automation networks, energy management systems, and weapon system components for reliability and uptime while maintaining "air gaps" to ensure nothing malicious could reach them. But air gaps eventually proved to be ineffective. In 2010, Stuxnet crossed the air gap at Iran's Natanz enrichment facility via a compromised USB drive, installing itself on the programmable logic controllers (PLCs) controlling centrifuges and causing physical destruction while feeding operators false sensor data. This incident demonstrated that adversaries could reach isolated systems without network connections.

At the same time, the business and mission benefits of IT/OT convergence—real-time operational visibility, predictive maintenance, remote monitoring, and the efficiency gains of industrial Internet of Things (IoT) integration—were becoming more compelling. What has replaced the air gap is a busy connectivity zone that adversaries enter and don't leave, extending dwell time often for months or years before detection, with consequences that extend beyond the enterprise into critical infrastructure, defense systems, and national security.

The scale of these once-quiet exploits has come into sharper focus. In 2024, roughly 70% of cyberattacks involved critical infrastructure, according to the House Homeland Security Committee. Facilities suffering cyber-driven physical disruptions jumped 146% year-over-year, from 412 sites in 2023 to 1,015 in 2024. China's cyber espionage operations targeting industrial and manufacturing sectors have recently surged 150%. And overall hacktivist campaigns against ICSs increased 51% into 2025. These attack spikes underscore the fragility of the nation's critical infrastructure.

To meet policy mandates, agencies will now implement zero trust—the cyber architecture already transforming federal IT—in OT systems where failure, in the worst cases, can have negative effects such as impaired warfighting readiness, contaminated water supplies, electric grid blackouts, or, in very rare instances, loss of life. The objective is not just compliance. It is to achieve cyber resilience that encompasses prevention, mitigation, and recovery from attacks, all to eliminate the possibility of catastrophic failure from malicious activities.

UNDERSTANDING ZERO TRUST

ZERO TRUST refers to an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources. A zero trust architecture grants no implicit trust to assets or user accounts. Key principles include: "Assume a breach;" "Never trust, always verify;" and "Allow only least-privileged access based on contextual factors."

To date, federal zero trust implementation efforts have focused primarily on IT environments. However, the threat landscape, the evolving federal policy framework, and technology advances have driven a natural progression of zero trust principles to OT.

OT is the broad category of programmable systems and devices that interact with the physical environment or manage devices that interact with this environment. These devices and systems thus influence "the real world." OT encompasses ICSs, weapon systems (which typically are OT with embedded IT components), building automation systems, transportation systems, communications systems, physical access control systems, and physical environment monitoring and measurement systems, among others.

Zero trust is an enabling framework, not an end in itself; resilient infrastructure is the ultimate goal.

Fragile Infrastructure, With Access Assumed

In a global networked environment where devices outnumber people two to one, new vulnerabilities are emerging. Everyday physical systems—HVAC, traffic lights, medical devices, emergency management systems—are now integrated with IT networks, expanding the attack surface. Recent cyberattacks against OT share key traits: they exploit this convergence of IT and OT, use legitimate credentials and native tools that defeat conventional detection, gain initial access through internet-facing devices with known vulnerabilities or default credentials, and establish dwell times from months to years. A 2024 report found that 98% of organizations experiencing cyberattacks said their IT security incidents also affected OT environments, indicating collapsed domain boundaries.

Volt Typhoon is the most well-known example of a modern OT attack that includes the use of pre-positioning. This Chinese state-sponsored operation has been active against U.S. OT infrastructure since at least 2021. In 2025, cyber firm Dragos disclosed the first confirmed Volt Typhoon compromise of the U.S. electric grid: an approximately 9-month intrusion at a Massachusetts utility in 2023 during which the adversary group exfiltrated data on OT operating procedures and grid layout to build knowledge for exploitation and impact later.

Rather than deploying custom malware, Volt Typhoon works mostly through living-off-the-land (LOTL) techniques involving PowerShell, Windows Management Instrumentation, and native system utilities that function the same as everyday administrative tools. Access comes through vulnerabilities in internet-facing network appliances, with traffic routed through compromised home routers to obscure its origin. Signature-based detection, a key element of conventional cyber tooling, fails against this approach by design. According to the MITRE ATT&CK for ICS framework, the observed techniques for Volt Typhoon's previous activity generally do not extend to the impact tactic.

