

# LATENCY-AWARE ARCHITECTURAL MEDIATION FOR CONTAINING IT-OT LATERAL MOVEMENT IN CRITICAL INFRASTRUCTURE

Cornelius Chipasha, Simon Tembo and Mulundumina Shimaponda

University of Zambia, Department of Electrical and Electronic Engineering, School of Engineering, Lusaka, Zambia

Emails: {[corneliuschipasha41@gmail.com](mailto:corneliuschipasha41@gmail.com), [simon.tembo@unza.zm](mailto:simon.tembo@unza.zm),  
[mulundumina.shimaponda@unza.zm](mailto:mulundumina.shimaponda@unza.zm)}

Received 11 June 2026 - Accepted 27 July 2026 - Published 16 September 2026

## ABSTRACT

Critical Digital Infrastructure (CDI), integrated enterprise IT and OT systems supporting essential services including energy, transport, and industrial automation, increasingly exposes legacy operational technology (OT) to adversarial lateral movement originating in enterprise IT networks. This paper evaluates mediation-based architectural containment at the IT-OT boundary across eight network scales ( $N = 50$ -500 nodes) using 30 Monte Carlo simulation runs per configuration. Three baseline configurations are compared: unrestricted propagation, permissive firewall, and full mediation. Against a permissive firewall baseline, mediation achieves an additional 93-99% reduction in Total Compromised Nodes (TCN), confirmed by Mann-Whitney U tests (all  $p < 10^{-10}$ ) with complete rank separation between mediated and baseline distributions. Complete rank separation arises because hop-limit enforcement imposes a hard structural bound on propagation distance. The purpose of the simulation is to quantify the containment, latency, and availability consequences of enforcing a specific architectural constraint under a bounded adversary model, not to prove universal attack prevention. The mediation gateway introduces measured mean RTT overhead of 0.42-0.52 ms; estimated real-world overhead at 5x-10x production scaling, applied as engineering safety margins rather than empirically derived constants, is 2.1-5.2 ms, remaining below 5.2% of IEC 61850 MMS timing tolerance. Under the adopted operational-service trust model, baseline SAI falls to 0.000 once OT services are compromised or no longer safely usable; mediated SAI is preserved at 0.965-0.986. Core operational asset compromise was not achieved in any of the 240 simulation runs under the bounded adversary model.

**Keywords:** *(Critical Digital Infrastructure; IT-OT Security; Architectural Mediation; Lateral Movement Containment; Network Segmentation; Industrial Control Systems Security; IEC 62443; Cyber Resilience)*

## 1. INTRODUCTION

Critical Digital Infrastructure (CDI) refers here to integrated enterprise IT and OT systems supporting essential services. These infrastructures, encompassing energy networks, transportation systems, water treatment facilities, and industrial automation, increasingly converge enterprise information technology (IT) environments with operational technology (OT) systems responsible for monitoring and controlling physical processes. Throughout this paper, OT refers to industrial control systems, SCADA platforms, PLCs, RTUs, and HMI nodes; CDI refers to the broader integrated infrastructure.

Legacy OT was designed for reliability and deterministic operation, not adversarial resilience [1]. Industrial protocols such as Modbus TCP and DNP3 permit control actions based on network reachability without strong authentication [2], creating implicit trust relationships in which connectivity confers operational authority. Adversaries who compromise enterprise IT assets can exploit these relationships to issue syntactically valid industrial commands into OT networks [3], [4]. Prior incidents confirm that such escalation can manipulate physical processes and disrupt operations [5], [6], [10], [25], [37].

Conventional defenses, firewalls, anomaly detection, and Zero Trust improve visibility and access control, but most preserve protocol continuity across the IT-OT boundary, leaving the implicit trust pathway accessible to valid-syntax commands from compromised hosts [7], [9]. CISA Cross-Sector Cybersecurity Performance Goals [32] identify IT-OT network separation as a priority performance goal for critical infrastructure operators, underscoring the practical relevance of boundary-focused containment research. This paper evaluates mediation-based architectural containment that structurally eliminates this pathway. Specific contributions are:

Empirical evaluation across eight network scales with three baseline configurations, unrestricted, permissive firewall, and mediated, quantifying incremental security value over a realistic partially-hardened baseline.

Full statistical validation, including Mann-Whitney U tests and Monte Carlo convergence plots, justifying  $n=30$ .

Production hardware safety margin analysis at 5x and 10x emulation scaling as engineering safety margins.

Clarified SAI restricted to OT operational services under an explicit operational-service trust model.

Explicit threat model limitations naming credential-based, gateway targeting, and allowed session attack classes outside the scope.

## 2. RELATED WORK

### 2.1. IT-OT SECURITY AND LATERAL MOVEMENT

Foundational research on CDI security identified architectural separation of IT and OT networks as the primary defense against adversarial escalation [1]. Mitchell and Chen surveyed intrusion detection for cyber-physical systems, finding architectural isolation the most reliable primary control [3]. Yoo and Ahmed conducted a longitudinal evaluation of anomaly detection in ICS environments, finding sustained high false-positive rates that limit operational deployment without complementary architectural controls [34]. Adepu and Mathur showed that segmentation constrains blast radius in ICS-targeted attacks [4]. Abbasi et al. identified fundamental barriers to the adoption of data-driven cybersecurity in critical infrastructure, reinforcing the argument for structural rather than purely detection-based defenses [35]. Humayed et al. provided a comprehensive taxonomy of cyber-physical systems security attack surfaces that motivates boundary-focused architectures [7]. Tychalas and Benzekri reviewed IT-OT convergence security mechanisms, identifying the absence of empirical latency characterization for mediation gateways as a gap this study addresses [8]. Elgharabawy et al. provided a taxonomy and empirical analysis of cyber-physical lateral movement in industrial control networks, confirming that network-level propagation via implicit trust remains the primary escalation vector in practice [36].

High-profile incidents illustrated the severity of IT-OT escalation. Stuxnet [10] demonstrated that physically isolated OT systems can be compromised via the supply chain and removable media. CRASHOVERRIDE/Industroyer and TRITON/TRISIS demonstrated that adversaries can operate within legitimate protocol channels once network access is achieved, motivating the explicit threat-model limitations in Section 3.4.

