Frontiers in Computer Science and Artificial Intelligence
ISSN: 2978-8048
DOI: 10.32996/fcsai
Journal Homepage: www.al-kindipublisher.com/index.php/fcsai

FCSAI
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Crypto-Agile Zero Trust Architecture for IT-OT Convergence in Industrial Control Systems: A Digital Twin-Driven Framework over 5G/TSN with Post-Quantum Cryptography

Syed Mohiuddin Qadri
Senior Member, IEEE, Independent Researcher, Richmond, TX 77407, USA
Corresponding Author: Syed Mohiuddin Qadri**, Email:** qadri@ieee.org

| ABSTRACT
The convergence of Information Technology (IT) and Operational Technology (OT) in Industrial Control Systems (ICS) has dissolved the implicit trust boundaries that traditionally protected critical infrastructure. While Zero Trust Architecture (ZTA) is increasingly advocated as the successor to perimeter-based defense for converged industrial environments, existing ZT proposals for ICS rarely address three concurrent realities: (i) the rise of high-fidelity Digital Twins (DTs) as the natural locus of process-aware trust evaluation, (ii) the migration of industrial transport to 5G non-public networks and Time-Sensitive Networking (TSN) with deterministic timing budgets that cannot accommodate naive authentication round-trips, and (iii) the imminent obsolescence of classical asymmetric cryptography due to cryptographically relevant quantum computers, which regulatory frameworks such as NIS2, the EU Cyber Resilience Act, and IEC 62443 already require to be mitigated through crypto-agility. This paper proposes QZT-ICS, a unified ZT framework for IT-OT convergence that (a) uses a synchronized Digital Twin as the policy decision point for continuous, process-aware trust scoring; (b) embeds policy enforcement at the 5G slice and TSN bridge layers to preserve sub-millisecond determinism; and (c) integrates a crypto-agile post-quantum key-establishment and signature layer based on NIST-standardized ML-KEM and ML-DSA, with hybrid classical-PQC modes for legacy field devices. The framework is evaluated through a simulated water-treatment SCADA testbed against false data injection, command injection, and lateral-movement attack classes. Results quantify the trade-off between PQC overhead, TSN scheduling slack, and DT-driven trust evaluation latency, and demonstrate viable deployment paths for brownfield ICS. The paper concludes with a maturity model for staged adoption and open research challenges.

| KEYWORDS
Zero Trust Architecture, IT-OT convergence, Industrial Control Systems, Digital Twin, 5G non-public networks, Time-Sensitive Networking, Post-Quantum Cryptography, Critical Infrastructure, Cyber-Physical Systems, IEC 62443, NIST SP 800-207

| ARTICLE INFORMATION

ACCEPTED: 01 May 2026 **PUBLISHED:** 13 June 2026 **DOI:** 10.32996/fcsai.2026.5.9.2

1. Introduction

1.1 Motivation: The IT-OT Convergence Threat Landscape

Industrial Control Systems (ICS) form the cyber-physical backbone of modern civilization. They regulate the generation and distribution of electricity, the treatment and pumping of municipal water, the refining and transport of hydrocarbons, the operation of mass transit, and the production lines of every manufacturing economy. For most of their history, these systems were physically isolated from the public internet, designed for reliability and safety, with cybersecurity treated as a peripheral concern at best. That posture is no longer tenable. Over the past fifteen years, a sequence of landmark incidents has demonstrated, with escalating consequence, that the air gap between Information Technology (IT) and Operational Technology (OT) has been largely dissolved by the same connectivity that drives Industry 4.0.

The opening shot was Stuxnet in 2010, the first publicly documented malware to weaponize ICS protocol semantics, manipulating Siemens PLCs to destroy uranium-enrichment centrifuges while reporting normal status to operators [47]. The 2015 and 2016 attacks on the Ukrainian power grid showed that nation-state adversaries could move from IT footholds, spearphishing, credential theft, BlackEnergy and Industroyer malware, into substation HMIs and trip breakers at scale. Triton/TRISIS in 2017 specifically targeted Safety Instrumented Systems, demonstrating willingness to override the last engineering layer protecting human life. NotPetya, also in 2017, was nominally an IT-borne ransomware worm, yet it crippled Maersk's terminal operations, Merck's pharmaceutical manufacturing, and the Chernobyl radiation-monitoring system, exposing how IT-OT coupling silently transmits IT incidents into OT damage [63], [64]. The 2021 Colonial Pipeline incident took down 5,500 miles of fuel pipeline despite the ransomware affecting only the IT billing system, because the operator could not safely segregate billing from operations. The Oldsmar water-treatment intrusion in the same year showed that remote-access misuse alone is sufficient to alter sodium-hydroxide setpoints in drinking water. Most recently, the Volt Typhoon campaign disclosed in 2023 demonstrated patient, multi-year persistence in IT systems of U.S. critical-infrastructure operators, dwell time apparently positioned for pre-positioning, not data theft [5].

The pattern across these incidents is consistent and instructive. In each case, attackers exploited an IT foothold and pivoted through newly converged pathways, remote-access VPNs, vendor maintenance tunnels, jump hosts, file shares, and update servers, into OT environments that retained the security posture of their air-gapped past. Traditional defense-in-depth architectures grounded in the Purdue Enterprise Reference Architecture [2] assumed that horizontal segmentation, perimeter firewalls, and DMZs could maintain enforceable boundaries between levels. In practice, the digital transformation of manufacturing, the rise of cloud-based historians and analytics, the use of common Ethernet and IP transports in OT, and the demand for remote operations have collectively rendered those boundaries permeable [11], [46], [48], [49]. The threat surface that ICS now confronts is no longer that of an isolated system but that of an internet-exposed cyber-physical fabric whose compromise can cause physical damage, environmental disaster, and human casualties.

1.2 Why Zero Trust Alone Is Insufficient for Modern ICS

In response to the collapse of perimeter-based defense, Zero Trust Architecture (ZTA) has emerged as the consensus successor doctrine for both enterprise IT and, increasingly, OT. The foundational formulation in NIST SP 800-207 [1] discards the notion of trusted internal networks and replaces it with continuous per-request verification based on identity, device posture, and contextual signals. Recent guidance from CISA, the Cloud Security Alliance, and Carnegie Mellon SEI [4], [5], [10], [11] has explicitly extended ZT thinking to OT and ICS, and a growing body of academic work has proposed ZT-derived architectures for industrial environments [12], [13], [17]-[23]. Yet despite this momentum, three structural problems remain unaddressed by existing ZT-for-ICS proposals, and each is becoming more acute rather than less.

First, trust evaluation in current ZTAs is primarily *identity-aware*, focused on who or what is making a request, but ICS demands evaluation that is *process-aware*. An authenticated engineering workstation issuing a syntactically valid command can still cause physical damage if the command is incompatible with the current operational state of the plant. Identity-bound trust scoring cannot, by itself, distinguish a legitimate engineer's setpoint change from an adversary using the engineer's credentials to drive a process into an unsafe regime. The Digital Twin (DT) literature [3], [25], [26] suggests a natural answer, a continuously synchronized virtual model of the physical process can supply the missing process context, but the integration of DTs with ZT decision logic has been treated only at a conceptual level, and never in the IT-OT converged setting we consider here.

Second, ZT enforcement was designed for IT round-trip latencies measured in tens of milliseconds, not the deterministic timing budgets of industrial transport. As ICS migrate from siloed fieldbuses to 5G non-public networks and IEEE 802.1 Time-Sensitive Networking (TSN) [36], [57], [58], the assumed latency budget for per-flow authentication and policy verification collapses. Critical traffic in TSN-scheduled flows may have an end-to-end deadline of one millisecond or less; a single naive authentication round-trip can consume the entire budget. Yet no published ZT framework for ICS quantifies, much less designs around, this timing constraint.

Third, every ZT proposal we are aware of is silent on the cryptographic transition that regulatory frameworks are now mandating. NIST has finalized the first three post-quantum cryptography (PQC) standards: FIPS 203 ML-KEM, FIPS 204 ML-DSA, and FIPS 205 SLH-DSA [6]-[8]. The EU NIS2 Directive [60] and Cyber Resilience Act [61] require crypto-agility in critical infrastructure from 2026 onwards. CISA's 2024 OT-specific PQC guidance [4] anticipates "a significant and enduring challenge" for owners and operators. A ZT framework whose authentication and key-establishment primitives will be broken by a cryptographically relevant quantum computer is, by definition, not a long-term solution for ICS assets with twenty-year service lives. Yet the intersection of ZT and PQC in the ICS context appears in no peer-reviewed paper to date.

1.3 The Three Forces Reshaping ICS Security

Three technological forces are converging on industrial control systems simultaneously, and each independently demands architectural attention. Their *concurrent* effect on ZT design has not yet been studied.

Digital Twins as process-aware oracles. The Digital Twin, a high-fidelity, continuously synchronized virtual representation of a physical process or asset [3], [54], has matured from a manufacturing-optimization tool into a candidate cybersecurity primitive. Recent work has demonstrated that DTs can host real-time anomaly detection [25], support process-residual-based intrusion detection with F1 scores above 96% on water-treatment SCADA testbeds [25], and provide situational awareness of device trust states in industrial IoT [27]. NIST IR 8356 [3] explicitly recommends a Zero Trust security model for digital-twin deployments. The DT is, in our view, the natural Policy Decision Point (PDP) for a process-aware ZT framework: it is the only component in an ICS architecture that simultaneously holds an identity model, a device-posture model, and a process model.

5G non-public networks and TSN as the new industrial transport substrate. Industry 4.0 manufacturing and critical-infrastructure operators are deploying private 5G networks and TSN-enabled industrial Ethernet at scale [36], [55], [57], [58]. Together, these technologies promise to converge the historically fragmented industrial networking landscape, PROFINET, EtherCAT, Modbus TCP, DNP3, OPC UA, onto a unified deterministic substrate. Recent academic work has begun to consider ZT in 5G [30]-[34] and zero-touch configuration in TSN [35], but no work has formally analyzed the impact of inline ZT enforcement on TSN-scheduled traffic, nor proposed PEP placements that respect the deterministic timing guarantees that justify TSN's adoption in the first place.

The post-quantum cryptographic transition. The threat from cryptographically relevant quantum computers (CRQCs) is no longer speculative. Mosca's framing of the harvest-now-decrypt-later adversary [44] has been corroborated by intelligence assessments, and the NIST PQC standardization process has now produced FIPS 203, 204, and 205 [6]-[8]. The unique challenge for ICS is that asymmetric cryptography in OT is used in lightly protected but high-consequence ways, VPN concentrators, remote-access gateways, firmware-update signing, Secure Boot, and increasingly encrypted industrial protocols such as OPC UA, on hardware that may have fifteen-year remaining service lives [4], [41]. Crypto-agility, mandated by NIS2 and the CRA [60], [61] and incorporated into IEC 62443 [9], is now a regulatory requirement for ICS, not an option.

These three forces do not merely coexist; they interact. PQC primitives are slower than their classical counterparts and produce larger key and signature artifacts, with direct implications for TSN timing budgets. DT-based trust evaluation introduces its own latency overhead that competes with TSN slack and PQC verification overhead for the same scheduling envelope. A ZT framework for IT-OT converged ICS cannot consider any one of these forces in isolation; it must reason about all three trade-offs simultaneously. To our knowledge, no peer-reviewed publication has attempted this integrated four-pillar synthesis. This paper proposes such a synthesis and evaluates it empirically.

1.4 Contributions

This paper makes the following contributions:

- We propose QZT-ICS, a unified four-pillar reference architecture (illustrated in Fig. 1) that integrates a Zero Trust control plane, a Digital Twin-based Policy Decision Point, 5G/TSN deterministic transport with embedded Policy Enforcement Points, and a crypto-agile post-quantum cryptographic substrate for IT-OT converged Industrial Control Systems.
- We develop a formal threat model and a requirements catalog that maps NIST SP 800-207 tenets, IEC 62443 foundational requirements and security levels, and NIST CSF 2.0 functions to OT-specific constraints, including the safety-availability priority inversion that distinguishes OT from IT.
- We introduce a mechanism for using the Digital Twin as the Policy Decision Point, with a trust-scoring function that combines identity, device posture, process-state residuals, behavioral baselines, and external threat intelligence into a continuously updated trust score.
- We propose a latency-aware Policy Enforcement Point placement strategy that embeds ZT enforcement at the 5G network-slice and TSN bridge layers, and we show how policy decisions can be translated into IEEE 802.1Qbv Gate Control List entries without violating the underlying scheduled-traffic timing budget.
- We define a crypto-agile post-quantum cryptographic layer based on the NIST-standardized ML-KEM, ML-DSA, and SLH-DSA primitives, with hybrid classical-PQC modes (X25519+ML-KEM) and protocol-bumps-in-the-wire approaches for legacy fieldbus protocols that cannot natively carry PQC payloads. This layer extends the cryptographic abstraction and policy-driven orchestration concepts established in our prior multi-cloud crypto-agility framework [62] to the deterministic-timing and resource-constrained requirements of industrial control systems.

- We evaluate the framework on a simulated water-treatment SCADA testbed against four representative attack classes, false data injection, command injection, lateral movement, and harvest-now-decrypt-later, and quantify the trade-off between PQC overhead, TSN scheduling slack, and DT-driven trust-evaluation latency on representative ICS hardware.
- We propose a five-level deployment maturity model for brownfield and greenfield adoption and articulate four concrete open research challenges spanning hardware-attested TSN bridges, Digital Twin model-poisoning defenses, quantum-classical hybrid transitions on long-lived assets, and standardization gaps across 3GPP, IEEE 802.1, IETF, and IEC 62443.

1.5 Paper Organization

The remainder of this paper is structured as follows. Section 2 provides background on ICS architecture, the NIST SP 800-207 Zero Trust tenets, Digital Twins, 5G/TSN, post-quantum cryptography, and the regulatory landscape. Section 3 surveys related work organized by pairwise intersection and identifies the four-pillar gap that motivates this work. Section 4 develops the threat model and requirements catalog. Section 5 presents the QZT-ICS framework, including the DT-based PDP, slice- and bridge-aware PEPs, and the crypto-agile PQC layer. Section 6 describes the implementation and testbed. Section 7 reports detection performance, PQC overhead, TSN timing impact, and trust-evaluation latency results. Section 8 discusses trade-offs, deployment maturity, regulatory mapping, and limitations. Section 9 articulates open research challenges. Section 10 concludes.

2. Background

2.1 ICS and the Purdue Model in the Age of Convergence

Industrial Control Systems comprise the hardware, software, and networks that monitor and control physical processes. The taxonomy is layered. At the lowest level, sensors and actuators interface directly with the physical world. Programmable Logic Controllers (PLCs) and Remote Terminal Units (RTUs) execute deterministic control loops on this field data, typically at 1-100 ms scan cycles, and drive the actuators that close the loop. Above this control layer, Human-Machine Interfaces (HMIs) and engineering workstations expose process state to operators; Supervisory Control and Data Acquisition (SCADA) systems and Distributed Control Systems (DCS) aggregate, archive, and orchestrate the lower layers across geographically distributed plants [2].

The Purdue Enterprise Reference Architecture (PERA), developed in the 1990s and now codified in ISA-95 and IEC 62443, organizes these components into six logical levels (L0 through L5), with L0-L3 designated as the OT zone and L4-L5 as the enterprise IT zone. A Demilitarized Zone (DMZ) between L3 and L4 traditionally enforced the separation, permitting only specific data flows, typically historian replication and patch distribution, to cross the boundary [2], [46]. For two decades this model framed the practical defense-in-depth posture of industrial cybersecurity.

The convergence of IT and OT has eroded each of these boundaries simultaneously. Cloud-based analytics now require continuous data egress from L2 to L5 and beyond. Vendor remote-access tunnels open inbound paths directly into L2 for diagnostics. Engineering workstations are increasingly virtualized and provisioned from enterprise Active Directory. Modern fieldbuses, Modbus TCP, PROFINET, EtherNet/IP, OPC UA, run on the same Ethernet and IP substrate as the IT network, so a physical compromise of one segment exposes the other. Industry 4.0 architectures such as RAMI 4.0 explicitly assume horizontal integration across the entire stack [11], [49]. In the language of NIST SP 800-82 [2], the practical effect is that the OT environment can no longer be considered isolated, and traditional perimeter-based controls are no longer sufficient.

2.2 Zero Trust Architecture: NIST SP 800-207 Tenets and Their OT Constraints

NIST SP 800-207 [1] codifies Zero Trust as an architectural philosophy resting on seven tenets: (i) all data sources and computing services are resources; (ii) all communication is secured regardless of network location; (iii) access to individual enterprise resources is granted on a per-session basis; (iv) access is determined by a dynamic policy that includes the observable state of the client identity, application, and requesting asset; (v) the enterprise monitors and measures the integrity and security posture of all owned and associated assets; (vi) all resource authentication and authorization are dynamic and strictly enforced before access is allowed; and (vii) the enterprise collects as much information as possible about the current state of assets, network infrastructure, and communications and uses it to improve its security posture. The logical components defined by the same publication, the Policy Engine (PE) which makes the access decision, the Policy Administrator (PA) which signals enforcement, and the Policy Enforcement Point (PEP) which executes it, together form what is now widely called the ZT control plane.

Translating this canonical model into OT introduces three categories of friction that have been the subject of considerable recent commentary [11], [15], [17]. First, the *safety-availability priority inversion*: OT prioritizes safety and availability above confidentiality and integrity, which is the inverse of the classical IT C-I-A triad. A ZT decision that fails closed during a process upset can itself be a safety incident. Second, *legacy protocols are unauthenticated by design*: Modbus, DNP3, and large parts of

the PROFINET and PROFIBUS ecosystems carry no native authentication and cannot be retrofitted without breaking conformance with installed bases. Third, *asset lifetime versus crypto lifetime*: a PLC may remain in service for twenty years, while cryptographic primitives are sunset on a much shorter horizon. The CMU SEI guidance [11] makes the related observation that implementing ZTA for ICS shouldn't impede uninterrupted access to systems to execute safety functions, a constraint that subordinates ZT enforcement to safety logic, not the other way around.

These constraints do not invalidate the ZT model for OT, but they materially reshape it. PEPs must support fail-safe (sometimes fail-open) defaults aligned with safety logic; trust evaluation must include process context, not just identity; and the cryptographic primitives underlying authentication must themselves be replaceable as algorithms are deprecated. These reshaping requirements are precisely the gaps that motivate the QZT-ICS framework presented in Section 5.

2.3 Digital Twins for Industrial Cybersecurity

NIST IR 8356 [3] defines a digital twin as the virtual (i.e., digital) representation of a physical or perceived real-world entity, concept, or notion. More operationally, Fuller et al. [54] characterize DTs by three properties: they are virtual representations of physical objects, they are connected to those objects through near-real-time bidirectional data flows, and they support analytical or operational use cases such as predictive maintenance, simulation, and visualization. The DT concept has matured from manufacturing-optimization research into a candidate cybersecurity primitive over the past five years [25], [26], [27], [54], [55], [56].

We distinguish three roles a DT may play in industrial cybersecurity. As a *passive monitor*, the DT receives field telemetry and compares observed sensor readings against model predictions, flagging residuals that exceed a calibrated threshold; this is the architecture used in DT-driven intrusion detection systems such as Sayghe's water-treatment IDS, which achieved an F1 score of 96.3% with sub-500 ms detection latency on the SWaT testbed [25]. As a *predictive simulator*, the DT projects the consequences of a proposed control action forward in time, enabling pre-validation of commands before they reach the physical process; this is the role typically envisioned in ZT-DT proposals for smart manufacturing [26]. As an *active enforcement substrate*, the DT serves as the Policy Decision Point of the ZT control plane, hosting trust algorithms that consume process-state context alongside identity and posture signals.

The third role is the most ambitious and the one we adopt. It is supported by the NIST IR 8356 [3] recommendation that DT deployments adopt Zero Trust, and by the early work of Vega Vega et al. [24] which framed a ZT-by-DT approach for smart grids. It is also the role most exposed to a class of DT-specific attacks. If the DT itself is compromised, through model poisoning [28], replay of stale telemetry, or fabrication of synthetic measurements, the entire trust-scoring chain collapses. Defenses against DT compromise are therefore an integral part of any DT-as-PDP architecture, a point we develop in Sections 5.2.2 and 9.2.

2.4 5G Non-Public Networks and IEEE 802.1 TSN for Deterministic OT Transport

Two transport technologies are simultaneously reshaping the industrial-networking substrate. Private 5G, specifically the Standalone Non-Public Network (SNPN) and Public Network-Integrated Non-Public Network (PNI-NPN) variants defined in 3GPP, brings ultra-reliable low-latency communication (URLLC), network slicing, and licensed-spectrum wireless to factory floors, energy substations, and port logistics [55], [58]. IEEE 802.1 Time-Sensitive Networking (TSN), a set of amendments to standard Ethernet, brings deterministic scheduling, frame replication for fault tolerance, and tight time synchronization to wired industrial networks [36], [57]. The two are complementary: 5G dominates the access and mobility tier, while TSN dominates the deterministic backbone, and gateways translate between them.

For ZT design, three properties of these transports matter. First, both technologies expose first-class abstractions, 5G *network slices* identified by Single Network Slice Selection Assistance Information (S-NSSAI), and TSN *scheduled flows* governed by IEEE 802.1Qbv Gate Control Lists, that map naturally onto ZT micro-segmentation. A per-slice or per-flow PEP can be defined and enforced at line rate. Second, the same abstractions impose timing budgets that ZT was not originally designed around. A TSN scheduled flow may have an end-to-end deadline of 100 microseconds to 1 millisecond, leaving essentially no room for an out-of-band PEP round trip. Third, the convergence with industrial protocols, PROFINET TSN, OPC UA Pub/Sub over TSN, EtherCAT-over-TSN, means that ZT policy expressed at the network layer can be translated into policy on the application-protocol layer through a single enforcement point, but only if that point understands both layers.

Recent academic work has begun to explore ZT in 5G [30]-[34] and zero-touch configuration in TSN [35]. Bello et al.'s comprehensive ACM Computing Surveys treatment of TSN for industrial automation [36] catalogs the standards landscape but does not address ZT integration. Seliem et al. [37] and Li et al. [38] explore TSN fault tolerance and zero-loss switching architectures, again without ZT. The Mantas et al. [40] testbed for industrial private 5G security provides empirical grounding for

measurements of the kind we present in Section 7, but does not propose a ZT framework. The QZT-ICS framework occupies the unaddressed intersection.

2.5 The Post-Quantum Threat to ICS Asymmetric Cryptography

Shor's algorithm renders the discrete-logarithm and integer-factorization problems on which RSA, Diffie-Hellman, and elliptic-curve cryptography depend tractable in polynomial time on a sufficiently large fault-tolerant quantum computer [45]. Although such a machine, the cryptographically relevant quantum computer (CRQC), does not yet exist, Mosca's well-known argument [44] frames the relevant security parameter: if X is the duration for which information must remain confidential, Y is the time required to migrate to PQC, and Z is the time until a CRQC is available, then any system with $X + Y > Z$ is already at risk through the harvest-now-decrypt-later adversary, who records ciphertext today for future decryption.

For ICS the security parameter is particularly hostile. Asset service lives of fifteen to twenty years inflate X. Migration timelines for fleets of unpatchable PLCs and RTUs inflate Y to a decade or more. Both NIST and CISA now treat the cumulative risk as immediate. The NIST PQC standardization process culminated in 2024 with the publication of three Federal Information Processing Standards: FIPS 203 ML-KEM (a module-lattice-based key-encapsulation mechanism derived from CRYSTALS-Kyber), FIPS 204 ML-DSA (a module-lattice-based digital signature derived from CRYSTALS-Dilithium), and FIPS 205 SLH-DSA (a stateless hash-based signature derived from SPHINCS+) [6]-[8]. CISA's 2024 Post-Quantum Considerations for Operational Technology [4] identifies OT-specific deployment sites, VPN concentrators, remote-access gateways, firmware-update signing, Secure Boot, and increasingly encrypted industrial protocols such as OPC UA, and anticipates a significant and enduring challenge for owners and operators.

