Volume 11, Issue 8, August – 2026 International Journal of Innovative Science and Research Technology
ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug612

Zero Trust Architecture for Industrial IoT Networks

Mohammed Ibrahim Abba[1]; Womeodu Blessing[2]; Raymond Ternenge Igbudu[3]

[1,2,3] Center for Cyberspace Studies
Nasarawa State University Keffi, Nigeria

Publication Date: 2026/08/31

Abstract: Industrial Internet of Things (IoT) networks connect sensors, programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) systems, enterprise platforms, and cloud services to support automation, monitoring, predictive maintenance, and data-driven decision-making. This connectivity also expands the attack surface of operational technology (OT), particularly where legacy equipment, remote access, heterogeneous protocols, and stringent availability and safety requirements coexist. This paper develops a Zero Trust Architecture (ZTA) tailored to IIoT environments. The proposed architecture combines strong device and user identity, continuous verification, least-privilege and attribute-based access control, micro-segmentation, secure gateways for legacy devices, encrypted communication, continuous monitoring, and risk-adaptive policy enforcement. A simulation-oriented evaluation framework is specified using a representative industrial network comprising field devices, PLCs, SCADA/HMI systems, an edge gateway, enterprise resources, and a remote maintenance user. The evaluation is designed to compare authorized and unauthorized access, lateral-movement attempts, malicious commands, and controller or gateway failures using security and performance metrics. The architecture is intended to contain compromise while preserving the availability and timing requirements of industrial processes. Importantly, the manuscript distinguishes the proposed evaluation framework from experimentally measured results: numerical performance values are not presented as empirical findings unless generated by an executable testbed. The study concludes that Zero Trust is most practical in IIoT when implemented incrementally, with OT safety and availability treated as first-class requirements and legacy assets protected through compensating controls rather than forced modernization.

Keywords: *Zero Trust Architecture; Industrial Internet of Things; Operational Technology; Micro-Segmentation; Access Control; PLC; SCADA; Network Security; Legacy Systems; Remote Access.*

How to Cite: Mohammed Ibrahim Abba; Womeodu Blessing; Raymond Ternenge Igbudu (2026) Zero Trust Architecture for Industrial IoT Networks. *International Journal of Innovative Science and Research Technology*,11(8), 2015-2020. https://doi.org/10.38124/ijisrt/26aug612

I. INTRODUCTION

The Industrial Internet of Things (IIoT) has transformed industrial operations by connecting sensors, actuators, PLCs, supervisory systems, engineering workstations, enterprise applications, and cloud services. The resulting convergence of information technology (IT) and operational technology (OT) supports real-time monitoring, automation, predictive maintenance, remote diagnostics, and analytics. However, connectivity also increases the number of pathways through which an adversary can reach systems that directly influence physical processes. NIST guidance emphasizes that OT security must account for performance, reliability, and safety requirements in addition to conventional cybersecurity objectives. [1]

Traditional perimeter-based security assumes that systems inside a protected network are relatively trustworthy. Zero Trust rejects this assumption. NIST SP 800-207 defines Zero Trust as a security approach in which implicit trust is not granted solely because a user or asset is located within a particular network boundary; authentication and authorization are treated as explicit functions before access to resources is established. [2] This principle is especially relevant to IIoT because industrial environments increasingly include remote users, vendor connections, cloud services, wireless devices, and legacy components.

The central challenge is not simply to place Zero Trust controls around an industrial network. Security mechanisms must be introduced without disrupting deterministic

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug612

communication, safety functions, availability, maintenance workflows, or legacy protocols. The proposed architecture therefore adopts a defense-in-depth model in which identity, policy, network segmentation, secure gateways, monitoring, and application-level controls reinforce one another.