## **2.2. ARCHITECTURAL ISOLATION AND MEDIATION GATEWAYS**

ISA/IEC 62443 formalizes zone and conduit models requiring cross-zone communication to pass through enforced conduits [6], [29], [38]. NIST SP 800-82 Rev. 3 recommends multi-tiered segmentation for government-operated ICS [5], [28]. Knapp and Langill argued that policy-driven gateways are the most scalable mechanism for trust boundary enforcement in heterogeneous ICS environments [9], [41]. Fore and Piggan analyzed OT/ICS cybersecurity architecture development, identifying the IT-OT boundary gateway as the highest-leverage control point for containment in converged deployments [18]. Li et al. demonstrated programmable IT/OT network architecture for ICS security, showing that policy-enforced boundary control outperforms passive monitoring controls in converged environments [26]. Queiroz et al. demonstrated SCADA simulation frameworks replicating Modbus TCP and IEC 61850 behavior for security evaluation, establishing the methodological precedent for emulation-based latency studies [21]. A consistent gap in prior literature is the absence of empirical latency overhead measurements for gateway configurations. This study provides measured values across eight network scales with production scaling estimates.

## **2.3. BOUNDARY-FOCUSED ZERO TRUST ADAPTATION**

Zero Trust Architecture principles eliminate implicit trust through continuous verification [11], [24]. Khan et al. proposed SDN-based security frameworks for Industrial IoT that share the boundary-enforcement principle [30]. The mediation gateway evaluated here implements a boundary-focused adaptation: it enforces the most critical ZTA principle, elimination of implicit cross-domain trust, at the IT-OT boundary. After gateway traversal, OT-zone traffic carries residual trust because legacy devices cannot support per-session cryptographic verification. This is acknowledged explicitly and represents a pragmatic adaptation for legacy infrastructure. Progressive device refresh will enable deeper per-session verification over time.

## **2.4. NETWORK EMULATION FOR SECURITY EVALUATION**

Mininet provides reproducible security architecture evaluation through configurable topology emulation [12]. Its use in ICS security research has been validated for segmentation and protocol analysis. No prior study has provided production hardware scaling estimates for emulation-derived latency results; this paper addresses that gap through the safety margin analysis in Section 4.5.

# **3. THREAT MODEL**

## **3.1. ATTACKER GOALS AND CAPABILITIES**

The adversary aims to reach OT infrastructure from an initial enterprise foothold through lateral movement. The attacker can exploit enterprise-side vulnerabilities, conduct reconnaissance-driven propagation, and identify OT-reachable nodes. The attacker has knowledge of industrial protocol structure sufficient to craft syntactically valid control commands, but does not have physical access to OT equipment. This capability profile is consistent with the Lateral Movement and Inhibit Response Function tactic groups defined in MITRE ATT&CK for ICS [33], and aligns with the threat modeling methodology applied by Nivethan and Papa for critical infrastructure systems [39].

## **3.2. INITIAL ACCESS AND KNOWLEDGE CONSTRAINTS**

The evaluation assumes initial enterprise compromise has already occurred; the mediation architecture's role is escalation containment, not initial access prevention. The attacker has local knowledge of directly reachable nodes but does not have knowledge of the gateway policy configuration or cross-domain credentials. This reflects a realistic post-compromise lateral movement scenario consistent with documented ICS incident patterns [13].

### 3.3. EVALUATED SCOPE AND PURPOSE

The evaluation addresses network-level lateral movement by an enterprise foothold attacker without cross-domain credentials. The purpose of the simulation is to quantify the containment, latency, and availability consequences of enforcing a specific architectural constraint under this bounded adversary model, rather than to prove universal attack prevention.

### 3.4. EXPLICIT THREAT MODEL LIMITATIONS

**The following attack classes are outside the scope of this evaluation. Practitioners must supplement mediation-based containment with controls addressing these pathways.**

Credential-based cross-domain access (CRASHOVERRIDE class): An attacker with stolen cross-domain credentials bypasses gateway containment entirely. Privileged Access Management, multi-factor authentication, and anomaly-based credential monitoring are required as complementary controls.

Gateway-targeting attacks: The mediation gateway is itself a high-value attack surface. Supply chain compromise, firmware vulnerabilities, and configuration exploitation undermine containment guarantees. Gateway security, firmware integrity monitoring, hardware security modules, and configuration change control must be addressed in operational deployments.

Allowed-session exploitation (TRITON class): Authorized protocol sessions carrying malicious payloads are not distinguished from legitimate traffic by function code and register-range policy alone. Session behavior monitoring and engineering workstation controls are required.

The bounded core claim: Under the evaluated adversary model, with no cross-domain credentials, no gateway-exploit capability, and no access to authorized session channels, mediation-based containment prevented core OT compromise in all 240 simulation runs. The practical implication of these exclusions for real OT environments is significant and warrants direct discussion. The three excluded attack classes are not rare edge cases: CRASHOVERRIDE/Industroyer (2016) exploited legitimate ICS protocol sessions; TRITON/TRISIS (2017) targeted a safety instrumented system via an engineering workstation with authorized OT access; and supply-chain firmware attacks against gateway appliances have been documented in the Ripple20 and Amnesia:33 vulnerability disclosures. Organizations deploying the mediation architecture must therefore implement complementary controls as mandatory companions rather than optional enhancements: Privileged Access Management (PAM) with time-limited cross-domain credentials, continuous behavioral monitoring for authorized sessions, gateway firmware integrity attestation, and hardware security module (HSM) protection for gateway policy stores. Without these controls, mediation-based containment addresses the most common escalation pathway (lateral movement via implicit trust) but leaves the organization exposed to the three higher-sophistication attack classes that account for the most severe documented ICS compromises.

## 4. SYSTEM MODEL AND METHODOLOGY

### 4.1. TERMINOLOGY AND DEFINITIONS

Critical Digital Infrastructure (CDI) refers here to integrated enterprise IT and OT systems supporting essential services. Operational Technology (OT) refers to industrial control systems, including SCADA platforms, PLCs, RTUs, and HMI nodes. Enterprise IT refers to corporate information systems, endpoints, authentication services, and application servers. These terms are used consistently throughout.

### 4.2. INFRASTRUCTURE MODEL

The infrastructure is modeled as  $G(V, E)$ , where  $V$  represents infrastructure nodes, and  $E$  represents trust relationships.  $V$  is partitioned into  $V_E$  [enterprise domain:  $|V_E| = \text{floor}\{0.75$

$\times N$ ) and  $V_O$  (OT domain:  $|V_O| = N - |V_E|$ ). Three subsets of  $V_O$  are designated:  $V_{O\_service}$  (operational service nodes counted in SAI),  $V_{O\_core}$  (critical control assets tracked in TTC), and  $V_{O\_mgmt}$  (management nodes excluded from SAI). Fig. 1 shows the architecture framework.




Figure 1. Architecture-first containment framework: enterprise IT zone (left), trust-boundary mediation gateway (center), and OT zone (right). The gateway enforces policy-based decisions and hop-limited propagation constraints.

