SYM-AM-26-074

# Excerpt from the Proceedings

of the

# Twenty-Third Annual Acquisition Research Symposium and Innovation Summit

## Wednesday, May 6, 2026 Sessions

## Volume I

## "Accelerating Warfighting Capabilities"

**Cyber Digital Twin-Informed Zero Trust: A Synergistic Framework for Securing Operational Technology in Defense Logistics Infrastructure**

**Published: April 30, 2026**

The research presented in this report was supported by the Acquisition Research Program, Graduate School of Defense Management at the Naval Postgraduate School.

**To request defense acquisition research, please contact:**

Acquisition Research Program
Department of Defense Management
Naval Postgraduate School
E: arp@nps.edu
**www.acquisitionresearch.net**

Copies of Symposium Proceedings and Presentations; and Acquisition Sponsored Faculty and Student Research Reports and Posters may be printed from the **NPS Defense Acquisition & Innovation Repository** at **https://dair.nps.edu/**.


ACQUISITION RESEARCH PROGRAM
DEPARTMENT OF ACQUISITION, FINANCE, AND MANPOWER
NAVAL POSTGRADUATE SCHOOL

# Cyber Digital Twin-Informed Zero Trust: A Synergistic Framework for Securing Operational Technology in Defense Logistics Infrastructure

**Dr. Barry A. Humphrey—**is the Program Manager for the Defense Logistics Agency's Research and Development Directorate. His program, Logistics Technology Research, aims at improving key business functions like data analytics, database security, supply chain risk management, and supply chain security through advanced technologies such as AI, Machine Learning, Blockchain, Augmented Reality, Quantum Computing, and Cybersecurity efforts to improve Warfighter support based on DLA's strategic guidance. Dr. Humphrey holds a PhD in Information Systems and Communications from Robert Morris University, an MA in Information Technology Management, Business and Organizational Security Management from Webster University, and a BS in Molecular Biology from Auburn University at Montgomery. He also completed Wharton Executive Education's Leadership Program in AI and Analytics and the Chief Technology Officer Program. Dr. Humphrey is co-founder and Director of Innovative Learning and Educational Technology at Ready Force Cyber, a non-profit organization providing cybersecurity, AI training and hands-on skill development, and career navigation in various technology fields. Dr. Humphrey is a retired military officer where he served over 22 years in the Air Force and Army.

## Abstract

The convergence of Information Technology (IT) and Operational Technology (OT) has exposed critical infrastructure to cyber-physical threats that perimeter-based security was never designed to handle. The consequences extend beyond data loss or equipment malfunction: compromised OT systems directly degrade military readiness, endanger both warfighter and civilian lives, and create national security vulnerabilities near-peer adversaries are actively probing. Legacy OT environments—the systems governing logistics, utilities, and manufacturing across military supply chains—operate under assumptions about isolation and trust that no longer hold. This research presents a security framework that integrates Cyber Digital Twins (CDT), Artificial Intelligence and Machine Learning (AI/ML), and a Zero Trust Architecture (ZTA) framework to provide an integrated defensive capability for OT cybersecurity. The approach centers on a high-fidelity virtual replica of the OT environment, training AI/ML models to recognize both normal operational signatures and simulated attack signatures within that replica, using the resulting risk intelligence to drive dynamic ZTA framework enforcement. The concept of the operational signature, the distinctive behavioral fingerprint of a device, process, or communication patterns is central to this framework: the CDT establishes baseline signatures, AI/ML models detect deviations from those signatures, and the ZTA framework enforces containment when anomalous signatures are identified.

**Keywords:** Cyber Digital Twin, Zero Trust Architecture, Operational Technology, AI/ML anomaly detection, operational signatures, defense logistics, national security, Industrial Control Systems

## Introduction

Anyone who has spent time inside a defense logistics operation understands the tension: these environments run on connectivity now—warehouse automation, predictive maintenance, enterprise resource planning integration—but the underlying Operational Technology (OT) systems were built in an era when cybersecurity meant locking the server room door. The digital transformation of defense logistics has delivered real operational gains, but it has also wired mission-critical Industrial Control Systems (ICS) into networks where sophisticated adversaries are establishing persistent, clandestine presence.

The Defense Logistics Agency (DLA) manages an extensive portfolio of OT systems that enable warfighter readiness: warehouse automation, environmental controls for material storage, fuel distribution infrastructure, and manufacturing equipment across the defense industrial base. These systems increasingly depend on network connectivity. That connectivity,

in turn, converts what were once air-gapped assets into potential entry points for nation-state actors and advanced persistent threats targeting military logistics (Cybersecurity and Infrastructure Security Agency, 2024). The national security implications are unambiguous. A compromised facility control system does not merely create a facility management problem—it can render temperature-sensitive munitions, pharmaceuticals, or electronics inventory unusable, directly degrading the readiness of deployed forces and placing both military personnel and the civilians they protect at risk.

The DLA Strategic Plan 2025–2030, "DLA Transforms: A Call to Action," confronts this reality head-on. The "Precision" imperative pushes for stronger digital interoperability and artificial intelligence (AI)/machine learning (ML)–powered decision advantage. The "Posture" imperative demands that the Agency illuminate and mitigate global supply chain risk to increase resiliency (DLA, 2025a). The Director's intent is blunt: to succeed in an environment where logistics are contested by adversaries across all domains and all levels of war, DLA must think, act, and operate differently. The CDT-ZTA Framework presented here is one answer to that mandate.

This research addresses these challenges through a Cyber Digital Twin–Informed Zero Trust Architecture (CDT-ZTA) Framework that emerged from multiple Small Business Innovation Research Phase I proof-of-concept projects conducted under the DLA R&D Directorate. The core argument is that integrating CDT simulation, AI/ML-driven signature analysis, and ZTA framework enforcement produces defensive capabilities none of the three can deliver alone—a security posture that is proactive, adaptive, and resilient rather than reactive and brittle.

**Operational Technology Challenges**

OT is the hardware and software that monitors and controls physical devices, processes, and events (Gartner, 2023). The critical distinction from Information Technology (IT) is the priority stack: where IT traditionally leads with confidentiality, OT leads with availability and safety. A crashed email server is an inconvenience. A crashed programmable logic controller (PLC) governing a physical process in a defense logistics facility is a safety event, a readiness event, and a national security event.

