CSAI Foundation | Cloud Security Alliance

CISA Zero Trust for Operational Technology

Applying Zero Trust Principles to Industrial Control Systems Without Disrupting Operations – 2026 Joint Guidance Analysis

2026-04-30

AI-assisted Rapid Research

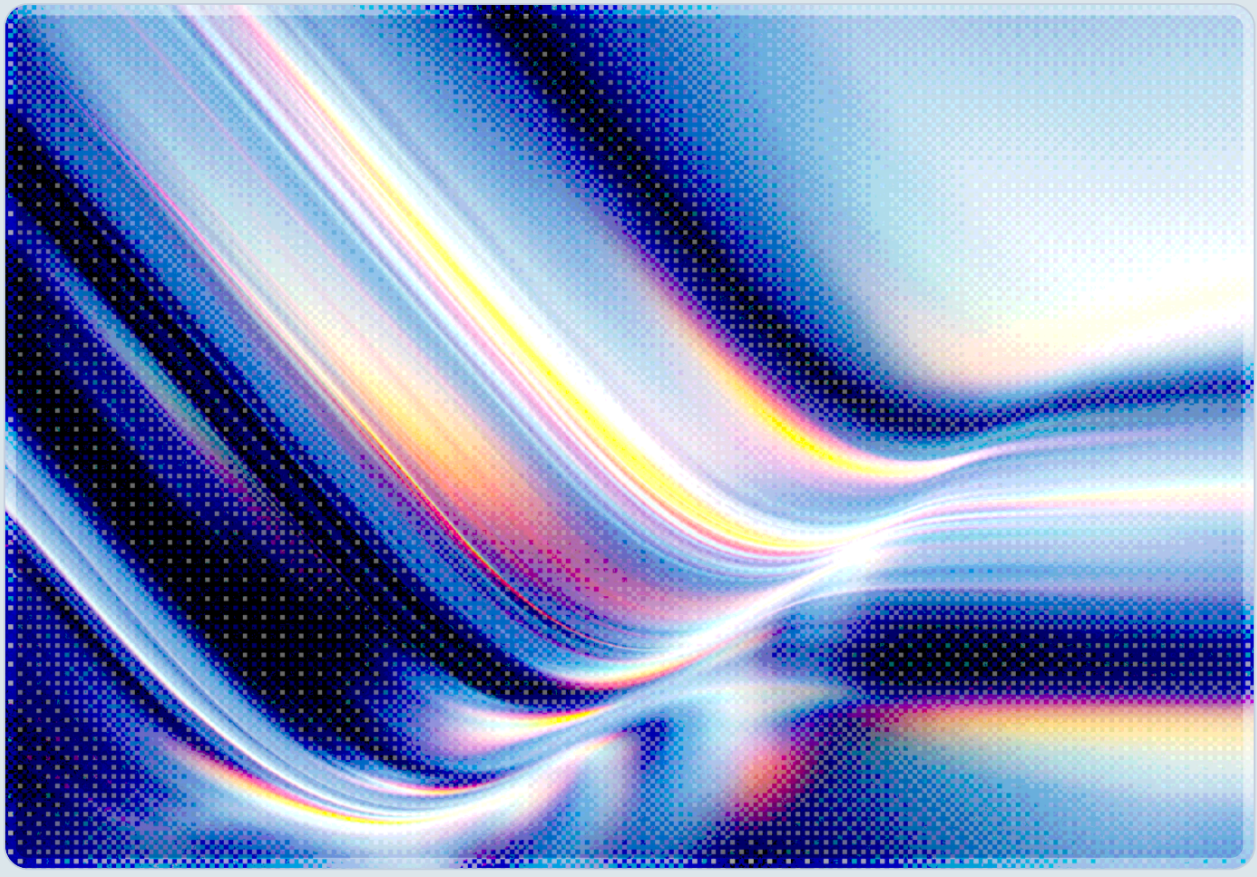

CSAI

CSA cloud security alliance®

Key Takeaways

- On April 29, 2026, CISA, the Department of War (DoW), DOE, FBI, and DOS jointly published the first comprehensive guidance for applying zero trust principles to operational technology (OT) environments [1].
- The guidance is explicit that IT-centric zero trust approaches cannot be applied wholesale to OT — availability imperatives, decades-long asset lifecycles, and legacy protocol constraints all demand tailored adaptation.
- State-sponsored actors, including those conducting activity tracked as Volt Typhoon, are actively prepositioning within critical infrastructure OT networks by exploiting IT/OT convergence pathways, Active Directory credential theft, and living off the land (LOTL) techniques [1].
- Supply chain risk management in OT remains immature; many legacy systems lack software bills of materials (SBOMs), and third-party vendor access remains one of the primary initial access vectors into OT environments.
- Passive network monitoring is the safe default for OT asset discovery and anomaly detection — active scanning risks disrupting legacy devices that were never designed to handle unexpected network queries.
- The document's six functional areas — Govern, Identify, Protect, Detect, Respond, and Recover — align with NIST Cybersecurity Framework (CSF) 2.0 and provide a structured, safety-conscious roadmap for OT zero trust implementation.