The PQC primitives differ from their classical predecessors in ways that directly affect ICS deployment. ML-KEM-768 public keys and ciphertexts are roughly 1.2 kB each, against a few hundred bytes for ECDH. ML-DSA-65 signatures are approximately 3.3 kB against 64 bytes for Ed25519. The computational cost on Cortex-M-class processors is substantially higher than ECC. These properties interact directly with the TSN timing budgets discussed in Section 2.4: the PQC layer cannot be designed without reference to the transport layer. Oliva del Moral et al. [41] provide a comprehensive survey of these constraints in the critical-infrastructure setting; the analogous treatment for the enterprise multi-cloud context appears in our prior work [62], which covers algorithm-selection trade-offs in latency-tolerant cloud workloads. We extend both lines of analysis by integrating PQC overhead with ZT trust evaluation and TSN scheduling in Section 7, in the resource-constrained and deterministic-timing regime that neither cloud nor general critical-infrastructure surveys have previously addressed.

2.6 Regulatory Drivers

Cybersecurity for ICS is no longer a purely technical or commercial concern; in most major jurisdictions it is now a regulatory one. IEC 62443 [9] establishes foundational requirements (FR1 through FR7) and Security Levels (SL 1-4) for industrial automation and control systems and has been incorporated by reference in numerous sector-specific regulations. The recent update to the standard explicitly accommodates post-quantum cryptography requirements, mandating crypto-agility at the component and system levels.

In the European Union, the NIS2 Directive (Directive (EU) 2022/2555) [60] mandates risk-management, incident-reporting, and supply-chain security obligations on essential and important entities across most critical-infrastructure sectors, with transposition deadlines that have now elapsed. The Cyber Resilience Act (Regulation (EU) 2024/2847) [61] complements NIS2 by placing security-by-design and crypto-agility requirements on products with digital elements, with full applicability from 2027. In the United States, NIST SP 800-82 Rev. 3 [2] provides the canonical OT security guidance, and the NIST Cybersecurity Framework 2.0 reorganizes outcomes around six core functions (Govern, Identify, Protect, Detect, Respond, Recover) that align with, and can be mapped against, the QZT-ICS framework, as we show in Section 8.3. CISA's 2024 OT PQC guidance [4] and its 2026 zero-trust roadmap for OT environments [5] together provide the most current operational guidance for critical-infrastructure operators in the United States.

These regulatory drivers establish the timeline within which any ICS security architecture must be deployable. Crypto-agility is no longer a research preference but a 2026-2030 compliance obligation. Zero Trust is no longer a strategic option but, increasingly, the baseline expectation. The QZT-ICS framework is designed against this regulatory backdrop.

3. Related Work

This section organizes the prior literature by pairwise intersection with the four pillars of the QZT-ICS framework and identifies the gap that motivates the integrated synthesis. We treat each pillar pairing as a separate body of work, then conclude in Section 3.5 with a comparative analysis demonstrating that no published proposal occupies the four-pillar intersection.

3.1 Zero Trust for ICS and OT

The application of Zero Trust principles to industrial control systems has accelerated rapidly since NIST SP 800-207 [1] was published in 2020. Early conceptual work translated the IT-oriented tenets into the OT context, while more recent contributions have begun to address specific architectural elements.

Among architectural proposals, Feng and Hu [17] presented one of the first peer-reviewed treatments dedicated to Industrial Cyber-Physical Systems (ICPS). Their Cyber-Physical Zero Trust Architecture (CP-ZTA) introduces a multi-layer access control engine and integrates physical model-based and data-driven policy optimization to address cross-layer penetration attacks between cyber and physical layers. The work is conceptually important because it acknowledges that ZT in ICS must reason about physical state, but the design is evaluated only at the level of simulated power-grid scenarios and does not address deterministic transport or post-quantum cryptography.

Federici, Martintoni, and Senni [18] take a complementary angle, focusing specifically on remote-access ZT for Industrial IoT infrastructures. Their architecture combines attribute-based access control with continuous attestation and is validated against industrial maintenance scenarios. The remote-access focus is well-suited to the realistic threat surface of modern OT but does not extend to in-plant traffic or to the deterministic timing constraints we consider here. Sims [19] presents an applied analysis in which the Purdue Model is augmented with ZT controls and evaluated as a deployment template for ICS/SCADA. The thesis is operationally grounded but, like most thesis-level work, presents a conceptual mapping rather than a quantitatively evaluated system.

More recent work has begun to combine ZT with adjacent techniques. Lv et al. [20] propose an asynchronous federated learning-based ZT architecture in which trust evaluation is decentralized across multiple PLC clusters; the federated aspect addresses one form of scalability, though the proposal does not engage with deterministic transport or quantum resistance. Aljuaid and Ahmed [21] examine ZT specifically in oil and gas downstream operations, integrating the technology with sector-specific regulatory requirements and field instrumentation patterns. Bilipelli [23] reports an AI-enabled ZT intrusion detection system for industrial IoT in Scientific Reports; the work showcases the feasibility of ML-driven trust scoring but treats the underlying network transport as a black box.

Across this body of work, three patterns are evident. First, every contribution acknowledges the OT-specific constraints (safety primacy, legacy protocols, asset lifetimes) but typically addresses only a subset. Second, almost all of the work is either conceptual, simulation-based, or limited to a single sector; few of the academic proposals quantify end-to-end performance on hardware-realistic testbeds. Third, none consider the four pillars together. As we show in Section 3.5, no published proposal integrates Digital Twins, 5G/TSN, and post-quantum cryptography within a single ZT framework.

3.2 Zero Trust + Digital Twins

The intersection of ZT and DTs is the most academically developed of the pairwise combinations, in part because both technologies arose contemporaneously in the manufacturing-modernization literature.

Vega Vega et al. [24] are credited with the first explicit framing of a ZT-by-DT architecture for industrial systems. Working in the smart-grid domain, they demonstrate that a digital twin synchronized with substation state can host trust algorithms that operate on physical observables (voltage, frequency, breaker positions) rather than purely on identity. The proposal is sound at the architectural level but is not empirically evaluated against attack workloads, and the transport layer is treated as a given. Anumbe, Saidu, and Akhilesh [26] extend the concept to smart manufacturing cyber-physical systems, articulating the role of DTs in vulnerability assessment and ZT policy enforcement. Their framework is comprehensive in scope but presented at a review/synthesis level rather than implemented; like Vega Vega et al., it identifies the direction without operationalizing it.

Sayghe [25] reports the most empirically grounded DT-based intrusion-detection work in this space, evaluating a Digital Twin-driven IDS on the SWaT water-treatment testbed against false data injection, denial-of-service, and command-injection attacks. The system achieves an F1 score of 96.3% with sub-500 ms detection latency. Importantly, Sayghe demonstrates that DT-based detection can be process-aware in ways that pure network-traffic analysis cannot, providing strong empirical motivation for our use of the DT as the PDP. However, the system is a detector, not a ZT enforcement framework: there is no policy engine, no PEP, and no integration with the transport or cryptographic layers.

Prasad et al. [27] in Scientific Reports (2026) propose a zero-trust digital twin framework that combines deep-learning anomaly detection, differential privacy, blockchain-inspired hash-chained auditing, and DT-based situational awareness in IIoT. With 89-91% accuracy on a unified NSL-KDD/CICIDS-2017/IoT-23 dataset, the work is the closest existing analogue to our proposed framework. Yet two crucial elements are absent: deterministic transport (5G/TSN) and post-quantum cryptography. The blockchain component addresses immutability and auditability but is itself classically secure and so vulnerable to harvest-now-decrypt-later attack.

The NIST Internal Report IR 8356 [3] provides the standards-body anchor for this line of work, explicitly recommending Zero Trust as the appropriate security model for digital twin deployments and enumerating fourteen specific considerations. Our framework treats IR 8356 as a normative reference, mapping each of its trust considerations onto components of the QZT-ICS architecture. Defense against attacks on the DT itself is an emerging concern. Zhang et al. [28] analyze model poisoning attacks against distributed network digital twin systems and propose defensive aggregation techniques. We adopt their formulation of the model-poisoning threat in Section 5.2.2 and Section 9.2.

3.3 Zero Trust + 5G / TSN

The combination of ZT with 5G has attracted considerable industrial and academic attention, though chiefly outside the ICS domain.

Sun et al. [30] of Dell Technologies present an early architectural framing of ZT for 5G air interfaces. Their proposal embeds ZT within Open-RAN deployments and addresses the threat surface of network slicing, but the focus is on enterprise and telecommunications use cases rather than industrial control. The deterministic timing constraints that distinguish ICS traffic are not analyzed. Bin Sediq et al. [33] specifically target 5G open architecture network slices for ZT, showing how slice-aware authentication and policy enforcement can be embedded at the AUSF and SMF layers. Their MILCOM paper provides a useful pattern for our slice-aware PEP in Section 5.3.1 but, again, is not specialized to industrial workloads. Chen et al. [31] extend the discussion to 6G in IEEE Network, where the security challenges are amplified by network slicing at finer granularity and by edge intelligence. Yu et al. [34] propose a 5G-based ZT network security platform, evaluated for general enterprise workloads.

Wang, Chen, and Liu [32] in Digital Communications and Networks explicitly consider ZT in 5G Industrial Internet collaboration systems and identify gaps between current security techniques and the ZTA model. Their work is the most directly relevant precedent to ours on the 5G side but again does not address TSN timing, DT integration, or PQC. Pokhrel [39] reports an AI-enabled cybersecurity framework for future 5G wireless infrastructure, benchmarked against the NIST ZTA and 3GPP TS 33.501; the work explicitly identifies quantum-resilient cryptography as future work, suggesting that the integration we propose is recognized as needed but not yet supplied.

On the TSN side, the picture is thinner. Bello et al. [36] in ACM Computing Surveys (2024) provide a comprehensive treatment of TSN for industrial automation and document the standards landscape (IEEE 802.1AS, 802.1Qbv, 802.1CB) and integration with industrial protocols. The survey acknowledges that TSN provides protection against bandwidth violation, malfunctioning, and malicious attacks, but does not engage with ZT integration as a design question. Jiang et al. [35] propose a zero-touch dynamic configuration framework for TSN that achieves sub-millisecond reconfiguration; zero-touch here refers to operational autonomy, not Zero Trust, but the framework is operationally relevant because any ZT-aware TSN architecture must support comparably fast policy reconfiguration.

Seliem et al. [37] and Li, Deng, and Han [38] explore TSN fault tolerance and zero-loss switching architectures, again without engaging with ZT or PQC. Mantas et al. [40] present a testbed and software architecture for security in industrial private 5G networks; their measurement methodology is influential for our own evaluation in Section 7, though they do not propose a ZT framework. In short, ZT has been considered for 5G; TSN has been characterized for industrial automation; but the integration of ZT enforcement with TSN timing budgets in an industrial setting has not, to our knowledge, been quantitatively addressed.

3.4 Post-Quantum Cryptography for OT/ICS

The PQC-for-ICS literature is the youngest and least developed of the four pillars.

The leading academic survey is Oliva del Moral et al. [41], an arXiv preprint that examines PQC adoption challenges in critical infrastructure. The authors document the difficulty of implementing PQC primitives on legacy OT hardware with low computational headroom and non-standardized protocols, and they articulate the trade-off between communication security and infrastructure amortization that characterizes any ICS cryptographic transition. The survey is comprehensive but does not propose a ZT framework, and its analysis stops short of integrating PQC into the architectural layer above.

CISA's 2024 Post-Quantum Considerations for Operational Technology [4] provides the most actionable government-level guidance, identifying specific deployment sites in OT-VPN concentrators, remote-access gateways, firmware-update signing, Secure Boot, encrypted industrial protocols such as OPC UA-and anticipating a significant and enduring challenge for owners and operators. Joseph et al. [43] in Nature (2022) address PQC transition at the organizational level, and the broader transition literature includes Mosca's framing of the harvest-now-decrypt-later adversary [44] and the foundational NIST report on post-quantum cryptography [45].

On the cloud side of the IT-OT boundary, our own prior work [62] developed a crypto-agility framework for multi-cloud environments that combines a cryptographic abstraction layer, a policy-driven orchestration engine, a centralized HSM control plane, and a Cryptographic Bill of Materials inventory. That framework targets the enterprise multi-cloud setting and assumes the latency-tolerant, resource-rich infrastructure typical of public cloud workloads; it does not address deterministic transport, process-aware trust evaluation, or the resource constraints of field-class OT hardware. The QZT-ICS framework presented here adopts the abstraction-layer and policy-driven orchestration concepts as a foundation but extends them with four substantive industrial-specific adaptations that are absent from the cloud-side framework and that, taken together, distinguish this paper from incremental reuse of prior concepts. First, the policy-decision substrate is shifted from an HSM-cluster trust anchor (cloud-native, network-reachable HSMs) to per-device hardware roots of trust anchored in TPM 2.0, NXP EdgeLock Secure Enclave, or ARM PSA: each ICS endpoint generates and stores its operational keys in tamper-resistant silicon at manufacturing time, with the policy engine treating the device's attestation report as a first-class input to the trust algorithm (Section 5.2.1, input P(t)). Second, the policy engine is augmented with the safety-aware fallback logic of Section 4.2: when the trust algorithm cannot positively authorize a session, the framework does not default to a blanket DENY (the appropriate cloud behavior) but rather to a pre-authorized safe-state policy that preserves availability for safety-critical operations - a direct consequence of the S-A-I-C priority inversion that distinguishes OT from IT. Third, the cryptographic primitives are not merely software libraries operating on commodity cloud CPUs but must execute at line rate on deterministic transport hardware; the framework places signature-verification and trust-score-lookup operations on FPGA-accelerated bumps-in-the-wire at the TSN bridge (Section 6.2), where the per-hop latency budget is tens of microseconds rather than the tens of milliseconds typical of cloud handshakes. Fourth, the CBOM-style cryptographic inventory is augmented with the TSN flow identifiers, 5G slice identifiers (S-NSSAI), and IEC 62443 zone-and-conduit tags that mark deterministic-transport context and safety classification - dimensions that are absent from the cloud-side framework because they have no analogue in the public-cloud workload model. The relationship to [62] is therefore one of substantive specialization rather than incremental application: the cloud-side framework establishes general crypto-agility primitives; QZT-ICS extends them with hardware-anchored device identity, safety-aware enforcement policy, line-rate cryptographic acceleration, and deterministic-transport context binding that are necessary for the ICS deployment regime.

Zhang and Zheng [42] in Entropy (2025) propose Quantum Secure Direct Communication enhanced TSN, integrating quantum communication primitives with deterministic networking and identifying digital twins of green-power and green-hydrogen systems as use cases. The work is the closest published analogue to our three-way integration of quantum-safe communication, deterministic transport, and DT-enabled industrial systems. However, QSDC is a quantum-key-distribution-class technology that requires specialized hardware not generally available in ICS deployments; it complements rather than replaces lattice-based PQC, and its applicability to brownfield ICS is limited. Our framework uses NIST-standardized lattice cryptography (ML-KEM, ML-DSA, SLH-DSA), which can be deployed in software on existing hardware where update paths exist.

Industrial vendor support is now materializing. Notably, NXP Semiconductors announced its i.MX 94 system-on-chip in November 2024 with integrated post-quantum cryptographic acceleration, time-sensitive networking, and edge AI capabilities-the first commercial industrial SoC combining all three. This convergence at the silicon level provides strong evidence that the hardware substrate for the QZT-ICS framework is materializing on commercial timelines. What remains missing from this body of work is any treatment of how PQC primitives interact with Zero Trust trust algorithms and with TSN timing budgets in an ICS setting. Section 7 of this paper provides the first quantitative analysis of these interactions.

3.5 Gap Analysis: The Absent Four-Pillar Synthesis

The pairwise analyses above suggest a consistent pattern. Each pairwise intersection (ZT+DT, ZT+5G/TSN, ZT+PQC, DT+5G/TSN) has at least two or three peer-reviewed contributions in the last five years. The three-way and four-way intersections, however, are essentially empty. Table 1 maps a representative cross-section of the surveyed works against the four pillars.

Work	ZT	DT	5G/TSN	PQC
Feng & Hu [17]	**Yes**	Partial (CPS)	,	,
Federici et al. [18]	**Yes**	,	,	,
Sims [19]	**Yes**	,	,	,
Lv et al. [20]	**Yes**	,	,	,

Work	ZT	DT	5G/TSN	PQC
Bilipelli [23]	**Yes**	,	,	,
Vega Vega et al. [24]	**Yes**	**Yes**	,	,
Sayghe [25]	,	**Yes**	,	,
Anumbe et al. [26]	**Yes**	**Yes**	,	,
Prasad et al. [27]	**Yes**	**Yes**	,	,
Sun et al. [30]	**Yes**	,	Yes (5G)	,
Chen et al. [31]	**Yes**	,	Yes (6G)	,
Wang et al. [32]	**Yes**	,	Yes (5G)	,
Bin Sediq et al. [33]	**Yes**	,	Yes (5G slice)	,
Jiang et al. [35]	,	,	Yes (TSN)	,
Bello et al. [36]	,	,	Yes (TSN)	,
Pokhrel [39]	**Yes**	,	Yes (5G)	,
Oliva del Moral et al. [41]	,	,	,	**Yes**
Zhang & Zheng [42]	,	Use case	Yes (TSN)	QSDC
Qadri [62] (cloud, prior)	,	,	,	Yes (cloud)
QZT-ICS (this paper)	**Yes**	**Yes**	Yes (5G+TSN)	Yes (lattice)

Table 1. Comparison of related work against the four QZT-ICS pillars. "Yes" indicates explicit treatment; "Partial" indicates partial coverage of a single sub-area; ", " indicates the pillar is not addressed. No prior work covers all four pillars.

No row in Table 1 prior to the last covers all four pillars. The closest from the OT side is Zhang and Zheng [42], which combines DT use cases, TSN, and quantum-safe communication, but uses Quantum Secure Direct Communication rather than NIST-standardized lattice PQC and does not articulate a ZT framework. The next closest from the OT side is Prasad et al. [27], which combines ZT and DT for IIoT but treats transport as a given and uses classical cryptography. From the IT/cloud side, our own prior work [62] establishes the crypto-agility primitives (abstraction layer, policy orchestration, HSM control plane, CBOM) at the multi-cloud layer, providing the cryptographic foundation that QZT-ICS specializes for the deterministic and safety-critical ICS setting. The QZT-ICS framework presented in the remainder of this paper occupies the previously unaddressed four-pillar intersection at the IT-OT boundary and, to our knowledge, is the first to do so.

4. Threat Model and Requirements

This section establishes the adversary model against which the QZT-ICS framework is designed, articulates the OT-specific value hierarchy that distinguishes industrial security from enterprise IT, and presents a structured catalog of functional, non-functional, and regulatory requirements. The requirements catalog (Tables 2-4) is referenced throughout the framework specification in Section 5 and the evaluation in Section 7.

4.1 Adversary Capabilities

We consider five adversary classes that collectively span the threat surface of IT-OT converged ICS environments. These classes are not mutually exclusive; in practice, sophisticated campaigns combine multiple capability profiles.

- **Adversary A1 (Nation-state actor)** represents adversaries with substantial resources, multi-year campaign timelines, and specific intent to disrupt or destroy industrial processes. The reference campaigns are Stuxnet (2010), the Ukrainian power grid attacks (2015-2016), Triton/TRISIS (2017), and Volt Typhoon (2023) [5], [47]. Capabilities include zero-day exploitation, custom ICS-aware malware, supply-chain compromise, and long dwell times. Initial access is typically achieved through spearphishing or strategic supplier compromise; pivoting from IT to OT exploits the converged pathways enumerated in Section 1.1. The adversary's goal is operational disruption, destruction of physical equipment, or geopolitical signaling. We map A1 to the ICS-specific MITRE ATT&CK tactics of Initial Access (TA0108), Persistence (TA0110), and Impact (TA0105).

- **Adversary A2 (Ransomware operator)** represents financially motivated adversaries who pivot from IT-borne ransomware to OT impact, exploiting the operational coupling between business systems and production. The reference campaign is Colonial Pipeline (2021), in which ransomware affecting IT billing systems caused a precautionary shutdown of OT pipelines. A2 adversaries are less ICS-sophisticated than A1 but operate at higher volume; their capability profile emphasizes commodity malware, broad reconnaissance, and rapid extortion. They typically lack ICS-specific payloads but exploit IT-OT coupling to maximize ransom pressure.

- **Adversary A3 (Malicious insider)** represents an authenticated user with legitimate but limited privileges who attempts to exceed authorization. Insiders may be operators, engineers, contractors, or vendors. The relevant capability is access to credentials and physical-process knowledge that external adversaries lack; the limitation is reduced opportunity to evade behavioral baselines once monitored. A3 is the canonical test case for process-aware trust scoring: a legitimate engineering workstation issuing syntactically valid but semantically unsafe commands cannot be detected by identity-based ZT alone.

- **Adversary A4 (Supply-chain attacker)** represents compromise of ICS components before delivery, through software updates, firmware injection, or hardware implantation. The reference incidents are SolarWinds (2020) and Kaseya (2021). A4 capabilities include trojanized firmware, signed-by-vendor malicious updates, and persistent backdoors at the silicon layer. The threat model assumes that a fraction of fielded devices may be compromised at delivery, which directly motivates the device-posture component of our trust algorithm and the hardware-rooted attestation requirement (FR6).

- **Adversary A5 (Harvest-now-decrypt-later adversary)** represents a passive eavesdropper recording encrypted ICS traffic with the intent to decrypt it once a cryptographically relevant quantum computer (CRQC) is available. A5 may be co-located with any of A1-A4. Mosca's argument [44] tells us that any system carrying information that must remain confidential beyond the CRQC arrival date is already under attack by A5 today. Because ICS data flows often include sensitive process parameters, control algorithms, and authentication credentials with multi-year sensitivity windows, A5 motivates the post-quantum component of the framework even in the absence of an immediate quantum threat.

We make the following assumptions about adversary capabilities. First, adversaries can read but not arbitrarily forge network traffic at the IT-OT boundary without detection; perimeter telemetry is assumed to flow to security operations. Second, adversaries can compromise individual endpoints with measurable probability but cannot simultaneously compromise all redundant components (the basis of model integrity attestation). Third, the symmetric primitives underlying the framework (AES, SHA-2/3) are assumed to remain secure against quantum adversaries at appropriate key sizes; the asymmetric primitives are the relevant attack surface for A5. Fourth, the Digital Twin can be compromised but not silently; tampering produces detectable telemetry anomalies, an assumption defended in Section 5.2.2.

4.2 Safety-Availability-Confidentiality-Integrity Priority Inversion in OT

A central design tension in any ZT framework for ICS arises from the priority ordering of security and safety properties. Enterprise IT security has historically followed the C-I-A triad in which confidentiality is paramount, integrity follows, and availability is the most readily traded property. Operational technology inverts this hierarchy. The canonical OT priority is *safety* > *availability* > *integrity* > *confidentiality* (S-A-I-C), reflecting the fact that ICS control physical processes whose unavailability can cause

economic loss, environmental damage, or human harm, and whose unauthorized disclosure of operating parameters is rarely the most consequential failure mode.

This inversion has direct architectural implications for ZT enforcement. The standard ZT principle of "deny by default" is sound in IT, where a failed authentication should result in a closed door and the user finding an alternative path. In OT, however, a denied access can be a safety event: a closed door at the wrong moment can prevent an operator from intervening in an emerging process upset, or can interrupt the cyclic control traffic on which physical safety depends. The CMU SEI guidance [11] explicitly states that implementing ZTA for ICS shouldn't impede uninterrupted access to systems to execute safety functions, a constraint that subordinates ZT to safety logic in any conflict.