A similar Chinese state-sponsored effort, Salt Typhoon, demonstrated the same emphasis on stealth, persistence, and strategic intelligence collection. In 2024, Salt Typhoon compromised at least nine major U.S. telecommunications providers, enabling the interception of senior government officials' communications in real time. While the campaign focused on telecom rather than OT, the objective—deep, long-term access to critical infrastructure for operational advantage—mirrored the pre-positioning behaviors seen in Volt Typhoon operations.

Similarly, Russia's Sandworm unit—the GRU operation behind the 2015–2016 Ukrainian power grid attacks—has expanded its targeting to NATO infrastructure. In early 2024, Sandworm operatives or proxies working behind hacktivist personas manipulated water treatment controls at a Texas facility. By December 2025, the group apparently had deployed DynoWiper malware against Polish combined heat-and-power plants and a renewable energy coordination system during peak winter demand. Defenders contained the attack before outages occurred, but the operation confirmed Russia's readiness to strike allied OT networks.

"IN A GLOBAL NETWORKED ENVIRONMENT WHERE DEVICES OUT NUMBER PEOPLE TWO TO ONE, NEW VULNERABILITIES ARE EMERGING."

And Iran is another adversary known to target critical infrastructure. In late 2023, its CyberAv3ngers, an Islamic Revolutionary Guard Corps-affiliated group, was found to have exploited default credentials on PLCs at a Pennsylvania water facility's booster station. With the launch of open hostilities with the United States, these activities have reportedly escalated, with the Cybersecurity and Infrastructure Security Agency (CISA) issuing a security advisory warning that the group continues to "exploit weaknesses in operational technology (OT) devices accessible from the internet."

Clarifying the threat picture, researchers have demonstrated that anyone with an internet connection can map infrastructure dependencies across military installations and the private sector using only publicly available information in just 20 to 30 hours per installation. Search engines such as Shodan and Censys freely index exposed OT devices; Metasploit, integrated with Shodan, provides ready-to-use exploit modules for common OT systems. Although DoW OT systems are rarely exposed directly to the internet, procurement records, contract award announcements, and other open-source data routinely surface sensitive infrastructure details that could inform targeting. And the private sector has similar exposure. As a result, security leaders should assume that adversaries already possess accurate information about OT infrastructure, including viable cyberattack paths.

PRE-POSITIONING refers to the practice by which adversaries—typically nation-state actors—infiltrate target networks months or years before any intended offensive action, establishing persistent, covert access that they can activate when geopolitical or operational circumstances align.

As they wait, pre-positioned actors map the environment, identify critical systems, and embed themselves deeply enough to establish persistence.

In this way, they maintain their ability to degrade, disable, or destroy critical systems at a strategic future moment.

Easy to Disrupt, Hard to Destroy—For Now

The most impactful OT security incidents don't always require adversaries to directly attack infrastructure. The 2021 Colonial Pipeline attack is one example. Ransomware compromised billing and invoicing systems, not OT. The pipeline remained functional, but operators shut it down voluntarily because they could no longer safely monitor or measure product flow. In a similar incident, cybercriminals gained access to a natural gas compression facility's IT network via a spearphishing link before pivoting to the OT network and deploying ransomware across both environments.

Only a handful of nation-states—mainly China and Russia—have demonstrated the ability to conduct truly destructive OT cyberattacks. By contrast, ransomware and hacktivist groups, despite growing OT activity, typically lack the engineering expertise and the patience required to cause catastrophic physical damage. In addition, causing this kind of damage to OT systems would likely undermine their financial objectives. However, these groups remain the most consequential threat most critical infrastructure and commercial operators face, responsible for some of the largest OT-related financial impacts on record, including the JBS Foods production halt and disruptions at Jaguar Land Rover in addition to the Colonial Pipeline shutdown.

AI is beginning to change these dynamics. Large language models accelerate analysis of exfiltrated operational data and lower the level of manpower and expertise needed for sophisticated attacks, making capabilities that once required teams of engineers more accessible.

Anthropic's Claude Mythos Preview—a frontier AI model with highly advanced capability at discovering and exploiting software vulnerabilities—offers an early glimpse of how AI will further reshape the cyber landscape. In testing, Mythos identified thousands of previously unknown flaws across every major operating system and browser, converting some into working exploits with little or no human guidance. The UK AI Security Institute found that Mythos succeeded at expert-level hacking challenges 73% of the time, and other frontier models could soon cross this threshold.