This priority difference creates security challenges that IT-centric teams routinely underestimate. ICS deployed 20 or 30 years ago were built for isolated operation using proprietary protocols. Nobody designed encryption, authentication, or logging because there was nothing to connect to. Many of these systems are still running—they are reliable, replacement costs are enormous, and the operational risk of upgrading a functioning system is hard to justify (Hassanzadeh et al., 2020). Industry analysis from 2026 confirms the pattern: many OT environments still run firmware predating modern cryptographic standards or no longer supported by manufacturers. Protection rather than replacement remains the prevailing strategy (B2B Cyber Security, 2026).

From a national security perspective, the risk calculus is fundamentally different from enterprise IT. When adversaries compromise OT systems in defense logistics, they do not merely steal data; they gain the ability to manipulate physical processes that sustain military operations. Fuel distribution disrupted during deployment. Warehouse environments degraded to destroy critical inventory. Manufacturing equipment sabotaged to introduce defects into the defense supply chain. These scenarios translate directly to degraded readiness, mission failure, and risk to the lives of service members and the civilian populations they defend. The old framing of OT vulnerability as primarily a safety concern understates the stakes: OT security in defense logistics is a national security imperative.

The convergence trend has accelerated as organizations chase operational efficiencies—connected to monitoring, predictive maintenance, and enterprise integration. But every new connection expands the attack surface. The Dragos 2026 OT/ICS Cybersecurity Report documented that organizations with comprehensive OT visibility detected and contained ransomware incidents in an average of 5 days, compared with the industry-wide average of 42 days (Dragos, 2026). The SANS Institute's analysis put it plainly: adversaries in 2025 showed patience and process awareness, investing in understanding how industrial environments work rather than developing novel malware. They were confident that long-standing architectural weaknesses would continue to deliver results (SANS Institute, 2026).

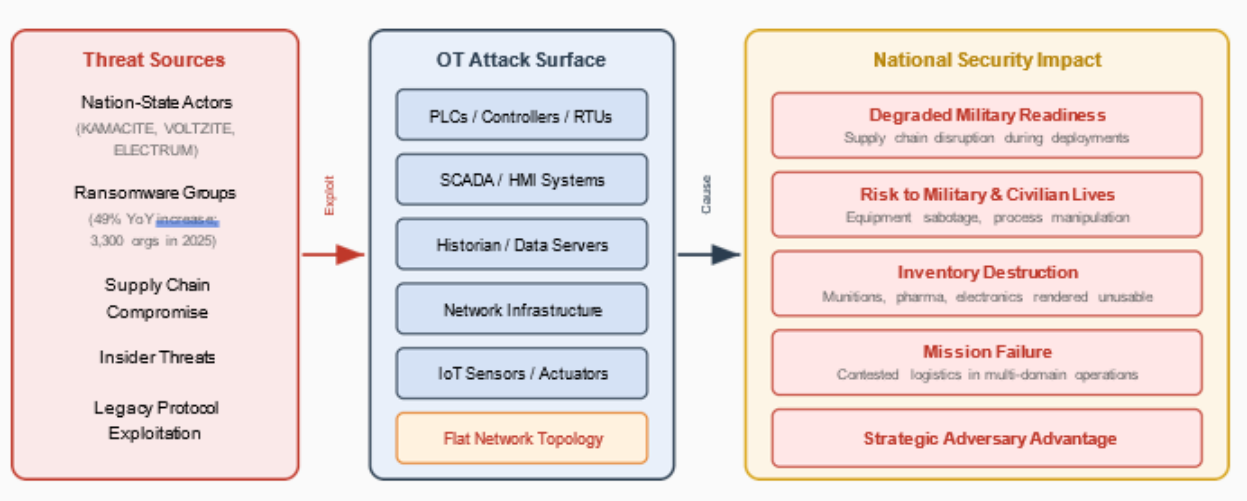


***Figure 1. The OT Threat Landscape Illustrating How Adversary Capabilities Target Defense Logistics OT Systems, with Consequences that Cascade from Facility-Level Disruptions to National Security Impacts Affecting Readiness and Lives***

## Zero Trust Architecture Framework Principles

The ZTA framework, at its core, is a rejection of the assumption that network location confers trustworthiness. Kindervag's (2010) "never trust, always verify" principle has matured into a full architectural framework codified in NIST SP 800-207 (Rose et al., 2020). The operating assumption is that adversaries may already be inside the network—an assumption that OT security teams, having watched the erosion of air gaps over the past decade, find entirely plausible. A systematic literature review covering 74 peer-reviewed articles from 2016–2025 confirmed that the ZTA framework has transitioned from theoretical paradigm to practical implementation, with authentication, authorization, and access control as the most consistently deployed components (Al-Sharif et al., 2025). A broader review analyzing over 1,700 articles documented growing attention to OT-specific implementations (Alqahtani & Almuhaideb, 2025).

The ZTA framework tenets that matter most for OT are continuous verification, least privilege access, microsegmentation, and comprehensive monitoring. Microsegmentation is particularly valuable because it limits the blast radius—when an adversary gets in, they hit a wall rather than an open floor plan. Segment the OT network into granular trust zones with explicit access controls between them, and you contain threats before they reach safety-critical systems. In signal processing terms, microsegmentation constrains the propagation of anomalous signatures: a compromised device's abnormal communication signature is contained within its segment rather than radiating across the entire network.

Implementing the ZTA framework in OT is harder than in enterprise IT, and anyone who says otherwise has not tried it. OT systems use legacy protocols with no authentication

mechanisms. They have real-time constraints where adding verification latency can disrupt physical processes. The Cloud Security Alliance's 2024 guidance on Zero Trust for Critical Infrastructure acknowledged these realities, offering a tailored five-step roadmap for industrial environments (Cloud Security Alliance, 2024). The Department of Defense (DoD) has moved aggressively on this front. In July 2025, DTM 25-003 directed all DoD components to achieve minimum Target Level Zero Trust across all systems, including OT and control systems (DoD, 2025a). In November 2025, the DoD CIO published "Zero Trust for Operational Technology Activities and Outcomes," defining 84 capability outcomes for target-level and 21 for advanced-level Zero Trust specific to OT (DoD, 2025b). Deadlines are set at fiscal 2030 for target-level and fiscal 2033 for advanced-level, with ZTA Strategy 2.0 expected in early 2026 (DefenseScoop, 2025). In January 2026, the NSA published a Zero Trust Implementation Guideline Primer providing phased direction aligned with the DoD ZTA Framework (National Security Agency, 2026).