This paper has four objectives: (1) identify the principal technical and operational challenges of applying Zero Trust to IIoT; (2) propose a multilayer architecture for device, network, and application protection; (3) define a reproducible simulation-based evaluation methodology for remote maintenance and attack-containment scenarios; and (4) establish security and performance metrics through which the architecture can be assessed. The work builds on NIST Zero Trust principles and OT-security guidance, as well as recent research on Zero Trust remote access and flexible IIoT architectures. [1]–[4]

II. BACKGROUND AND ZERO TRUST SECURITY PRINCIPLES

Zero Trust is a security paradigm centered on explicit verification, least privilege, continuous evaluation, and resource protection rather than implicit trust based on network location. NIST SP 800-207 identifies logical functions such as the Policy Engine, Policy Administrator, and Policy Enforcement Point as central components of a Zero Trust Architecture. [2]

➢ *Core Principles*

- Verify explicitly: each access request should be evaluated using available identity, device, resource, environmental, and risk information.
- Use least privilege: users and devices should receive only the permissions required for a defined task and for the minimum necessary duration.
- Assume breach: the architecture should limit the consequences of compromise through segmentation, constrained communication paths, monitoring, and rapid revocation.
- Continuous monitoring: device posture, behavior, network flows, and policy compliance should be monitored throughout a session.
- Protect resources rather than relying on network location: access should be tied to the specific resource and operation requested.

➢ *Zero Trust in OT and IIoT*

Applying these principles to OT requires adaptation. NIST SP 800-82 Rev. 3 notes that OT systems interact with the physical environment and therefore have security requirements that differ from many enterprise IT systems, particularly with respect to safety, reliability, availability, and timing. [1] A Zero Trust design for IIoT should consequently avoid indiscriminate inspection or authentication mechanisms that could interrupt control loops. Where a legacy device cannot perform modern authentication, compensating controls such as secure gateways, protocol-aware proxies, network segmentation, and monitoring can provide a security boundary around the device.

III. INDUSTRIAL IOT SECURITY CHALLENGES

➢ *Device Diversity and Resource Constraints:*

IoT environments contain sensors, PLCs, RTUs, HMIs, gateways, servers, and specialized controllers with widely different processing, memory, operating-system, and cryptographic capabilities. Some legacy assets cannot support contemporary endpoint agents or strong cryptographic protocols.

➢ *Real-Time, Availability and Safety Requirements:*

Industrial communication may have strict timing constraints, while unexpected downtime can interrupt production or create safety consequences. Security controls must therefore be engineered around process requirements rather than added without timing analysis.

➢ *Legacy and Proprietary Protocols:*

Protocols such as Modbus and older proprietary interfaces may lack contemporary authentication or encryption. Replacing these assets is often expensive or operationally impractical.

➢ *IT/OT Convergence and Remote Access:*

Remote maintenance, enterprise connectivity, cloud analytics, and vendor access introduce additional identities and pathways into OT environments. These pathways require strong authentication, authorization, monitoring, and time-bound access.

➢ *Lateral Movement:*

A compromised engineering workstation, gateway, or device can become a stepping stone toward other assets if the network is flat. Micro-segmentation is therefore important for reducing blast radius.

➢ *Policy and Lifecycle Complexity:*

A Zero Trust deployment requires device inventories, identity lifecycle management, policy governance, certificate renewal and revocation, logging, exception management, and coordination between IT, OT, engineering, safety, and security teams.

IV. RELATED AND EMPIRICAL LITERATURE REVIEW

NIST SP 800-207 provides the foundational Zero Trust architecture model, but it is intentionally abstract and not an IIoT-specific implementation guide. [2] For OT environments, NIST SP 800-82 Rev. 3 provides complementary guidance emphasizing the unique performance, reliability, and safety requirements of operational technology. [1] The combination of these documents provides a useful conceptual basis: Zero Trust

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug612

supplies the access-control paradigm, while OT security guidance constrains how those controls should be deployed.