These developments have direct implications for OT. Enterprises have long relied on proprietary firmware, niche protocols like Modbus and DNP3, and physical isolation, yet these embedded systems rarely undergo the systematic security analysis Mythos automates. Many can't be patched without halting operations, creating conditions ripe for long-dwell, high-consequence intrusions. While Anthropic's defensive coalition, Project Glasswing, extends access to major IT vendors, observers note the apparent omission of OT. Meanwhile, stockpiled zero-day exploits now depreciate faster as AI-driven discovery accelerates, incentivizing earlier deployment before defenders can respond.

"ONLY A HANDFUL OF NATION-STATES—MAINLY CHINA AND RUSSIA—HAVE DEMONSTRATED THE ABILITY TO CONDUCT TRULY DESTRUCTIVE OT CYBERATTACKS."

Government and commercial leaders can frame OT risk in these two ways:

- The near-term threat, now open to a broad range of adversaries, is disruption: loss of visibility, operational shutdown, extended recovery, and eroded public confidence. Even seemingly contained incidents can produce business or mission failure.
- The longer-term, higher-consequence threat, is physical destruction, currently concentrated in nation-state actors but expanding as AI lowers barriers.

Both warrant investment, even as resource competition between OT security and traditional enterprise priorities (e.g., control system upgrades vs. new aircraft) creates difficult decisions.

Tracking Federal Policy Mandates

The federal policy framework has been moving in a clear direction, although the pace and specificity of OT-focused guidance have lagged behind the IT side. NIST SP 800-207 established zero trust architectural principles; NIST SP 800-82 Revision 3 expanded OT security guidance and specified integrated zero trust, with caveats; and DoW's 2022 Zero Trust Strategy set a target-level implementation deadline of the end of Fiscal Year 2027, later codified by DTM 25-003 in July 2025.

As mentioned, the *Zero Trust for Operational Technology Activities and Outcomes guidance*, issued by the DoW CIO's Zero Trust Portfolio Management Office in November 2025, is another consequential OT-specific policy. It defines 105 zero trust activities (84 target-level and 21 advanced) across seven pillars: users, devices, applications and workloads, data, networks and environments, automation and orchestration, and visibility and analytics. It distinguishes between an operational layer and a process control layer. This guidance empowers DoW asset owners to tailor these requirements to their systems and environments. Compliance against all 105 activities is neither required nor expected in most cases.

More recently, CISA and interagency partners—including DoW—released a joint guide on *Adapting Zero Trust Principles to Operational Technology*. Its recommendations, including asset visibility, segmentation, strong identity, and continuous monitoring, closely align with and reinforce the operational and architectural themes discussed in this article and the earlier requirements issued by DoW. The guide also elevates supply chain risk management as a core zero trust consideration. Our analysis complements this new guidance by incorporating DoW policy requirements as well as commercial implementation realities.

An updated DoW Zero Trust Strategy is expected in the near future, with additional guidance for weapons systems and defense critical infrastructure.

A GLOBAL PERSPECTIVE ON ZERO TRUST FOR OT

In January 2026, CISA and the UK's National Cyber Security Centre (NCSC-UK), in collaboration with the FBI and international partners, released *Secure Connectivity Principles for Operational Technology,* a joint framework comprising eight principles for designing, securing, and managing connectivity into OT environments.

Developed with agencies in Australia, Canada, Germany, the Netherlands, and New Zealand, the principles are framed as goals rather than minimum compliance requirements.

The principles address risk-opportunity balance, connectivity exposure limitation, network connection standardization, secure protocol adoption, OT boundary hardening, compromise impact limitation, comprehensive logging and monitoring, and isolation planning. The guidance addresses operators of essential services across critical infrastructure sectors and has no implementation deadline.

This guidance and DoW's *Zero Trust for Operational Technology Activities and Outcomes* share basic zero trust principles—least-privilege access, network segmentation, continuous monitoring, and strong authentication—but differ in other ways. The CISA/NCSC-UK guidance is advisory and applicable across both public and private critical infrastructure sectors worldwide.

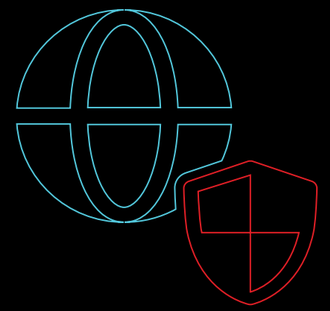

The Complexity of OT Zero Trust

As the federal government focuses on translating zero trust principles into OT-native implementations that preserve operational uptime, agencies will need different tools, timelines, and risk calculations compared with zero trust for IT, particularly when the goal is detecting and neutralizing access that may already exist. While IT cybersecurity often focuses on preventing data breaches, OT protocols focus on keeping plants and facilities up and running.