**Cyber Digital Twins for Security**

Digital twin technology—creating virtual replicas of physical systems for simulation and analysis—is well established in manufacturing and engineering. Extending the concept to cybersecurity means building CDTs that replicate not just functional behavior but security characteristics, vulnerabilities, attack surfaces, and critically, the operational signatures of every device, process, and communication flow in the environment (Eckhart & Ekelhart, 2019). The term "signature" is used deliberately. Just as radar systems identify aircraft by their electromagnetic signatures and sonar systems classify vessels by their acoustic signatures, a CDT captures the behavioral signatures of an OT environment: the timing patterns of PLC polling cycles, the characteristic data payloads of industrial protocol requests, the process response curves of physical equipment under load, and the network traffic profiles that constitute normal operations. These signatures become the baseline against which all deviations are measured.

The field has matured considerably. An ACM Computing Surveys review aggregating 201 publications mapped digital twins for security operations against the NIST Cybersecurity Framework, concluding that the technology sits at a convergence point of technology, cybersecurity, and industrial requirements (Vielberth et al., 2025). Repetto (2026) pushed further, proposing architectures for Cybersecurity Digital Twins that bridge fragmented security operations across multi-ownership digital service chains. The World Economic Forum highlighted CDTs as a transformative cybersecurity technology, noting that AI/ML has already cut breach detection times by 33% and containment times by 43% in security operations centers, and that CDTs build on those gains by enabling SOCs to simulate attacks and optimize defenses in real time (World Economic Forum, 2025). Industry has followed: Trend Micro announced enterprise-grade CDT technology for proactive cybersecurity in July 2025, powered by NVIDIA computing, spanning on-premises to cloud and IT to OT (Trend Micro, 2025).

For OT security specifically, the CDT's value is deeply practical. CDTs let you run attack simulations, test defensive measures, and validate security configurations without risking production systems. El Hachem et al. (2025) demonstrated that traditional threat modeling struggles with the dynamic, interconnected nature of ICS and that CDTs fill that gap with operational-context-aware security analysis throughout the system life cycle. By 2026, CDTs as virtual replicas for testing, simulation, and resilience training have become a defining trend in OT security, with organizations implementing application-aware microsegmentation and deep device signature fingerprinting within twin environments as standard practice (B2B Cyber Security, 2026). The DLA R&D portfolio has invested directly in these capabilities, with funding supporting CDT projects for process analysis and smart warehouse modernization (DLA, 2025b).

### Artificial Intelligence and Machine Learning for Anomaly Detection

AI/ML has fundamentally changed what is possible in cybersecurity analysis—automated pattern recognition and anomaly detection at volumes no human team can match. In OT environments, where normal operations produce consistent, predictable patterns of network traffic and process behavior, AI/ML models can learn the signature of what "normal" looks like and flag deviations that may indicate attack or malfunction (Zhou et al., 2021). The concept of the signature is again central: AI/ML models are, at their core, signature recognition systems. They learn the multidimensional signature of normal OT operations—comprising network traffic signatures, process variable signatures, device communication signatures, and timing signatures—and identify when observed behavior deviates from those learned baselines.

A systematic review published in Electronics in January 2026, analyzing 89 studies from 2021–2025, categorized AI/ML-based ICS anomaly detection across five data types: network traffic, operational data, simulation data, hybrid data, and auxiliary data. The key finding: network data catches cyber-layer attack signatures (reconnaissance, denial-of-service, man-in-the-middle), while operational data catches physical-layer anomaly signatures (process disturbances, false data injection, stealth deviations; Kim et al., 2026). This directly validates the CDT-ZTA Framework's approach of combining both data types for comprehensive signature coverage.

The data scarcity problem in OT is real—production ICS environments do not experience enough attacks to train effective models from operational data alone. Recent work on Semi-Supervised Generative Adversarial Networks addressed this by using adversarial training to generate synthetic operational data that supplements limited real-world examples (Processes, 2025). This parallels the way the CDT-ZTA Framework uses its CDT to generate synthetic attack signatures. Birihanu and Lendák (2025) advanced explainable anomaly detection using multivariate Gaussian models that identify anomalous correlations between sensors and actuators while providing interpretable results—addressing the growing demand for Explainable AI/ML in critical infrastructure. At industry scale, Fortinet's 2025 report confirmed the payoff: organizations on integrated OT security platforms with AI/ML-driven detection achieved up to 93% fewer cyber incidents compared to flat networks (Fortinet, 2025). The synergy between AI/ML detection and CDTs is straightforward: the CDT provides a safe environment for generating the training data—the attack and normal signatures—that AI/ML models need, while models deployed to both CDT and production enable continuous signature comparison and validation (Xu et al., 2022).

## Literature Review

The current literature reveals a clear trajectory: CDT technology, AI/ML, and the ZTA framework have each matured individually to the point where their convergence for OT security is not just feasible but increasingly necessary. The systematic reviews published between 2025 and 2026 tell a consistent story across independent research streams. Al-Sharif et al. (2025) and Alqahtani and Almuhaideb (2025) document the ZTA framework's transition from conceptual model to deployable architecture, with growing attention to OT-specific implementations that account for legacy protocols, real-time constraints, and availability requirements that enterprise IT deployments never face. Vielberth et al. (2025) mapped 201 publications on CDTs for security operations against the NIST Cybersecurity Framework, concluding that the technology has reached a maturity level where it functions as an operational convergence point rather than a research curiosity. Kim et al. (2026), analyzing 89 studies on AI/ML-based ICS anomaly detection, found that hybrid approaches combining network traffic signatures and operational data signatures—the approach the CDT-ZTA Framework employs—represent the most effective detection strategy for industrial environments.

What the literature has not yet delivered is a validated integration architecture that ties these three capabilities together into a unified defensive system purpose-built for defense logistics OT—one where signature analysis is the connective thread. That gap is significant because the individual technologies create dependencies on each other that the literature acknowledges but has not resolved in practice. AI/ML-driven anomaly detection in OT environments is constrained by a well-documented training data problem: production ICS environments do not generate enough attack signatures to train effective models, and generating synthetic attack signatures on live systems risks operational disruption (Kim et al., 2026; Processes, 2025). CDTs solve that problem by providing safe environments for attack signature simulation and labeled dataset generation, but without an enforcement mechanism, detection insights remain informational rather than operational. ZTA framework microsegmentation provides that enforcement, but static segmentation policies degrade as environments change and adversaries adapt limitation that AI/ML-driven continuous risk assessment directly addresses.