Federici, Martintoni, and Senni proposed a Zero Trust architecture for remote access in IIoT infrastructures. Their work is particularly relevant to industrial maintenance because it considers multi-level authorization and enforcement at both network and edge/device levels. The authors describe an approach in which remote access is not equivalent to unrestricted network membership; access can instead be constrained to specific industrial resources and operations. [3]

Zanasi, Russo, and Colajanni proposed a flexible Zero Trust architecture for industrial IoT infrastructures that emphasizes micro-segmentation and software-defined enforcement. Their approach is relevant to IIoT because dynamic network policy can help adapt security controls to heterogeneous environments while reducing reliance on static perimeter boundaries. [4]

Basta et al. examined evaluation considerations for Zero Trust micro-segmentation, highlighting the trade-off between security granularity and operational/network complexity. [5] This trade-off is particularly important in IIoT because excessively coarse segmentation may permit lateral movement, whereas excessively fine segmentation can increase policy complexity and management overhead.

Køien examined Zero Trust principles for legacy components, reinforcing the importance of compensating controls where older assets cannot directly participate in modern security mechanisms. [6] Industry frameworks such as the Industrial Internet Consortium Trustworthiness Framework and the U.S. Department of Defense Zero Trust Reference Architecture further illustrate the convergence of identity, segmentation, monitoring, and continuous risk assessment in modern security programs. [7], [8]

The literature therefore indicates a consistent research direction: Zero Trust can be adapted to IIoT, but practical implementation depends on legacy integration, policy scalability, network segmentation, remote-access governance, and careful management of OT performance and safety. The present work extends this direction by integrating these elements into a single simulation-oriented architecture and by explicitly defining an evaluation methodology rather than treating illustrative numerical values as measured experimental results.

V. PROPOSED ZERO TRUST ARCHITECTURE FOR IOT

The proposed architecture follows a multilayer defense-in-depth model. It separates the Zero Trust control plane from industrial data/control paths while deploying enforcement points at gateways, network boundaries, and selected endpoints. The architecture contains six principal functional layers: identity and trust foundation; continuous authentication and posture assessment; policy decision and administration; policy enforcement and micro-segmentation; secure communication and legacy-device protection; and continuous monitoring and analytics.