To start, a significant complicating factor is that many of the OT environments essential to federal missions fall outside direct federal control. National security and defense operations rely on a mix of government-owned and managed systems, government-owned but contractor- or vendor-managed systems, and fully vendor-owned or -operated infrastructure. Civilian power, water, gas, and telecommunications networks that feed defense installations often sit under municipal, regional, or private governance. This creates a fragmented security landscape in which federal agencies must secure mission outcomes without full authority over underlying OT assets.

Legacy device constraints are another immediate barrier to quick zero trust implementation. OT devices in many federal and commercial environments often have lifecycles of 15 years to more than 30 years. These aging PLCs, remote terminal units (RTUs), and legacy SCADA systems were engineered for reliability and availability, not cybersecurity. Many cannot run authentication software, accept patches without operational risk, or participate in identity-based access control. Adversaries have already mapped these attack surfaces.

OT also brings a stringent requirement for availability. Where IT security treats patching and rebooting as routine, any control that causes even brief downtime on a power grid, water system, or weapons platform is likely to be unacceptable. Zero trust controls for OT must be deployed non-disruptively.

Asset visibility is another challenge. NIST has identified incomplete asset inventory as a major barrier to zero trust adoption. The challenge is harder in OT environments where devices may be geographically dispersed, decades-old, and uncatalogued. Many also communicate through specialized industrial protocols—such as Modbus, DNP3, and BACnet—that were not designed for security and are often incompatible with standard IT monitoring and security tools, creating significant interoperability gaps.

Also, organizations may be factoring in air gaps that no longer exist. Seemingly isolated OT systems can be exposed to insider threats, temporary device connections, removable media, and hardware or software with undocumented external connectivity. For example, threat detection analytics by Booz Allen found that 60% of cyber incidents of unknown origin within a global manufacturing firm were attributable to USB devices carried by insiders, requiring proactive steps to reduce this threat.

Finally, IT-OT governance silos create exploitable gaps. In 2023, an estimated 37% of ransomware attacks on industrial organizations impacted both IT and OT environments, up from 27% in 2021. Wholly separate governance structures, budgets, risk frameworks, and reporting chains for the two domains should be rethought.

FEDERAL LEADERS GRAPPLING WITH REAL CHALLENGES

According to our exclusive survey of over 100 federal IT and cyber leaders, the biggest barriers they face implementing zero trust within their OT environment include (top two responses):

Legacy OT systems that can't support modern security controls **56%**	Fragmented governance across IT, OT, and external operators **44%**	Shortage of staff with both OT operations and cybersecurity expertise **37%**
Compliance timelines or policy mandates that do not align with technical or operational realities **29%**	Operation uptime requirements that limit ability to implement changes **19%**	Limited visibility into OT assets and communication pathways **15%**

A Technical Toolkit

The capabilities below translate DoW zero trust pillars into OT-specific implementations that work within real-world constraints such as legacy devices, requirements for non-disruptive deployment, proprietary protocols, and the need for continuous operations. Together, they illustrate how a layered architecture—grounded in device identity, segmentation, behavioral analytics, and coordinated response—can materially reduce attacker dwell time and harden mission-critical OT environments.