Repetto (2026) argued that restricting security operations to individual administrative domains is inadequate against multi-step attacks that propagate across boundaries, while El Hachem et al. (2025) demonstrated that traditional threat modeling fails to capture the dynamic, interconnected nature of ICS environments. Literature, in other words, has identified the pieces and documented why each one alone is insufficient. What has been missing is the integration framework that leverages the CDT to generate and catalog operational and attack signatures, AI/ML to detect deviations from those signatures, and the ZTA framework to contain what the AI/ML inevitably misses, creating the kind of closed-loop, self-improving defensive architecture that the current threat landscape demands.

The threat intelligence literature reinforces the urgency from a national security standpoint. Dragos tracked 26 threat groups globally, including three newly identified, and reported a 49% year-over-year increase in ransomware groups hitting industrial organizations, affecting 3,300 organizations in 2025 (Dragos, 2026). KAMACITE conducted systematic reconnaissance of U.S. industrial devices throughout 2025, mapping entire control loops down to HMIs, variable frequency drives, and metering modules—essentially cataloging the operational signatures of American industrial infrastructure. ELECTRUM expanded beyond Ukrainian targets to European energy infrastructure with intent to disrupt operations. For defense logistics, translate those capabilities to ammunition production lines, fuel distribution networks, or warehouse control systems during a contested deployment, and the national security implications—degraded readiness, endangered military and civilian lives, strategic adversary advantage—are stark.

## Methodology

### Research Design

This research uses design science methodology. Design science focuses on architecting solutions to identified problems while generating knowledge that transfers beyond immediate implementation (Hevner et al., 2004). The CDT-ZTA Framework is the designed artifact; the problem is OT security in the defense logistics enterprise. The research team includes personnel from DLA's R&D department, cybersecurity division, and OT subject-matter experts, along with several trusted vendors currently under DLA R&D contract. This effort includes case-study analysis of framework deployment within DLA environments, allowing rigorous evaluation of framework effectiveness while capturing the contextual and organizational factors that determine whether a technically sound framework works in practice.

## Framework Architecture

The CDT-ZTA Framework integrates three components: the CDT environment, the AI/ML anomaly detection system, and the ZTA framework microsegmentation enforcement layer. The architecture defines distinct interfaces between them—data flows from physical systems to the CDT, from the CDT to AI/ML models, and from AI/ML risk assessments to ZTA framework policy enforcement. The concept of the operational signature is the connective tissue: the CDT captures signatures, AI/ML models learn and evaluate signatures, and the ZTA framework responds when anomalous signatures are detected.

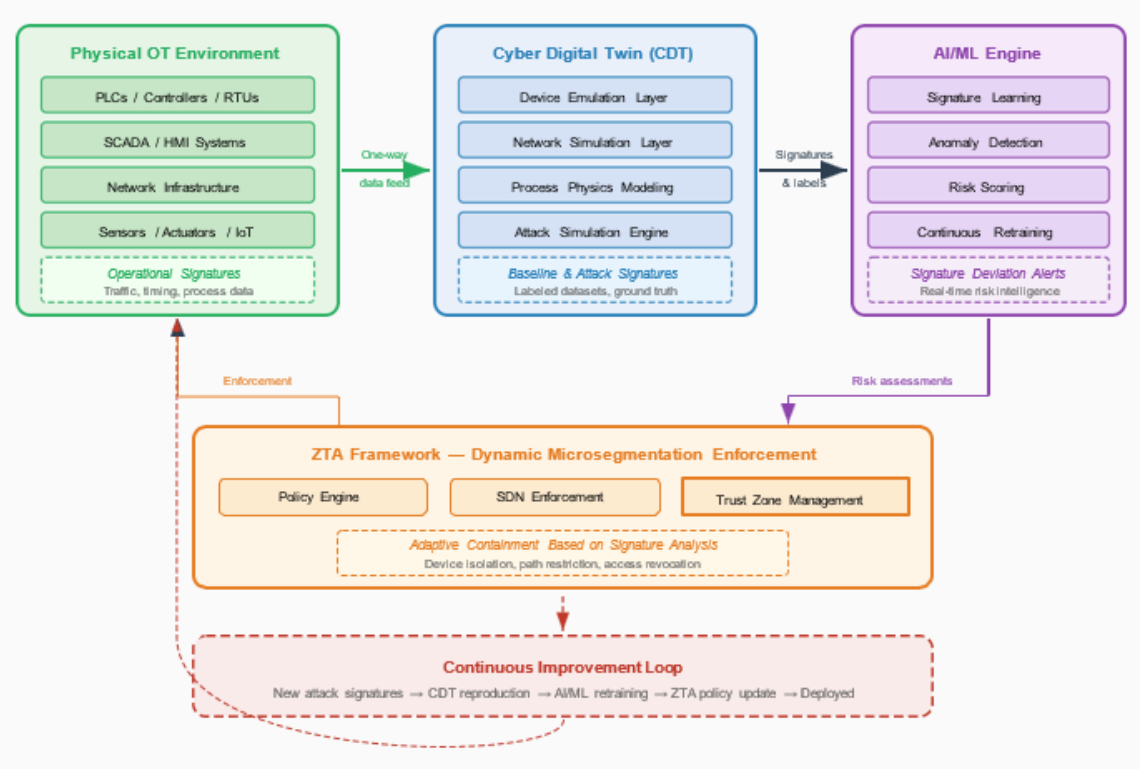


**Figure 2. CDT-ZTA Framework Architecture Showing the Signature-Driven Data Flow from Physical OT Systems Through the CDT, AI/ML Engine, and ZTA Framework Enforcement, with the Continuous Improvement Loop that Drives Adaptive Defense.**

The CDT is a virtual replica of the target OT environment, including network topology, device emulations, and process simulations. It ingests real-time operational data from sensors and network-monitoring points in the physical environment, keeping synchronized with the actual system state. Critically, this data flow is one-directional—the CDT reflects current conditions while remaining isolated from the production network to prevent bidirectional propagation of risk. The CDT captures and catalogs the full spectrum of operational signatures: device communication signatures (protocol-specific traffic patterns, polling intervals, data payload structures), process signatures (physical response curves, setpoint trajectories, control loop behaviors), and network signatures (traffic volume baselines, flow directionality, service discovery patterns). These baseline signatures become the ground truth against which all future behavior is measured.