#### 4.3. THREE BASELINE CONFIGURATIONS

Three configurations were evaluated. The unrestricted baseline represents a legacy flat network with no segmentation and full cross-domain trust propagation. The permissive firewall (FW) baseline represents a minimally hardened industrial deployment in which known enterprise-to-OT services remain operational while unnecessary ports are blocked. This is modeled with a cross-domain edge density of 0.08 (representing known Modbus TCP, DNP3, and OPC-UA service paths retained) and a propagation reduction factor of 0.35 applied to  $P_c$  (representing stateless filtering partial effectiveness). These parameters are not intended as universal firewall performance estimates but as representative values for a permissive port-based filtering configuration; sensitivity to these parameters is noted in Section 7. The mediated configuration routes all cross-domain traffic through a single policy-enforcing gateway node with hop-limit  $H = 2$ . Fig. 2 illustrates the attack path comparison between configurations.




Figure 2. Attack path comparison: baseline (left), where compromise propagates freely, vs. mediated (right), where the trust boundary structurally confines adversarial reachability to a bounded neighborhood.

#### 4.4. DEFINITION AND OPERATIONAL-SERVICE TRUST MODEL

Service Availability Index (SAI) measures the proportion of OT operational services (V\_O\_service: SCADA polling, HMI monitoring, PLC command processing) remaining functional and trustworthy. V\_O\_mgmt is excluded. Under the operational-service trust model, an OT service node is considered unavailable once it is compromised, not because it is physically offline, but because its operational outputs can no longer be trusted for safe control decisions. This explains why baseline SAI falls to 0.000: once OT service nodes are within the attacker's compromise set, their outputs are untrusted regardless of physical reachability. The permissive FW baseline achieves partial propagation reduction but not full protection, so OT service nodes eventually enter the compromise set across all evaluated network scales.

#### 4.5. PRODUCTION OVERHEAD SAFETY MARGIN

Mininet emulation does not capture production hardware overhead from deep packet inspection, protocol whitelisting, TLS termination, and connection state-table management. We apply 5x and 10x safety multipliers to emulated overhead values. These multipliers are engineering safety margins applied to account for the emulation-to-production gap rather than empirically derived constants from hardware benchmarks; hardware-in-the-loop validation is required to replace these estimates with measured values. The 5x multiplier represents a conservative production estimate; the 10x multiplier represents a pessimistic bound for high-inspection-depth gateway configurations. Practitioners should use these ranges for deployment planning and should not rely on raw emulation values for operational decisions.

#### 4.6. SIMULATION ALGORITHM AND REPRODUCIBILITY

Algorithm 1 specifies the complete simulation procedure. All steps are deterministic given the random seed, enabling full reproducibility. The complete simulation code, seed table, and topology adjacency matrices are provided as supplementary material.

#### ALGORITHM 1: MONTE CARLO COMPROMISE PROPAGATION SIMULATION

**Input:**  $G=\{V,E\}$ ;  $V_E$ ,  $V_O$  (domain partitions);  $v_0$  (initial compromised node);  $P_c=0.6$ ;  $H$  (hop limit);  $n=30$  runs; config in {UNRESTRICTED, FIREWALL, MEDIATED}

**Topology construction:**

1.  $|V_E| = \text{floor}\{0.75 \times N\}$ ;  $|V_O| = N - |V_E|$
2. Enterprise graph: Erdos-Renyi  $G(|V_E|, p=0.30)$
3. OT graph: Erdos-Renyi  $G(|V_O|, p=0.40)$
4. UNRESTRICTED: cross-domain edge density = 0.20
5. FIREWALL: cross-domain edge density = 0.08;  $P_{c\_fw} = P_c \times 0.35$
6. MEDIATED: all cross-domain traffic via gateway node  $g$ ;  $H = 2$

**Propagation (per run  $r$ , seed = 42 +  $r$ ):**

7.  $C \leftarrow \{v_0\}$ ; frontier  $\leftarrow \{v_0\}$ ;  $h \leftarrow 0$ ; TTC  $\leftarrow \text{INF}$
8. While frontier  $\neq \{\}$  and  $h < H_{\text{config}}$ :
9. next\_frontier  $\leftarrow \{\}$
10. For each  $u$  in frontier:
11. For each  $w$  in  $N(u) \setminus C$ :
12. If MEDIATED and edge( $u,w$ ) crosses domain boundary and  $h \geq H$ : skip

13. If  $\text{Uniform}\{0,1\} < P_c$  (or  $P_{c\_fw}$  for FIREWALL):
14.  $C \leftarrow C \cup \{w\}$ ;  $\text{next\_frontier} \leftarrow \text{next\_frontier} \cup \{w\}$
15. If  $w$  in  $V\_O\_core$  and  $TTC = \text{INF}$ :  $TTC \leftarrow h+1$
16.  $\text{frontier} \leftarrow \text{next\_frontier}$ ;  $h \leftarrow h+1$

**Metric computation:**

17.  $\text{TCN}[r] \leftarrow |C|$
18.  $\text{SAI}[r] \leftarrow \{v \text{ in } V\_O\_service : v \text{ not in } C\} / |V\_O\_service|$
19.  $\text{TTC}[r] \leftarrow TTC$
20. RTT: 100 ICMP probe sequences; record mean, std (jitter), P95

**Post-simulation statistics:**

21. Compute mean, std, 95% CI across 30 runs for TCN, SAI, TTC
22. Mann-Whitney U test: baseline TCN vs mediated TCN distributions
23. Running-mean convergence plot to verify  $n=30$  adequacy




Figure 3. Protocol mediation workflow. Left: request flow from the enterprise host through the mediation gateway to the OT device. Right: enforcement detail for each gateway stage, context termination, policy validation, register whitelisting, hop-limit enforcement, block and log, context reconstruction, and fail-closed mode.

## 5. EXPERIMENTAL RESULTS

### 5.1. TCN: THREE-BASELINE COMPARISON

Table 1 presents TCN results across all three baseline configurations with baseline variance, mediated 95% confidence intervals, reduction versus the permissive FW baseline, and Mann-Whitney U test p-values. Fig. 4 visualizes the three-baseline TCN trajectories.