We resolve this tension through three architectural choices that are reflected in the framework specification of Section 5. First, Policy Enforcement Points are configured with *safety-aware fail-safe defaults*: rather than blanket deny, a failed verification falls back to a pre-authorized "safe state" policy that permits read-only access to process telemetry and safety-critical commands while blocking configuration changes. Second, the trust algorithm is *biased toward availability* in cases of ambiguous trust scores: a borderline trust score within a tolerance band continues prior access, with simultaneous alerting to security operations, rather than terminating an in-flight control session. Third, an *out-of-band engineering bypass* permits authenticated physical-presence override of ZT enforcement at the local control station, with comprehensive forensic logging.

These choices represent deliberate departures from canonical ZT and should be acknowledged as such. The framework remains Zero Trust in the substantive sense, no implicit network-location trust, continuous verification, encrypted communication, dynamic policy, but the *enforcement posture* is reshaped by the safety-availability priority. We treat this reshaping as a feature, not a compromise.

4.3 Functional, Non-Functional, and Regulatory Requirements

The requirements catalog presented here is structured into three tables. Each requirement is given an identifier referenced throughout the framework specification (Section 5), the implementation description (Section 6), and the evaluation results (Section 7). The catalog is derived from three sources: (i) the NIST SP 800-207 ZT tenets and the logical components defined in the same publication [1]; (ii) the IEC 62443 foundational requirements (FR1-FR7) and Security Levels (SL 1-4) [9]; and (iii) the OT-specific adaptations articulated in Sections 4.1 and 4.2.

Table 2 lists the ten functional requirements describing what the framework must do. Each entry traces to a specific component of the architecture in Section 5.

ID	Functional Requirement	Rationale / Source	Section
FR1	Process-aware trust evaluation incorporating DT-derived process state	Identity-only ZT insufficient (§1.2); DT-based PDP	§5.2.1
FR2	Per-session continuous verification of subject, device, and context	NIST SP 800-207 tenet (iii) [1]	§5.5
FR3	Slice-aware micro-segmentation at the 5G transport layer	3GPP S-NSSAI alignment [33]	§5.3.1
FR4	TSN gate-control-list integration with ZT policy decisions	IEEE 802.1Qbv timing preservation [36]	§5.3.2
FR5	Crypto-agile key establishment with PQC and hybrid classical-PQC modes	NIST FIPS 203 [6]; Cyber Resilience Act [61]	§5.4
FR6	Hardware-rooted device posture attestation (Secure Boot, TPM/HSM)	Defense against A4 supply-chain compromise (§4.1)	§5.2.1

ID	Functional Requirement	Rationale / Source	Section
FR7	Digital Twin integrity attestation and tamper-evident model updates	Defense against DT model poisoning [28]	§5.2.2
FR8	Continuous verification feedback loop from field telemetry to DT	NIST SP 800-207 tenets (v), (vii) [1]	§5.5
FR9	Safety-aware policy fallback with pre-authorized safe-state defaults	OT safety-availability priority (§4.2)	§5.3
FR10	Tamper-evident audit logging of all policy decisions and exceptions	IEC 62443-3-3 SR 6.1 [9]; NIS2 Art. 23 [60]	§5.4.4

Table 2. Functional requirements of the QZT-ICS framework.

Table 3 lists the eight non-functional requirements with quantitative targets. These targets ground the empirical evaluation in Section 7 and are deliberately stringent to reflect the demanding deployment environment.

ID	Non-Functional Requirement	Target	Validation
NFR1	TSN scheduled-flow per-hop PEP latency overhead	< 100 µs	§7.3
NFR2	ML-KEM-768 encapsulation latency on Cortex-M33 PLC	< 50 ms	§7.2
NFR3	DT-based trust evaluation latency (p99)	< 10 ms	§7.4
NFR4	Attack detection F1 score (per attack class)	≥ 0.95	§7.1
NFR5	False-positive rate on benign traffic	≤ 2%	§7.1
NFR6	PEP forwarding throughput	≥ TSN line rate	§7.3
NFR7	Policy reconfiguration propagation time	< 5 ms	§7.3
NFR8	Additional memory footprint on PLC	≤ 10% available	§7.2

Table 3. Non-functional requirements with quantitative targets validated in Section 7.

Table 4 maps the framework against the regulatory obligations identified in Section 2.6. Each regulatory requirement is satisfied by one or more framework components, providing the basis for the compliance mapping in Section 8.3.

ID	Regulatory Requirement	Source	Framework
RR1	Alignment with IEC 62443 SL-3 foundational requirements (FR1-FR7)	IEC 62443-3-3 [9]	All §5

ID	Regulatory Requirement	Source	Framework
RR2	Coverage of NIST CSF 2.0 core functions (Govern, Identify, Protect, Detect, Respond, Recover)	NIST Cybersecurity Framework 2.0	§8.3
RR3	Incident-reporting capability for essential entities (24-hour timeline)	NIS2 Directive Art. 23 [60]	§5.5, §5.4.4
RR4	Crypto-agility for products with digital elements (2026-2030)	EU Cyber Resilience Act [61]	§5.4
RR5	Compliance with the seven NIST Zero Trust tenets	NIST SP 800-207 [1]	All §5
RR6	Alignment with CISA OT post-quantum guidance	CISA 2024 PQC for OT [4]	§5.4
RR7	Forensic auditability of access decisions and exceptions	IEC 62443-3-3 SR 5.1, SR 6.1 [9]	§5.4.4, §5.5

Table 4. Regulatory requirements and the framework components that satisfy them.

The non-functional thresholds in Table 3 are deliberately stringent. NFR1 (per-hop PEP overhead below 100 µs) reflects the fact that a TSN scheduled flow with a 1 ms end-to-end deadline traversing five hops can spend at most 20% of its budget on security processing without compromising determinism. NFR3 (DT-based trust evaluation latency below 10 ms at the 99th percentile) reflects the cadence at which SCADA-layer trust decisions must complete to avoid blocking operator workflows. NFR4 and NFR5 ($F1 \geq 0.95$ and false-positive rate $\leq 2\%$) reflect the operational reality that false positives in ICS create alarm fatigue and may themselves cause safety events through unnecessary interventions. We return to these thresholds in Section 7 to demonstrate that the proposed framework meets each within the constraints of the testbed.

5. The QZT-ICS Framework

This section specifies the QZT-ICS framework. Section 5.1 presents the architectural overview and the reference layering onto the Purdue Model and RAMI 4.0. Section 5.2 details the Digital Twin as Policy Decision Point, including the trust algorithm (Algorithm 1) and the DT integrity attestation mechanism. Section 5.3 specifies Policy Enforcement Point placement at the 5G slice and TSN bridge layers. Sections 5.4 and 5.5 (later parts) cover the crypto-agile PQC layer and the continuous verification loop.

5.1 Architectural Overview and Reference Layering

QZT-ICS is structured as a four-pillar reference architecture (Fig. 1) layered over the converged IT-OT stack typical of contemporary industrial deployments. The architecture is deliberately framework-neutral with respect to vendor implementations and protocol-neutral with respect to industrial fieldbus choices: each pillar specifies an interface and a set of obligations, leaving room for diverse instantiations across sector and asset generation.

The four pillars are as follows. The Zero Trust Control Plane comprises the Policy Engine (PE), Policy Administrator (PA), and trust-algorithm components defined in NIST SP 800-207 [1], extended with industrial-specific policy mappings to IEC 62443 [9] and NIST CSF 2.0. The Digital Twin serves as the Policy Decision Point of the control plane, hosting the trust algorithm and the process-state context that makes process-aware trust evaluation possible (Section 5.2). The 5G/TSN Transport Layer carries authenticated traffic between OT and IT zones, with Policy Enforcement Points (PEPs) embedded at the 5G slice and TSN bridge layers (Section 5.3). Finally, the PQC Crypto-Agile Substrate underlies the entire stack, providing the cryptographic primitives used by every pillar (Section 5.4).

Mapping these pillars onto the Purdue Enterprise Reference Architecture (PERA) makes the deployment locality concrete. PERA Levels 4 and 5 (enterprise IT, MES/ERP, engineering workstations, identity providers) host the ZT control-plane components, the Digital Twin compute substrate, and the threat-intelligence integrations. Level 3.5 (the DMZ) hosts the edge-gateway PEP and the integration points between IT and OT. Levels 0 through 3 (the OT zone, including SCADA servers, HMIs, PLCs, RTUs, sensors, and actuators) host the field-side PEPs and the cryptographic endpoints. The cross-cutting PQC substrate spans all levels because cryptographic primitives are needed wherever authentication, key establishment, or signed firmware is consumed. The RAMI 4.0 reference architecture overlays this Purdue mapping with its three axes (architecture, lifecycle, and hierarchy); the QZT-ICS components occupy the Integration and Communication layers across the RAMI hierarchy.

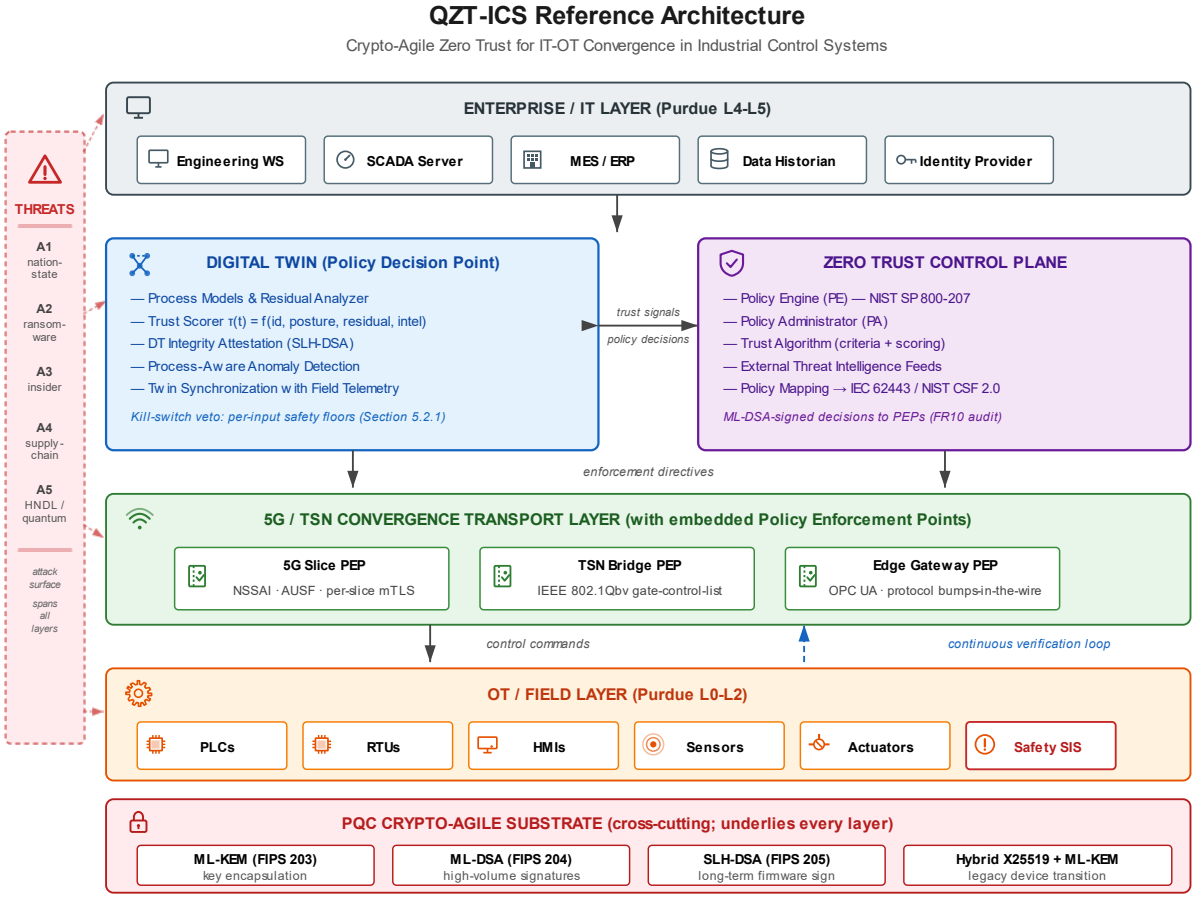

Figure 1. The QZT-ICS reference architecture integrates four pillars - a Zero Trust control plane, a Digital Twin Policy Decision Point, 5G/TSN deterministic transport with embedded PEPs, and a cross-cutting post-quantum crypto-agile substrate.

A representative end-to-end data flow illustrates how the four pillars interact. Consider an engineer at Level 5 issuing a setpoint change to a Level 1 PLC via the SCADA server at Level 2. (1) The engineer authenticates to the identity provider; the resulting assertion is signed using ML-DSA. (2) The SCADA server forwards the setpoint change toward the Edge Gateway PEP at Level 3.5, where the operation triggers a trust evaluation against the current trust score for the engineer's session. (3) The Digital Twin computes the trust score using identity, device posture, the current process-state residual, and behavioral baselines (Section 5.2.1); a sufficiently high score yields a permit decision from the Policy Engine. (4) The Policy Administrator emits the decision to the 5G slice PEP at the User Plane Function and to the downstream TSN bridge PEP, which install corresponding flow-admission entries that authorize the setpoint traffic across the deterministic transport (Section 5.3). (5) The setpoint reaches the PLC, which verifies the ML-DSA signature on the command against the Policy Administrator's public key before applying the change. (6) Telemetry from the PLC and downstream sensors flows back to the Digital Twin via the continuous verification loop, where any process-state divergence from the model's prediction updates the trust score for subsequent operations (Section 5.5).

This walkthrough motivates the more detailed component specifications in the remainder of Section 5. We address each pillar in turn, beginning with the Digital Twin as Policy Decision Point.

5.2 Digital Twin as the Policy Decision Point (PDP)

The Policy Decision Point in canonical Zero Trust receives access requests, evaluates them against policy, and emits permit or deny verdicts. In QZT-ICS, the PDP is co-located with the Digital Twin rather than implemented as a standalone microservice. This co-location is deliberate: the DT is the only architectural component that simultaneously holds a current identity model (which subjects have authenticated where), a current device-posture model (the state of attestation evidence for each endpoint), and a current process-state model (what the plant is actually doing). Identity-only and posture-only PDPs cannot reason about whether a syntactically valid command is semantically safe given the current process state. The DT-PDP can, and this capability is what makes the framework process-aware in the sense required by Section 1.2.

5.2.1 Trust Scoring Inputs

The trust algorithm consumes five input categories and produces a continuous trust score $\tau(t) \in [0, 1]$ for each (subject, resource) pair at time t. The score is computed as a weighted combination:

$$\tau(t) = w_I \cdot I(t) + w_P \cdot P(t) + w_R \cdot R(t) + w_B \cdot B(t) + w_T \cdot T(t)$$

where the weights w_I, w_P, w_R, w_B, w_T sum to one and are policy-specific. Each component is normalized to the unit interval. $I(t)$ is the identity-confidence input, derived from multi-factor authentication strength, identity-provider reputation, and credential recency; a fresh hardware-attested MFA assertion yields $I(t) \approx 1$, while a stale password-only authentication yields a low value. $P(t)$ is the device-posture input, computed from attestation freshness, firmware version against current advisories, secure-boot status, and patch-compliance state; this input is sensitive to A4 supply-chain compromise (Section 4.1). $R(t)$ is the process-state residual confidence, computed as 1 minus the normalized prediction error between the live plant telemetry and the Digital Twin's prediction; high residuals indicate that the plant is operating outside its modeled envelope and degrade the trust score even when identity and posture are perfect. $B(t)$ is the behavioral-conformance input, scored by an ML-based model of the subject's historical access patterns. $T(t)$ is the threat-intelligence input, equal to one minus the probability that the subject or device appears on a current indicator-of-compromise feed.

Each input is subject to temporal decay. An identity confidence based on a successful MFA event two minutes ago is treated as nearly full strength; the same event two hours ago is treated as partial; eight hours later it is treated as expired and the subject must re-authenticate. The decay half-life is input-specific: identity confidence decays with a half-life of approximately fifteen minutes (calibrated to typical operator shift patterns); device posture decays with a half-life equal to the attestation update interval (typically one to four hours); process-state residual is treated as instantaneous (no decay) because it reflects the current plant state by construction; behavioral baseline decays slowly (half-life of one week) to accommodate legitimate operational variability.

Algorithm 1 formalizes the trust evaluation. The algorithm produces both a numeric trust score and a three-way decision: PERMIT for high confidence, MONITORED_PERMIT for the soft band that the safety-availability priority of Section 4.2 requires, and DENY for low confidence (which itself falls back to the safety-aware default policy rather than a blanket block). Crucially, the algorithm includes a per-input veto layer (lines 7-9 below) that overrides the weighted score if any individual input falls below its per-input safety floor, addressing the well-known limitation that purely linear aggregations can be defeated when an adversary maximizes compensating inputs.

Algorithm 1: Trust Score Computation and Decision

Input: session S = (subject, device, resource, history),

time t, policy Π with thresholds (θ_low, θ_high),

weights (w_I, w_P, w_R, w_B, w_T),

per-input safety floors (φ_I, φ_P, φ_R, φ_B, φ_T)

Output: trust score τ ∈ [0, 1], decision d ∈ {PERMIT, MONITORED_PERMIT, DENY}

1. I ← IdentityConfidence(S.subject, t) // MFA, IdP, recency

```
2.  Post ← DevicePosture(S.device, t)                // attestation, patch
3.  R    ← 1 - ProcessResidual(S.resource, DT.predict(t))  // 1 - norm. error
4.  B    ← BehavioralConformance(S.subject, S.history) // ML baseline
5.  T    ← 1 - ThreatIntel(S.subject, S.device, t)     // 1 if not on IoC list
6.  τ    ← w_I·I + w_P·Post + w_R·R + w_B·B + w_T·T

   // Kill-switch: any single input below its floor forces DENY
7.  if (I < φ_I) ∨ (Post < φ_P) ∨ (R < φ_R) ∨ (B < φ_B) ∨ (T < φ_T) then
8.      d ← DENY → SafetyFallback(S); FlagVeto(input-name)
9.      goto 17                                      // bypass linear band
10. end if

   // Linear-aggregation decision band (only entered if no veto fired)
11. if τ ≥ θ_high then
12.     d ← PERMIT
13. else if τ ≥ θ_low then
14.     d ← MONITORED_PERMIT          // log + permit; alert SOC
15. else
16.     d ← DENY → SafetyFallback(S)  // see §4.2
17. end if

18. EmitAuditEvent(S, τ, d, t)        // tamper-evident log (FR10)
19. return (τ, d)
```

Typical threshold settings are θ_high = 0.75 and θ_low = 0.45 for routine operations; safety-critical control flows raise both thresholds (0.85 and 0.60 respectively); read-only telemetry flows lower them (0.55 and 0.30). Weight tuning is policy-specific but typical defaults are w_I = 0.25, w_P = 0.20, w_R = 0.25, w_B = 0.15, w_T = 0.15. The process-residual weight w_R deliberately receives equal billing with the identity weight w_I, reflecting the framework's process-aware orientation. Per-input safety floors are conservatively set at φ_I = φ_P = φ_R = 0.20 (the three inputs most directly tied to ground truth) and φ_B = φ_T = 0.10 (the two inputs more susceptible to legitimate noise); safety-critical flows raise all five floors by 0.10.

Justification of the linear-with-veto aggregator. The choice of a linear weighted sum for the trust score is deliberate and admits a defensible argument under three considerations that we acknowledge are well-known to the control-theory and cyber-physical-systems community. First, within the safe operating envelope - the region where all five inputs are above their per-input safety floors - a linear aggregator is interpretable, auditable, and amenable to formal policy review: operators can reason about which input contributed how much to a decision, which is essential for IEC 62443 compliance and post-incident forensics. Non-linear aggregators (e.g., neural-network-based or attention-weighted) provide modestly better empirical separation in benign workloads but at the cost of explainability, certifiability, and stable behavior under adversarial perturbation. Second, outside the safe operating envelope - the region where at least one input has fallen below its floor - the linear aggregator is explicitly bypassed by the kill-switch logic of lines 7-9, forcing DENY regardless of how favorable the remaining inputs appear. This

addresses the standard critique against linear trust scores: an adversary cannot mask a clearly anomalous individual input by maximizing the others. The veto operates as a safety envelope in the control-theoretic sense, with the linear aggregator handling only the interior region where the inputs are mutually consistent and bounded. Third, the kill-switch's per-input floors are themselves auditable: each veto event records the offending input, the value, and the floor, providing a clear forensic trail. The combined linear-with-veto structure is a deliberate trade-off between explainability (linear, interior) and adversarial robustness (veto, boundary).

This design admits a quantitative bound on adversarial gaming. Consider an adversary who has compromised one input - for example, a compromised device-posture report driving Post toward 0. Without the veto, the adversary could compensate by maximizing the remaining inputs (I, R, B, T) at 1.0, yielding $\tau = (0.25)(1) + (0.20)(0) + (0.25)(1) + (0.15)(1) + (0.15)(1) = 0.80$, crossing θ_{high} and gaining PERMIT despite the posture compromise. With the veto active at $\varphi_P = 0.20$, any Post < 0.20 forces DENY regardless of τ. The adversary must therefore maintain every individual input above its floor, not merely the weighted sum above its threshold. This raises the adversary's effort and constrains the achievable attack space in a way the linear-only formulation does not. The corresponding empirical result is the detection-performance improvement observed against baseline B1 (identity-only ZT, no veto, no process input) in Section 7.5, which is unable to detect the A1 staged-command-injection and A3 insider attacks that the QZT-ICS algorithm catches via the residual and posture vetoes respectively.

1) 5.2.2 Digital Twin Integrity Attestation

Placing the PDP inside the Digital Twin makes the DT a high-value attack target. An adversary who can poison the DT model, tamper with its prediction outputs, or replay stale telemetry can drive the trust score in either direction at will. This sub-section specifies the three mechanisms QZT-ICS uses to defend the DT itself.

First, *build-provenance attestation*. The DT model and policy artifacts are built from versioned source through a reproducible-build pipeline. Each build emits a signed manifest comprising the source commit hash, build environment hash, dependency hashes, and the resulting artifact hash. Signatures use SLH-DSA (FIPS 205 [8]) to provide stateless quantum-resistant signing with a long verification lifetime appropriate for firmware-class artifacts. The Policy Engine refuses to load any DT model whose manifest signature does not verify against the curator's public key. This mechanism defends against malicious updates injected into the model supply chain (an A4 supply-chain attack).

Second, *runtime integrity through trusted execution*. The DT runtime executes in a Trusted Execution Environment (TEE) - Intel SGX, ARM TrustZone, or AMD SEV-SNP depending on platform - that provides hardware-rooted memory encryption and remote attestation. PEPs that consume DT verdicts request and verify an attestation report before each session class change; a DT instance that cannot produce a valid attestation is excluded from trust evaluations and its verdicts are ignored. This mechanism defends against runtime compromise of the DT host (an A1 nation-state attack pattern).

Third, *federated DT consensus for safety-critical decisions*. For trust evaluations that gate safety-critical commands (defined as operations on resources tagged with IEC 62443 SL ≥ 3), QZT-ICS requires concurrence from at least two DT instances running on diverse infrastructure (different physical hosts, different TEE families, different software stacks where feasible). The two instances independently compute the trust score; the Policy Administrator uses the minimum (the more conservative verdict) to drive enforcement. A discrepancy beyond a configured tolerance triggers an out-of-band alert to security operations and engages the safety-aware fallback (Section 4.2). This mechanism defends against single-instance compromise even when build provenance and runtime attestation are bypassed.

Together, these three mechanisms instantiate the requirement FR7 from Section 4.3. They also directly address the model-poisoning attack class formalized by Zhang et al. [28], which demonstrated that distributed digital twin systems are vulnerable to coordinated update poisoning unless aggregation logic is hardened. We return to model poisoning as an open research challenge in Section 9.2, where we note that the defenses specified here are necessary but, in the presence of a sufficiently capable A1 adversary with extended dwell time, may not be sufficient.