The AI/ML component runs multiple detection models trained on signature data generated within the CDT. Nominal operational data establishes behavioral baselines—the full set of normal signatures. Synthetic attack data generated through controlled threat simulation

within the CDT provides labeled examples of malicious signatures that supervised models need. The resulting models deploy to both the CDT (for continued refinement) and the production environment (for real-time monitoring). In signal processing terms, the AI/ML models function as matched filters: they are tuned to the known signatures of the environment and alert when incoming signals deviate from expected patterns beyond defined thresholds.

The ZTA framework microsegmentation component consumes continuous risk assessments from the AI/ML models and dynamically adjusts access control policies based on detected signature anomalies and predicted threat states. Software-defined networking enables granular enforcement down to individual devices. When the AI/ML system flags an anomalous signature from a specific device, microsegmentation policies automatically restrict that device's communication paths—containing the threat before lateral movement begins.

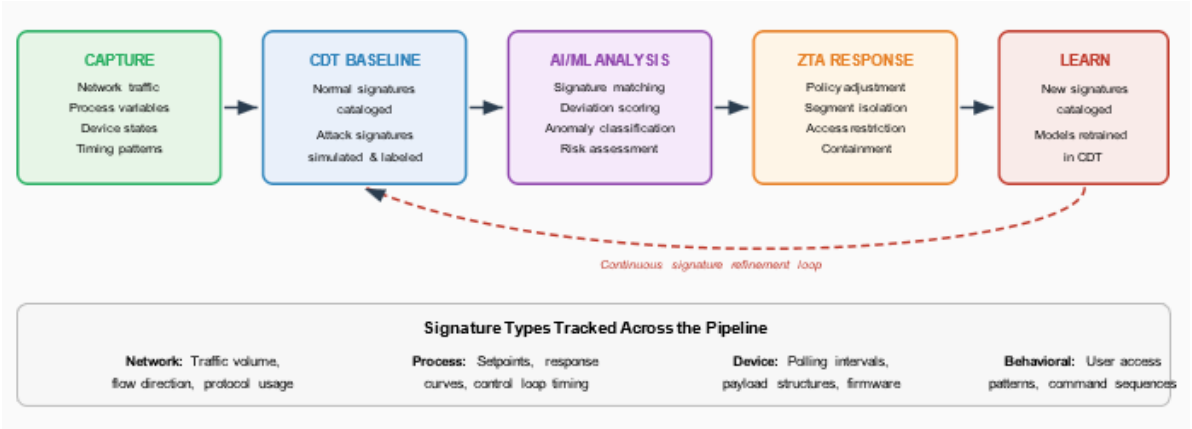


**Figure 3. The Signature Analysis Pipeline Showing How Operational Signatures Flow From Data Capture Through CDT Baselining, AI/ML Analysis, ZTA Framework Enforcement, and Back into the Learning Loop**

### Data Collection and Analysis

Data comes from multiple sources: network traffic captures, process variable logs, device state information, and security event data. Inside the CDT, comprehensive logging captures all system interactions during both normal operations and simulated attacks. This controlled environment produces labeled datasets with known ground truth—known normal signatures and known attack signatures—because the CDT generated both. Production data collection uses existing network monitoring infrastructure augmented with sensors positioned to capture OT-specific protocols and device communication signatures. Data anonymization and aggregation protect sensitive operational information while preserving the traffic patterns and behavioral signatures needed for analysis. Integration with DLA security operations centers provides additional threat intelligence and incident response context.

Analysis methods include statistical baseline characterization of operational signatures, AI/ML model performance evaluation (detection accuracy, false positive rates, detection latency, evasion resilience), and security effectiveness assessment through red team exercises. Red team evaluations pit skilled adversaries against systems protected by the framework, with results compared to baseline configurations without CDT-ZTA protections. The signature-centric approach enables precise measurement: how reliably do AI/ML models detect known attack signatures, how quickly do they identify novel signatures that deviate from baselines, and how effectively does the ZTA framework contain threats once anomalous signatures are flagged?

### Implementation Environment

Framework implementation is centered on a representative DLA facility OT environment.

In accordance with DLA CUI and security policies, specific system identifiers, facility names, and detailed equipment configurations are not disclosed in this paper. The test environment is representative of the Facility Related Control Systems (FRCS) common across DLA installations—systems that govern physical processes essential to maintaining the environmental, operational, and safety conditions required for mission-critical inventory storage and logistics operations. The U.S. Army Engineering and Support Center's Cybersecurity Systems Program has awarded contracts for logical and physical cybersecurity inventories of these systems at DLA locations, providing the foundational asset visibility that the CDT-ZTA Framework builds upon (DVIDS, 2024).

The test environment includes programmable logic controllers governing critical facility processes; supervisory control and monitoring systems coordinating operational parameters across multiple zones; local controllers managing individual subsystems; human-machine interfaces providing operators with real-time visibility into equipment status and conditions; and historian servers logging operational data for trend analysis and compliance reporting. These components are characteristic of the legacy OT architectures found across the DLA's logistics infrastructure—systems designed for reliability and availability that now require cybersecurity protections commensurate with their national security significance.

The CDT will be constructed on commercial virtualization platforms augmented with process simulation capabilities. Device emulations replicate the communication signatures and control logic of physical controllers. Network simulation accurately reproduces traffic signatures, including the periodic polling cycles, setpoint adjustments, and alarm conditions characteristic of the facility's OT operations. Process simulations model the physical dynamics of the controlled environment—so when an attack scenario manipulates a setpoint or disables a controller, the CDT reflects the downstream impact on process conditions and cascading equipment responses. This fidelity matters for signature analysis: the CDT must reproduce signatures with sufficient accuracy that AI/ML models trained within it perform reliably against real-world signatures in the production environment. An adversary who subtly adjusts a critical process setpoint may not trigger a conventional alarm, but the deviation in the operational process signature—detected by AI/ML models tuned to the CDT-established baseline—would flag the anomaly before the resulting conditions compromise mission-critical inventory or endanger personnel.

## Results and Current Progress

### Phased Implementation Status

The CDT-ZTA Framework is being implemented in phases, and it is important to present this work as it stands rather than as a finished product. Current efforts are concentrated on establishing the ZTA framework and microsegmentation foundation—the enforcement architecture that will eventually receive dynamic risk intelligence from the AI/ML and CDT components. This sequencing is deliberate: you cannot build adaptive, AI/ML-informed access control on top of a network that has no segmentation to adapt. The enforcement layer must exist before intelligence can drive it.