Background

Operational technology encompasses the programmable systems and devices that interact directly with physical processes: industrial control systems (ICS), building automation systems, transportation infrastructure, energy distribution networks, water treatment facilities, and the countless other systems on which modern society depends. Unlike information technology, where the consequences of a security failure are typically measured in data loss or service unavailability, OT failures can produce immediate

physical consequences — equipment damage, service disruptions, environmental harm, and risks to human life. This fundamental asymmetry has shaped how the OT security community has historically approached cybersecurity, often prioritizing isolation over integration.

That posture is now under sustained pressure. OT environments that were once operationally isolated or entirely manually controlled are becoming increasingly interconnected, digitally monitored, and remotely operated as organizations pursue operational efficiency and predictive maintenance [1]. Industry 4.0 digitization, cloud-connected sensors, and AI-driven process optimization are accelerating this convergence between IT and OT networks — and with it, the expansion of the attack surface. Pathways that once required physical access now exist as logical connections, and threat actors have adapted accordingly.

The threat landscape facing OT has matured significantly. Nation-state actors have demonstrated purpose-built OT attack capabilities through malware families including CrashOverride, Havex, BlackEnergy 2 and 3, Trisis, and Incontroller — each tailored to manipulate or destroy specific classes of industrial equipment [1]. More recently, actors conducting activity tracked as Volt Typhoon have demonstrated a patient, low-and-slow approach: compromising IT environments, stealing Active Directory credentials, and using those credentials to move laterally into adjacent OT networks while relying on LOTL techniques to evade detection. The goal appears to be prepositioning for potential future disruption rather than immediate exploitation.

It is against this backdrop that five U.S. government agencies — CISA, the Department of War (DoW), the Department of Energy (DOE), the Federal Bureau of Investigation (FBI), and the Department of State (DOS) — collaborated to produce the first joint guidance on applying zero trust principles specifically to OT environments. Published on April 29, 2026, the 28-page document represents both a technical implementation guide and a conceptual reframing of how organizations should approach security in environments where the wrong configuration change can have immediate physical consequences [1].

Security Analysis

Why Standard Zero Trust Falls Short in OT

Zero trust's foundational premise — assume breach has occurred, continuously validate access, minimize implicit trust — translates differently in OT than in enterprise IT. The guidance is candid about this: "The blanket application of traditional information technology (IT)-focused ZT capabilities to OT is neither reasonable nor feasible" [1]. Understanding why reveals the deeper challenge this document grapples with throughout its pages.

OT systems operate under an availability imperative that IT environments rarely face with equivalent strictness. A programmable logic controller (PLC) managing a refinery process or a power substation cannot simply reboot for a security patch without structured maintenance procedures, engineering reviews, and often regulatory approval. Many OT components run continuously for years — sometimes decades — between planned maintenance windows, and near-constant availability requirements mean that opportunistic or emergency patching is rarely an option. The guidance notes that some critical infrastructure components have lifespans exceeding fifty years, meaning assets deployed before modern networking concepts existed may remain in service well into the future.

Legacy OT systems compound this challenge by relying on proprietary protocols — Modbus, DNP3, EtherNet/IP — that were never designed with authentication or encryption in mind. These protocols often cannot be actively scanned without risking the destabilization of connected devices. They typically produce minimal logs, limiting the forensic and detection capabilities that zero trust implementations depend upon. They often cannot support agent-based security software due to constrained hardware, warranty restrictions, or tight coupling of application and operating system layers. And implementing zero trust controls across these environments demands genuine partnership between OT engineers, IT architects, and cybersecurity professionals — without which well-intentioned security changes can cause unplanned outages or safety incidents.

Governing Zero Trust in OT Environments

The guidance opens its functional recommendations with governance, recognizing that technical controls cannot succeed without organizational structures that account for OT's unique stakeholder landscape. Effective OT governance must clearly delineate the roles of asset owners — plant managers, energy managers — alongside operators, engineers, system integrators, vendor support personnel, and OT security teams. The document emphasizes that personnel with cross-disciplinary fluency, capable of understanding both cybersecurity threat models and OT safety constraints, are especially valuable for avoiding accidental disruptions during zero trust implementation.