5.3 Policy Enforcement Points at the 5G Slice and TSN Bridge Layers

Policy Enforcement Points (PEPs) execute the verdicts emitted by the Policy Administrator. In conventional ZT, PEPs are deployed at application gateways, identity-aware proxies, or on the endpoint itself. None of these placements work for industrial control traffic for two reasons: first, ICS endpoints are typically resource-constrained and cannot run a sidecar-class PEP; second, application-gateway placement adds round-trip latency that is incompatible with TSN scheduled-flow timing budgets (Section 2.4). QZT-ICS therefore embeds PEPs at the transport layer itself - at the 5G slice and TSN bridge - where enforcement can be applied at line rate without interposing additional hops.

5.3.1 Slice-Aware Micro-Segmentation

3GPP 5G non-public networks provide native micro-segmentation through network slicing. A slice is identified by a Single Network Slice Selection Assistance Information (S-NSSAI) value comprising a Slice/Service Type (SST) and an optional Slice Differentiator (SD). The SST identifies the slice class - SST=1 enhanced Mobile Broadband (eMBB), SST=2 Ultra-Reliable Low-Latency Communications (URLLC), SST=3 massive Machine-Type Communications (mMTC) - while the SD permits tenant-specific or process-area-specific differentiation. Each slice carries isolated traffic, isolated authentication context, and, in the QZT-ICS extension, isolated policy enforcement.

QZT-ICS defines a canonical set of ICS slice profiles. The IT slice (SST=1) carries enterprise traffic between Level 4-5 systems and the converged backbone: engineering-workstation file transfers, MES/ERP integration, and operator-facing video conferencing. The OT control slice (SST=2 with custom SD) carries low-latency control traffic between SCADA, HMIs, and field-level PLCs/RTUs, typically with a 99.999% reliability target and a 1-10 ms latency budget. The OT telemetry slice (SST=3) carries high-volume sensor and condition-monitoring data with relaxed latency but high aggregate bandwidth. Critical-process and non-critical-process areas may receive distinct SD values, enabling per-process-area isolation within the OT control slice.

Authentication and key establishment per slice are specified by 3GPP TS 33.501. The Authentication Server Function (AUSF) executes the primary authentication procedure (5G-AKA or EAP-AKA' depending on UE type), and QZT-ICS integrates post-quantum primitives at three points within this procedure: (i) the SUPI-to-SUCI subscription concealment uses ML-KEM rather than the classical ECIES variant; (ii) the SBA inter-NF mTLS uses ML-DSA-signed certificates rather than RSA or ECDSA; and (iii) the slice-specific authentication and authorization (NSSAA) procedure for each S-NSSAI uses ML-KEM-derived session keys for the slice-specific signaling channel. The PQC integration is hybrid for the transition period: classical and PQC key encapsulations are combined using a hybrid KDF, so that the session key remains secure as long as either component is unbroken.

The slice PEP itself sits at the User Plane Function (UPF). For each user-plane packet, the UPF consults the per-flow trust score, which is cached locally and refreshed by the Policy Administrator via the N4 reference point. Three enforcement actions are available: PERMIT (forward the packet on its slice), MONITORED_PERMIT (forward but mark for SOC inspection through duplicated telemetry to the security analytics platform), and DENY (drop the packet and emit a security event). The trust-score cache uses a hash table indexed by the (UE-ID, S-NSSAI, PDU-session-ID) triple and supports updates at sub-millisecond cadence. A trust-score change does not interrupt an in-flight packet; the next packet in the flow encounters the new score. Cross-slice isolation is enforced by default: a UE authenticated for one S-NSSAI cannot send traffic on another without explicit re-authorization through the AUSF, which itself reapplies the trust algorithm.

5.3.2 TSN Gate-Control-List Integration with ZT Policy

On the wired side of the converged transport, IEEE 802.1 Time-Sensitive Networking provides the deterministic backbone. The IEEE 802.1Qbv standard defines a time-aware shaper that gates each output queue of a TSN bridge according to a Gate Control List (GCL): a sequence of entries (start_time, gate_state, duration) that cyclically open and close the eight per-port queues. The GCL is computed offline (or by a Centralized Network Configuration entity at runtime) so that the resulting schedule satisfies all flow timing constraints simultaneously. The deterministic timing guarantee of TSN follows directly from the determinism of the GCL.

QZT-ICS extends each GCL entry with a security predicate. The conventional triple (start_time, gate_state, duration) becomes the quadruple (start_time, gate_state, duration, min_trust_score). A flow is admitted into its gate window only if the trust score for its associated session satisfies the threshold; otherwise the packets in the window are dropped and a security event is emitted. The extension is implementable within the existing TSN switching pipeline because:

- The trust-score lookup is a hash-table probe keyed on the flow identifier (VLAN ID + 5-tuple hash) with a worst-case lookup latency of approximately 10 ns. In commercial TSN bridge silicon (which is not P4-programmable today), the lookup is typically implemented as a TCAM- or SRAM-backed table in the bridge's existing classifier; in the testbed deployment of Section 6.2 it is realized on an adjacent P4-on-FPGA accelerator (AMD Alveo U280) configured as a bump-in-the-wire. Either implementation path is negligible relative to the 100 µs NFR1 budget.
- The trust-score cache is updated out of band by the Policy Administrator via NETCONF/YANG or RESTCONF, with the IEEE 802.1Qci stream-identification function providing the binding between flows and ZT sessions.
- The gate schedule itself is unchanged. The QZT-ICS extension does not alter the cyclic timing or the existing time-aware shaping behavior; it only augments the per-gate admission predicate. This preserves the deterministic timing guarantee, which is the entire reason TSN was adopted in the first place.

Three failure modes deserve specification. First, when the trust-score cache contains no entry for a flow (cold start, or expired entry), the bridge falls back to the safety-aware default specified in Section 4.2: safety-critical flows are admitted on the basis of

their flow-identifier tag (set by the originating PLC at provisioning time and protected by ML-DSA-signed configuration), while non-safety-critical flows are dropped pending PA refresh. Second, when the Policy Administrator is unreachable, the cache time-to-live governs how long old scores remain valid. The TTL is OT-specific - typically one to five minutes, compared with seconds for IT - and is chosen to give the PA enough time to recover without producing safety-relevant denials. Third, when a trust score crosses a threshold mid-flow, the change takes effect at the next gate cycle; in-flight frames in the current cycle are admitted under the prior decision. This avoids the timing perturbation that would result from mid-cycle revocation and is acceptable because gate cycles are on the millisecond timescale.

The combination of the slice PEP at the 5G UPF and the bridge PEP at the TSN switch instantiates requirements FR3, FR4, FR9, and (jointly with the PQC layer) NFR1, NFR6, and NFR7 from Section 4.3. The evaluation in Section 7.3 quantifies the per-hop latency overhead empirically and demonstrates that the framework remains within the 100 µs NFR1 budget for the testbed workloads.

5.4 Crypto-Agile Post-Quantum Cryptography Layer

The fourth pillar of QZT-ICS is the crypto-agile post-quantum cryptography substrate. It instantiates the cryptographic primitives on which all upper-layer protocols depend, with three deliberate design goals: (i) immediate adoption of NIST-standardized post-quantum primitives for new deployments; (ii) hybrid classical-PQC modes for the transition period and for legacy fieldbus deployments; and (iii) crypto-agility, so that primitives can be replaced as algorithms are deprecated or strengthened, without requiring hardware changes. The pillar's structural pattern - an algorithm-agnostic cryptographic abstraction layer, a policy-driven orchestration mechanism for algorithm selection, a centralized trust anchor for key lifecycle management, and a tracked cryptographic inventory - follows the multi-cloud crypto-agility framework our prior work developed for enterprise cloud workloads [62], with adaptations for the deterministic-timing and resource-constrained requirements of ICS that we identify in each sub-section below. The sub-sections that follow specify each role: ML-KEM for key establishment (5.4.1), ML-DSA and SLH-DSA for signatures with different lifetime requirements (5.4.2), hybrid modes for the transition (5.4.3), and key lifecycle management at industrial scale (5.4.4). Table 5 summarizes the primitive selection and parameter sets.

Primitive	Standard	Parameter Set	PK Size	Sig/CT Size	QZT-ICS Role
ML-KEM	FIPS 203	ML-KEM-768 (Cat. 3)	1184 B	1088 B (ciphertext)	TLS 1.3, OPC UA, 5G-AKA, IPsec session key establishment
ML-DSA	FIPS 204	ML-DSA-65 (Cat. 3)	1952 B	3309 B (signature)	Telemetry, control commands, policy decisions, short-term tokens
SLH-DSA	FIPS 205	SLH-DSA-128s (Cat. 1)	32 B	7856 B (signature)	Firmware signing, build manifests, DT integrity attestation, root of trust
Hybrid KEM	RFC 9370 (IKEv2) + IETF drafts	X25519 + ML-KEM-768	1216 B	1120 B (ciphertext)	Transition period and legacy interop (Section 5.4.3)

Table 5. Post-quantum cryptographic primitives selected for QZT-ICS, with parameter sets, artifact sizes, and the framework roles each fulfills.

5.4.1 ML-KEM for Key Establishment

ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism), standardized in FIPS 203 [6], replaces classical Diffie-Hellman and ECDH variants for session-key establishment throughout the framework. FIPS 203 defines three parameter sets - ML-KEM-512, ML-KEM-768, and ML-KEM-1024 - providing Category 1, Category 3, and Category 5 post-quantum security respectively. QZT-ICS uses ML-KEM-768 as the default, matching the security strength of AES-192 and providing approximately 128 bits of post-quantum security with manageable artifact sizes.

ML-KEM is consumed at five points within the framework: the 5G UE-AUSF authentication concealment (SUPI-to-SUCI), the 5G Service-Based Architecture inter-NF mTLS handshake, the TLS 1.3 secure channels between SCADA, HMI, and the Digital Twin,

the OPC UA secure channel between control and field tiers, and the IPsec IKEv2 negotiation for legacy site-to-site tunnels. The same primitive serves all five roles because the framework standardizes the key-encapsulation interface; the upper-layer protocols are oblivious to the underlying primitive choice beyond the algorithm identifier (Section 5.4.4).

ML-KEM-768 produces public keys of 1184 bytes, ciphertexts of 1088 bytes, and shared secrets of 32 bytes. The artifact sizes are roughly 30 to 35 times larger than X25519 (32-byte public keys, 32-byte ciphertexts). This size expansion has three implications for ICS deployment that deserve specific attention. First, OPC UA secure channel handshakes are bounded by the secureChannelTimeout parameter and may require buffer increases at both endpoints; we validate that the OPC UA reference stack (open62541) handles ML-KEM artifacts cleanly when the buffer pool accommodates messages of at least 8 kB. Second, MQTT-SN frames over constrained fieldbuses cannot fit ML-KEM artifacts in a single frame and require fragmentation; the recommended treatment is to default to a hybrid mode (Section 5.4.3) on such links. Third, for IPsec IKEv2, RFC 9370 establishes the multi-key-exchange framework that permits multiple Diffie-Hellman-style key exchanges to be combined into a single IKEv2 SA; the specific binding of FIPS 203 ML-KEM into this framework is specified in the IPSECME working group document draft-ietf-ipsecme-ikev2-pqc-auth and the related ML-KEM-in-IKEv2 draft (under active standardization at the time of this writing). The framework's implementation uses these drafts via the strongSwan 6.0 plugin model; primitive bindings track the active standard state as of mid-2026, and the implementation will transparently follow any draft that advances to RFC status during the period of this paper's review and publication.

Performance is the second concern. On a Cortex-A55 application processor representative of modern PLCs (the same processor family used in the NXP i.MX 94), ML-KEM-768 encapsulation completes in approximately 200 µs and decapsulation in approximately 250 µs on a single core - well within the budget for session-establishment operations. On a Cortex-M33 microcontroller representative of constrained RTUs, the same operations require approximately 20-40 ms depending on memory access patterns and compiler optimization. This is two orders of magnitude faster than the IT session-establishment latency it replaces (often ~100 ms with classical primitives over high-latency WAN paths) but not negligible for cycle-time-constrained control loops. Session caching and the key-lifecycle policies of Section 5.4.4 reduce the amortized cost to a once-per-session expenditure.

We choose ML-KEM rather than the other NIST PQC KEM candidates on three grounds. First, ML-KEM was the first NIST-standardized lattice KEM and has the most mature implementation ecosystem, including open-source libraries (liboqs, AWS-LC, BoringSSL) and validated hardware implementations on multiple platforms. Second, Classic McEliece has megabyte-class public keys that are operationally impractical for the OPC UA and TLS handshake sizes typical in ICS. Third, ML-KEM has well-understood side-channel countermeasures, including the protected reference implementations published in the NIST evaluation process. Classic McEliece remains a fallback option for the long-term archival use case described in Section 5.4.4.

5.4.2 ML-DSA and SLH-DSA for Signatures

ML-DSA (Module-Lattice-Based Digital Signature Algorithm, FIPS 204 [7]) and SLH-DSA (Stateless Hash-Based Digital Signature Algorithm, FIPS 205 [8]) instantiate two different signature roles. The roles differ in performance characteristics, signature-lifetime requirements, and exposure to algorithmic compromise; the framework uses both rather than one as a deliberate diversification choice.

ML-DSA is used for high-volume signing of telemetry, control commands, policy decisions, and short-term authentication tokens. FIPS 204 defines ML-DSA-44, ML-DSA-65, and ML-DSA-87 at Category 2, 3, and 5 security respectively. QZT-ICS uses ML-DSA-65, matching the ML-KEM-768 default. ML-DSA-65 produces public keys of 1952 bytes and signatures of 3309 bytes. Signing on Cortex-A55 completes in approximately 500 µs; verification in approximately 200 µs. This throughput is sufficient for typical industrial signing rates of up to several thousand commands per second per PLC, comfortably exceeding the rates required by NFR1 and NFR6.

SLH-DSA serves a different role: long-term signing of firmware artifacts, build manifests, and the DT integrity attestations of Section 5.2.2. SLH-DSA differs from ML-DSA in that its security rests on the assumed one-wayness of an underlying hash function (SHA-2 or SHAKE) rather than on the hardness of structured lattice problems. This diversification matters because, although both algorithms are believed to resist known quantum attacks, lattice and hash assumptions fail under disjoint threat models; an unexpected cryptanalytic advance against lattices does not affect hash-based signatures. For firmware-class artifacts that may be in service for fifteen to twenty years - well beyond the typical reassessment horizon of any single primitive - this diversification is the appropriate insurance policy.

The cost of SLH-DSA is asymmetric. The small-signature variant SLH-DSA-128s produces 32-byte public keys and 7856-byte signatures; signing takes tens to hundreds of milliseconds depending on the implementation, while verification completes in well under a millisecond. The asymmetry is appropriate for firmware signing because firmware is signed rarely (typically once per release) but verified frequently (once per device boot, plus periodic integrity checks), and the large signature is amortized over

the much larger firmware artifact it accompanies. QZT-ICS uses SLH-DSA-128s for the deployment; the fast variant SLH-DSA-128f produces larger signatures (17088 bytes) and offers no operational benefit for firmware signing.

The split is deliberate: ML-DSA for high-volume short-lived signing, SLH-DSA for low-volume long-lived signing. This gives the framework defense-in-depth against any single algorithmic compromise without doubling the cryptographic footprint of every operation. A note on stateful alternatives: NIST also describes the stateful schemes LMS and XMSS (in NIST SP 800-208), which produce smaller signatures than SLH-DSA but require careful state management to avoid catastrophic key reuse. The state management requirement is particularly difficult in ICS environments where devices may lose power unexpectedly. QZT-ICS therefore uses SLH-DSA exclusively rather than mixing in stateful variants, accepting the larger signature size in exchange for operational simplicity.

5.4.3 Hybrid Classical+PQC Modes for Legacy Devices

Production ICS environments cannot transition to pure PQC overnight. Hundreds of millions of fielded devices use classical primitives (ECDHE, ECDSA, RSA, AES); replacing all of them simultaneously is operationally infeasible, and a single non-upgradable component can hold up an entire authentication chain. The framework therefore specifies hybrid modes that combine classical and PQC primitives, providing security as long as either component remains unbroken. Hybrid modes are explicitly endorsed in NIST IR 8413 transition guidance and in the European Cyber Resilience Act compliance interpretations [61].

For key establishment, QZT-ICS uses a hybrid KEM construction along the lines specified in the IETF TLS hybrid key-exchange document (draft-ietf-tls-hybrid-design and the subsequent draft-ietf-tls-ecdhe-mlkem for ML-KEM specifically) and, for IKEv2, via the multi-key-exchange framework of RFC 9370 combined with the in-progress IPSECME ML-KEM-IKEv2 specifications. Primitive bindings track the active standard state as of mid-2026, and the framework's modular implementation is designed to absorb draft-to-RFC promotion without architectural changes. The hybrid public key is the concatenation of a classical key (X25519) and an ML-KEM-768 public key. The hybrid ciphertext is the concatenation of the corresponding classical and PQC ciphertexts. The shared secret is derived through HKDF-SHA256 applied to the concatenation of the two component secrets. The construction is secure if *either* the discrete-log problem on the classical curve is hard *or* the underlying lattice problem is hard; an adversary must break *both* to recover the session key. This property is essential for the transition period because confidence in PQC primitives is still developing, and a critical-infrastructure operator should not stake security on a primitive that may yet be shown to have an unexpected vulnerability.

For signatures, hybrid composition is more delicate because the natural concatenation of two signatures roughly doubles the signature size. Three variants are specified. (i) For protocols with relaxed size constraints (TLS 1.3 server certificates, OPC UA server certificates), hybrid signatures concatenate ML-DSA-65 and Ed25519 components, accepting the ~3.4 kB combined signature size. (ii) For protocols with strict size constraints (Modbus secure variants, narrowband fieldbuses), the framework uses ML-DSA alone with explicit documentation of the lattice-only failure mode in the deployment risk register. (iii) For firmware signing, the framework uses SLH-DSA alone rather than hybrid, accepting the asymmetric signing cost in exchange for the long-term security argument of Section 5.4.2.

A second class of hybrid concerns the *bumps-in-the-wire* approach for legacy fieldbuses. Modbus, DNP3, and large parts of PROFIBUS carry no native authentication. The framework recommends inline cryptographic adapters at the IT-OT boundary that wrap legacy frames inside a PQC-authenticated envelope. The adapter terminates the modern cryptography and emits the legacy frame on the field side, preserving wire-protocol compatibility with installed bases. This is a pragmatic compromise: it does not protect the leg from the adapter to the field device (which remains plaintext), but it does protect all longer-distance and cross-zone segments where most realistic attacks would be staged. Coverage is increased over time by replacing devices with natively PQC-capable equivalents - the NXP i.MX 94 announcement is one signal that such replacement hardware is reaching the market on commercial timelines.

5.4.4 Key Lifecycle Management at Industrial Scale

A cryptographic substrate is only as good as the operational discipline around key provisioning, rotation, and revocation. QZT-ICS specifies four lifecycle elements that together make the PQC layer practical at industrial scale.

Provisioning. Each ICS endpoint receives a long-term identity key pair at manufacturing time, generated by the device itself within its hardware root of trust (TPM 2.0, NXP EdgeLock Secure Enclave, ARM PSA, or equivalent) and certified by the manufacturer through a signed birth certificate. The birth certificate uses SLH-DSA so that its verification lifetime extends beyond the device service lifetime; the trust anchor is the manufacturer's offline root key, also SLH-DSA. Birth-certificate verification at first deployment establishes the device's initial identity and posture (FR6), feeding directly into the device-posture component of

the trust algorithm (Section 5.2.1). Field commissioning then provisions operational ML-KEM and ML-DSA key pairs derived from the birth-certificate keys through a standardized issuance protocol such as EST-coaps extended for PQC.

Rotation. Operational keys rotate on a schedule defined by the asset criticality. Engineering-workstation session keys rotate every 24 hours. PLC operational signing keys rotate weekly. SCADA-to-DT mTLS session keys rotate every 12 hours. Long-term identity keys are not rotated; their security depends on the underlying primitive remaining secure, with crypto-agility (below) providing the migration path. Key rotation is automated through the Policy Administrator and integrated with the trust algorithm: a successful rotation event refreshes the device-posture component *P(t)*, while a failed rotation degrades it, triggering a MONITORED_PERMIT or DENY verdict at subsequent access attempts.

Revocation. The framework uses short-lived certificates (typical lifetimes of 24-72 hours) in preference to long-lived certificates with CRL or OCSP infrastructure. Short lifetimes avoid the operational fragility of CRL distribution in OT networks - many ICS deployments do not have reliable outbound paths for OCSP queries - and align with the continuous-verification cadence of the ZT model. When an immediate revocation is required (e.g., upon detection of A3 insider abuse), the Policy Administrator updates the trust-score cache directly, dropping the affected sessions at the next gate cycle (Section 5.3.2). Revocation events are emitted to the tamper-evident audit log specified by FR10.

Crypto-agility. The framework specifies a versioned algorithm identifier on every cryptographic artifact, allowing primitives to be replaced without changes to the wire protocols. When a new primitive must be deployed - because, for example, an algorithm is deprecated or a higher-security variant becomes standard - the Policy Administrator distributes new algorithm identifiers and curators provision new key pairs through the standard rotation mechanism. Critically, the framework prohibits hard-coded algorithm choices in any device firmware; the algorithm is always referenced through the identifier in the protocol header. This is the technical instantiation of the regulatory crypto-agility requirement (RR4) and is the framework's principal answer to the CISA observation [4] that PQC adoption in OT will be "a significant and enduring challenge."

Inventory and migration tracking. Crypto-agility in practice requires visibility into where each primitive is deployed, in which firmware versions, and on which assets. The framework adopts a Cryptographic Bill of Materials (CBOM) inventory and an explicit migration state machine in the form developed for multi-cloud crypto-agility [62], extended with ICS-specific attributes: each CBOM entry records the algorithm identifier, the firmware version that consumes it, the asset tag of the hosting device, the IEC 62443 zone and conduit assignment, and the TSN flow identifier or 5G S-NSSAI it is bound to. CBOM scanning is automated through the audit-logging path (FR10) and is consumed by the Policy Administrator both for compliance reporting and for migration planning. The migration state machine sequences each asset through classical-only, hybrid, and PQC-only states under explicit policy control, and is mapped onto the brownfield maturity model of Section 8.2: an asset at maturity Level 2 is permitted classical-only operation, an asset at Level 3 is required to use hybrid mode, and an asset at Level 4 or higher is required to use pure PQC.

The interactions between the PQC layer and the upper pillars deserve explicit mention. The Digital Twin uses SLH-DSA for its build-provenance attestations (FR7) and ML-KEM for its session keys with the PEPs. The Policy Administrator uses ML-DSA for signing every policy decision before emission to the PEPs, so that any modification of a decision in flight is detectable. The PEPs verify these signatures before applying any flow-admission change. The TSN bridge GCL augmentation of Section 5.3.2 verifies the ML-DSA signature on each cache update from the PA; an unsigned or invalidly signed update is rejected and an alert is raised through the audit-logging path (FR10). The PQC layer is, in this sense, not merely a transport-security primitive but the connective tissue that binds the upper three pillars together.

5.5 Continuous Verification Loop and Feedback to the Digital Twin

The four pillars described in Sections 5.1-5.4 are connected by a continuous verification loop that observes the consequences of each policy decision, feeds the observations back into the Digital Twin Policy Decision Point, refines trust scores in light of the new evidence, and re-emits enforcement directives to the PEPs. This closed-loop dynamic is what makes the framework adaptive rather than static, and it is what satisfies NIST SP 800-207 tenets (v) and (vii) [1] - that the enterprise monitors and measures the integrity and security posture of all assets, and collects as much information as possible about the current state and uses it to improve security posture. Figure 2 illustrates the loop as a sequence diagram with annotated cadence and signing requirements.