CAPABILITY	HOW IT WORKS	KEY TECHNICAL DETAILS	WHY IT MATTERS IN OT
OT Asset Inventory and Device Identity (Devices pillar)	Establish a centralized, continuously updated inventory of every device in the OT environment before any other zero trust control	• Passive discovery only • PKI/X.509 certificates deployed where technically feasible • Flagging of legacy devices incapable of certificate support	Organizations can't enforce least-privilege access, segment a network, or detect behavioral anomalies on assets they can't see
Identity and Access Management for Machines and People (Users pillar)	Extend authentication and authorization to non-person entities; OT context includes machine service accounts and automated processes	• Attribute-based access control (ABAC) governs access • Privileged access management (PAM) enforces least-privilege	Legacy OT was built on implicit trust within the perimeter; extending identity controls to machines closes the most common lateral movement path
Network Segmentation and Microsegmentation (Networks and Environments pillar)	Macro-separate IT and OT environments, then progressively isolate critical OT functions through software-defined architectures	• Software-defined networking (SDN) and software-defined perimeters (SDP), consistent with NIST SP 800-207 • Policy Enforcement Points to isolate process control systems, safety instrumented systems (SISs), and HMIs	Limits blast radius by constraining lateral movement across the IT/OT stack, directly countering the persistence-then-pivot model used by Volt Typhoon and similar actors
Behavioral Monitoring and Anomaly Detection (Visibility and Analytics pillar)	Continuously baseline normal device and communication behavior across the OT environment	• Passive detection of anomalous OT device communication patterns, out-of-window firmware changes, and out-of-parameter control commands • Feeds a security data lake or risk operations center	Defeats living-off-the-land (LOTL) tactics that evade signature-based detection; extended attacker dwell time in OT creates intervention windows that behavioral analytics can surface
OT-Aware Data Protection (Data pillar)	Protect process data, configuration files, and engineering workstation data at rest and in transit	• Encryption enforced across the IT/OT boundary • Access controls limiting who can read or modify PLC logic • Exfiltration monitoring at the IT/OT demarcation point	Control logic and safety configurations are high-value targets; theft or manipulation of PLC ladder logic can enable physical attacks
Automation, Orchestration, and Response (Automation and Orchestration pillar)	Coordinate policy enforcement, threat response, and access decisions to reduce dependence on manual intervention	• All response actions validated in testbed environments • Policy updates implemented after testing to ensure no operational impact	OT environments are understaffed relative to their attack surface; automation closes the gap between detection and response

Table 1. Zero trust capabilities for OT environments, mapped to DoW's *Zero Trust for Operational Technology Activities and Outcomes* (November 2025). Together, these capabilities limit attacker movement, detect pre-positioned access, and address the legacy device constraints that make traditional patching-based security insufficient.

Note: The standard zero trust pillar for Applications & Workloads is not presented as a standalone capability in this OT-focused model. In most OT environments, application functionality is tightly coupled to physical devices (e.g., PLCs, HMIs, and control systems), and associated protections are implemented through device identity, network segmentation, and data-centric controls.

Taken together, these capabilities show that achieving zero trust in OT environments is less about isolated controls and more about building a coordinated, layered architecture designed for adversaries that may already be inside. OT resilience depends on visibility into every device, restrictions on how those devices communicate, continuous monitoring that exposes subtle behavioral shifts, and carefully orchestrated response actions that avoid operational disruption.

The Road to Resilience

Knowing how to secure an OT environment is not the same as being able to implement those protections at scale, especially when systems cannot be taken offline. Achieving resilience depends on disciplined execution: turning architectural principles into sustained operational change through effective planning, engineering, and monitoring. The phases below outline the practical steps organizations must take to close the gap between zero trust intent and measurable, mission-aligned resilience:

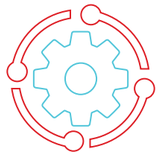

ASSESS AND ROADMAP
Establish a defensible starting point by assessing OT cybersecurity maturity, identifying gaps, and hunting for threats already inside the environment. Then build a prioritized roadmap that sequences investment and guides the transformation ahead.

DESIGN AND ENGINEER
Translate assessment findings into the architecture by defining target-state environments, selecting OT-appropriate tools, and embedding secure-by-design principles. This engineers security into the OT systems rather than retrofitting them after deployment.

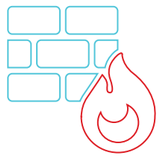

IMPLEMENT AND REMEDIATE
Execute the designed architecture by deploying segmentation; hardening networks, firewalls, and individual systems and devices; and progressively reducing attack surface without disrupting operational uptime.

DETECT AND RESPOND
Move from passive defense to active resilience by standing up threat reduction programs, deploying managed detection and response capabilities, and establishing OT-specific incident response protocols that neutralize threats.

Use Case: Deterrence in Action

Recently, a large industrial company undertook a sweeping modernization of its OT security, beginning with an assessment that revealed blind spots across thousands of devices. From there, the organization launched a focused effort to strengthen foundational controls—rapidly identifying OT assets, segmenting critical networks, and deploying hardened endpoint protections across its plants. These early moves brought previously unmonitored systems into view and quickly reduced risk without disrupting production.

As the program matured, the company expanded to global-scale deployments, introduced stricter removable-media governance, and built a dedicated OT threat-detection capability that sharply cut malicious activity. Just as important, a culture of cyber awareness took root, with widespread adoption of secure practices and a growing community of OT security champions. Soon the organization had established a resilient, zero trust–aligned OT posture that could sustain itself long after the initial transformation.