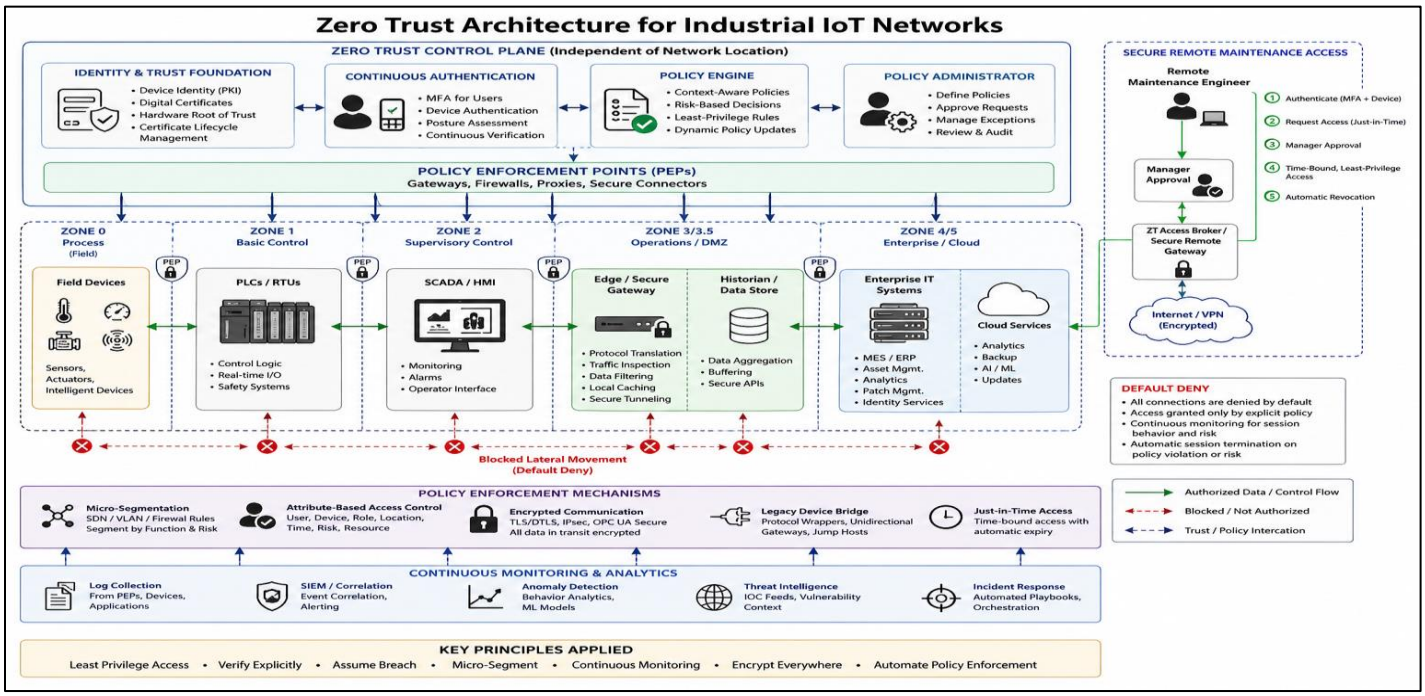

Fig 1. Proposed Zero Trust Architecture for Industrial IoT Networks.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug612

➢ *Device Identity and Authentication*

Each capable IIoT device should receive a unique cryptographic identity, preferably managed through a certificate-based public-key infrastructure. X.509 certificates, secure credential storage, hardware security modules, or trusted platform modules may be used according to device capability. Human users should use strong authentication, including multi-factor authentication for remote and privileged access. Devices that cannot implement these mechanisms should be placed behind trusted gateways that perform compensating authentication and policy enforcement.

➢ *Micro-Segmentation and Network Isolation*

The industrial environment should be divided into security zones according to function, process dependency, consequence, and communication requirements. A practical segmentation model can distinguish field devices, basic control, supervisory control, operations/DMZ, enterprise systems, and cloud services. Communication between zones should be explicitly allowed by policy rather than implicitly permitted. Software-defined networking may provide dynamic rule enforcement where the underlying infrastructure supports it. Segmentation should complement, rather than replace, established OT zoning and conduit practices.

➢ *Attribute-Based and Risk-Adaptive Access Control*

Access decisions should consider attributes such as user identity, device identity, role, resource, requested operation, location, time, device posture, process state, and current risk. For example, a maintenance engineer may be permitted to read diagnostic registers during a scheduled maintenance window but prohibited from changing control logic unless that specific operation has been authorized.

➢ *Secure Communication and Data Protection*

Where technically and operationally appropriate, communications should use authenticated and encrypted protocols. TLS or DTLS can protect supported communications, while secure industrial protocols such as OPC UA security mechanisms may be used where applicable. Legacy unencrypted protocols should be isolated and protected through secure gateways or protocol-aware tunnels. Encryption should not be introduced blindly into deterministic control paths; its processing and latency implications should be tested before deployment.

➢ *Legacy Device Integration*

Legacy PLCs and controllers should be protected through compensating controls rather than assumed to be trustworthy. A secure gateway can authenticate remote users, enforce operation-level policies, inspect traffic, log commands, and restrict communication to approved destinations. This approach creates a security boundary around devices that cannot themselves support modern identity mechanisms.

➢ *Continuous Monitoring and Anomaly Detection*

Telemetry from network flows, gateways, devices, identity systems, and policy decisions should feed a monitoring and security analytics capability. Baselines can be established for normal communication patterns, and deviations can trigger alerts, additional verification, session restriction, or revocation. Monitoring should support forensic investigation through protected and time-synchronized logs.

VI. METHODOLOGY

Because the supplied manuscript describes a simulation-based evaluation but does not provide the executable testbed, raw logs, repetitions, confidence intervals, or machine-generated measurements required to substantiate numerical experimental findings, the revised methodology treats the evaluation as a reproducible simulation protocol. Numerical values in the original draft that were presented as measured latency or security outcomes have therefore not been retained as empirical measurements. They should only be reported after execution of the testbed and analysis of its output.

➢ *Research Design*

A controlled simulation study should compare a baseline industrial network with the proposed Zero Trust configuration. The baseline represents conventional perimeter-oriented access, while the experimental condition implements device identity, policy enforcement, micro-segmentation, secure gateway controls, and continuous monitoring. The same traffic workloads and attack scenarios should be executed under both conditions.