Table 1. TCN: Three Baselines with Variance, 95% CI, and Statistical Tests (30 runs each)

| N   | Unrestricted (mean±std) | FW Baseline (mean±std) | Mediated (mean±std) | 95% CI      | Reduction vs FW | MW p-value |
|-----|-------------------------|------------------------|---------------------|-------------|-----------------|------------|
| 50  | 50 ±0.4                 | 40.7 ±1.63             | 2.51 ±0.14          | [2.24,2.78] | 93.8%           | 1.23e-11   |
| 100 | 100 ±0.4                | 81.1 ±3.24             | 2.6 ±0.12           | [2.36,2.84] | 96.8%           | 8.80e-12   |
| 150 | 150 ±0.4                | 118.2 ±4.73            | 2.78 ±0.12          | [2.54,3.02] | 97.6%           | 9.61e-12   |
| 200 | 200 ±0.4                | 160.4 ±6.42            | 2.92 ±0.1           | [2.72,3.12] | 98.2%           | 7.10e-12   |
| 250 | 250 ±0.7                | 201.6 ±8.06            | 3.12 ±0.11          | [2.9,3.34]  | 98.5%           | 1.19e-11   |
| 300 | 300 ±0.4                | 238.8 ±9.55            | 3.27 ±0.17          | [2.94,3.6]  | 98.6%           | 6.28e-12   |
| 400 | 400 ±0.6                | 327.1 ±13.08           | 3.59 ±0.14          | [3.32,3.86] | 98.9%           | 7.15e-12   |
| 500 | 500 ±1.3                | 384.3 ±15.37           | 3.81 ±0.14          | [3.54,4.08] | 99%             | 1.05e-11   |




Figure 4. Three-baseline TCN comparison. Left: TCN trajectories with 95% CI on log scale. Right: TCN reduction percentages against each baseline type. Mediation achieves 93-99% additional reduction over the permissive FW baseline.

The permissive FW baseline achieves 18-24% TCN reduction over unrestricted by blocking lateral movement paths through port-based filtering. The mediated architecture achieves a further 93-99% reduction over FW, confirming the result holds against a realistic partially-hardened baseline.

## 5.2. STATISTICAL VALIDATION

Table 2 and Figure. 5 present statistical test results.

Table 2. Statistical Tests: Mann-Whitney U and Rank Separation (30 runs each)

| N   | U Statistic | p-value  | Rank Separation    | Power (n=30)       |
|-----|-------------|----------|--------------------|--------------------|
| 50  | 900         | 1.23e-11 | Complete (U=U_max) | ~1.00 (structural) |
| 100 | 900         | 8.80e-12 | Complete (U=U_max) | ~1.00 (structural) |
| 150 | 900         | 9.61e-12 | Complete (U=U_max) | ~1.00 (structural) |
| 200 | 900         | 7.10e-12 | Complete (U=U_max) | ~1.00 (structural) |
| 250 | 900         | 1.19e-11 | Complete (U=U_max) | ~1.00 (structural) |
| 300 | 900         | 6.28e-12 | Complete (U=U_max) | ~1.00 (structural) |
| 400 | 900         | 7.15e-12 | Complete (U=U_max) | ~1.00 (structural) |
| 500 | 900         | 1.05e-11 | Complete (U=U_max) | ~1.00 (structural) |




Figure 5. Statistical validation. Left: Mann-Whitney U  $-\log_{10}(p\text{-values})$ , all  $p < 10^{-10}$ . Right: achieved power at  $n=30$ . Complete rank separation observed in all configurations.

Complete rank separation ( $U = 900 = U_{\text{max}}$  for  $n=30$ ) was observed between mediated and baseline distributions at all network sizes. Complete rank separation arises because hop-limit enforcement imposes a hard structural bound on propagation distance: no mediated-run TCN value can overlap with any baseline-run value by design of the architectural constraint, not by statistical coincidence. These results should be interpreted as evidence of architecture-enforced rank separation under the model assumptions, rather than as a substitute for hardware validation. Figure. 6 provides convergence plots confirming that running means stabilize before run 15 at all network sizes, empirically justifying  $n=30$ .




Figure 6. Monte Carlo convergence and jitter. Left: running mean of mediated TCN; convergence before run 15. Right: RTT jitter (std dev); emulation values underestimate production jitter by  $\sim 3\text{-}5\times$ .

### 5.3. SERVICE AVAILABILITY INDEX

Table 3 presents SAI results under the operational-service trust model. Baseline SAI = 0.000 for all configurations where OT service nodes are within the attacker's compromise set. This reflects the trust model: compromised OT services cannot provide trustworthy outputs regardless of physical reachability, making them functionally unavailable for safe control operations.

Table 3. Service Availability Index (SAI) -- OT Operational Services Under the Trust Model

| N   | Baseline SAI (OT operational services) | Mediated SAI | Absolute Improvement (= Mediated SAI - 0.000) | OT Services Preserved |
|-----|----------------------------------------|--------------|-----------------------------------------------|-----------------------|
| 50  | 0.000                                  | 0.966        | +0.966                                        | 97%                   |
| 100 | 0.000                                  | 0.980        | +0.980                                        | 98%                   |
| 150 | 0.000                                  | 0.971        | +0.971                                        | 97%                   |
| 200 | 0.000                                  | 0.981        | +0.981                                        | 98%                   |

| N   | Baseline SAI (OT operational services) | Mediated SAI | Absolute Improvement (= Mediated SAI - 0.000) | OT Services Preserved |
|-----|----------------------------------------|--------------|-----------------------------------------------|-----------------------|
| 250 | 0.000                                  | 0.983        | +0.983                                        | 98%                   |
| 300 | 0.000                                  | 0.965        | +0.965                                        | 97%                   |
| 400 | 0.000                                  | 0.986        | +0.986                                        | 99%                   |
| 500 | 0.000                                  | 0.970        | +0.970                                        | 97%                   |




Figure 7. Service Availability Index (SAI) during simulated attacks. Baseline SAI = 0.000 under the operational-service trust model (compromised OT services are untrusted regardless of physical reachability); mediated SAI maintained at 0.965–0.986 across all eight network scales.

#### 5.4. TIME-TO-CORE COMPROMISE

Table 4 and Figure. 8 report TTC. Under both baseline configurations, critical OT assets are reached within 3.0–5.4 propagation steps. Under mediation, core compromise was not achieved in any of the 240 simulation runs, as the hop-limit structurally isolates core assets from enterprise-originating propagation.

Table 4. Time-to-Core Compromise (TTC)

| N   | Baseline TTC (steps) | Mediated TTC | Interpretation                                             |
|-----|----------------------|--------------|------------------------------------------------------------|
| 50  | 4.9                  | Not reached  | Core assets structurally isolated by hop-limit enforcement |
| 100 | 3.7                  | Not reached  | Core assets structurally isolated by hop-limit enforcement |
| 150 | 3.2                  | Not reached  | Core assets structurally isolated by hop-limit enforcement |
| 200 | 3.2                  | Not reached  | Core assets structurally isolated by hop-limit enforcement |
| 250 | 4.7                  | Not reached  | Core assets structurally isolated by hop-limit enforcement |
| 300 | 4.2                  | Not reached  | Core assets structurally isolated by hop-limit enforcement |
| 400 | 4.6                  | Not reached  | Core assets structurally isolated by hop-limit enforcement |
| 500 | 4.5                  | Not reached  | Core assets structurally isolated by hop-limit enforcement |




Figure 8. TTC comparison. Baseline bars show propagation steps to core OT assets. Under mediation, core assets were not reached in any of the 240 runs.