Governance structures must also establish acceptable risk tolerances before incidents occur. In environments where security controls could impact system availability or safety, escalation procedures and pre-approved compensating controls need to exist in advance, not be improvised under pressure. The guidance frames procurement as an underutilized governance lever: newer OT components increasingly support security logging, secure communication protocols, and identity management capabilities that legacy systems lack entirely, allowing strategic procurement decisions to gradually advance zero trust posture without disruptive mid-lifecycle retrofits.

Supply chain risk management receives significant attention. While SBOMs are increasingly available for newer control systems, legacy OT products rarely provide such documentation. Organizations should assess vendor maturity by evaluating existing SBOMs, vulnerability management practices, CVE Numbering Authority status, and whether component security capabilities match their connectivity scenarios [1]. Third-party access — often granted broadly for monitoring or maintenance — requires explicit authorization, strictly limited scope, and persistent logging. Technical controls are preferred over procedural ones, though the guidance acknowledges that procedural controls may be the only available option for many legacy environments.

Identifying Assets and Mapping Risk

Before any zero trust architecture can be designed, organizations need accurate visibility into what exists in their OT environment. The guidance treats comprehensive asset inventory as foundational — an "organized, regularly updated list of an organization's systems, hardware, and software" that enables understanding of the devices, communications, and functions that compose the environment [1]. Achieving this in OT requires different tooling and approaches than IT asset discovery: active scanning that would be routine in an enterprise network can knock legacy OT devices offline.

Passive monitoring is therefore the recommended default for OT asset discovery. SPAN ports and network TAPs allow traffic observation without generating queries that could disrupt control system processes. Organizations are advised to understand network topologies deeply enough to recognize that automated tools will produce only partial results in highly segmented or air-gapped environments — and to plan sensor placement accordingly, prioritizing visibility into the largest or highest-risk network segments rather than attempting comprehensive coverage at prohibitive cost.

Risk assessment in OT must extend beyond the likelihood-and-impact calculations standard in IT to encompass cyber-physical consequences. The guidance emphasizes threat modeling that explicitly accounts for how IT/OT convergence might expand the attack surface, maps lateral movement paths an adversary could exploit, and considers the downstream physical consequences of a successful attack. This includes cascading failures across interconnected industrial systems and the potential for a cyber incident to produce outcomes — equipment damage, environmental harm, personnel safety incidents — that dwarf the direct cost of the breach itself [1].

Protecting OT Systems Without Breaking Them

Network segmentation is described as "one of the most foundational and effective security controls in OT environments, often serving as the primary line of defense" [1]. This is not new guidance for OT practitioners — the concept of security zones and controlled conduits has been embedded in standards

like ISA/IEC 62443 for years [10]. What the document adds is a reframing of segmentation as a dynamic, continuously enforced security policy rather than a static architectural decision. Air gaps can be bridged. VLANs can be misconfigured. Without continuous validation and active monitoring within segmented boundaries, even well-designed architectures degrade over time.

The guidance draws a sharp distinction between IT and OT segmentation philosophies. Where IT segmentation primarily aims to contain breaches, protect data confidentiality, and enforce compliance, OT segmentation must prioritize the protection of physical processes, system availability, and operational safety. The tools and approaches reflect this: data diodes enforcing hardware-controlled unidirectional communication, jump hosts as hardened single-entry points for remote administration, and secure gateways as controlled interfaces for data transfer between network zones.

Identity, credential, and access management (ICAM) in OT presents distinct challenges. Many OT systems predate modern ICAM capabilities, and directly connecting IT and OT identity systems — for instance, sharing an Active Directory domain — is explicitly discouraged as it increases the attack surface and creates a path for lateral movement [1]. Emergency access cannot be an afterthought; OT environments require pre-defined break-glass procedures with limited-lifespan accounts, stringent auditing, and physical safeguards to ensure that availability and safety are never compromised by overly restrictive access controls.

The encryption guidance reflects OT's prioritization of integrity and availability over confidentiality. In OT environments where latency is safety-critical and where adding full encryption could disrupt time-sensitive communications, the guidance recommends focusing encryption efforts on protecting the semantic structure of communications — addresses and function codes that reveal operational intent — rather than encrypting high-frequency real-time process values. Where secure protocol extensions are available (Modbus over TLS, BACnet/SC, CIP Security, DNP SAv5), organizations should enable them; where they are not, wrapping legacy protocols in TLS-enabled gateways provides authentication and integrity protection outside the control devices themselves [1].