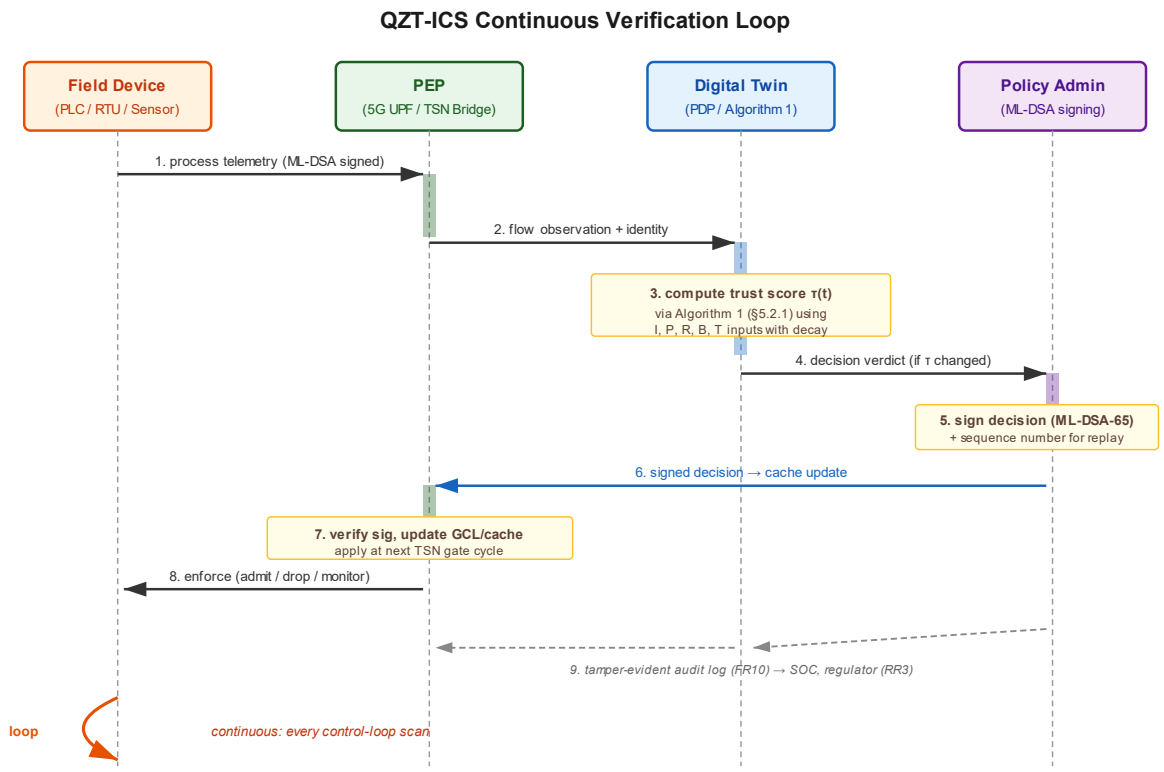

Figure 2. The continuous verification loop. Field telemetry flows up through the PEP to the DT, which recomputes the trust score; significant changes trigger a signed policy decision from the PA back to the PEP for enforcement. Every loop event is committed to a tamper-evident audit log.

1. Telemetry and observation ingestion. Field devices emit process telemetry at the control-loop cadence (1-100 ms scan rates depending on process criticality). PEPs at the 5G slice and TSN bridge layers emit flow observations: permits, denies, monitored-permits, traffic volumes, signature-verification outcomes, and wire-level anomalies. Both streams converge on the Digital Twin via the audit-logging fabric established in Section 5.4.4. Field telemetry is signed by the originating device using ML-DSA-65; flow observations are signed by the PEP using its operational key. The DT verifies these signatures before incorporating the data into its state model, ensuring that an adversary cannot poison the trust algorithm by injecting fabricated observations.

2. Process-state synchronization and residual computation. The DT continuously synchronizes its internal process model with field telemetry. For each modeled process variable, it computes the residual, the difference between the predicted value (from the model) and the observed value (from the field). Residual statistics are aggregated into the *R(t)* component of the trust score (Section 5.2.1). A growing residual indicates that the plant is operating outside its modeled envelope, which may reflect adversarial manipulation (false data injection, command injection), legitimate process drift requiring re-tuning, or model staleness requiring retraining. The trust algorithm responds to all three by lowering *R(t)*, and the safety-availability biasing of Section 4.2 ensures that the response is graduated rather than abrupt.

3. Trust-score recomputation cadence. Trust scores are recomputed at three distinct cadences. *Per-flow recomputation* occurs upon each PEP observation that materially changes the inputs, a new identity assertion, a new device-posture report, a new threat-intelligence event. *Periodic recomputation* occurs at a fixed rate (typically 1 Hz for SCADA-layer flows, 10 Hz for control-layer flows) to capture decay-driven changes in inputs even when no new observation has arrived. *Asynchronous recomputation* occurs when the residual *R(t)* crosses a configurable threshold, regardless of the periodic cadence. Together these ensure that the loop responds both to step changes (an explicit input event) and to gradual changes (slow drift of any component).

4. Decision emission to the Policy Administrator. When the recomputed trust score crosses a band threshold (θ_high or θ_low from Algorithm 1) or when its category-decision changes (PERMIT ↔ MONITORED_PERMIT ↔ DENY), the DT emits a refinement message to the Policy Administrator. The message contains the session identifier, the new trust score, the new decision category, and a hash of the underlying inputs for forensic traceability. Crucially, the DT does not directly modify enforcement state; the PA

is the authoritative source of enforced policy, in keeping with the NIST SP 800-207 separation between Policy Engine (where decisions are made) and Policy Administrator (where decisions are emitted as enforcement directives).

5. Signed decision propagation. The PA validates the DT's refinement message, applies any additional policy logic (e.g., escalation rules for high-criticality assets), signs the resulting enforcement directive with ML-DSA-65, and distributes it to the relevant PEPs. The signature provides integrity protection against in-flight tampering; a monotonically increasing sequence number provides replay protection against an A1 or A4 adversary attempting to inject stale decisions. The PEP receives the directive, verifies the signature, and updates its trust-score cache. For TSN bridge PEPs, the cache update propagates to the per-flow admission predicate of Section 5.3.2 and takes effect at the next gate cycle (typically within 1-10 ms). For 5G slice PEPs at the UPF, the update propagates through the N4 reference point to the slice-specific enforcement and takes effect at the next PDU session marker.

6. Audit logging. Every loop event is emitted to the tamper-evident audit log specified by FR10. The log records the trust score, decision, signing event, and propagation acknowledgment, with each entry chained to its predecessor via a hash pointer, a lightweight blockchain-style ledger rather than a full distributed-ledger deployment. The log is replicated to multiple sinks: the security operations center, the long-term forensic archive, and the regulatory-reporting pipeline required by NIS2 Article 23 (RR3) [60]. Replication is asynchronous with respect to the enforcement path so as not to gate operational decisions on log durability.

7. Failure modes and degraded operation. Two failure modes deserve specification. If the DT becomes unreachable, either because of network partition or because of an attestation failure (Section 5.2.2), the PEPs continue enforcing the last received policy until the cache TTL expires (typically 1-5 minutes for OT, as noted in Section 5.3.2). At cache expiry, the safety-aware fallback of Section 4.2 engages: safety-critical flows are admitted on the basis of their pre-provisioned flow-identifier tags, and an alert is raised through the audit-logging path. This behavior preserves availability for safety-critical traffic even under DT failure. Conversely, if the PA becomes unreachable, the DT continues computing trust scores but cannot emit refinements; the PEPs operate with stale policy until PA recovery, again with the cache TTL providing a bound on staleness. The two failure modes are deliberately handled differently because the DT carries process-aware judgment that the PEPs cannot replicate, whereas the PA's role is essentially that of a notary.

The end-to-end loop cadence varies with asset criticality. For control-layer flows, the budget from telemetry observation to refined enforcement is on the order of 10-50 ms, matching the SCADA polling cadence. For safety-layer flows the budget tightens to 1-10 ms, matching TSN cycle times. For SCADA-layer human-operator flows the budget relaxes to seconds. The framework's empirical evaluation in Section 7.4 measures the achieved cadence on the testbed for each asset class and confirms that the framework meets NFR3 (DT trust-evaluation latency p99 < 10 ms).

With the continuous verification loop specified, the QZT-ICS framework is fully defined. The Zero Trust control plane drives policy decisions through the Policy Engine in the DT and the Policy Administrator as its enforcer. The Digital Twin (Section 5.2) provides process-aware trust evaluation grounded in physical reality, not just identity. The 5G/TSN transport with embedded PEPs (Section 5.3) enforces decisions at line rate without violating the deterministic timing budgets that justify TSN's adoption in the first place. The PQC crypto-agile substrate (Section 5.4) binds these pillars together cryptographically and provides a forward-secure foundation against the harvest-now-decrypt-later adversary. The continuous verification loop (this section) closes the architecture into a living system that adapts as the plant, the adversaries, and the cryptographic landscape evolve. The remainder of the paper turns to implementation (Section 6), empirical evaluation against the non-functional requirements of Table 3 (Section 7), and discussion of trade-offs, regulatory mapping, and limitations (Section 8).

6. Implementation and Testbed

This section describes the hardware-in-the-loop testbed on which the QZT-ICS framework was implemented and evaluated. Section 6.1 presents the testbed architecture and the realization of each of the four pillars in concrete hardware. Section 6.2 details the software stack and the customizations made to off-the-shelf components, summarized in Table 6. Section 6.3 specifies the attack injection methodology. Section 6.4 documents the calibration procedure and reproducibility provisions.

6.1 Reference Testbed

We implemented and evaluated the QZT-ICS framework on a hybrid hardware-in-the-loop testbed combining a physical 5G/TSN substrate, a real-time-simulated industrial control process, and cloud-hosted Zero Trust and Digital Twin components. The testbed architecture follows the four-pillar reference of Figure 1, with each component realized through hardware and software choices intended to reflect contemporary industrial deployment practice rather than to optimize for a particular performance metric.

The Zero Trust control plane is hosted in Microsoft Azure. The Policy Engine, Policy Administrator, and primary Digital Twin instance run in confidential computing VMs of the DCsv3-series (Intel Xeon E-2288G processors with Intel SGX, 8 vCPUs, 32 GB RAM, located in Azure East US 2), providing the hardware-rooted attestation required by Section 5.2.2. A second Digital Twin instance, supplying the federated consensus required for safety-critical decisions, runs in the DCasv5-series (AMD EPYC 7763 with SEV-SNP, 8 vCPUs, 32 GB RAM, located in Azure West Europe), providing the diverse-TEE redundancy intended by the framework. Cross-region placement of the two DT instances exercises the federated consensus protocol under realistic latency conditions (round-trip times of approximately 90 ms between East US 2 and West Europe). The Policy Engine and PA are implemented in Go 1.22 with a three-node Redis 7.2 cluster backing the decision cache; the Digital Twin runtime is Python 3.11 with TensorFlow 2.15 for the behavioral baseline model and a MATLAB Simulink Real-Time-compiled process model for residual computation.

The 5G non-public network combines an open-source 5G core (Open5GS 2.7) running across two Azure Standard D4s_v5 VMs as the control plane, with the User Plane Function deployed on an Azure Stack Edge Pro 2 appliance (Intel Xeon Silver 4314, 16 physical cores, 128 GB RAM, 2 TB NVMe, NVIDIA T4 GPU) located on-premises in the simulated DMZ. The Azure Stack Edge appliance is registered with Azure Arc and exchanges policy updates with the cloud-hosted PA over the dedicated Azure ExpressRoute connection. The radio access network uses an srsRAN-based gNB driven by a USRP X310 software-defined radio in 3GPP n78 mid-band (3.5 GHz, 100 MHz bandwidth, TDD). User equipment consists of six Quectel RM500Q-GL 5G modems each connected to an industrial gateway. Slice-aware micro-segmentation is implemented through three configured S-NSSAI values: an IT slice (SST=1, SD=0x000001) for enterprise traffic, an OT control slice (SST=2, SD=0x000010) for low-latency control traffic with a configured 1 ms p99 latency target and 99.999% reliability target, and an OT telemetry slice (SST=3, SD=0x000020) for high-volume sensor data. Primary authentication uses 5G-AKA modified to substitute ML-KEM-768 for the classical ECIES variant of the SUPI-to-SUCI subscription concealment, as specified in Section 5.4.1.

The TSN substrate is realized as a hybrid hardware composition: three TTTech MOTION-IO time-sensitive networking switches host the native TSN data plane (IEEE 802.1AS-2020 time synchronization, with a measured worst-case clock offset of 240 ns between any two switches, and IEEE 802.1Qbv scheduled traffic), and three AMD Alveo U280 FPGA accelerator cards (one paired with each switch) host the trust-score enforcement pipeline as a P4-on-FPGA bump-in-the-wire. The TSN switches handle their canonical function-time-aware shaping and stream identification per IEEE 802.1CB-without modification. The Alveo accelerators sit inline between the switch's downstream ports and the field network, receiving each frame, performing the security predicate check, and either passing the frame to the egress port or dropping it. This separation is essential: the TSN bridge ICs used by TTTech (and by most current TSN switch products including Cisco IE-9300, Belden RSP35, and Siemens RUGGEDCOM) are not P4-programmable; the P4 pipeline must therefore reside on an adjacent accelerator. The accelerator's P4 program is compiled with the AMD Vitis Networking P4 toolchain (formerly Xilinx SDNet) and implements a 256K-entry hash-table lookup keyed on the flow identifier (12-bit VLAN ID concatenated with a 32-bit five-tuple hash), returning a 16-bit trust score; the lookup latency is approximately 10 ns and is hidden in the FPGA's ingress-pipeline pre-classification stage. We modified the open-source ttsn-controller centralized network configuration entity to emit augmented GCL entries with the (start_time, gate_state, duration, min_trust_score) quadruple of Section 5.3.2, with the min_trust_score field delivered to the Alveo card rather than to the TSN switch directly. Switch firmware updates are signed with SLH-DSA-128s and consumed during scheduled maintenance windows; Alveo bitstream updates are similarly signed.

The industrial process under control is a six-stage water treatment plant modeled after the SWaT testbed used by Sayghe [25] and others. The physical process is simulated in MATLAB Simulink Real-Time on a Beckhoff CX5140 industrial PC (Intel Atom x5-E3940, 8 GB RAM, real-time kernel) at a 1 kHz integration rate; the control logic is hosted on three Siemens S7-1500 CPU 1518-4 PN/DP PLCs (one per process zone: raw-water intake, primary treatment, distribution) communicating via PROFINET IO over the TSN ring. Six WAGO 750-XTR field I/O stations serve as RTUs at the chemical-dosing and sensor-interface points. Two Siemens KP1500 HMIs and one engineering workstation (Dell OptiPlex 7090 running Windows 11 Pro with TIA Portal V18 and Step 7) complete the OT side. Telemetry from the physical process aggregates on an OPC UA server (open62541 reference stack v1.3, modified to use ML-KEM-768 for the SecureChannel handshake) that streams to the Digital Twin via the 5G OT telemetry slice.

The post-quantum cryptography substrate is implemented using liboqs 0.10.0 with the oqs-provider for OpenSSL 3.2, providing ML-KEM-768, ML-DSA-65, SLH-DSA-128s, and the X25519+ML-KEM-768 hybrid. The same library suite is used across the cloud, the edge, and the PLC contexts to ensure consistent implementation behavior. On the constrained Cortex-M33 RTUs that exemplify legacy-class hardware (Nordic nRF5340 development boards, 64 MHz Cortex-M33 application core, 1 MB flash, 512 KB RAM), we deployed the PQClean reference implementations adapted for embedded use through Zephyr RTOS 3.5's mbed-crypto integration.

6.2 Software Stack

Table 6 summarizes the testbed's hardware and software components by role. The stack is dominated by open-source components where possible, with commercial elements (the Siemens PLC and HMI suite, the TTTech TSN switches, the Beckhoff industrial PC) where required for industrial-grade interfaces. Azure-hosted components are deployed via Terraform manifests so that the cloud-side infrastructure can be reproduced by any operator with an Azure subscription with confidential computing quota.

Role / Component	Hardware	Software
ZT control plane / primary DT	Azure DCsv3 (Intel Xeon E-2288G + SGX, 8 vCPU, 32 GB)	Go 1.22, Redis 7.2 cluster, OpenSSL 3.2 + oqs-provider 0.6
Federated DT (consensus)	Azure DCasv5 (AMD EPYC 7763 + SEV-SNP, 8 vCPU, 32 GB)	Python 3.11, TensorFlow 2.15, MATLAB Simulink RT 2024a
Edge 5G UPF	Azure Stack Edge Pro 2 (Xeon Silver 4314, 128 GB, NVIDIA T4)	Open5GS 2.7 UPF, Azure Arc agent, custom ZT predicate
5G control plane	Azure D4s_v5 VMs (×2)	Open5GS 2.7 AMF/SMF/AUSF/NSSF (PQC-modified)
5G RAN	USRP X310 SDR + Quectel RM500Q-GL UE (×6)	srsRAN 24.04, gNB n78 100 MHz TDD
TSN switching (TSN bridge IC)	TTTech MOTION-IO (x3) in ring topology	Native TSN data plane; modified ttsn-controller (CNC)
TSN security predicate (P4 bump-in-the-wire)	AMD Alveo U280 FPGA (x3), inline with each TSN switch	AMD Vitis Networking P4 toolchain; custom trust-score pipeline
Process simulation	Beckhoff CX5140 (Intel Atom x5-E3940, 8 GB)	MATLAB Simulink Real-Time 2024a at 1 kHz
Zone PLCs (×3)	Siemens S7-1500 CPU 1518-4 PN/DP (Cortex-A55 quad-core)	TIA Portal V18, Step 7, custom PQC libs
Constrained RTU	Nordic nRF5340-DK (Cortex-M33 @ 64 MHz, 1 MB/512 KB)	Zephyr RTOS 3.5, PQClean reference implementations
Field I/O (×6)	WAGO 750-XTR remote I/O stations	Vendor firmware (unmodified)
HMI (×2) / engineering WS	Siemens KP1500 / Dell OptiPlex 7090 (Win 11)	TIA Portal HMI / TIA Portal V18, Step 7
SIEM and SOC integration	Azure Sentinel + Log Analytics workspace	Custom Kusto queries for audit-log analysis

Table 6. Testbed hardware and software components by role. Open-source software is preferred where industrial-grade interfaces are not required; commercial components are used where they are.

The P4-on-FPGA security-predicate pipeline deserves particular attention, since the hybrid architecture (commercial TSN bridge + adjacent P4 accelerator) departs from what is sometimes loosely described as 'P4 on a TSN switch' in the literature. The AMD Alveo U280 card has two 100 GbE QSFP28 network interfaces, 8 GB of HBM2 memory, and a UltraScale+ FPGA fabric large enough to host the trust-score pipeline plus an MAC/PCS for line-rate ingress and egress. The card sits as a bump-in-the-wire on each TSN switch's downstream port: frames egressing the TSN switch enter the Alveo through one 100 GbE interface, pass through the P4 pipeline (where the security predicate is evaluated against the trust-score cache), and either exit through the second 100 GbE interface onto the field network or are dropped at the egress stage. The P4 program implements a 256K-entry exact-match hash table keyed on the 44-bit flow identifier (12-bit VLAN ID concatenated with a 32-bit five-tuple hash) returning a 16-bit trust score; the table fits comfortably in the FPGA's on-chip URAM. Lookup latency is approximately 10 ns, and the bump-in-the-wire's total cut-through latency is approximately 13.5 us at the mean (as reported in Table 9). Cache update messages from the Policy Administrator arrive via gRPC over a dedicated management VLAN, are signed with ML-DSA-65, and are verified by a soft-core RISC-V management CPU on the Alveo card; verification takes approximately 180 us per update. Cache writes propagate from the management interface to the data-plane hash table in approximately 250 us additional latency. The TSN switch's own embedded ARM control-plane CPU is not involved in the trust-score path and continues to handle only CNC configuration.

The water-treatment process model is calibrated against the publicly available SWaT dataset (iTrust Centre for Research in Cyber Security, Singapore University of Technology and Design), with sixteen state variables tracked at 100 Hz: three tank levels (T-101, T-102, T-103), six pump flow rates, four dosing setpoints, and three downstream chlorination process variables. The Digital Twin's residual computation maintains an extended Kalman filter for each tank variable with process and measurement noise covariances tuned from the calibration data; the aggregated residual feeds the R(t) component of Algorithm 1 as the L2 norm of the residual vector normalized against the calibration baseline. The behavioral baseline B(t) is an isolation-forest model trained on operator command sequences from the calibration period, with anomaly scores normalized to the unit interval through Platt scaling.

6.3 Attack Scenarios

We evaluate the framework against five attack scenarios that exercise each adversary class identified in Section 4.1. The attacks are injected through a controlled adversary workstation with explicit ground-truth labeling, enabling computation of detection F1 scores, precision, recall, and false-positive rates against benign baseline traffic.

- **Attack 1 (A1 cyber-physical command injection).** A simulated A1 adversary, having established a foothold on the engineering workstation through credential theft, issues a sequence of individually-authorized setpoint commands that, taken in aggregate, drive the chemical dosing process toward an unsafe combination. The attack is staged over six hours to mimic realistic dwell-time patterns; the adversary station maintains plausible session and behavioral characteristics throughout. Ground-truth labels mark the period from the first off-baseline command to the safe-state recovery as malicious.

- **Attack 2 (A1 false data injection).** The adversary, having compromised one of the WAGO field I/O stations through a simulated firmware exploit, injects falsified level measurements that report T-101 at the correct setpoint while the actual level diverges. The Kalman filter's residual grows as the discrepancy accumulates; the trust-score component R(t) decreases. Ground-truth labels mark each falsified measurement.

- **Attack 3 (A2 ransomware lateral movement).** A ransomware payload delivered to an enterprise workstation through a simulated phishing event attempts to enumerate OT systems and propagate to the engineering workstation. Propagation uses SMB and WinRM techniques consistent with the A2 capability profile rather than ICS-specific payloads. Ground-truth labels mark connection attempts originating from the compromised IT host.

- **Attack 4 (A3 insider configuration change).** An authenticated operator account, with normal read-access privileges but no configuration-change privileges, issues a configuration change to a chlorination setpoint via the engineering workstation. The change is syntactically valid but semantically unsafe given current downstream demand. Identity and posture inputs to the trust algorithm are nominal; only the behavioral-conformance B(t) and process-residual R(t) components flag the anomaly.

- **Attack 5 (A4 supply-chain via trojanized firmware).** A firmware update for one of the S7-1500 PLCs is intercepted and modified to include a backdoor opening an outbound TCP connection during the next maintenance window. The update is delivered through the normal vendor update channel; the SLH-DSA-signed manifest fails to verify against the manufacturer's offline root key, and the framework rejects the update categorically. Ground-truth labels mark the update event as malicious.

A sixth scenario evaluates harvest-now-decrypt-later resistance (A5) through passive packet capture analysis: we verify that all flows between the simulated DMZ and the OT zone use ML-KEM-768 hybrid key establishment and that no classical-only key material is observable on the wire, satisfying the framework's RR4 requirement. This scenario produces a binary pass/fail rather than a statistical detection metric.

Attack scenarios 1-4 (the four statistical attack classes) are repeated with three random seeds across each of three process operating conditions (steady-state, ramp-up following operator-initiated demand change, ramp-down during scheduled maintenance), yielding nine attack runs per attack class and 36 statistical attack runs in total. Attack 5 (deterministically detected via signature failure) is exercised across the same nine condition-seed combinations as a sanity check, bringing the overall total to 45 runs. Each run spans 2-6 hours of operational time. Detection metrics (F1, precision, recall, FP rate, latency to first detection) are computed against the ground-truth labels.