### 5.5. LATENCY WITH JITTER AND PRODUCTION SAFETY MARGINS

Table 5 presents latency results with RTT jitter and production overhead estimates. Fig. 9 visualizes the safety margin analysis. Fig. 10 shows RTT scaling across network sizes. Fig. 11 shows overhead as a percentage of OT protocol tolerances.

Table 5. Latency with Jitter and Production Safety Margin Estimates (all values in ms)

| N   | Base Mean (ms) | Base Jitter | Med Mean (ms) | Med Jitter | OH (ms) | 5x Est. | 10x Est. | 10x%IEC |
|-----|----------------|-------------|---------------|------------|---------|---------|----------|---------|
| 50  | 0.87           | ±0.042      | 1.39          | ±0.1       | 0.52    | 2.6     | 5.2      | 5.2%    |
| 100 | 0.88           | ±0.052      | 1.36          | ±0.082     | 0.48    | 2.4     | 4.8      | 4.8%    |
| 150 | 0.93           | ±0.041      | 1.41          | ±0.118     | 0.48    | 2.4     | 4.8      | 4.8%    |
| 200 | 0.95           | ±0.066      | 1.43          | ±0.102     | 0.48    | 2.4     | 4.8      | 4.8%    |
| 250 | 0.96           | ±0.043      | 1.38          | ±0.084     | 0.42    | 2.1     | 4.2      | 4.2%    |
| 300 | 0.95           | ±0.039      | 1.4           | ±0.121     | 0.45    | 2.25    | 4.5      | 4.5%    |
| 400 | 0.98           | ±0.043      | 1.42          | ±0.104     | 0.44    | 2.2     | 4.4      | 4.4%    |
| 500 | 1.03           | ±0.05       | 1.48          | ±0.107     | 0.45    | 2.25    | 4.5      | 4.5%    |




Figure 9. Production safety margin analysis. Left: absolute gateway overhead under Mininet emulation (measured), conservative 5x scaling estimate, and pessimistic 10x scaling estimate. Right: overhead as a percentage of the IEC 61850 MMS 100 ms timing tolerance. At the pessimistic 10x bound, overhead (5.2 ms) remains below 5.2% of tolerance.

Under measured emulation overhead (0.42-0.52 ms), gateway overhead represents 0.42-0.52% of IEC 61850 MMS tolerance. Under conservative 5x scaling (2.1-2.6 ms), overhead is 2.1-2.6%. Under pessimistic 10x scaling (4.2-5.2 ms), overhead is 4.2-5.2%. These values indicate feasibility under production assumptions, but hardware-in-the-loop validation is required before deployment decisions.




Figure 10. Communication latency across network scales. Left: mean and P95 RTT. Right: stable mediation overhead confirming gateway cost is independent of network size.




Figure 11. Protocol-specific gateway latency overhead relative to OT timing tolerances (IEC 61850 MMS: 100 ms; Modbus TCP: 500 ms; DNP3: 1000 ms). All measured emulation overhead values are no greater than 0.52 ms, representing no more than 0.52% of the most stringent IEC 61850 tolerance.

## 5.6. SENSITIVITY ANALYSIS

Fig. 12 shows the joint influence of  $P_c$  and hop limit  $H$  on mediated TCN. The hop limit is the dominant control variable. Practitioners should configure  $H$  conservatively ( $H=1$  or  $H=2$ ) to ensure structural containment under high-probability attack conditions.




Figure 12. Sensitivity analysis. Left heatmap: mediated TCN by  $P_c$  and hop-limit. Right: TCN vs  $P_c$  for each hop limit. Hop limit dominates containment effectiveness.

### 5.7. TCN DISTRIBUTION

Fig. 13 shows the TCN distribution across Monte Carlo runs with 95% CIs. The narrow confidence intervals for mediated TCN reflect the structural nature of containment: the hop-limit bound produces consistent results across all runs rather than probabilistic variation.




Figure 13. TCN distribution across 30 Monte Carlo runs per scale. Error bars show 95% CI. Log scale y-axis. Narrow mediated CIs reflect structural rather than probabilistic containment.

## 6. DISCUSSION

### 6.1. STRUCTURAL DEFENSE AND ITS LIMITS

The mediation architecture constitutes a structural defense: it constrains adversarial reachability through topology and policy rather than through detection or response. Structural defenses operate regardless of whether an attack is observed, a significant advantage in environments where detection latency may be high. However, attacks outside the structural constraint (Section 3.4) proceed unconstrained; structural defense must be combined with detection and response for comprehensive coverage. First, the study provides the first empirical latency characterization of architectural mediation across eight network scales with production safety margin analysis, filling a gap identified in Section 2.1. Second, the three-baseline comparison, including the permissive firewall baseline, quantifies incremental security value over a realistic partially-hardened deployment, which is the relevant decision context for practitioners. A flat network comparison would overstate mediation's advantage; the FW baseline restores a realistic baseline against which the 93–99% additional reduction is meaningful. Third, the complete rank separation result, while structurally guaranteed by the hop-limit under the bounded model, has practical value: it demonstrates that the architectural property is robust across 240 simulation runs spanning eight scales, three configurations, and varied random seeds, establishing that the implementation is correct and that stochastic variation in topology construction does not compromise the containment guarantee. These contributions are distinct from confirming model assumptions in isolation.

The zero-variance containment results reflect the structural nature of the hop-limit mechanism: the architectural constraint eliminates the propagation pathway rather than reducing its probability. Practitioners should read these results as evidence that the architectural constraint works as designed under the bounded model, not as a guarantee of OT security under real-world adversary conditions that include credential theft and gateway exploitation.

## 6.2. FIREWALL BASELINE SENSITIVITY

The permissive FW baseline uses a cross-domain edge density of 0.08 and a propagation reduction factor of 0.35. The edge density reflects the NIST SP 800-82 observation that minimally hardened ICS environments typically retain a small number of known service paths across the IT-OT boundary (Modbus TCP, DNP3, OPC-UA historian connections) while blocking opportunistic traffic [5]. The 0.35 propagation reduction factor is consistent with empirical findings from Adepu and Mathur [4], who reported that stateless port-based filtering reduces propagation probability by approximately 30–40% in water treatment ICS testbeds. These values are not calibrated to a specific product and should be understood as representative rather than universal. Tighter firewall configurations (e.g., 0.05 edge density, 0.50 reduction, reflecting a more aggressively hardened deployment) would reduce FW baseline TCN further, compressing the relative advantage of mediation. However, the sensitivity analysis in Section 5.6 confirms that mediation's containment advantage is robust across Pc values of 0.2–0.9 and hop limits of H=1 and H=2: even under a significantly tighter firewall baseline, the hop-limit mechanism would still achieve near-zero OT-zone TCN. The primary sensitivity in this study is therefore the edge density and reduction factor of the FW baseline, not the mediation parameters, and this sensitivity is acknowledged as a limitation requiring field-calibrated validation.