Patch management in OT requires structured coordination with operational staff and cannot follow the rapid cadence common in IT environments. The guidance recommends the Stakeholder-Specific Vulnerability Categorization (SSVC) decision tree framework to evaluate the tradeoff between downtime risk and exploitation risk for specific vulnerabilities [1]. When patching cannot be promptly performed, the recommended response is to isolate vulnerable systems, increase monitoring, restrict communication pathways, and deploy virtual patches through network defenses — reducing exposure while preserving operational continuity.

Detecting Threats in Environments with Minimal Logging

Detection in OT is complicated by the limited logging and forensics capabilities that characterize many legacy environments. The guidance highlights two primary detection approaches suited to OT's relatively static operational baseline. Baseline-based detection uses statistical models to identify deviations from normal behavior patterns — an approach well-suited to OT environments where authorized communications follow consistent, predictable patterns, but which requires carefully defined learning windows that capture all operational modes before tuning detection thresholds. Specification-based detection defines the range of valid behaviors explicitly, akin to safety system design, and robustly captures known failure states and operationally viable conditions without requiring probabilistic models [1].

CISA maintains an open-source SIEM tool, Malcolm, which includes Zeek parsers for common OT protocols and supports deep traffic analysis of industrial communications [1]. Passive monitoring using SPAN ports or network TAPs remains the preferred approach at OT boundaries — these provide traffic visibility without requiring reconfiguration of switch infrastructure or generating load on monitored systems.

Endpoint detection and response (EDR) in OT faces significant barriers: legacy embedded components may not support standard agents below the HMI layer, warranty policies often prohibit software additions, and cloud-connected EDR solutions are impractical in isolated or air-gapped networks. The guidance recommends a lightweight approach for embedded systems — monitoring CPU and memory telemetry, new processes, and configuration changes — recognizing that embedded OT components should rarely be modified under normal conditions. LOTL techniques deserve particular attention: attackers can leverage legitimate engineering tools and protocols to conduct malicious activity while blending into normal operations, and the prevalence of unauthenticated protocols in OT networks makes this approach especially effective [7].

Responding to Incidents Without Causing Additional Harm

Zero trust architecture's assumption that a breach has already occurred makes incident response planning an integral part of ZT implementation rather than a separate exercise. The guidance recommends that OT IR plans build upon existing emergency response procedures, risk assessments, and business continuity plans, then layer cyber-specific additions: isolation procedures with pre-authorized personnel, decision matrices and playbooks for common OT threat scenarios, and clearly defined communication protocols for internal and external stakeholders.

Threat containment in OT is fundamentally different from IT containment. Aggressively isolating or air-gapping OT systems during an incident can itself cause the operational disruptions and safety incidents that the attacker sought to produce — security teams must collaborate with process engineers to determine which functions can continue under heightened segmentation, and for how long. The guidance introduces the concept of "soft segmentation" — a carefully defined demarcation point within the network — as an appropriate intermediate containment step when full isolation would be disruptive [1]. The document is explicit that OT incident response often requires the difficult balance of aggressive containment against maintaining safe operations.

Coordinating incident response for critical infrastructure requires engagement with the broader ecosystem: DoW, the intelligence community, and Sector Risk Management Agencies. The guidance calls for clear information-sharing protocols, well-defined joint IR roles and responsibilities, and regular exercises to maintain readiness.

Recovering with OT-Specific Continuity Planning

Recovery in OT requires backup procedures that differ substantially from IT. OT-specific devices — controllers, PLCs, safety instrumented systems — require backups capturing engineering logic, I/O lists, startup values, and operational specifications, not just data and configuration. Backup capabilities vary significantly across devices; some support only offline backups and some lack backup functionality entirely [1]. Retaining detailed engineering documentation — cause-and-effect matrices, control narratives, logic printouts — is critical for restoration in environments where interoperability challenges can extend recovery timelines dramatically. Business continuity plans must integrate cybersecurity considerations explicitly, with backup files tested regularly via file hashing and checksum validation before they are needed.

Recommendations

Immediate Actions

Organizations operating OT environments should treat passive asset inventory as an urgent priority. Without knowing what exists on OT networks, no meaningful zero trust architecture can be designed. Organizations lacking passive monitoring infrastructure should deploy network TAPs at key boundaries — particularly where OT connects to IT or external networks — and begin constructing the comprehensive asset inventory the guidance treats as foundational. Jump hosts should be assessed immediately: the

guidance strongly recommends a hardened, MFA-enforced jump host as the sole remote access entry point into OT environments. Any organization currently permitting direct remote access to OT systems from vendor or operator workstations should treat this as an urgent remediation target.