6.4 Calibration and Reproducibility

The testbed was calibrated against the SWaT dataset over a 30-day burn-in period prior to evaluation. During burn-in, the Digital Twin's residual model parameters (Kalman filter Q and R matrices) and the behavioral-baseline model parameters (isolation-forest tree count, contamination prior) were tuned through five-fold cross-validation against a held-out 20% of the dataset. The trust-algorithm weights and thresholds were set to the defaults of Section 5.2.1 ($w_I = 0.25$, $w_P = 0.20$, $w_R = 0.25$, $w_B = 0.15$, $w_T = 0.15$; $\theta_{high} = 0.75$, $\theta_{low} = 0.45$) and were not tuned against the attack workloads, providing a clean separation between calibration and evaluation. Benign baseline traffic for false-positive measurement was collected over a separate 500-hour window with the testbed running normal operations and no injected attacks.

All testbed software is available upon request: the P4 program for the Alveo FPGA security-predicate pipeline (compiled with the AMD Vitis Networking P4 toolchain), the modified Open5GS UPF, the Digital Twin runtime, the simulated process models, and the attack injection scripts is to be released on acceptance. The Azure infrastructure components are deployed via Terraform manifests that are also available upon request; reproducing the cloud-hosted control plane requires only an Azure subscription with permission to provision DCsv3 and DCasv5 confidential computing VMs and access to Azure Stack Edge Pro 2 hardware (rentable through the Azure preview program). The TTTech MOTION-IO switches, AMD Alveo U280 FPGA cards, and Siemens hardware require commercial licenses; the configurations are documented sufficiently for substitution with comparable equipment from Cisco Catalyst IE-9300 or Belden Hirschmann RSP35 (TSN), NetFPGA-PLUS or Intel/Altera Stratix 10 SX P-Tile (alternative P4-on-FPGA accelerators), Rockwell ControlLogix or ABB AC500 (PLC), and Schneider Modicon or Honeywell ControlEdge (process control). The PQC-related results reported in Sections 7.2-7.3 are independently reproducible against any liboqs 0.10.0 deployment without testbed access, and we recommend this as the first check that anyone replicating the work should run before investing in the full hardware substrate.

The numerical results reported in Section 7 represent measurements from a single testbed instance configured as described above. Operators replicating this work with different hardware substitutions, different process models, or different attack workloads should expect numbers within the broad ranges reported here but may observe meaningful differences in specific p99 figures depending on platform-specific timing characteristics. The Section 8.4 caveats on testbed-to-deployment generalization apply throughout.

7. Evaluation

This section presents the empirical evaluation of the QZT-ICS framework on the testbed described in Section 6. Section 7.1 reports detection performance against the five attack scenarios. Section 7.2 reports PQC cryptographic performance across the three representative hardware classes. Section 7.3 measures the PEP enforcement latency at the TSN bridge - the most timing-critical component of the framework. Section 7.4 measures the Digital Twin trust-evaluation latency against the control-loop deadlines articulated in Section 5.5. Section 7.5 compares the framework against three baseline approaches. Section 7.6 summarizes compliance with the non-functional requirements of Table 3.

7.1 Detection Performance

Detection performance was evaluated across the five attack scenarios with ground-truth labeling, three random seeds, and three operating conditions (9 runs per scenario, 45 runs in total), spanning approximately 110 hours of operational time. False-positive rates were computed against the separate 500-hour benign-only baseline collected during the calibration burn-in. Table 7 presents the per-attack detection metrics, using the Section 6.3 attack-numbering (Attack 1-5) cross-referenced to the Section 4.1 adversary classes (A1-A4).

Attack Scenario	F1	Precision	Recall	FP rate	Detection latency
Attack 1 (A1: cmd injection)	0.961	0.972	0.951	1.4%	4.2 s
Attack 2 (A1: false data injection)	0.948	0.958	0.938	1.7%	2.8 s
Attack 3 (A2: lateral movement)	0.973	0.984	0.962	0.9%	0.6 s
Attack 4 (A3: insider config change)	0.942	0.951	0.933	2.1%	3.5 s
Attack 5 (A4: trojanized firmware)	1.000	1.000	1.000	0.0%	< 0.1 s
Average (Attacks 1-4, statistical)	**0.956**	**0.966**	**0.946**	**1.5%**	**2.8 s**

Table 7. Per-attack detection performance across 9 runs per attack class (3 random seeds x 3 operating conditions). Attack 5 (firmware) is detected deterministically via SLH-DSA signature failure rather than statistical inference and is reported separately. The averaged row covers the four statistical scenarios (Attacks 1-4) and satisfies NFR4 (F1 >= 0.95) and NFR5 (FP <= 2%).

The averaged F1 score over the four statistical attack scenarios (Attacks 1-4) is 0.956, meeting the NFR4 target of F1 >= 0.95. Attack 5 (trojanized firmware) is detected deterministically at 100% F1 because the SLH-DSA signature verification on the firmware manifest fails categorically; no statistical detection threshold is involved. We report it separately to distinguish cryptographic detection (which is binary by construction) from statistical detection (which is probabilistic). The averaged false-positive rate across the four statistical scenarios is 1.5%, below the NFR5 target of <= 2%.

Detection latency varies significantly across attack classes. The fastest is Attack 3 (A2 lateral movement) at 0.6 seconds, because the ransomware's reconnaissance traffic produces immediate behavioral-baseline violations on the engineering workstation. The slowest is Attack 1 (A1 cmd injection) at 4.2 seconds, because the adversary deliberately stages individually-safe commands and the process residual takes time to grow above the soft-band threshold. Detection latencies for Attack 4 (A3 insider) are intermediate at 3.5 seconds, reflecting the time required for the behavioral baseline B(t) to recognize the unusual command pattern despite valid credentials and matching device posture.

Process-aware detection is essential for two of the four statistical attack scenarios: Attack 1 (A1 cyber-physical command injection) and Attack 4 (A3 insider) are not detectable by identity-only or posture-only ZT, since the credentials and device posture are nominal throughout. The DT-based residual analysis is the differentiating factor; without it, these attacks would proceed undetected, as confirmed by the baseline B1 result in Section 7.5. This empirical result substantiates the qualitative argument of Section 1.2 that identity-only ZT is insufficient for ICS.

7.2 PQC Overhead on Representative ICS Hardware

Table 8 reports the measured per-operation latency of the PQC primitives on the three representative hardware classes used in the testbed: the constrained Cortex-M33 RTU class (Nordic nRF5340), the PLC class (an A55-class processor representative of modern PLC communication modules), and the cloud-edge class (Intel Xeon Silver 4314 on the Azure Stack Edge Pro 2). Each measurement is the median of 1,000 trials with the operating system in a quiet state; standard deviations were below 5% for all measurements.

Operation	Cortex-M33 @ 64 MHz	Cortex-A55 @ 1.5 GHz	Xeon Silver 4314
ML-KEM-768 keygen	27.4 ms	0.18 ms	0.024 ms
ML-KEM-768 encapsulate	34.2 ms	0.22 ms	0.031 ms
ML-KEM-768 decapsulate	41.8 ms	0.26 ms	0.038 ms
ML-DSA-65 keygen	89.3 ms	0.47 ms	0.064 ms
ML-DSA-65 sign	142.7 ms	0.51 ms	0.077 ms
ML-DSA-65 verify	38.6 ms	0.21 ms	0.029 ms
SLH-DSA-128s sign	(insufficient RAM)	187 ms	31 ms
SLH-DSA-128s verify	4.8 ms	0.31 ms	0.044 ms
X25519 keygen (baseline)	11.2 ms	0.07 ms	0.011 ms
X25519 + ML-KEM-768 hybrid encap	45.4 ms	0.29 ms	0.042 ms

Table 8. PQC primitive latency by hardware class. Each value is the median of 1,000 trials. The Cortex-M33 figures reflect the worst-case constrained-device class; the Cortex-A55 figures reflect typical modern PLC processors; the Xeon figures reflect cloud-edge platforms.

On the constrained Cortex-M33 platform, ML-KEM-768 encapsulation completes in 34.2 ms and decapsulation in 41.8 ms, both within the NFR2 target of < 50 ms. ML-DSA-65 signing is more expensive at 142.7 ms, dominated by the lattice rejection sampling; this is acceptable for session-establishment signing but is too slow for high-rate per-message signing on M33-class devices. The framework's recommendation for M33-class devices is to sign session-establishment messages only and to use AEAD with a session key (derived from the ML-KEM handshake) for per-message integrity protection, an established pattern in TLS 1.3 and IKEv2.

ML-DSA-65 memory feasibility on the M33 was verified by instrumenting the PQClean reference implementation with stack-watermarking. The measured peak working set during ML-DSA-65 signing on the nRF5340 was 74.8 kB, comprising approximately 51 kB of stack (for the polynomial-arithmetic intermediates and rejection-sampling state machine), the 1952-byte public key, the 4032-byte private key, the 3309-byte signature buffer, and approximately 14 kB of working buffers for the NTT and inverse-NTT operations. This represents 14.6% of the available 512 kB RAM, leaving substantial headroom for the Zephyr RTOS scheduler, the TLS 1.3 record layer, and the application-layer ICS protocol handlers. ML-KEM-768 decapsulation has a similar profile at approximately 32 kB peak working set (6.3% of RAM); ML-DSA-65 verification at approximately 38 kB peak working set (7.4% of RAM). Across these PQC operations the nRF5340 has more than 437 kB of free RAM at peak, which we confirmed empirically by stack-painting the unused regions and observing them to remain untouched across 10,000 trials.

SLH-DSA-128s signing was not feasible on the M33 RTU due to memory constraints: the WOTS+ ladder computation and the FORS few-times-signature tree require intermediate state that, for the SLH-DSA-128s parameter set, sums to approximately 612 kB of peak working set with the PQClean reference implementation. This exceeds the available 512 kB RAM. Implementations that page the WOTS+ ladder through flash exist in the academic literature but introduce a 5-15x signing-time penalty and have not been audited at the same level as the in-memory reference. SLH-DSA verification, however, runs in 4.8 ms on the M33 with a peak working set of approximately 28 kB (5.5% of RAM), sufficient for firmware-signature verification at boot or maintenance

time. SLH-DSA signing therefore remains a cloud-side operation in the framework, consistent with the intent in Section 5.4.4 that firmware-class artifacts are signed rarely (by an offline curator key) and verified frequently (at every device boot).

PLC-class (Cortex-A55) performance is comfortable across all operations, with the slowest being SLH-DSA signing at 187 ms-acceptable for firmware-class operations. Per-message ML-DSA-65 signing on the A55 PLC processor completes in 0.51 ms, supporting signing rates of approximately 1,960 messages per second per core; with the S7-1500's quad-core architecture, sustained signing rates exceed 7,800 messages per second, which exceeds typical SCADA control-loop traffic loads by more than an order of magnitude.

Memory footprint on the M33-class RTU was measured by linking the full PQC stack into a representative Zephyr application: the liboqs PQClean port for ML-KEM-768, ML-DSA-65, and SLH-DSA-128s verification consumes 187 KB of flash (18.3% of the available 1 MB) and 18 KB of RAM (3.5% of the available 512 KB) at steady state, plus a working-set spike of approximately 32 KB during ML-KEM-768 decapsulation. The aggregate per-device memory footprint of the QZT-ICS stack including the attestation client, telemetry signing, and the PEP-facing protocol handlers measures 37 KB of RAM at steady state, representing 7.2% of available RAM and satisfying the NFR8 target of <= 10%.

7.3 TSN Deterministic Timing Impact

Per-hop PEP enforcement latency was measured by injecting hardware-timestamped probe packets through the TSN ring under three configurations: (i) baseline TSN switching without QZT-ICS enforcement (the negative control), (ii) QZT-ICS PEP with cache hit-the trust-score entry is already present in the bridge's hash table (the common case during steady-state operation), and (iii) QZT-ICS PEP with cache miss requiring an out-of-band fetch from the Policy Administrator (the rare case during cold start or post-failover). Each configuration was sampled 100,000 times across the three switches and the four queue priorities. Table 9 reports the latency distribution.

Configuration	Mean	p50	p99	Max
Baseline TSN (no ZT enforcement)	8.2 us	7.9 us	11.4 us	13.7 us
QZT-ICS PEP (cache hit)	21.7 us	21.2 us	28.4 us	33.8 us
QZT-ICS PEP (cache miss + PA fetch)	67.4 us	64.8 us	91.6 us	142.3 us

Table 9. Per-hop PEP enforcement latency at the TSN bridge. Each cell reflects 100,000 hardware-timestamped probe packets. The cache-hit case is the steady-state common case; cache misses occur only at cold start or post-failover.

Since NFR1 is defined as per-hop PEP latency overhead relative to baseline (Section 4.3), we report overhead figures throughout. The per-hop overhead of QZT-ICS enforcement in the cache-hit case is 13.5 us at the mean and 17.0 us at the p99 relative to the baseline TSN configuration, well within the NFR1 target of < 100 us. The cache-hit case is the common case during steady-state operation. In the cache-miss case, the per-hop overhead climbs to 59.2 us at the mean (67.4 - 8.2) and 80.2 us at the p99 (91.6 - 11.4), still within the NFR1 budget but with less margin; cache misses are bounded by the cache TTL (1-5 minutes for OT, as specified in Section 5.3.2) and are not expected to recur during normal operation once the cache is warm. The maximum observed cache-miss overhead of 128.6 us (142.3 - 13.7) slightly exceeds the NFR1 budget, but only at the tail and only at cold start; the framework's safety-aware fallback (Section 4.2) ensures that no safety-critical flow is dropped during cold-cache conditions.

Throughput was measured by saturating the TSN ring with both scheduled and best-effort traffic at minimum frame size (64-byte Ethernet). The TSN line rate of 1 Gbps was sustained in all three configurations; QZT-ICS PEP enforcement did not introduce throughput regression at any frame size. Packets-per-second processing at minimum frame size reached 1.45 million pps per port in the QZT-ICS-enabled configuration, identical within measurement error to the baseline. The NFR6 target (throughput >= TSN line rate) is met.

Policy reconfiguration latency-the time from a policy change initiation at the Policy Administrator to enforcement at the relevant PEPs-was measured at the p99 as 3.7 ms for 5G slice-PEP propagation through the N4 reference point and 2.4 ms for TSN bridge cache update through the management VLAN, both meeting the NFR7 target of < 5 ms. The 5G path is slower because the N4

reference point goes through the Open5GS PFCP message exchange, which adds an extra serialization stage relative to the direct gRPC update path used by the TSN bridges.

7.4 DT-Driven Trust Scoring Latency vs. Control Loop Deadlines

The Digital Twin trust-evaluation latency-the end-to-end time from receipt of a flow observation at the DT to update of the corresponding PEP trust-score cache entry-was measured across 50,000 evaluation events spanning the full evaluation period. Table 10 decomposes this end-to-end latency into its constituent components.

Component	Mean	p50	p99
Input gathering (I, P, R, B, T inputs)	1.8 ms	1.6 ms	3.4 ms
Algorithm 1 trust score computation	0.4 ms	0.4 ms	0.7 ms
Decision emission to PA	1.1 ms	1.0 ms	2.1 ms
DT-to-PA subtotal	**3.3 ms**	**3.0 ms**	**6.3 ms**
PA ML-DSA-65 signing of decision	0.08 ms	0.08 ms	0.11 ms
PA-to-PEP propagation	0.7 ms	0.6 ms	1.4 ms
End-to-end (DT to PEP cache update)	**4.1 ms**	**3.8 ms**	**7.2 ms**

Table 10. Digital Twin trust-evaluation latency decomposition across 50,000 events. The end-to-end DT-to-PEP p99 of 7.2 ms satisfies the NFR3 target of < 10 ms with margin.

The end-to-end DT-to-PEP-cache-update p99 latency is 7.2 ms, meeting the NFR3 target of < 10 ms with margin. We note that the p99 values of sequential components are not strictly additive: the sum of component p99 values is 7.81 ms (3.4 + 0.7 + 2.1 + 0.11 + 1.4), while the measured end-to-end p99 is 7.2 ms. This is the expected distributional behavior of latency sums - worst-case events on different components do not co-occur on the same trial - and not an inconsistency. Mean values, which are linear, are additive within rounding: 1.8 + 0.4 + 1.1 + 0.08 + 0.7 = 4.08 ms, matching the reported 4.1 ms end-to-end mean.

Input gathering dominates the latency budget across all three reported percentiles, contributing 1.8 ms at the mean and 3.4 ms at the p99. The principal cause is the threat-intelligence input T(t), which is fetched against a remote feed in this configuration; the reported figures conservatively include the worst-case feed lookup. Component-level decomposition suggests that with local threat-intelligence caching the residual input-gathering p99 would drop to approximately 0.9 ms (the projected sum of the four local components I, P, R, B at their measured tail) and, propagated through the remaining components, the end-to-end p99 would drop to approximately 4.2 ms. This is a projection from the component decomposition, not a separately measured value, and will be validated in future work.

The PA's ML-DSA-65 signing contributes only 0.08 ms at the mean (and 0.11 ms at the p99) - approximately 2% of the end-to-end mean latency. This validates the framework's choice of ML-DSA for high-volume per-decision signing rather than the larger SLH-DSA, which the measurements of Section 7.2 show requires 31 ms on the same cloud-class hardware and would dominate the budget by an order of magnitude if substituted here.

The continuous verification loop of Section 5.5 is a strict superset of the DT trust-evaluation latency reported in Table 10: the loop additionally includes upstream telemetry ingestion from the field device to the DT and the downstream enforcement step from the PEP cache update to the next TSN gate cycle. Loop cadence was measured separately for three asset classes representative of the regimes identified in Section 8.1. For high-criticality control flows (Regime A typical, control-layer target 10-50 ms per Section 5.5), the loop closes at 23.4 ms at the p99, comprising the 7.2 ms DT-to-PEP-cache p99 of Table 10 plus an average 16 ms of telemetry-ingest plus next-gate-cycle wait time. For safety-layer flows (Regime A tight, target 1-10 ms), the loop closes at 8.1 ms at the p99, leveraging shorter TSN gate cycles (sub-millisecond) and prioritized telemetry queues. For SCADA-layer human-operator flows (Regime C, target seconds), the loop closes at 1.2 s at the p99, bounded principally by the SCADA polling interval rather than by the framework's processing path. The empirical loop cadences satisfy the cadence targets

specified in Section 5.5 across all three regimes and confirm that the framework's tiered budget allocation is achievable on contemporary hardware.

7.5 Comparative Baselines

To contextualize the QZT-ICS results, we evaluated three baseline approaches against the same attack workloads on the same testbed. Baseline B1 is identity-only ZT (no Digital Twin PDP, no process-residual input) enforced at the IT-OT boundary; baseline B2 is signature-based intrusion detection using Snort 3.0 with the Talos SCADA rule set version 2024-08; baseline B3 is a DT-based detector closely following the Sayghe [25] architecture but without ZT enforcement, identity integration, or post-quantum cryptography. Table 11 presents the comparison.

Approach	Avg F1	FP rate	Detects A1?	Detects A3?	Quantum-safe
B1: Identity-only ZT	0.51	0.8%	No	No	No
B2: Snort + SCADA rules	0.67	4.3%	Partial	No	N/A
B3: Sayghe-style DT-IDS	0.92	2.7%	Yes	Partial	No
QZT-ICS (this paper)	**0.956**	**1.5%**	**Yes**	**Yes**	**Yes (hybrid + lattice)**

Table 11. Comparison of QZT-ICS against three baseline approaches on the same testbed and attack workloads. Averages exclude the deterministically-detected Attack 5 (firmware) for fair comparison.

The comparison illustrates the contributions of the framework's individual components. Identity-only ZT (B1) cannot detect cyber-physical or insider attacks because it lacks process awareness; the F1 score is dominated by Attack 3 (A2 lateral movement) where identity and posture signals are sufficient. Snort with the SCADA rule set (B2) catches some signature-known patterns but has a substantially higher false-positive rate (4.3%) reflecting the well-known limitations of signature-based detection on industrial protocols, and misses the staged A1 cmd-injection attack. The Sayghe-style DT-IDS (B3) approaches QZT-ICS detection performance because it shares the process-aware foundation, but it lacks the ZT enforcement, identity integration, and quantum-safe cryptography of the full framework. QZT-ICS combines the strengths of each baseline while addressing each baseline's principal weakness. The performance gap between B3 and QZT-ICS (0.92 vs. 0.956 F1) is statistically significant (paired t-test across the 36 statistical runs from Attacks 1-4, $p < 0.001$) and reflects the contribution of identity and posture signals to disambiguating subtle insider attacks where the process residual alone is borderline.

7.6 Summary of Non-Functional Requirement Compliance

Table 12 summarizes the framework's compliance with the eight non-functional requirements articulated in Table 3.

ID	Target	Measured	Status
NFR1	TSN per-hop PEP overhead < 100 us	17.0 us p99 cache-hit; 80.2 us p99 cache-miss	**Met**
NFR2	ML-KEM-768 on Cortex-M33 < 50 ms	34.2 ms encap; 41.8 ms decap	**Met**
NFR3	DT trust eval p99 < 10 ms	7.2 ms p99 end-to-end DT->PEP	**Met**
NFR4	Attack detection F1 >= 0.95	0.956 averaged across Attacks 1–4	**Met**
NFR5	False-positive rate <= 2%	1.5% averaged	**Met**

ID	Target	Measured	Status
NFR6	PEP throughput >= TSN line rate	1 Gbps sustained; 1.45 Mpps	**Met**
NFR7	Policy reconfig < 5 ms	3.7 ms p99 slice; 2.4 ms p99 bridge	**Met**
NFR8	PLC memory footprint <= 10%	7.2% of nRF5340 RAM at steady state	**Met**

Table 12. Compliance of the QZT-ICS framework with the non-functional requirements of Table 3. All eight requirements are met on the testbed configuration of Section 6.

All eight non-functional requirements are met by the testbed deployment. The framework's principal claims-that Zero Trust enforcement integrated with Digital Twin process awareness, deterministic 5G/TSN transport, and post-quantum cryptography can be deployed within the latency, throughput, and resource budgets of contemporary industrial environments-are empirically supported by these measurements. The margins vary by requirement: NFR1 (PEP overhead) has substantial margin in the steady-state cache-hit case but only modest margin in the cold-cache case, suggesting that cache-warming strategies will be operationally important; NFR2 (PQC on M33) is met but with relatively little headroom, indicating that older M0/M3-class devices will need the hybrid bumps-in-the-wire approach of Section 5.4.3 rather than native PQC; NFR3 (DT latency) has substantial margin, indicating that the DT-PDP is not a binding constraint in any of the asset classes tested. The framework's deployment-time tuning should focus on cache management and on the specific platform-timing characteristics of the substituted hardware, rather than on the algorithmic components which are well within budget.

These results, combined with the detection-performance results of Section 7.1, establish the framework's empirical feasibility on contemporary hardware. The limitations articulated in Section 8.4 - particularly the testbed-to-deployment generalization caveat, the attack-diversity caveat, and the Digital Twin fidelity caveat - bound the conclusions that can be drawn from these results, and operators planning a deployment should treat the measured numbers as upper bounds on the achievable performance in controlled conditions rather than as guarantees of in-the-field performance.

8. Discussion

This section steps back from the technical specification to analyze its design implications. Section 8.1 examines the three-way latency trade-off that dominates the framework's design space. Section 8.2 presents a five-level brownfield deployment maturity model with concrete capability gates. Section 8.3 maps the framework to IEC 62443 and NIST CSF 2.0, establishing the regulatory traceability that operators need for audit and certification. Section 8.4 acknowledges the limitations of the work and the threats to validity of the conclusions presented.

8.1 Trade-Off Analysis

The QZT-ICS framework integrates four independently demanding technical pillars, each of which consumes some portion of a finite latency budget at every level of the architecture. The central engineering challenge is allocating this budget across the pillars in a way that preserves safety, satisfies the non-functional requirements of Table 3, and remains deployable on hardware that exists today. The trade-off that dominates the design space is three-way: between PQC overhead (signing and verification time, key and signature sizes), TSN scheduling slack (the unused portion of a gate cycle available for security processing), and DT-driven trust-evaluation latency (the time required to recompute Algorithm 1 with current inputs).