The current initiative centers on ZTA framework microsegmentation deployment. The work involves defining granular trust zones around individual OT subsystems—isolating controllers governing different physical processes from one another, separating supervisory management interfaces from local controllers, and enforcing explicit access policies between zones. Software-defined networking infrastructure enables policy enforcement at the device level, and initial segmentation rules are being validated against baseline traffic signatures to ensure that security policies do not inadvertently block legitimate operational communications.

The CDT and AI/ML anomaly detection components are planned for subsequent phases. The CDT environment will be constructed once the microsegmentation architecture is stable and validated, providing both a simulation platform for AI/ML model training and a pre-deployment testing environment for policy refinement. AI/ML models will then be trained within the CDT using both nominal operational signatures and synthetic attack signatures generated through controlled simulation, with the resulting detection capabilities layered onto the existing ZTA framework enforcement architecture.

**Microsegmentation and ZTA Framework Progress**

Early results from the microsegmentation deployment are encouraging, though incomplete. Initial segmentation of the OT test environment has demonstrated the feasibility of device-level policy enforcement without disrupting operational processes. Communication between controllers and their supervisory systems has been restricted to only the specific data exchanges required for coordinated operations, eliminating the broad network access that previously allowed any device on the OT network to communicate with any other device. From a national security perspective, this is foundational: an adversary who compromises a single controller in a non-critical zone is now contained to that zone rather than gaining access to process-critical controllers, supervisory management interfaces, or historian servers—assets whose compromise could cascade into the kind of operational failures that destroy mission-critical inventory and degrade military readiness.

However, several challenges have emerged. Legacy controllers running outdated firmware have limited compatibility with modern SDN-based policy enforcement, necessitating workarounds that add management complexity. Some OT service discovery mechanisms rely on broadcast traffic that device-level segmentation restricts, requiring careful policy exceptions validated individually to avoid creating exploitable gaps. These are not insurmountable problems, but they illustrate why OT microsegmentation cannot be treated as a straightforward IT network exercise—every exception is a potential attack path, and every restriction is a potential operational disruption.

The broader industry data provides context for what mature microsegmentation should deliver. Fortinet's 2025 report found that organizations on integrated OT security platforms with device-level segmentation experienced up to 93% fewer cyber incidents than those on flat networks (Fortinet, 2025). Dragos documented that organizations with comprehensive OT visibility detected and contained ransomware incidents within 5 days, compared to the 42-day industry average (Dragos, 2026). These benchmarks represent the target state the CDT-ZTA implementation is building toward. The significance becomes concrete when mapped to documented threats. Dragos elevated VOLTZITE to Stage 2 of the ICS Cyber Kill Chain after observing the group manipulate engineering workstation software to extract configuration files and probe for operational conditions triggering process shutdowns (Dragos, 2026). In a DLA OT context, that translates to an adversary moving laterally from a compromised workstation to the controllers governing critical facility processes—exactly the kind of movement that device-level microsegmentation is designed to block.

**Anticipated Operational Impact**

Because the project remains active in implementation, a final operational impact assessment is premature. What can be assessed are the anticipated impacts—both positive and negative—based on current progress, the test environment characteristics, and the broader evidence base.

If the framework performs as designed through full implementation, the anticipated benefits are substantial. Microsegmentation alone fundamentally changes the security posture by eliminating the flat network topology that allows unrestricted lateral movement. For a DLA

facility storing sensitive material—pharmaceuticals, electronics, munitions components—containment means the difference between a localized security incident and a facility-wide operational disruption that renders critical inventory unusable, with cascading consequences for military readiness and, ultimately, for the safety of military and civilian lives that depend on an unbroken logistics chain.

When the CDT and AI/ML layers are integrated, the anticipated capability uplift is significant. The CDT will enable pre-deployment validation of segmentation policy changes through signature analysis, eliminating trial-and-error on live systems. AI/ML-driven anomaly detection trained on OT-specific operational signatures should identify subtle attacks—a gradual process setpoint manipulation producing an anomalous process signature, an unauthorized polling pattern producing an anomalous network signature—that rule-based monitoring would miss entirely. The combination of rapid AI/ML signature deviation detection feeding dynamic microsegmentation enforcement should compress the window between intrusion and containment from hours or days to seconds.

The risks of getting this wrong are not abstract. A misconfigured microsegmentation policy that blocks a critical communication path between a controller, and its supervisory system does not just generate an alert can disrupt the physical process that the controller governs, with potential consequences for both inventory integrity and personnel safety. False positives from the AI/ML detection layer, once integrated, present a related risk. If AI/ML models are not adequately trained on the full range of legitimate operational signature seasonal variations, maintenance cycling, demand response events, equipment degradation patterns, they will flag normal operations as suspicious. There is also the risk that adversaries specifically target the framework's integration points, manipulating what the CDT sees to corrupt baseline signatures or trigger inappropriate ZTA framework responses—the kind of patient, process-aware adversary behavior that Dragos documented in KAMACITE's systematic mapping of U.S. industrial control loops throughout 2025 (Dragos, 2026).