Short-Term Mitigations

Supply chain risk management should be elevated as a governance priority. Organizations should begin assessing vendor maturity through SBOM availability, CVE Numbering Authority status, and vulnerability management practices. Third-party remote access should be audited to ensure that access is explicitly authorized, narrowly scoped, and comprehensively logged. On the detection side, organizations should evaluate CISA's open-source Malcolm SIEM for OT protocol visibility and invest in OT-specific passive monitoring tools capable of understanding industrial protocols. Incident response plans should be reviewed and updated to include OT-specific isolation procedures, with pre-authorized personnel and decision matrices in place before they are needed.

Strategic Considerations

The guidance frames procurement as the long-term path to sustainable zero trust in OT. Building ZT-aligned security requirements — logging, secure protocol support, identity and access management — into procurement specifications for new OT equipment leverages natural asset replacement cycles rather than requiring disruptive mid-lifecycle retrofits. Developing personnel with genuine cross-disciplinary fluency — security professionals who understand OT constraints and OT engineers who understand threat modeling — may ultimately deliver more security value than any specific tool. Business continuity plans should explicitly integrate cybersecurity considerations, with cyber incident response procedures tested alongside operational continuity plans at least annually.

CSA Resource Alignment

This guidance connects directly to several CSA frameworks and publications that organizations can use to operationalize zero trust in OT contexts. CSA's *Zero Trust Guidance for Critical Infrastructure* (2024) provides a complementary 65-page implementation roadmap specifically addressing OT and ICS environments, including a five-step process — Define Protect Surface, Map Operational Flows, Build ZTA, Create ZT Policy, Monitor and Maintain — that aligns well with the CISA guidance's functional structure [3]. CSA's *Zero Trust Guidance for Achieving Operational Resilience* addresses the business continuity

and resilience dimensions that the CISA document's Recover section emphasizes. CSA's *CCZT (Certificate of Competence in Zero Trust)* curriculum provides training resources for developing the cross-disciplinary fluency that both documents identify as essential.

At the framework level, the CISA guidance's alignment with NIST CSF 2.0 makes it directly applicable within organizations already using CSA's Cloud Controls Matrix (CCM), which references the NIST framework extensively. CSA's AI Controls Matrix (AICM) is also relevant as OT environments increasingly incorporate AI-enabled process optimization and predictive maintenance: the AICM's guidance on AI system access controls, monitoring, and supply chain security applies directly to AI components deployed within or adjacent to OT networks. MAESTRO, CSA's threat modeling framework for agentic AI systems, is particularly worth attention as autonomous AI agents begin appearing in industrial contexts — the same assumptions about lateral movement and supply chain compromise that MAESTRO addresses in AI agent deployments apply with heightened physical-consequences stakes in OT environments.

Organizations seeking assurance validation for their zero trust OT programs can leverage CSA's STAR (Security, Trust, Assurance, and Risk) program. STAR's continuous monitoring tier provides a mechanism for demonstrating ongoing compliance with zero trust principles — an approach that aligns with the CISA guidance's emphasis on continuous monitoring and validation as core zero trust behaviors rather than point-in-time assessments.

References

[1] CISA, Department of War, Department of Energy, FBI, Department of State. "Adapting Zero Trust Principles to Operational Technology." CISA, April 29, 2026.

[2] CISA. "CISA and U.S. Government Partners Unveil Guide to Accelerate Zero Trust Adoption in Operational Technology." CISA Press Release, April 29, 2026.

[3] Cloud Security Alliance. "Zero Trust Guidance for Critical Infrastructure." CSA Zero Trust Working Group, 2024.

[4] NIST. "Cybersecurity Framework 2.0." NIST, February 2024.

[5] Keith Stouffer et al. "Guide to Operational Technology (OT) Security, NIST SP 800-82r3." NIST Special Publication, September 2023.

[6] CISA. "Zero Trust Maturity Model Version 2.0." CISA, April 2023.

[7] CISA, NSA, FBI, et al. "PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure." Joint Cybersecurity Advisory, February 2024.

[8] CISA, FBI, NSA, et al. "Identifying and Mitigating Living Off the Land Techniques." Joint Guidance, February 2024.

[9] MITRE. "MITRE ATT&CK for ICS." MITRE Corporation, continuously updated.

[10] ISA/IEC. "ISA/IEC 62443 Series of Standards for Industrial Automation and Control Systems Security." ISA, continuously updated.

[11] NIST. "Zero Trust Architecture, NIST SP 800-207." NIST Special Publication, August 2020.