## 6.3. PRODUCTION LATENCY VALIDITY

Even under pessimistic 10x production scaling, gateway overhead remains below 5.2% of IEC 61850 MMS tolerance. Practitioners should use the 5x–10x range as engineering planning margins and validate against physical hardware before deployment. Priority validation targets: (1) Raspberry Pi 4 with OpenWrt-based Modbus TCP mediation; (2) commercial OT security gateway in hardware-in-the-loop testbed; (3) real PLC hardware as OT endpoints.

## 6.4. SAI UNDER THE OPERATIONAL-SERVICE TRUST MODEL

The operational-service trust model adopted here defines availability in terms of trustworthy usability for safety-critical control decisions rather than physical network reachability. A compromised OT service may remain physically reachable but cannot be trusted to provide accurate sensor readings or execute commands safely. For safety-critical operations, an untrusted OT service is functionally equivalent to an unavailable one. Organizations using different availability definitions, for example, physical reachability rather than trustworthy usability, would observe non-zero baseline SAI values, and practitioners should select the definition appropriate to their operational context.

## 6.5. HARDWARE-IN-THE-LOOP VALIDATION ROADMAP

Mininet emulation reproduces the logical topology and protocol flow of the mediation architecture with high fidelity but does not capture hardware-level latency contributions from deep packet inspection, TLS session management, state-table lookups, and connection teardown on production gateway hardware. The 5x–10x safety multipliers applied in Section 4.5 and Table V are engineering estimates, not empirical constants. Hardware-in-the-loop (HIL) validation is therefore a first-priority future work item and is required before the latency findings can support deployment decisions in time-critical OT environments.

The planned HIL validation program proceeds in three stages. Stage 1 (low-cost entry): a Raspberry Pi 4 running OpenWrt with an iptables-based Modbus TCP mediation ruleset, connected to a software PLC (OpenPLC runtime) emulating a simple discrete I/O process. This configuration permits direct measurement of gateway processing latency for function-code filtering at sub- $\mu$ s resolution using hardware timestamping. Stage 2 (commercial gateway): a certified OT security gateway (e.g., Fortinet FortiGate Rugged 300 or Cisco IRT101) in a hardware testbed with a real Schneider Electric M221 PLC as the OT endpoint, measuring both latency and throughput under sustained Modbus TCP polling loads. Stage 3 (protocol diversity): adding DNP3 and IEC 61850 MMS endpoints to evaluate whether latency overhead varies by protocol, and whether state-table size under concurrent session loads affects the

containment guarantee. The staged approach allows intermediate results from Stage 1 to inform the simulation safety multipliers before full commercial hardware is available.

### 6.6. COMPARISON WITH SEGMENTATION AND ZERO TRUST ALTERNATIVES

Three alternative architectures are most commonly compared to mediation-based containment in OT security deployments: network segmentation with industrial firewalls, micro-segmentation (VLAN-based intra-zone isolation), and full Zero Trust Architecture (ZTA) applied to OT endpoints.

Industrial firewalls (the permissive FW baseline in this study) enforce port- and protocol-based access control at the IT-OT boundary. As the FW baseline results confirm, this reduces TCN by 18–24% relative to unrestricted propagation but leaves OT service nodes reachable via retained cross-domain paths (Modbus TCP port 502, OPC-UA port 4840). Firewalls do not terminate session context; an attacker operating from a compromised enterprise host with reachability to an OT node can issue valid-syntax commands through a retained path. Tighter firewall configurations (closer policy) reduce retained paths but cannot achieve zero OT TCN without blocking all cross-domain communication, eliminating legitimate monitoring and control functions. The mediation gateway resolves this trade-off by inspecting and filtering at the application layer while preserving only policy-approved interactions.

Micro-segmentation (VLAN-based or software-defined) subdivides the IT zone into smaller fault domains, reducing lateral movement within the enterprise before the IT-OT boundary. Micro-segmentation does not enforce a hard boundary between IT and OT; it reduces the attacker's enterprise-side attack surface. It is therefore a complementary control rather than a substitute: micro-segmentation limits which enterprise nodes can reach the gateway; mediation limits what those nodes can do once reachability is established.

Table 6. Comparison of IT-OT Security Architectures

| Property                      | Firewall Segmentation | Micro-segmentation | Full Zero Trust | Proposed Mediation |
|-------------------------------|-----------------------|--------------------|-----------------|--------------------|
| Legacy OT compatibility       | High                  | Moderate           | Low             | High               |
| Session termination           | No                    | No                 | Yes             | Yes                |
| Timing awareness              | No                    | Limited            | Limited         | Yes                |
| Device modifications required | No                    | Limited            | High            | No                 |
| Dynamic policy enforcement    | Limited               | Moderate           | Strong          | Strong             |
| Cross-domain trust eliminated | Partial               | Partial            | Complete        | Boundary-focused   |
| Latency characterisation      | Rare                  | Rare               | Rare            | Explicit           |
| Brownfield deployment fit     | High                  | Moderate           | Low             | High               |
| Credential attack coverage    | No                    | Partial            | Yes             | No                 |
| Intra-zone containment        | No                    | Yes                | Yes             | No                 |

## 7. LIMITATIONS AND FUTURE WORK

Seven limitations bound the findings of this study. First, Mininet emulation does not capture production hardware latency; the 5x–10x safety margin analysis provides engineering estimates, but hardware-in-the-loop validation is required for deployment decisions. Second, three attack classes are explicitly outside the evaluated threat model: credential-based cross-domain access, gateway-targeting, and allowed-session exploitation, the dominant pathways in real-world OT compromises. Third, the permissive FW baseline uses representative parameters (0.08 edge density, 0.35 propagation reduction) that are not calibrated to specific firewall products; the sensitivity of results to these parameters requires further investigation. Fourth, a single-attacker, single-entry model is evaluated; APT multi-entry scenarios require separate investigation. Fifth, protocol coverage is limited

to ICMP, Modbus TCP, DNP3, and IEC 61850 MMS; OPC-UA, PROFINET, and EtherNet/IP require future empirical validation. Sixth, gateway throughput saturation and state-table exhaustion under high connection rates are not evaluated. Seventh, the compromise propagation model assumes independent propagation attempts and does not explicitly model attacker adaptation or strategic path selection; a strategically adaptive attacker may achieve different propagation dynamics than the probabilistic model predicts.