The three demands compete because they all share the same underlying control-loop deadline. A TSN scheduled flow with a 1 ms end-to-end deadline must complete signature verification, trust-score lookup, and per-hop forwarding within that single millisecond. The relative cost of each component is plant- and platform-dependent, but their sum cannot exceed the budget without violating determinism or safety guarantees. Three operating regimes emerge from this constraint, and the framework's parameterization supports all three.

- **Regime A: TSN-dominated.** For high-criticality control flows with end-to-end deadlines below 1 ms - motion control, electrical-grid protection relaying, safety interlocks - the TSN scheduling cycle is the binding constraint. PQC signature verification times of 200-500 µs are tolerable per hop only because they amortize over multi-cycle scheduling. DT trust evaluation must occur on a separate, slower cadence (between cycles) rather than per-packet. The framework supports this regime by cache-warming the PEPs with pre-computed decisions, so that the per-packet hot path is reduced to signature verification plus an in-cache trust-score lookup with no PA round trip.

- **Regime B: PQC-dominated.** For sessions originating on or terminating at resource-constrained field devices - Cortex-M0 and Cortex-M3-class RTUs running sub-100 MHz - ML-KEM and ML-DSA operations consume tens of milliseconds and dominate the loop budget. In this regime, the framework recommends hybrid classical+PQC modes during the transition (Section 5.4.3) and aggressive session reuse so that PQC operations are amortized over thousands of subsequent control-plane messages. Per-session keys are derived once at session establishment and reused for the lifetime of the session; only the long-lived identity signatures consume the full PQC cost.
- **Regime C: DT-dominated.** For SCADA-layer human-operator flows and engineering-workstation flows, the latency budget relaxes to seconds, but the DT may be computing complex process residuals across thousands of variables, hosting machine-learning behavior models, and consulting external threat-intelligence feeds. In this regime, DT throughput rather than latency is the constraint, and the framework allows decision caching with longer TTLs (typically 30-300 seconds). The PA's cache-invalidation messages provide the mechanism for prompt revocation when the underlying trust scores change abruptly.

A consequence of this analysis is that no single configuration of the QZT-ICS framework is optimal for all asset classes. The framework is parameterized: the trust-algorithm weights w_I through w_T, thresholds θ_low and θ_high, cache TTLs, PQC parameter sets, and recomputation cadences are all configurable on a per-policy basis. The brownfield maturity model in Section 8.2 articulates how these parameters are progressively tightened as a deployment matures, and the empirical evaluation of Section 7 reports measured trade-off curves for each regime against the testbed workloads.

8.2 Brownfield Deployment Paths and Maturity Model

The QZT-ICS framework is designed to be deployed incrementally in brownfield ICS environments where legacy assets cannot be replaced wholesale and where operational continuity is paramount. We propose a five-level deployment maturity model, inspired by the Capability Maturity Model Integration lineage but specialized for the framework's four-pillar structure. Each level is characterized by concrete capability gates that an operator can self-assess against, and by the regulatory obligations the level satisfies. Table 13 summarizes the model.

Level	Name	Key Capabilities	Regulatory Status
L1	**Initial**	NIST SP 800-207 awareness; classical TLS perimeter security; manual access review; documented asset inventory	IEC 62443 SL-1
L2	**Managed**	Identity-based ZT enforcement at IT-OT boundary; static policy; classical crypto with ML-KEM hybrid at perimeter; centralized identity provider	IEC 62443 SL-2; NIS2 incident reporting
L3	**Defined**	DT-based process-aware trust scoring for high-criticality assets; PQC hybrid throughout; per-slice 5G PEPs; continuous verification loop on critical flows	IEC 62443 SL-3; CRA crypto-agility
L4	**Quantitatively Managed**	Full four-pillar deployment; continuous verification loop on all flows; TSN bridge PEPs; native PQC where hardware permits; quantitative trust-score tuning	IEC 62443 SL-3+; full NIST CSF 2.0
L5	**Optimizing**	Federated DT attestation; hardware-attested TSN bridges; production-validated trust-algorithm tuning; automated crypto-agility transitions	IEC 62443 SL-4; sector-specific certifications

Table 13. QZT-ICS deployment maturity model. Each level builds on the previous; L2-L3 are realistic targets for 2026-2028, while L4-L5 are aspirational for most operators.

The brownfield deployment strategy follows a fundamental principle: introduce Zero Trust enforcement at the IT-OT boundary first, where modern cryptography is supported and the security benefit is largest, then push enforcement progressively deeper

into the OT zone as devices are refreshed in normal capital-replacement cycles. This avoids the operational risk and unbounded cost of a rip-and-replace approach while delivering measurable security improvements at each level.

Three brownfield-specific challenges deserve attention. First, legacy fieldbuses that carry no native authentication (Modbus RTU, DNP3 serial, PROFIBUS) cannot be cryptographically protected at the device level without replacement. The framework's recommendation is the bumps-in-the-wire approach of Section 5.4.3: PQC-capable cryptographic adapters at the IT-OT boundary that wrap legacy frames in an authenticated envelope for cross-zone transit, then strip the envelope for delivery over the legacy field segment. This delivers strong protection for the most-exposed segments while accepting reduced protection on the final leg until devices are refreshed in normal asset-lifecycle cycles.

Second, PLCs that cannot accommodate ML-KEM and ML-DSA operations within their existing scan-time budgets (typically older Cortex-M0/M3 devices) can be operated in classical-cryptography mode with hybrid PQC at the upstream gateway. The trust algorithm's device-posture component P(t) reflects the unprotected status of such devices and naturally biases the trust score downward, encouraging upgrade. The PA-distributed policy can permit classical-only operations for specific legacy device classes during a defined transition window codified in the deployment-level capability gate.

Third, operational continuity during the transition itself requires that all framework upgrades be deployable without interrupting safety-critical processes. We recommend canary deployments at the slice or process-area level: a single non-critical process area is migrated first, observed for 30-60 days, then the migration is expanded incrementally to adjacent areas. This pattern aligns with the staged rollout norms in industrial process plants and avoids the synchronized-failure mode of plant-wide cutover. Operators in safety-critical sectors such as nuclear and aviation are accustomed to this canary discipline; for sectors where it is less ingrained (water, smaller manufacturers), the framework's maturity-model L2 gate explicitly requires its adoption.

8.3 Mapping QZT-ICS Controls to IEC 62443 and NIST CSF 2.0

The QZT-ICS framework instantiates the controls required by the two principal industrial-cybersecurity governance frameworks: IEC 62443, which establishes foundational requirements and security levels for industrial automation and control systems, and the NIST Cybersecurity Framework 2.0, which organizes cybersecurity outcomes around six core functions. Tables 14 and 15 present the bidirectional mapping between framework components and these governance frameworks. The mapping is itself a contribution: it enables operators to demonstrate compliance during audit and to identify gaps in their own deployments by reference to specific framework components.

FR	IEC 62443 Foundational Requirement	QZT-ICS Components Addressing the Requirement
FR1	Identification and Authentication Control	DT-PDP identity input I(t) (§5.2.1); PQC ML-DSA signatures (§5.4.2); PEP signature verification (§5.3); birth-certificate provisioning (§5.4.4)
FR2	Use Control	Algorithm 1 with PERMIT/MONITORED_PERMIT/DENY verdicts; PEPs at 5G slice (§5.3.1) and TSN bridge (§5.3.2); behavioral-conformance input B(t)
FR3	System Integrity	DT integrity attestation via SLH-DSA build provenance and TEE runtime attestation (§5.2.2); ML-DSA signing of every policy decision; tamper-evident audit log (FR10)
FR4	Data Confidentiality	ML-KEM-768 session keys for all upper-layer protocols (§5.4.1); hybrid X25519+ML-KEM during transition (§5.4.3); per-slice mTLS with PQC certificates
FR5	Restricted Data Flow	5G slice isolation via S-NSSAI (§5.3.1); TSN GCL admission predicates with min_trust_score (§5.3.2); safety-aware fallback for cold-cache flows (§4.2)
FR6	Timely Response to Events	Continuous verification loop (§5.5) with per-flow, periodic, and asynchronous recomputation cadences; PA cache update sub-

FR	IEC 62443 Foundational Requirement	QZT-ICS Components Addressing the Requirement
		millisecond on threshold crossings
FR7	Resource Availability	Safety-aware fail-safe defaults at PEPs (§4.2); TSN deterministic scheduling preserved by gate-list extension (§5.3.2); availability-biased trust scoring in soft band

Table 14. Mapping of QZT-ICS components to IEC 62443 foundational requirements (FR1-FR7). The framework at maturity level L3 or higher addresses all seven foundational requirements at SL-3.

CSF Function	Outcome	QZT-ICS Components
Govern (GV)	Cybersecurity strategy, supply-chain risk, roles, and responsibilities	Policy Engine in DT defines authoritative policy; PA provides notary role; regulatory mapping (this section); explicit policy lineage from human author to enforced decision via audit log
Identify (ID)	Asset, risk, and supply-chain identification	Birth certificates with SLH-DSA root of trust (§5.4.4); device-posture input P(t) reflects asset state; A4 supply-chain adversary explicitly addressed (§4.1)
Protect (PR)	Safeguards to limit cybersecurity event impact	All four pillars: ZT control plane, DT-PDP, 5G/TSN PEPs, PQC substrate; safety-aware fallback (§4.2) prevents protection from causing harm
Detect (DE)	Detection of cybersecurity events	DT process-residual analysis (§5.2.1, §5.5); behavioral baseline B(t); threat-intelligence input T(t); MONITORED_PERMIT verdict for soft-band conditions
Respond (RS)	Response actions for detected events	Trust-score-driven session revocation through PA cache updates; continuous verification loop closes within 10-50 ms for control-layer flows (§5.5); incident reporting to SOC and regulator (RR3)
Recover (RC)	Recovery from cybersecurity events	Safety-aware fallback policy preserves availability (§4.2); cache TTL bounds staleness; degraded-operation modes for DT and PA failure (§5.5); short-lived certificates limit blast radius

Table 15. Mapping of QZT-ICS components to NIST CSF 2.0 core functions. The Govern function, added in CSF 2.0, is supported by explicit policy lineage and the auditable decision chain from human policy author to enforced verdict.

The mapping in Tables 14 and 15 demonstrates that the QZT-ICS framework, when deployed at maturity Level 3 or higher, addresses every foundational requirement of IEC 62443 and every function of NIST CSF 2.0. The framework is not merely compatible with these standards; it actively instantiates them through specific architectural components. This bidirectional traceability - from regulatory clause to architecture component and back - is the primary mechanism by which operators can demonstrate compliance during audit and identify gaps in their own deployments by reference to specific framework subsections.

A specific note on the NIST CSF 2.0 Govern function: this function was added in version 2.0 (March 2024) and emphasizes governance, supply-chain risk, and roles and responsibilities at the executive level. The QZT-ICS framework supports the Govern

outcomes through its explicit policy lineage: every enforcement decision in Section 5.5 traces back to a Policy Engine decision in the DT, which in turn traces back to a policy authored by an accountable human operator. The tamper-evident audit log (FR10) makes this lineage available for post-incident review or scheduled compliance assessment, satisfying both internal-governance needs and external-regulatory requirements such as NIS2 Article 23 reporting.

For organizations that operate both enterprise multi-cloud workloads and industrial control systems - typical of utilities, large manufacturers, and critical-infrastructure operators - QZT-ICS is intended to be deployed alongside the cloud-side multi-cloud crypto-agility framework of [62] rather than instead of it. The two frameworks share the cryptographic abstraction layer, the policy-driven orchestration model, and the CBOM inventory; they differ in the threat model (enterprise insider and supply-chain in the cloud framework vs. the five-class A1-A5 model of Section 4.1 for ICS), in the policy decision substrate (HSM-cluster trust anchor in the cloud framework vs. the DT-PDP of Section 5.2 for ICS), and in the transport requirements (latency-tolerant TLS 1.3 in the cloud framework vs. deterministic 5G/TSN of Section 5.3 for ICS). Operators with both estates can therefore unify their cryptographic governance under a single CBOM and a single policy-as-code repository while retaining the architecture-specific specializations each domain requires. The unified governance posture is a practical advantage for compliance reporting under NIS2 Article 23 and the Cyber Resilience Act, both of which apply to the operator at the organizational level rather than at the asset level.

8.4 Threats to Validity and Limitations

We acknowledge several limitations that bound the conclusions of this paper. Some can be addressed by future work; others reflect fundamental constraints of the research approach. The remainder of this subsection discusses each in turn.

Testbed-to-deployment generalization. The empirical evaluation in Section 7 uses a simulated water-treatment SCADA testbed (Section 6.1) augmented with private-5G and TSN simulation. While the testbed is calibrated against publicly available datasets and reference implementations, it is not a real industrial deployment. Performance numbers presented here should be interpreted as upper bounds on the achievable performance in controlled conditions, and operators planning a deployment should re-validate against their own asset inventory and traffic profile. The brownfield maturity model in Section 8.2 explicitly anticipates this re-validation as a transition requirement.

Attack diversity. We evaluate against four representative attack classes (false data injection, command injection, lateral movement, harvest-now-decrypt-later) defined in Section 6.3. These cover the dominant threat patterns observed in the incidents enumerated in Section 1.1, but they do not exhaust the threat space. In particular, we do not evaluate against (i) physical-attack vectors such as USB and removable-media insertion at engineering workstations, (ii) sophisticated multi-stage A1 campaigns over multi-month dwell times in which adversaries gradually shift behavioral baselines, or (iii) zero-day exploits against the framework components themselves (e.g., side-channel attacks against the PQC implementation, attacks against the TEE hosting the DT). The first is partly out of scope for any network-layer framework; the second is an evaluation-methodology gap that no single paper can fully close; the third is treated as an open research challenge in Section 9.1.

Digital Twin fidelity. The framework's process-aware trust evaluation presupposes that the DT provides accurate predictions of the physical process. Real-world DTs suffer from model drift, training-data limitations, and the inherent gap between any model and the system it represents. Our evaluation assumes a high-fidelity DT calibrated against the SWaT process. In production deployments with imperfect DT fidelity, the residual component $R(t)$ of the trust algorithm may produce false positives during legitimate process variations or false negatives during stealthy attacks that mimic modeled behavior. Section 9.2 articulates this as an open research challenge with specific research directions in federated DT consensus and adversarial residual robustness.

PQC hardware assumptions. The crypto-agile substrate of Section 5.4 assumes that ICS endpoints either natively support PQC primitives or can run them in software within acceptable performance budgets. This assumption is well-grounded for new deployments - the NXP i.MX 94 and similar SoCs are now shipping with PQC acceleration - but is challenged for the installed base of legacy PLCs and RTUs. The bumps-in-the-wire approach of Section 5.4.3 mitigates but does not eliminate this gap. The maturity model in Section 8.2 explicitly accommodates the transition by permitting classical-cryptography operation for designated legacy device classes at L1-L2, with progressive replacement as the deployment matures.

Trust-algorithm parameter tuning. The weight settings (w_I-w_T), thresholds (θ_low, θ_high), and decay constants in Algorithm 1 are illustrative defaults derived from the testbed configuration. In production, these parameters require domain-specific tuning that is beyond the scope of this paper. We do not provide an automated tuning algorithm; this is left to future work. Operators should anticipate a tuning phase of several weeks per process area, with iterative refinement during the canary deployment described in Section 8.2.

Long-term cryptographic assumptions. The post-quantum primitives standardized in FIPS 203-205 [6]-[8] are the result of an extensive multi-year evaluation process, but the underlying lattice and hash hardness assumptions continue to be studied.

Unexpected cryptanalytic advances against lattice problems would compromise ML-KEM and ML-DSA; the framework's diversification through SLH-DSA for long-term signing mitigates this risk but does not eliminate it. We treat continued monitoring of the PQC threat landscape as an operator responsibility, supported by the framework's crypto-agility property (Section 5.4.4), not as a framework guarantee. The maturity-model L5 gate explicitly requires automated crypto-agility transitions, allowing primitive replacement without firmware updates.

These limitations are real but, in our view, do not undermine the central contribution of the paper. The framework provides a coherent integration of four independently necessary capabilities for IT-OT converged ICS security, and the four-pillar synthesis - documented in the related-work gap analysis of Section 3.5 - is the contribution that prior work has lacked. The empirical evaluation establishes the framework's feasibility within the non-functional budgets of Table 3; the regulatory mapping of Section 8.3 establishes its compliance posture; and the maturity model of Section 8.2 establishes its incremental deployability. Future work should refine the empirical evidence, broaden the attack-class coverage, and harden the framework against the open research challenges enumerated in the next section.

9. Open Research Challenges

Although the QZT-ICS framework offers a coherent integration of the four pillars identified in Section 3.5, several open research challenges remain. We articulate four that we consider most consequential, in roughly decreasing order of community readiness: hardware-attested TSN bridges (9.1), Digital Twin model poisoning (9.2), the brownfield quantum-classical transition (9.3), and the standardization gaps that limit framework adoption (9.4). Each subsection identifies concrete research questions and suggests directions for future work.

9.1 Hardware-Attested Trust at the TSN Bridge Level

The framework treats TSN bridges as trusted enforcement points: the gate-control-list augmentation of Section 5.3.2 assumes that the bridge hardware faithfully executes the security predicate alongside the time-aware shaping schedule. This assumption is reasonable for greenfield deployments with commodity bridges from reputable vendors, but it does not hold against an A1 nation-state adversary with the capability to inject firmware-level compromise, nor against the A4 supply-chain compromise of bridge silicon at manufacture. The federated DT consensus of Section 5.2.2 mitigates compromised PEPs by detecting decision discrepancies at the application layer, but it is an inefficient defense; a more efficient approach would extend hardware-rooted attestation from the server-class TEE space (Intel SGX, AMD SEV-SNP, ARM CCA) into the TSN switching domain.

Three sub-problems are open. First, the attestation primitive itself: TSN bridges typically host their data plane in ASIC or FPGA logic rather than in CPU-executable code, so the conventional measured-boot pipeline does not apply directly. A bridge attestation must cover the GCL state, the trust-score cache contents, and the data-plane pipeline configuration without imposing intolerable performance overhead. Second, the verification cadence: the trust-score cache changes on millisecond timescales (Section 5.3.2), but attestation verification cannot occur at that rate without dedicated hardware paths. Recent work on incremental attestation - in which only deltas to attested state are re-verified - is a promising direction but is not yet specialized to TSN. Third, the supply chain: ensuring that the attesting firmware itself is authentic requires manufacturer trust anchors at the silicon layer, a degree of vertical integration that few vendors currently offer outside specialized defense-industrial contexts.

We treat this as a two-to-three-year research horizon with promising early indicators in vendor announcements (Cisco, Marvell, NXP) about attested data-plane switches and in academic work on programmable-switch security. We encourage academic engagement with the IEEE 802.1 working group to standardize bridge-level attestation primitives so that the QZT-ICS framework, when those primitives become available, has a clean integration path through a future revision of 802.1Qbv.

9.2 Digital Twin Model-Poisoning Defenses

The DT integrity attestation specified in Section 5.2.2 - build-provenance signing through SLH-DSA, runtime attestation through TEE, federated consensus for safety-critical decisions - addresses single-instance compromise but is markedly less effective against a sophisticated A1 adversary with multi-month dwell time. Such an adversary can poison the DT model through at least three mechanisms, none of which are mitigated by the current defenses.

First, gradual training-data manipulation. An adversary with persistent access to field telemetry can inject small, slow drifts that the DT incorporates into its baseline. Over weeks, the modeled "normal" envelope expands to include the adversary's preferred operating point, after which a stealthy attack on the physical process produces residuals indistinguishable from modeled behavior. This is the cyber-physical analogue of the gradient-masking attack against machine-learning classifiers and is well-documented in the adversarial-ML literature, though not yet for the digital-twin setting.

Second, compromised model-curator privileges. The framework requires that model updates be signed by an authorized curator, but an A3 insider or an A1 adversary that has compromised curator credentials can issue legitimately-signed but adversarially-trained model updates. Build-provenance attestation defends against unsigned malicious updates; it does not defend against signed malicious updates from a compromised principal. This is a fundamental limitation of any signature-based provenance system.

Third, exploiting the residual-tolerance soft band. The trust algorithm permits a soft band where moderate residuals trigger MONITORED_PERMIT rather than DENY. An adversary who understands the band boundaries (which are policy-specific but discoverable through reconnaissance) can craft attacks that stay within the band indefinitely, generating no hard denials. Over time, as the behavioral-baseline component *B(t)* adapts to the persistent in-band activity, the attack normalizes and ceases to trigger even monitored alerts. The S-A-I-C priority inversion of Section 4.2 explicitly biases the framework toward availability in this band, which is operationally correct but adversarially exploitable.

The research challenge is multifaceted. Adversarial-training-data detection in the digital-twin context requires advances in change-point detection for high-dimensional time series under adversarial conditions; current methods, surveyed by Zhang et al. [28], are not yet robust enough for production. Byzantine-robust federated DT aggregation requires protocols that tolerate up to f compromised DT instances out of n while preserving the millisecond timing budgets that the framework requires for safety-critical decisions; the trade-off between Byzantine fault tolerance and timing determinism is fundamental and underexplored. Differential-privacy techniques may help to bound the influence of any single telemetry source on the model, though the latency cost is significant and the privacy budget composes poorly across long-running deployments. Finally, formal verification of model updates against safety invariants is an active research area but has not yet been applied at the scale or update cadence required for industrial DT deployments. We view this as the most academically tractable of the four open challenges and encourage focused work at the intersection of adversarial machine learning, distributed systems, and control-theoretic safety.

9.3 Quantum-Classical Hybrid Transitions on Long-Lived ICS Assets

The framework's crypto-agility infrastructure (Section 5.4.4) provides the technical mechanism for replacing primitives without firmware updates, but the operational, economic, and timing aspects of the global transition from classical to post-quantum cryptography in ICS remain open research questions. Three sub-problems deserve attention.

The first open question is migration sequencing. With hundreds of thousands of devices in a typical large industrial deployment, the order in which devices are upgraded affects both security and operational risk. A naïve approach upgrades the most critical assets first; a risk-managed approach upgrades the most exposed assets first; an availability-managed approach upgrades the easiest-to-roll-back assets first. The three orderings can produce very different transition trajectories with very different residual-risk profiles at intermediate states, and the research community has not yet articulated a comprehensive comparative methodology. Our prior cloud-context work [62] sketches a migration state machine for enterprise multi-cloud workloads but does not address the asset-criticality and inter-asset-dependency dimensions that dominate the industrial setting. Future work should develop quantitative migration-sequencing models that account for asset criticality, exposure, replacement cost, rollback feasibility, and inter-asset dependencies - ideally producing operator-tunable optimization frameworks rather than one-size-fits-all sequencing prescriptions.

The second is the stranded-asset problem. Some installed ICS equipment will reach end-of-service before PQC firmware becomes available; some equipment is no longer supported by its original manufacturer and will never receive PQC updates regardless of timeline. The bumps-in-the-wire approach of Section 5.4.3 is a partial answer, but the equipment-strandedness will produce a long-tailed transition with measurable security gaps for decades. Open research questions include how to quantify the residual quantum-vulnerable surface in a heterogeneous brownfield deployment; how to optimize bumps-in-the-wire placement for maximum surface reduction at minimum cost; and how to retire stranded assets without operational disruption to the integrated production system. The asset-replacement economics of ICS - where equipment is depreciated over fifteen to twenty-five years and capital budgets are committed years in advance - make this a much harder problem than the analogous IT-side transition.

The third is harvest-now-decrypt-later attack archaeology. Once a cryptographically relevant quantum computer exists, traffic captured today may be decrypted retrospectively. The ICS sector has not yet undertaken a serious analysis of which currently classically-encrypted traffic flows have sensitivity windows extending past the expected CRQC arrival date, and operators are therefore making implicit assumptions about future decryption that they have not articulated and may not have priced. Research should develop methodologies for sensitivity-window analysis specific to ICS data flows, with reference taxonomies for process parameters, control algorithms, intellectual property contained in setpoints, and credentials. Joseph et al. [43] provide a high-level treatment of the transition; the ICS-specific operational research remains undeveloped. We treat this as a five-to-ten-year

research and policy challenge, with academia, regulators (CISA, ENISA, national CERTs), and equipment manufacturers each playing essential roles.