➢ *Simulation Environment*

The proposed testbed comprises field sensors and actuators, PLC/RTU nodes, an HMI/SCADA layer, an edge security gateway, historian and enterprise servers, a remote maintenance workstation, and an SDN-capable virtual network. Mininet can emulate programmable network topology and OpenFlow control, while ns-3 can be used where packet-level simulation is required. Containerized services or virtual machines can represent industrial applications and security components.

➢ *Remote Maintenance Use Case*

- An authorized maintenance engineer authenticates using user credentials and an additional factor, while the originating device is evaluated.
- A maintenance request identifies the target industrial asset, requested operations, expected duration, and justification.
- A designated manager or authorized approver reviews the request.
- The Policy Engine evaluates identity, role, asset, time, process state, device posture, and risk attributes.
- If approved, a time-limited authorization token is issued and enforcement points create the minimum required communication path.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug612

- The secure gateway performs a second authorization check before forwarding commands.
- Network behavior and individual operations are monitored throughout the session.
- At expiration or upon policy violation, the token is revoked and temporary flows are removed.

➢ *Threat and Test Cases*

- Authorized maintenance access to the approved PLC.
- Invalid or expired credentials.
- Forged or expired access token.
- Access to a PLC outside the approved scope.
- Lateral scanning from a compromised engineering workstation.
- Unauthorized PLC write or control-logic modification.
- Anomalous traffic during an authorized session.
- Temporary failure of the policy service or SDN controller.
- Compromise of a legacy device protected by a secure gateway.

Table 1 Evaluation Metrics

Dimension	Metric	Interpretation
Security	Unauthorized access success rate	Lower is better.
Security	Blocked lateral-movement attempts	Higher containment is better.
Security	Policy-violation detection rate	Measures monitoring effectiveness.
Security	Mean time to revoke	Time from confirmed violation to termination.
Performance	Authentication/policy latency	Overhead from Zero Trust checks.
Performance	Flow provisioning time	Time to create authorized network paths.
Performance	Command latency	Delay added to industrial operations.
Availability	Continuity during control-plane failure	Measures resilient enforcement.
Scalability	Policy decision throughput	Capacity under increasing requests.

VII. EXPECTED EVALUATION OUTCOMES AND INTERPRETATION

The architecture is expected to reduce unauthorized access and lateral movement by replacing broad network trust with resource-specific authorization and segmentation. Authorized maintenance should remain possible while access to unrelated devices is denied. A compromised endpoint should have a smaller blast radius because communication with other segments is not automatically permitted.

The expected performance trade-off is additional processing and policy-decision overhead. This overhead should be measured rather than assumed to be negligible. Control-loop traffic should be evaluated separately from less time-sensitive maintenance and monitoring traffic.

The evaluation should also determine whether the architecture continues enforcing existing policies when the central policy service becomes unavailable. Cached policies, redundant controllers, and carefully defined fail-safe behavior can reduce the risk of a security component becoming a single point of failure.

VIII. DISCUSSION

The principal contribution of the proposed architecture is the integration of identity, segmentation, policy enforcement, legacy protection, and monitoring into an IIoT-oriented Zero Trust model. Zero Trust security cannot be reduced to authentication alone; a valid identity is insufficient if it can communicate with every device in the plant.

Micro-segmentation is central to containment, but segmentation policy must be based on actual industrial workflows. Excessive segmentation can create operational complexity, while insufficient segmentation can leave pathways for lateral movement. Policy design should be reviewed jointly by cybersecurity personnel, OT engineers, process owners, and safety specialists.

Legacy equipment remains a strong argument for a gateway-based deployment strategy. A secure gateway can provide an external enforcement layer, but gateways become high-value assets and should be hardened, redundantly deployed where necessary, monitored, patched, and incorporated into incident-response plans.