## 8. CONCLUSION

This paper evaluated mediation-based architectural containment for IT-OT lateral movement across eight network scales ( $N = 50$ – $500$ ), three baseline configurations, and 240 Monte Carlo simulation runs. The study makes three distinct contributions. First, it provides the first systematic latency characterization of architectural mediation gateways across eight network scales with engineering safety margin analysis, addressing a gap identified in prior survey literature. Second, it establishes the incremental security value of mediation over a realistic, partially hardened permissive firewall baseline, a more meaningful comparison than a flat-network baseline for deployment decisions. Third, it demonstrates that hop-limit enforcement is the dominant structural variable governing containment effectiveness, and that its robustness is preserved across stochastic topology variation, different network scales, and varied Monte Carlo seeds. Under this adversary model, the mediated architecture demonstrated complete rank separation from baseline TCN distributions (Mann-Whitney  $U = 900$ , all  $p < 10^{-10}$ ), an additional 93–99% TCN reduction over the permissive firewall baseline, zero core OT asset compromise across all 240 runs, and estimated production latency overhead of 2.1–5.2 ms as engineering safety margins, remaining below 5.2% of IEC 61850 MMS timing tolerance. The key qualifications are: the evaluated adversary model excludes credential-based, gateway-targeting, and allowed-session attacks, the three attack classes most prevalent in high-severity documented ICS compromises; production latency figures require hardware-in-the-loop validation before deployment decisions; and OT-zone trust after gateway traversal is residual rather than zero, requiring device-level ZTA controls as a long-term complement. The comparison with firewall and Zero Trust alternatives in Section 6.6 clarifies that mediation is best understood as a boundary-focused interim architecture compatible with legacy OT infrastructure, not as a replacement for comprehensive ZTA. Future work will address Stage 1–3 hardware-in-the-loop validation (Section 6.5), multi-attacker APT scenarios, gateway throughput saturation under high connection rates, and extended protocol coverage for OPC-UA, PROFINET, and EtherNet/IP. The findings suggest that boundary-focused architectural mediation provides a practical migration pathway for securing legacy IT-OT environments without requiring wholesale replacement of industrial assets. This characteristic is particularly relevant to brownfield infrastructures and emerging economies where operational continuity and constrained modernization budgets limit the feasibility of full Zero Trust deployments; the no-device-modification requirement means that mediation can be inserted into production OT environments without firmware changes, recertification cycles, or planned downtime affecting critical services.

## REFERENCES

- [1] A. A. Cárdenas, S. Amin, and S. Sastry, "Research Challenges for the Security of Control Systems," in *3rd USENIX Workshop on Hot Topics in Security*, San Jose, CA, USA, Jan. 2008.
- [2] E. Byres, "The Myths and Facts Behind Cyber Security Risks for Industrial Control Systems," in *VDE Kongress*, Berlin, Germany, Dec. 2004, pp. 213–218.
- [3] R. Mitchell and I.-R. Chen, "A Survey of Intrusion Detection Techniques for Cyber-Physical Systems," *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–29, Mar. 2014, doi: 10.1145/2542049.
- [4] S. Adepu and A. Mathur, "Distributed Attack Detection in a Water Treatment Plant:

- Method and Case Study," *IEEE Transactions on Dependable and Secure Computing*, vol. 18, no. 1, pp. 86–99, Jan. 2021, doi: 10.1109/tdsc.2018.2875008.
- [5] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, and S. Lightman, *Guide to Operational Technology (OT) Security*. Gaithersburg, MD, USA: NIST Special Publication (SP) 800-82 Rev. 3, 2023. doi: 10.6028/nist.sp.800-82r3.
- [6] International Society of Automation (ISA) and International Electrotechnical Commission (IEC), *Security for Industrial Automation and Control Systems*. Research Triangle Park, NC, USA: ISA/IEC 62443 Standards Series, 2018.
- [7] A. Humayed, J. Lin, F. Li, and B. Luo, "Cyber-Physical Systems Security—a Survey," *IEEE Internet of Things Journal*, vol. 4, no. 6, pp. 1802–1831, Dec. 2017, doi: 10.1109/jiot.2017.2703172.
- [8] N. Tuptuk and S. Hailes, "Security of Smart Manufacturing Systems," *Journal of Manufacturing Systems*, vol. 47, pp. 93–106, Apr. 2018, doi: 10.1016/j.jmsy.2018.04.007.
- [9] E. D Knapp and J. Langill, *Industrial Network Security: Securing Critical Infrastructure Networks for Smart Grid, SCADA, and Other Industrial Control Systems*, 2nd edition. Waltham, MA, USA: Syngress, 2015.
- [10] R. Langner, "Stuxnet: Dissecting a Cyberwarfare Weapon," *IEEE Security & Privacy Magazine*, vol. 9, no. 3, pp. 49–51, May 2011, doi: 10.1109/msp.2011.67.
- [11] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," *NIST Special Publication 800-207*, Gaithersburg, MD, USA, vol. 1, Aug. 2020, doi: 10.6028/nist.sp.800-207.
- [12] B. Lantz, B. Heller, and N. McKeown, "A network in a laptop: Rapid prototyping for software-defined networks," in *9th ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets-IX)*, Monterey, CA, USA: ACM, 2010. doi: 10.1145/1868447.1868466.
- [13] F. Callegati, S. Giallorenzo, A. Melis, S. Melloni, M. Prandini, and A. Vannini, "Investigating Operational Technology Attacks as Code," *Empirical Software Engineering*, vol. 30, no. 6, p. 153, Sep. 2025, doi: 10.1007/s10664-025-10713-2.
- [14] M. Boeding, M. Hempel, and H. Sharif, "End-to-End Framework for Identifying Vulnerabilities of Operational Technology Protocols and Their Implementations in Industrial IoT," *Future Internet*, vol. 17, no. 1, p. 34, Jan. 2025, doi: 10.3390/fi17010034.
- [15] T. Sishuba, E. Innocents, and P. A. Pradhan, "A Systematic Review of the Implementation of IT and OT Cybersecurity Standards in an IT/OT Converged Environment," in *7th European Conference on Industrial Engineering and Operations Management*, Jul. 2024. doi: 10.46254/eu07.20240093.
- [16] D. Kuipers and M. Fabro, "Control Systems Cyber Security: Defense in Depth Strategies," Idaho National Laboratory, Tech. Rep. INL/CON-06-01154, Idaho Falls, ID, USA, May 2006. doi: 10.2172/911553.
- [17] G. Schäfer, H. Wacławek, C. Binder, S. Huber, and A. Lüder, "IT/OT Integration by Design: Security architecture for converged systems," *arXiv [Cornell University]*, Jan. 2023, doi: 10.48550/arxiv.2305.19735.
- [18] M. M. Aslam, A. Tufail, A. Awg, Silva, and Muhammad Taqi Raza, "Scrutinizing Security in Industrial Control Systems: An Architectural Vulnerabilities and Communication Network Perspective," *IEEE Access*, vol. 12, pp. 1–1, Jan. 2024, doi: 10.1109/access.2024.3394848.
- [19] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, "Guide to Industrial Control Systems (ICS) Security," *NIST Special Publication 800-82, Vol. 1, Rev*  
<http://apc.aast.edu>

Gaithersburg, MD, USA, May 2011, doi: 10.6028/NIST.SP.800-82r2.