Beyond the Perimeter: The Whole-of-Nation Challenge

Federal agencies can mandate zero trust for systems they own, operate, or fund—but, as mentioned, much of the critical infrastructure that underpins national defense sits outside federal authority, creating a need for cross-sector collaboration. In many cases, these systems and networks are operated by municipal utilities, investor-owned companies, and vendors whose decisions on segmentation, authentication, or patching remain independent of federal control. While these organizations often maintain advanced cybersecurity protocols, they don't necessarily align with federal requirements.

Closing this gap requires mechanisms and collaborations that extend influence beyond ownership boundaries while respecting the decentralized nature of U.S. infrastructure. Federal procurement can incentivize higher OT security standards across the commercial market. A more empowered CISA—equipped with clearer authorities, scalable funding, and the ability to set specific baseline requirements for high-risk sectors—can shift voluntary coordination toward measurable outcomes. Finally, harmonizing federal expectations with widely adopted commercial frameworks and industry-specific standards can reduce compliance friction, accelerate adoption, and create a shared national baseline. Together, these changes would transform today's advisory model into a whole-of-nation zero trust posture capable of securing the interconnected civilian-military environment.

Conclusion: Accelerating Beyond Compliance

The need to extend zero trust to OT is now foundational to national resilience. DoW has set the direction and will now move to validate controls, reduce implicit trust, and rigorously close pathways before they can ever be exploited.

For other federal agencies, OT weaknesses rarely stay isolated given interdependent missions, shared infrastructure, and cross-sector dependencies. Agencies must align policy, funding, procurement, and incident response in ways that elevate zero trust to a mission-readiness requirement, strengthening their own systems and the broader federal ecosystem.

For the private sector, the challenge is to treat zero trust as both a business imperative and a contribution to national security. This means confronting difficult tradeoffs: modernizing aging systems, prioritizing visibility and segmentation, and investing ahead of disruption. Organizations that act early will maintain continuity under pressure.

Across all sectors, the real test is execution. Deeply understanding the architecture is the starting point; resilience comes from sustained, disciplined implementation. Beyond compliance milestones, progress will center on continually strengthening the ability to detect, contain, and neutralize threats before they affect operations. Meeting that challenge will determine how rapidly the nation's critical infrastructure can adapt to and overcome the latest OT threats in an increasingly contested environment.

David Forbes and Brandon Grimes are leaders in Booz Allen's National Cyber business supporting government missions, and Kyle Miller and Pia Capra are leaders in the company's Commercial Cyber business.

COMMERCIAL OT SECURITY: WHAT LEADERS NEED TO KNOW

- In the commercial OT world, a projected $5 trillion wave of U.S. industrial modernization and new construction—driven by reshoring, geopolitical pressures, and tariff realignment—presents an opportunity to build security into critical systems from the ground up rather than retrofit later at far greater cost.
- Ransomware remains the primary risk for most commercial operators, while nation-state actors increasingly target energy and oil and gas sectors. Like government environments, many commercial industrial verticals are subject to cyber regulations, such as the electric sector, chemical facilities, and pipelines.
- Commercial control systems carry 20-to-30-year lifecycles and cannot be taken offline for patching. Vendor and automation contractor constraints limit security controls. And the talent pool of professionals who understand both industrial operations and cybersecurity is usually thin.
- For CISOs navigating this environment, it's important to establish comprehensive asset visibility to understand your environment; implement network segmentation to create defensible barriers between IT and OT environments; and identify the operational "crown jewels" to guide risk-based investment decisions.
- Security measures that threaten uptime or production continuity will lose institutional support. The strongest programs position security as an enabler of critical operations, secure funding through business leadership rather than IT budgets, and pair outside expertise for initial buildout with smaller in-house teams.

Survey Methodology

Market Connections and Booz Allen partnered to design an online survey of over 100 federal government decision makers/influencers, involved with IT and cybersecurity, regarding AI's impact on government cybersecurity practices. The survey consisted of 14 questions and was fielded in April of 2026. Due to rounding, responses may not add up to exactly 100%.

SPEED READ

The convergence of IT and operational technology (OT) is eliminating once-protective air gaps, allowing nation-state adversaries to embed themselves inside critical infrastructure long before disruption is detected.

While few actors can yet execute truly destructive OT attacks, AI enabled techniques are accelerating disruptive campaigns—forced shutdowns, loss of visibility, and ransomware—narrowing the space between disruption and destruction.

Given the extensive challenges in implementing zero trust across critical infrastructure, organizations should view policy deadlines as starting points for sustained resilience, not finish lines for compliance.

Booz Allen®