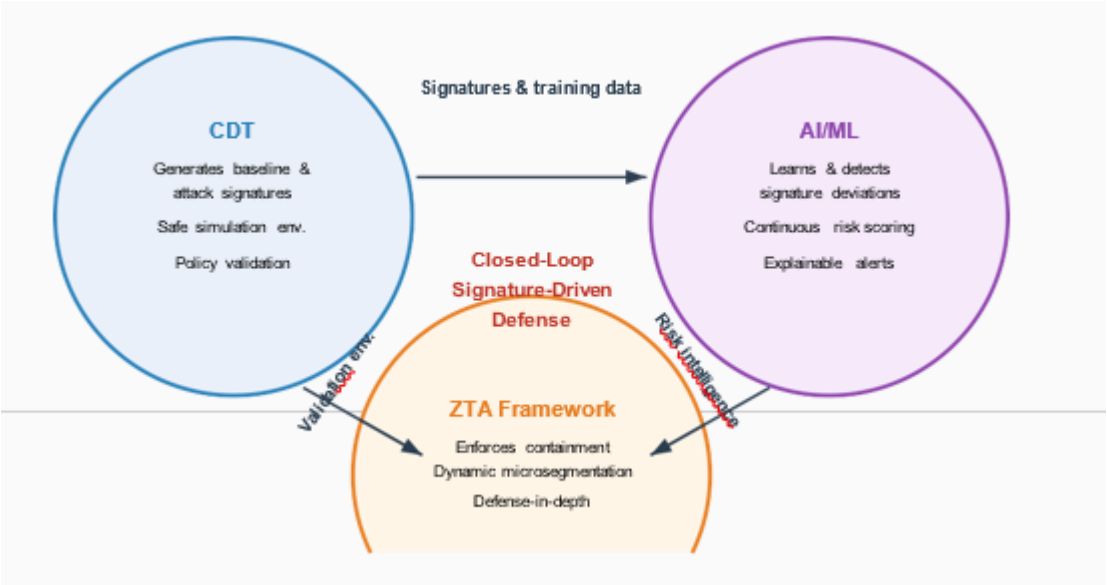


***Figure 4. The Synergistic Integration Model Shows How CDT, AI/ML, and the ZTA Framework Components Depend On One Another. The CDT Generates Signatures for AI/ML Training, AI/ML Provides Risk Intelligence to the ZTA Framework, and the ZTA Framework Validates Policies Within the CDT—Creating a Closed-Loop, Signature-Driven Defensive Architecture.***

## Conclusions

This research is in early stages, and saying otherwise would misrepresent where the work stands. What can be said is that the CDT-ZTA Framework's foundational phase—establishing ZTA framework microsegmentation within a representative DLA OT environment—has demonstrated that device-level policy enforcement in legacy OT environments is achievable without compromising the availability that mission-critical facility control systems demand. That is a necessary first step; it is a promising one.

The early results validate the core architectural premise: build the enforcement layer first, then add signature-based intelligence on top of it. Microsegmentation is establishing the trust zone boundaries and access controls that will eventually receive dynamic risk assessments driven by AI/ML signature analysis from the CDT components. The challenges encountered so far—legacy firmware incompatibilities, OT broadcast dependencies, the operational complexity of managing device-level policy exceptions—are real but solvable and documenting them provides the implementation lessons that other DoD organizations will need as they pursue their own OT ZTA framework mandates under DTM 25-003 and the November 2025 OT-specific guidance.

The national security stakes elevate this work beyond technical exercise. The threat landscape is not waiting for the research to conclude. Dragos documented adversary groups actively mapping industrial control loop signatures across U.S. infrastructure throughout 2025, with KAMACITE conducting systematic reconnaissance and VOLTZITE escalating to Stage 2 of the ICS Cyber Kill Chain (Dragos, 2026). Ransomware groups targeting industrial organizations surged 49% year-over-year, impacting 3,300 organizations. For DLA facilities where OT systems are critical for mission success, these are not theoretical risks—they are active adversary capabilities being refined against the same classes of systems this framework is designed to protect. Compromised OT systems in defense logistics mean degraded readiness, endangered military and civilian lives, and a strategic advantage surrendered to adversaries. The CDT-ZTA Framework is designed to prevent that outcome.

The DLA Strategic Plan 2025–2030 calls for the Agency to think, act, and operate in new ways—and doing so responsibly means validating each phase before building on it, being transparent about what works and what does not, and maintaining the intellectual honesty to adjust course as results and threats evolve. The CDT-ZTA Framework, with its signature-centric approach to detection and containment, is on a trajectory to deliver the integrated defensive capability that DLA's OT environments need. The framework aligns with DTM 25-003, November 2025 OT-specific ZTA framework guidance, and the NSA's January 2026 Zero Trust Implementation Guideline Primer, positioning DLA to meet compliance deadlines while building capabilities that improve security rather than merely satisfying checklists. The start is promising. The work continues.

## Future Steps

The DLA R&D Directorate will continue this research through subsequent phases that expand the framework's capabilities as the technology landscape evolves. Near-term priorities include constructing the CDT of the OT test environment and initiating AI/ML model training using both nominal operational signatures and synthetic attack signatures generated within the CDT. The signature library, the catalog of known normal and known malicious behavioral fingerprints—will grow with each phase, becoming an increasingly powerful asset for detection and response.

Several emerging technology integrations merit focused R&D attention. Large Language Model (LLM) integration into the CDT ecosystem represents a near-horizon opportunity. Recent

work has demonstrated that LLMs can automate interpretation of complex security telemetry, assist real-time threat analysis, and generate human-readable explanations that make sophisticated signature detection outputs accessible to facility operators who are not data scientists (Ferraro et al., 2025). Post-quantum cryptographic protocols for securing OT communications channels will become critical as quantum computing matures. Federated learning approaches that enable AI/ML models to improve across multiple DLA installations without centralizing sensitive operational signature data address both the data scarcity problem and data sovereignty requirements. Autonomous response orchestration that compresses the decision loop between signature deviation detection and containment is another area under evaluation.

Scalability requires candid assessment. CDT demands significant computational resources and specialized expertise for accurate OT modeling. Organizations need to decide whether to build that capability in-house, leverage commercial platforms, or pursue a hybrid approach. The AI/ML component requires ongoing care: models trained today will degrade as adversaries evolve their attack signatures and systems change their operational signatures. Continuous retraining, integration of current threat intelligence, and model validation are operational requirements for the life of the system. Enterprise-scale microsegmentation introduces network management complexity addressable only through automation and orchestration.

Organizational change management remains essential. OT operations personnel need to understand how security controls affect their systems and have clear channels for identifying legitimate communications that policies might inadvertently block. Security personnel need to develop OT-specific expertise—investigating alerts in an industrial environment is fundamentally different from investigating alerts in an enterprise IT network. The DLA Strategic Plan's emphasis on organizational agility and mission-driven skill development supports these management requirements. Technical capability without organizational readiness is expensive shelfware.

As the DoD develops additional ZTA framework guidance for weapon systems and defense-critical infrastructure (DefenseScoop, 2025), extending the CDT-ZTA Framework to those domains—adapting the architecture to the unique operational signatures and constraints of weapons platforms and mission-critical infrastructure—is a natural and necessary evolution. Broader ambition is a framework that does not just keep pace with adversary evolution but anticipates it—one that learns new signatures, adapts its models, and extends to new DLA environments and technology domains as the defense logistics mission demands. Warehouse automation, fuel distribution infrastructure, and manufacturing equipment across the defense industrial base—these all present variations of the same fundamental challenge: legacy OT systems connected to modern networks, operating under availability and national security constraints that make traditional IT security approaches inadequate and sometimes dangerous.