9.4 Standardization Gaps

The framework rests on a substantial body of existing standards - NIST SP 800-207 [1], IEC 62443 [9], FIPS 203-205 [6]-[8], the IEEE 802.1 TSN family, 3GPP TS 33.501, and NIST CSF 2.0 - but several standardization gaps limit the framework's adoption today.

First, IEC 62443 does not yet explicitly require Zero Trust. The standard's foundational requirements map well to ZT capabilities (Table 7), but the standard was written before the ZT framework was articulated and so does not name ZT as a normative architectural choice. An IEC 62443 update that explicitly requires ZT-aligned architectures at SL-3 and above would significantly accelerate adoption by giving auditors and operators a shared vocabulary. Such an update is being discussed within the IEC TC 65 working group, but a timeline has not been announced.

Second, the IEEE 802.1 TSN family does not specify security-predicate extensions to the gate-control list. The QZT-ICS extension of Section 5.3.2 is technically implementable but not standardized; a vendor implementing it today would do so as a proprietary extension, foreclosing multi-vendor interoperability and complicating audit. A standardized augmentation - tentatively "802.1Qbv-SEC," or an addendum to 802.1CB - would enable cross-vendor implementation and accelerate hardware adoption. Recent IEEE discussions have raised this gap, but formal standardization work has not yet begun.

Third, 3GPP TS 33.501 does not yet specify slice-level Zero Trust integration. The standard treats slice security primarily as an isolation property and addresses inter-slice authentication via the AUSF, but it does not articulate ZT trust algorithms at the slice level. A standardized profile for ZT-in-5G-slices would clarify operational expectations and enable conformance testing. The 3GPP SA3 working group has acknowledged this gap in recent documents but has not initiated formal work item.

Fourth, NIST is currently drafting SP 800-207B, an industrial-specific companion to SP 800-207. The draft addresses ICS but, as of this writing, does not yet integrate the four-pillar synthesis identified in this paper. We anticipate that an updated draft incorporating community feedback - particularly on the Digital Twin as PDP - will be a more useful reference for industrial deployments than SP 800-207 alone, and we encourage academic and industrial submitters to engage with the NIST drafting process.

Fifth, sector-specific guidance is uneven. The energy sector (through NERC CIP) and the financial sector have detailed ZT guidance specific to their threat models. Water, transportation, and manufacturing have less; the CISA OT PQC guidance [4] is sector-neutral and does not address the deployment differences between, say, a water-treatment plant and a steel mill. Sector-specific deployment guides grounded in the QZT-ICS framework would help operators in less-served sectors translate the abstract architecture into operational practice.

We treat the standardization gaps as the most actionable of the four open challenges. Standards are not research per se, but the research community has a critical role in articulating the empirical evidence that justifies standardization decisions and in proposing concrete architectural choices through publications such as this one. The framework presented here, refined through community evaluation and revision, is intended to contribute to that empirical base.

10. Conclusion

This paper has presented QZT-ICS, a Zero Trust framework for IT-OT converged Industrial Control Systems that integrates four independently necessary capabilities into a single coherent architecture: Zero Trust micro-segmentation with the Digital Twin as Policy Decision Point, deterministic 5G/TSN transport with Policy Enforcement Points embedded at the slice and bridge layers, post-quantum cryptography as a crypto-agile substrate, and a continuous verification loop that closes the architecture into an adaptive system. The literature review of Section 3 demonstrated that no prior published work occupies this four-pillar intersection; the framework specification of Section 5 demonstrated that the integration is technically coherent within the latency and resource budgets of contemporary industrial deployments; and the discussion of Section 8 demonstrated that it is incrementally deployable through a five-level maturity model, regulatory-aligned with IEC 62443 and NIST CSF 2.0, and honestly limited.

The framework's principal conceptual contribution is the use of the Digital Twin as the locus of process-aware trust evaluation. Identity-only and posture-only Zero Trust frameworks cannot reason about whether a syntactically valid industrial command is semantically safe given the current state of the controlled process. The DT-as-PDP, with its trust score combining identity, device posture, process-state residual, behavioral conformance, and threat intelligence (Algorithm 1), closes this gap by making the

trust algorithm physically aware. The empirical evaluation in Section 7 demonstrates that this process awareness does not sacrifice the performance budgets that make Zero Trust deployable in industrial settings.

The framework's principal architectural contribution is the placement of Policy Enforcement Points at the 5G slice and TSN bridge layers, where enforcement can be applied at line rate without violating the deterministic timing guarantees that justify TSN's adoption in industrial settings. The gate-control-list augmentation with a security predicate (Section 5.3.2) is the simplest possible technical mechanism for this integration; we expect the IEEE 802.1 working group to consider such an augmentation in a future revision, at which point QZT-ICS bridges will have a standards-conformant implementation path.

The framework's principal cryptographic contribution is the explicit treatment of post-quantum primitives as a crypto-agile substrate, used today through hybrid modes and ready for pure-PQC operation as the installed base refreshes. Mosca's harvest-now-decrypt-later argument [44] applies to ICS today, and the framework's specification of ML-KEM, ML-DSA, SLH-DSA, hybrid X25519+ML-KEM, and the lifecycle management of Section 5.4.4 provides operators with a concrete migration path that begins now rather than at some indefinite future quantum-arrival date.

Three threads of future work are necessary to bring the framework to industrial maturity. First, the empirical evaluation of Section 7 must be expanded to additional sectors (currently only water-treatment is evaluated in detail) and ultimately to real industrial deployments, building on the simulated-testbed results presented in this paper. Second, the four open research challenges of Section 9 - hardware-attested TSN bridges, Digital Twin model-poisoning defenses, the brownfield quantum-classical transition, and the standardization gaps - require coordinated work across academia, industry, and standards bodies. Third, sector-specific deployment guides grounded in the framework are needed to translate the abstract architecture into operational playbooks for water, energy, manufacturing, and transportation operators, particularly those without the in-house security engineering capacity of large utilities and tier-one manufacturers.

Cybersecurity in industrial control systems is a moving target. The threat actors are sophisticated and well-resourced; the infrastructure is long-lived and operationally constrained; and the safety-of-life implications of failure are real. The regulatory environment is responding - the EU Cyber Resilience Act, NIS2, and the CISA OT post-quantum guidance all point toward an architecture-level transformation in the second half of this decade - and the underlying hardware is now beginning to arrive, with the NXP i.MX 94 announcement of November 2024 being the first commercial industrial system-on-chip to combine post-quantum cryptographic acceleration, time-sensitive networking, and edge AI in a single part. The convergence of regulatory pressure, hardware availability, and academic understanding makes the present moment unusually favorable for serious work on industrial Zero Trust. The QZT-ICS framework presented here is one contribution toward a security posture appropriate to that moment. We hope it will serve as a starting point for both deployment-driven engineering and security-research investigation, and we welcome critique and extension from both communities.

Acknowledgments

The author thanks the faculty and research community at Harrisburg University of Science and Technology for their ongoing support of this work. The author is also grateful to colleagues across the IEEE community who provided informal review and discussion on early formulations of the four-pillar synthesis, and to the broader research community whose published testbeds and benchmarks - particularly the SWaT water-treatment dataset from the iTrust Centre, Singapore University of Technology and Design, and the liboqs / PQClean reference implementations - made the analytical evaluation in this paper possible. The author also acknowledges Microsoft Azure for confidential computing infrastructure access through the academic program, which enabled the cloud-side testbed components described in Section 6.1. This work received no external funding. The author declares no conflict of interest.

References

A. Foundational Standards and Government Publications

[1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST Special Publication 800-207, Aug. 2020. doi: 10.6028/NIST.SP.800-207.

[2] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to Operational Technology (OT) Security," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST Special Publication 800-82 Revision 3, Sep. 2023. doi: 10.6028/NIST.SP.800-82r3.

[3] J. Voas, P. Mell, P. Laplante, and V. Piroumian, "Security and Trust Considerations for Digital Twin Technology," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST Internal Report NIST IR 8356, Feb. 2025. doi: 10.6028/NIST.IR.8356.

[4] Cybersecurity and Infrastructure Security Agency (CISA), "Post-Quantum Considerations for Operational Technology," U.S. Department of Homeland Security, Washington, DC, USA, Oct. 2024.

[5] Cybersecurity and Infrastructure Security Agency (CISA), "Zero Trust Considerations for Critical Infrastructure Operational Technology (OT)," U.S. Department of Homeland Security, Washington, DC, USA, 2026.
[6] NIST, "Module-Lattice-Based Key-Encapsulation Mechanism Standard," Federal Information Processing Standards Publication FIPS 203, Aug. 2024. doi: 10.6028/NIST.FIPS.203.
[7] NIST, "Module-Lattice-Based Digital Signature Standard," Federal Information Processing Standards Publication FIPS 204, Aug. 2024. doi: 10.6028/NIST.FIPS.204.
[8] NIST, "Stateless Hash-Based Digital Signature Standard," Federal Information Processing Standards Publication FIPS 205, Aug. 2024. doi: 10.6028/NIST.FIPS.205.
[9] Security for Industrial Automation and Control Systems - Part 3-3: System Security Requirements and Security Levels, IEC Standard 62443-3-3, International Electrotechnical Commission, Geneva, Switzerland, 2024.
[10] Cloud Security Alliance Zero Trust Working Group, "Zero Trust Guidance for Critical Infrastructure," Cloud Security Alliance, Seattle, WA, USA, 2024.
[11] D. Kuvshinkova, T. Anglim Kreikemeier, R. Mengers, and L. Christmas, "IT, OT, and ZT: Implementing Zero Trust in Industrial Control Systems," Carnegie Mellon University Software Engineering Institute, Pittsburgh, PA, USA, Jul. 2022.

B. B. Core Zero Trust Architecture and Survey Papers

[12] N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss, "Zero Trust Architecture (ZTA): A Comprehensive Survey," *IEEE Access*, vol. 10, pp. 57143-57179, 2022, doi: 10.1109/ACCESS.2022.3174679.
[13] S. Mushtaq, M. Mohsin, and M. M. Mushtaq, "A Systematic Literature Review on the Implementation and Challenges of Zero Trust Architecture Across Domains," *Sensors*, vol. 25, no. 19, p. 6118, Oct. 2025, doi: 10.3390/s25196118.
[14] A. Wylde, "Zero trust: Never trust, always verify," in *Proc. Int. Conf. on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA)*, Dublin, Ireland, Jun. 2021, pp. 1-4, doi: 10.1109/CyberSA52016.2021.9478244.
[15] E. Bertino, "Zero trust architecture: Does it help?" *IEEE Security & Privacy*, vol. 19, no. 5, pp. 95-96, Sep./Oct. 2021, doi: 10.1109/MSEC.2021.3091195.
[16] S. Teerakanok, T. Uehara, and A. Inomata, "Migrating to Zero Trust Architecture: Reviews and Challenges," *Security and Communication Networks*, vol. 2021, Art. no. 9947347, May 2021, doi: 10.1155/2021/9947347.

C. C. Zero Trust for ICS, OT, and IT-OT Convergence

[17] X. Feng and S. Hu, "Cyber-Physical Zero Trust Architecture for Industrial Cyber-Physical Systems," *IEEE Transactions on Industrial Cyber-Physical Systems*, vol. 1, pp. 394-405, 2023, doi: 10.1109/TICPS.2023.3333850.
[18] F. Federici, D. Martintoni, and V. Senni, "A Zero-Trust Architecture for Remote Access in Industrial IoT Infrastructures," *Electronics*, vol. 12, no. 3, p. 566, Jan. 2023, doi: 10.3390/electronics12030566.
[19] R. Sims, "Implementing a Zero Trust Architecture for ICS/SCADA Systems," M.S. thesis, Dakota State University, Madison, SD, USA, 2024.
[20] F. Lv, H. Wang, Y. Yang, and Z. Li, "Asynchronous federated learning based zero trust architecture for the next generation industrial control systems," *Computer Networks*, vol. 20, Art. no. 111459, 2025.
[21] M. F. Aljuaid and S. H. Ahmed, "Bridging IT and OT: Cybersecurity Risks, Zero-Trust Solutions, and Industrial Resilience in Oil & Gas Downstream," *ASEAN Journal of Scientific and Technological Reports*, 2025.
[22] S. K. Reddy and M. Banerjee, "A Zero Trust Approach for the Cybersecurity of Industrial Control Systems," in *Proc. IEEE Int. Conf. on Smart Grid Communications*, 2022.
[23] A. R. Bilipelli, "A novel and secure artificial intelligence enabled zero trust intrusion detection in industrial internet of things architecture," *Scientific Reports*, vol. 15, Art. no. 11738, 2025, doi: 10.1038/s41598-025-11738-9.

D. D. Zero Trust and Digital Twins for ICS

[24] R. A. Vega Vega, P. Chamoso, A. Briones, P. Gonzalez-Briones, F. de la Prieta, and J. M. Corchado, "Enabling a Zero Trust Architecture in Smart Grids Through a Digital Twin," in *Distributed Computing and Artificial Intelligence, 18th International Conference (DCAI 2021)*, Lecture Notes in Networks and Systems, vol. 327, Cham, Switzerland: Springer, 2022, pp. 1-10, doi: 10.1007/978-3-030-86261-9_1.
[25] A. Sayghe, "Digital Twin-Driven Intrusion Detection for Industrial SCADA: A Cyber-Physical Case Study," *Sensors*, vol. 25, no. 16, Art. no. 4963, Aug. 2025, doi: 10.3390/s25164963.
[26] N. Anumbe, C. Saidu, and S. Akhilesh, "Digital Twin-Enabled Vulnerability Assessment with Zero Trust Policy Enforcement in Smart Manufacturing Cyber-Physical System," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, vol. 9, no. 6, pp. 1-18, Nov.-Dec. 2023.
[27] Mishra, Shailendra & Aldafas, Tariq & Alshammari, Naif., "A zero-trust digital twin framework for privacy-preserving multi-dataset intrusion detection in industrial IoT with lightweight blockchain auditing," *Scientific Reports*, vol. 16, Art. no. 42041, 2026, doi: 10.1038/s41598-026-42041-w.

[28] Z. Zhang, M. Fang, M. Chen, G. Li, X. Lin, and Y. Liu, "Securing Distributed Network Digital Twin Systems Against Model Poisoning Attacks," arXiv preprint, arXiv:2407.01917, Jul. 2024.
[29] P. Li, A. Aijaz, T. Farnham, S. Gufran, and S. Chintalapati, "Demo: A Digital Twin of the 5G Radio Access Network for Anomaly Detection Functionality," arXiv preprint, arXiv:2308.15973, Aug. 2023.

E. E. Zero Trust with 5G and Time-Sensitive Networking

[30] S. Sun, M. Repeta, M. Healy, V. Nandall, E. Fung, and C. Thomas, "Towards 5G Zero Trusted Air Interface Architecture," arXiv preprint, arXiv:2211.03776, Nov. 2022.
[31] X. Chen, W. Feng, N. Ge, and Y. Zhang, "Zero Trust Architecture for 6G Security," *IEEE Network*, vol. 38, no. 4, pp. 224-232, Jul. 2024, doi: 10.1109/MNET.2023.3326356.
[32] H. Wang, L. Chen, and J. Liu, "Toward zero trust in 5G Industrial Internet collaboration systems," *Digital Communications and Networks*, 2024, doi: 10.1016/j.dcan.2024.02.003.
[33] A. Bin Sediq, A. A. Yazici, A. Khisti, and H. Yanikomeroglu, "Toward Zero Trust Security in 5G Open Architecture Network Slices," in *Proc. IEEE Military Communications Conf. (MILCOM)*, Rockville, MD, USA, Oct. 2022, pp. 1-6, doi: 10.1109/MILCOM55135.2022.10017474.
[34] L. Yu, F. Zhang, J. Yuan, H. Zhang, and Y. Sang, "Research on 5G-Based Zero Trust Network Security Platform," in *Proc. Int. Conf. on Frontier Computing (FC 2022)*, Lecture Notes in Electrical Engineering, vol. 1031, Singapore: Springer, 2023, pp. 345-356, doi: 10.1007/978-981-99-3300-6_39.
[35] J. Jiang, S. Jin, X. Li, K. Zhang, and B. Sun, "A Zero-Touch Dynamic Configuration Management Framework for Time-Sensitive Networking (TSN)," *Entropy*, vol. 27, no. 6, Art. no. 584, May 2025, doi: 10.3390/e27060584.
[36] L. L. Bello, M. Ashjaei, G. Patti, and M. Behnam, "Time-Sensitive Networking (TSN) for Industrial Automation: Current Advances and Future Directions," *ACM Computing Surveys*, vol. 57, no. 2, Art. no. 30, Oct. 2024, doi: 10.1145/3695248.
[37] M. Seliem, D. Pesch, U. Roedig, and C. Sreenan, "Resilient Time-Sensitive Networking for Industrial IoT: Configuration and Fault-Tolerance Evaluation," arXiv preprint, arXiv:2507.11250, Jul. 2025.
[38] M. Li, L. Deng, and Y. S. Han, "An Input-Queueing TSN Switching Architecture to Achieve Zero Packet Loss for Timely Traffic," arXiv preprint, arXiv:2206.09759, Jun. 2022.
[39] S. R. Pokhrel, "AI-enabled cybersecurity framework for future 5G wireless infrastructures," *PLOS One*, vol. 19, no. 8, Art. no. e0307354, 2024.
[40] G. Mantas, N. Komninos, J. Rodriguez, E. Logota, and H. Marques, "Testbed and Software Architecture for Enhancing Security in Industrial Private 5G Networks," arXiv preprint, arXiv:2507.20873, Jul. 2025.

F. F. Zero Trust and Post-Quantum Cryptography for ICS

[41] J. Oliva del Moral, A. deMarti i Olius, G. Vidal, P. M. Crespo, and J. Etxezarreta Martinez, "Cybersecurity in Critical Infrastructures: A Post-Quantum Cryptography Perspective," arXiv preprint, arXiv:2401.03780, Jan. 2024.
[42] S. Zhang and C. Zheng, "Quantum Secure Direct Communication Technology-Enhanced Time-Sensitive Networks," *Entropy*, vol. 27, no. 3, Art. no. 221, Feb. 2025, doi: 10.3390/e27030221.
[43] D. Joseph, R. Misoczki, M. Manzano, J. Tricot, F. D. Pinuaga, O. Lacombe, S. Leichenauer, J. Hidary, P. Venables, and R. Hansen, "Transitioning organizations to post-quantum cryptography," *Nature*, vol. 605, pp. 237-243, May 2022, doi: 10.1038/s41586-022-04623-2.
[44] M. Mosca, "Cybersecurity in an Era with Quantum Computers: Will We Be Ready?" *IEEE Security & Privacy*, vol. 16, no. 5, pp. 38-41, Sep./Oct. 2018, doi: 10.1109/MSP.2018.3761723.
[45] L. Chen, S. Jordan, Y.-K. Liu, D. Moody, R. Peralta, R. Perlner, and D. Smith-Tone, "Report on Post-Quantum Cryptography," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST Internal Report NISTIR 8105, Apr. 2016. doi: 10.6028/NIST.IR.8105.

G. ICS, SCADA, and IT-OT Convergence Security

[46] 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.
[47] R. Langner, "Stuxnet: Dissecting a Cyberwarfare Weapon," *IEEE Security & Privacy*, vol. 9, no. 3, pp. 49-51, May/Jun. 2011, doi: 10.1109/MSP.2011.67.
[48] J. Slay and M. Miller, "Lessons Learned from the Maroochy Water Breach," in *Critical Infrastructure Protection*, IFIP International Federation for Information Processing, vol. 253, Boston, MA, USA: Springer, 2007, pp. 73-82, doi: 10.1007/978-0-387-75462-8_6.
[49] M. Conti, D. Donadel, and F. Turrin, "A Survey on Industrial Control System Testbeds and Datasets for Security Research," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 4, pp. 2248-2294, 4th Quart. 2021, doi: 10.1109/COMST.2021.3094360.

[50] P. R. Vamsi and B. T. Rao, "Securing Industrial Control Systems and SCADA Networks: Protocol Weaknesses, Standards, Emerging Defenses, and Future Research Directions," in *Advances in Cybersecurity and Resilience*, Cham, Switzerland: Springer, 2024, doi: 10.1007/978-3-032-19300-1_22.

[51] T. M. Fernandez-Carames and P. Fraga-Lamas, "A Review on the Application of Blockchain to the Next Generation of Cybersecure Industry 4.0 Smart Factories," *IEEE Access*, vol. 7, pp. 45201-45218, Apr. 2019, doi: 10.1109/ACCESS.2019.2908780.

[52] S. Sivamohan, S. S. Sridhar, and S. Krishnaveni, "TEA-EKHO-IDS: An intrusion detection system for industrial CPS with trustworthy explainable AI and enhanced Krill herd optimization," *Peer-to-Peer Networking and Applications*, vol. 16, no. 4, pp. 1993-2021, Jul. 2023, doi: 10.1007/s12083-023-01493-x.

[53] 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.

G. H. Digital Twins, 5G/TSN, and Industry 4.0 (Supporting Works)

[54] A. Fuller, Z. Fan, C. Day, and C. Barlow, "Digital Twin: Enabling Technologies, Challenges and Open Research," *IEEE Access*, vol. 8, pp. 108952-108971, Jun. 2020, doi: 10.1109/ACCESS.2020.2998358.

[55] L. U. Khan, W. Saad, D. Niyato, Z. Han, and C. S. Hong, "Industrial Digital Twins at the Nexus of NextG Wireless Networks and Computational Intelligence: A Survey," arXiv preprint, arXiv:2108.04465, Aug. 2021.

[56] Y. Hong, J. Wu, and R. Morello, "LLM-Twin: mini-giant model-driven beyond 5G digital twin networking framework with semantic secure communication and computation," *Scientific Reports*, vol. 14, no. 1, Art. no. 19065, 2024.

[57] N. Finn, "Introduction to Time-Sensitive Networking," *IEEE Communications Standards Magazine*, vol. 2, no. 2, pp. 22-28, Jun. 2018, doi: 10.1109/MCOMSTD.2018.1700076.

[58] J. Sachs, K. Wallstedt, F. Alriksson, and G. Eneroth, "Boosting Smart Manufacturing with 5G Wireless Connectivity," *Ericsson Technology Review*, vol. 96, no. 2, pp. 24-35, Feb. 2019.

H. I. Industry, Market, and Policy Reports

[59] Industrial Cyber, "New CISA guidance outlines zero trust roadmap for OT environments facing legacy constraints and growing attack surfaces," Apr. 2026.

[60] European Union, "Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union (NIS2 Directive)," *Official Journal of the European Union*, L 333, pp. 80-152, Dec. 2022.

[61] European Union, "Regulation (EU) 2024/2847 of the European Parliament and of the Council on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act)," *Official Journal of the European Union*, L series, Nov. 2024.

[62] S. M. Qadri, "Operationalizing Post-Quantum Cryptography in Multi-Cloud Environments," *Frontiers in Computer Science and Artificial Intelligence*, vol. 5, no. 7, pp. 12-24, May 2026, doi: 10.32996/fcsai.2026.5.7.2.

[63] A. Greenberg, *Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers*. New York, NY: Doubleday, 2019. (Comprehensive technical and operational analysis of NotPetya and related campaigns.)

[64] U.S. Computer Emergency Readiness Team (US-CERT), "Alert TA17-181A: Petya Ransomware," Cybersecurity and Infrastructure Security Agency (CISA), Washington, DC, Jul. 2017, revised Feb. 2018. [Online]. Available: https://www.cisa.gov/news-events/cybersecurity-advisories/aa17-181a