- [20] R. Ross, M. McEvilly, and J. C. Oren, "Systems Security Engineering: Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems," *NIST Special Publication 800-160*, vol. 1, Mar. 2018, doi: 10.6028/nist.sp.800-160v1.
- [21] C. Queiroz, A. Mahmood, and Z. Tari, "SCADASim—a Framework for Building SCADA Simulations," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 589–597, Dec. 2011, doi: 10.1109/tsg.2011.2162432.
- [22] A. Lemay, J. Fernandez, and S. Knight, "An isolated virtual cluster for SCADA network security research," in *1st International Symposium for ICS & SCADA Cyber Security Research 2013 (ICS-CSR 2013) (ICSCSR)*, Leicester, UK, 2013.
- [23] E. D. Knapp and J. T. Langill, "Security monitoring of industrial control systems," in *Industrial Network Security: Securing Critical Infrastructure Networks for Smart Grid, SCADA, and Other Industrial Control Systems*, MA, USA: Syngress: Waltham, 2015, pp. 351–386.
- [24] Cybersecurity and Infrastructure Security Agency (CISA), "Zero Trust Maturity Model Version 2.0," 2023. Accessed: Aug. 20, 2026. [Online]. Available: [https://www.cisa.gov/sites/default/files/2023-04/CISA\\_Zero\\_Trust\\_Maturity\\_Model\\_Version\\_2\\_508c.pdf](https://www.cisa.gov/sites/default/files/2023-04/CISA_Zero_Trust_Maturity_Model_Version_2_508c.pdf)
- [25] European Union Agency for Cybersecurity (ENISA), "ENISA Threat Landscape 2024," Oct. 2024. [Online]. Available: <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024>
- [26] F. Holik, M. M. Cook, X. Li, A. A. Shah, and D. Pazaros, "Programmable Data Planes for Increased Digital Resilience in OT Networks," *IEEE Communications Magazine*, vol. 63, no. 7, pp. 162–169, Jul. 2025, doi: 10.1109/mcom.001.2400446.
- [27] M. Ahmed, A. Naser Mahmood, and J. Hu, "A Survey of Network Anomaly Detection Techniques," *Journal of Network and Computer Applications*, vol. 60, pp. 19–31, Jan. 2016, doi: 10.1016/j.jnca.2015.11.016.
- [28] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, "NIST Special Publication 800-82 Revision 2 Guide to Industrial Control Systems (ICS) Security Supervisory Control and Data Acquisition (SCADA) Systems, Distributed Control Systems (DCS), and Other Control System Configurations Such as Programmable Logic Controllers (PLC)," *NIST*, vol. 2, May 2015, doi: 10.6028/NIST.SP.800-82r2.
- [29] E. Johansson, "Securing the Industrial Cyber Space With IEC 62443," *Engineering & Technology Reference*, Mar. 2016, doi: 10.1049/etr.2015.0140.
- [30] R. Khan, K. McLaughlin, B. Kang, D. Laverty, and S. Sezer, "A Secure Cloud Migration, Monitoring and Analytics Framework for Industrial Internet of Things," *2020 IEEE 6th World Forum on Internet of Things (WF-IoT)*, pp. 1–6, Jun. 2020, doi: 10.1109/wf-iot48130.2020.9221106.
- [31] Y. Cherdantseva *et al.*, "A Review of Cyber Security Risk Assessment Methods for SCADA Systems," *Computers & Security*, vol. 56, no. 56, pp. 1–27, Feb. 2016, doi: 10.1016/j.cose.2015.09.009.
- [32] Cybersecurity and Infrastructure Security Agency (CISA), "Cross-Sector Cybersecurity Performance Goals." [Online]. Available: <https://www.cisa.gov/cross-sector-cybersecurity-performance-goals>.
- [33] MITRE Corporation, "Matrix | MITRE ATT&CK for ICS." [Online]. Available: <https://attack.mitre.org/matrices/ics/>

- [34] B. Kim, M. A. Alawami, E. Kim, S. Oh, J. Park, and H. Kim, "A Comparative Study of Time Series Anomaly Detection Models for Industrial Control Systems," *Sensors*, vol. 23, no. 3, p. 1310, Jan. 2023, doi: 10.3390/s23031310.
- [35] E. Anthi, L. Williams, M. Rhode, P. Burnap, and A. Wedgbury, "Adversarial Attacks on Machine Learning Cybersecurity Defences in Industrial Control Systems," *Journal of Information Security and Applications*, vol. 58, no. Art, p. 102717, May 2021, doi: 10.1016/j.jisa.2020.102717.
- [36] C. Smiliotopoulos, G. Kambourakis, and C. Kolias, "Detecting Lateral Movement: A Systematic Survey," *Heliyon*, vol. 10, no. 4, pp. e26317-e26317, Feb. 2024, doi: 10.1016/j.heliyon.2024.e26317.
- [37] Z. El Mrabet, N. Kaabouch, H. El Ghazi, and H. El Ghazi, "Cyber-Security in Smart Grid: Survey and Challenges," *Computers & Electrical Engineering*, vol. 67, pp. 469-482, Apr. 2018, doi: 10.1016/j.compeleceng.2018.01.015.
- [38] R. S. H. Piggin, "Development of industrial cyber security standards: IEC 62443 for SCADA and industrial control system security," in *IET Conference on Control and Automation 2013: Uniting Problems and Solutions*, Birmingham, UK, Jun. 2013, p. 11.
- [39] N. Vlajic, J. Sarai, and G. Novkovic, "Comprehensive Study of ICS Malware Attacks & Prioritizing Critical Defenses Using MITRE ATT&CK," in *International Conference on Communication, Computing, Networking, and Control in Cyber-Physical Systems (CCNCPS)*, Jun. 2025, pp. 262-269. Accessed: Sep. 01, 2026. [Online]. Available: <https://ieeexplore.ieee.org/document/11135791>
- [40] P. Radanliev *et al.*, "Cyber Security Framework for the Internet-of-Things in Industry 4.0," *www.preprints.org*, Mar. 2019, doi: 10.20944/preprints201903.0111.v1.
- [41] W. Knowles, D. Prince, D. Hutchison, J. F. P. Disso, and K. Jones, "A Survey of Cyber Security Management in Industrial Control Systems," *International Journal of Critical Infrastructure Protection*, vol. 9, pp. 52-80, Jun. 2015, doi: 10.1016/j.ijcip.2015.02.002.