Remote maintenance deserves particular attention because third-party access can create a direct path into industrial environments. Just-in-time, time-limited access with explicit approval and operation-level restrictions is preferable to persistent VPN credentials or unrestricted remote accounts. Federici et al. demonstrate the relevance of multi-level authorization and gateway enforcement for this problem. [3]

Artificial intelligence and machine learning may improve anomaly detection and adaptive risk scoring, but they should be supporting mechanisms rather than sole authorization authorities. False positives, explainability, model drift, adversarial manipulation, and the consequences of automated denial in safety-critical operations must be evaluated.

The proposed framework is architecture- and simulation-oriented and therefore cannot establish real-world effectiveness without deployment or validated experimental data. Industrial

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug612

networks vary substantially across sectors, so laboratory results should not automatically be generalized to operational plants.

IX. PRACTICAL DEPLOYMENT ROADMAP

- Inventory assets, identities, communications, dependencies, safety functions, and remote-access pathways.
- Classify assets by criticality and risk; identify the highest-consequence and most exposed resources.
- Establish identity management for users and capable devices, including certificate lifecycle management.
- Introduce segmentation at high-risk boundaries and validate communication dependencies before enforcement.
- Deploy secure gateways around critical legacy assets that cannot support modern authentication.
- Implement just-in-time remote access with MFA, approval, least privilege, logging, and automatic expiration.
- Centralize telemetry and establish behavioral baselines and incident-response playbooks.
- Test policy enforcement, failover, recovery, and safety interactions in a controlled environment before production rollout.
- Expand Zero Trust incrementally and continuously review policies, exceptions, identities, vulnerabilities, and asset changes.

X. CONCLUSION

This paper presented a Zero Trust Architecture tailored to Industrial IoT networks. The proposed design combines cryptographic device identity, strong user authentication, continuous verification, least-privilege access, attribute-based policy decisions, micro-segmentation, secure gateways for legacy systems, protected communication, and continuous monitoring. The architecture addresses the central tension between stronger cybersecurity and the operational requirements of industrial systems.

A reproducible simulation methodology was defined around a remote-maintenance scenario and a set of attack and resilience tests. The revised manuscript deliberately avoids presenting unsupported numerical latency or security results as empirical findings. Such values should be generated from repeated testbed execution and reported with appropriate statistical analysis.

Zero Trust can provide a practical security direction for IIoT when implemented incrementally and adapted to OT realities. The most effective approach is to combine Zero Trust principles with OT-aware segmentation, safety engineering, legacy compensating controls, resilient enforcement, and continuous monitoring. Future research should validate the architecture in larger heterogeneous testbeds, investigate AI-assisted adaptive policy enforcement, evaluate post-quantum migration strategies, and establish sector-specific implementation profiles.

REFERENCES

[1]. K. Stouffer et al., “Guide to Operational Technology (OT) Security,” NIST Special Publication 800-82 Rev. 3, National Institute of Standards and Technology, Gaithersburg, MD, USA, Sep. 2023. doi: 10.6028/NIST.SP.800-82r3.

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

[3]. F. Federici, D. Martintoni, and V. Senni, “A Zero-Trust Architecture for Remote Access in Industrial IoT Infrastructures,” Electronics, vol. 12, no. 3, Art. no. 566, 2023. doi: 10.3390/electronics12030566.

[4]. C. Zanasi, S. Russo, and M. Colajanni, “Flexible Zero Trust Architecture for the Cybersecurity of Industrial IoT Infrastructures,” Ad Hoc Networks, vol. 156, Art. no. 103414, 2024.

[5]. N. Basta et al., “Towards a Zero-Trust Micro-segmentation Network Security Strategy: An Evaluation Framework,” in Proc. IEEE/IFIP Network Operations and Management Symposium (NOMS), 2022.

[6]. G. M. Køien, “Zero-Trust Principles for Legacy Components,” Wireless Personal Communications, vol. 121, pp. 1169–1186, 2021.

[7]. Industrial Internet Consortium, “The Industrial Internet of Things Trustworthiness Framework Foundations,” Industrial Internet Consortium, 2019.

[8]. U.S. Department of Defense Chief Information Officer, “Department of Defense Zero Trust Reference Architecture,” Version 2.0, Jul. 2022.