## References


Ahmad, Z., et al. (2025). A new era for digital twins: Progress and industry adoption. *International Journal of Computer Integrated Manufacturing*. https://doi.org/10.1080/27525783.2025.2555877

Al-Sharif, Z. A., et al. (2025). A systematic literature review on the implementation and challenges of Zero Trust architecture across domains. *Sensors, 25*(19), 6118. https://doi.org/10.3390/s25196118

Alqahtani, A., & Almuhaideb, A. (2025). *Comprehensive review of Zero Trust research landscape* [Reference details].

B2B Cyber Security. (2026, February). *Digital twins for modern OT resilience*. https://b2b-cyber-security.de/en/digitale-zwillinge-fuer-moderne-ot-resilienz/

Birihanu, E., & Lendák, I. (2025). Explainable correlation-based anomaly detection for Industrial Control Systems. *Frontiers in Artificial Intelligence, 7*, 1508821. https://doi.org/10.3389/frai.2024.1508821

Cloud Security Alliance. (2024). *Zero Trust guidance for critical infrastructure*.

Cybersecurity and Infrastructure Security Agency. (2024). *Industrial control systems cyber emergency response team year in review 2023*. U.S. Department of Homeland Security.

Defense Logistics Agency. (2025a). *DLA transforms: A call to action—Strategic plan 2025–2030*. https://www.dla.mil/Info/Strategic-Plan/

Defense Logistics Agency. (2025b). *Defense Logistics Agency FY2026 research, development, test & evaluation budget justification*. https://comptroller.war.gov/

DefenseScoop. (2025, December 9). *Pentagon plans to publish Zero Trust strategy 2.0 in early 2026*. https://defensescoop.com/2025/12/09/dod-zero-trust-strategy-2-0-expected-early-2026/

Dickson, P. (2025). *IDC assessment of digital twin technology for cybersecurity exposure management*. IDC.

DoD. (2025a). *DTM 25-003: Implementing the DoD Zero Trust strategy*. https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dtm/DTM%2025-003.PDF

DoD. (2025b). *Zero Trust for operational technology activities and outcomes*. DoD CIO Zero Trust Portfolio Management Office. https://dodcio.defense.gov/Portals/0/Documents/Library/ZTOperationalTechnologyActivitiesOutcomes.pdf

Dragos. (2026). *2026 OT/ICS cybersecurity report and year in review*. https://www.dragos.com/ot-cybersecurity-year-in-review

DVIDS. (2024). Cybersecurity program providing vital services for National Guard, Defense Logistics Agency. https://www.dvidshub.net/news/489056/

Eckhart, M., & Ekelhart, A. (2019). Digital twins for cyber-physical systems security: State of the art and outlook. In *Security and quality in cyber-physical systems engineering* (pp. 383–412). Springer.

El Hachem, J., et al. (2025). Leveraging digital twins for advanced threat modeling in cyber-physical systems cybersecurity. *International Journal of Information Security*. https://doi.org/10.1007/s10207-025-01043-x

El Ouedrhiri, Z., et al. (2025). An explainable Markov chain–machine learning sequential-aware anomaly detection framework for industrial IoT systems based on OPC UA. *Sensors, 25*(19), 6122.

Ferraro, A., et al. (2025). Enabling cyber security education through digital twins and generative AI. *arXiv preprint*, arXiv:2507.17518.

Fortinet. (2025). *2025 State of operational technology and cybersecurity report*.

Gartner. (2023). *Information technology glossary: Operational technology*.

Hassanzadeh, A., Burkett, R., & Xu, S. (2020). A cyber-physical testbed for detecting attacks on industrial control systems. *Proceedings of the 2020 ACM SIGSAC Conference on*

*Computer and Communications Security*, 2211–2213.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–105.

Kim, J., et al. (2026). AI-based anomaly detection in industrial control and cyber-physical systems: A data-type-oriented systematic review. *Electronics, 15*(1), 20. https://doi.org/10.3390/electronics15010020

Kindervag, J. (2010). *Build security into your network's DNA: The Zero Trust network architecture*. Forrester Research.

National Institute of Standards and Technology. (2023). *Guide to operational technology security* (NIST Special Publication 800-82 Rev. 3). U.S. Department of Commerce.

National Security Agency. (2026). *Zero Trust implementation guideline primer* (U/OO/102936-26). https://media.defense.gov/2026/Jan/08/2003852320/-1/-1/0/CTR_ZERO_TRUST_IMPLEMENTATION_GUIDELINE_PRIMER.PDF

Processes. (2025). *Semi-supervised generative adversarial networks for synthetic OT data generation*. [Reference details].

Repetto, M. (2026). Cybersecurity digital twins: Concept, blueprint, and challenges for multi-ownership digital service chains. *Journal of Information Security and Applications, 96*, 104299. https://doi.org/10.1016/j.jisa.2025.104299

Resnick, J. (2025). *Zero Trust Portfolio Management Office commentary on OT requirements* [Reference details].

Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero trust architecture* (NIST Special Publication 800-207). U.S. Department of Commerce.

SANS Institute. (2026). T*op takeaways from the Dragos 2026 OT cybersecurity report*. https://www.sans.org/blog/top-takeaways-from-the-dragos-ot-cybersecurity-report-2026

Stouffer, K., Pease, M., Tang, C., Zimmerman, T., Pillitteri, V., & Lightman, S. (2023). *Guide to operational technology security* (NIST Special Publication 800-82 Rev. 3). National Institute of Standards and Technology.

Trend Micro. (2025, July 31). Trend Micro reinvents proactive security with digital twin technology. https://newsroom.trendmicro.com/2025-07-31

Vielberth, M., et al. (2025). Digital twins in security operations: State of the art and future perspectives. *ACM Computing Surveys*. https://doi.org/10.1145/3746279

World Economic Forum. (2025, March). How digital twin technology can enhance cybersecurity. https://www.weforum.org/stories/2025/03/how-digital-twin-technology-can-enhance-cyber-security/

Xu, J., Wen, H., Li, M., & Chen, X. (2022). Digital twin-enabled anomaly detection for industrial control systems. *IEEE Internet of Things Journal, 9*(15), 13585–13595.

Zhou, Y., Zhang, X., Li, Y., & Yang, Q. (2021). A survey on deep learning approaches for network intrusion detection systems in industrial control systems. *IEEE Access, 9*, 85379–85402.




Acquisition Research Program
Department of Acquisition, Finance, and Manpower
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net