

TrustPulse-IoT: Context-Aware Adaptive Access Control for Heterogeneous IoT Networks

Pawar Pawan Appasaheb¹, Mr. R. R. Tajane², Prof. J. R. Mahajan³

Student, Department of Artificial Intelligence and Data Science¹

Assistant Professor, Department of Artificial Intelligence and Data Science²

HOD, Department of Artificial Intelligence and Data Science³

Adsul Technical Campus, Chas, Ahilyanagar, India

Abstract: *IoT deployments increasingly combine constrained sensors, gateways, legacy controllers, cloud services, and vendor-managed interfaces. This heterogeneity makes a single perimeter or static access rule insufficient because a device that was legitimate at enrollment may later exhibit an unsafe communication pattern. This paper proposes TrustPulse-IoT, an adaptive security architecture that converts device identity, behavioral stability, protocol compliance, asset criticality, and telemetry freshness into a continuously refreshed trust value at the edge. The trust value is used to select one of four actions—permit, permit-with-restrictions, quarantine, or deny—while a lightweight behavior profile is maintained locally to reduce dependence on centralized raw telemetry. The design is motivated by security gaps identified in the supplied IoT literature, particularly weaknesses involving device identity, access control, segmentation, insecure interfaces, data protection, and limited attention to continuous security state. Unlike a conventional checklist of controls, the proposed model links observation to policy enforcement and feedback. A formal trust function, decision algorithm, threat model, prototype plan, and evaluation protocol are presented. No experimental performance values are claimed; the paper defines a reproducible path for implementation and measurement. The proposed architecture is intended as an incremental research contribution that can be evaluated against a baseline IoT network using detection quality, unauthorized-access prevention, response latency, overhead, and false-positive rate.*

Keywords: adaptive access control, anomaly detection, edge security, Internet of Things, trust management

I. INTRODUCTION

The Internet of Things (IoT) has moved from isolated sensing systems to interconnected environments in which devices exchange measurements, commands, status information, and software updates across organizational and vendor boundaries. Security consequently depends on more than confidentiality of transmitted data. The identity of a device, the services it is expected to contact, the timing and frequency of its communication, and the consequences of a compromised endpoint all influence the risk of a request.

The three papers supplied as the reference material expose complementary aspects of this problem. The automation-focused work describes how IoT devices can create communication paths that do not fit traditional defense-in-depth assumptions and emphasizes device/API hardening and segmentation. The broader IoT review identifies recurring concerns such as integrity, encryption, privacy, and lifecycle/update management. The survey of security issues highlights the breadth of confidentiality, integrity, availability, authorization, access control, authenticity, non-repudiation, and interoperability concerns. [1]–[3]

The common weakness across these perspectives is not the absence of individual security controls; it is the difficulty of deciding how those controls should interact when the security state of a device changes. A device can possess valid credentials while simultaneously generating abnormal traffic. A network segment can be correctly configured while an



authorized device misuses an exposed service. A device can be authenticated correctly while its firmware or configuration has changed. These cases motivate a security model that evaluates each access decision using multiple pieces of evidence rather than identity alone.

This paper introduces TrustPulse-IoT, a context-aware adaptive access-control architecture. Its central research question is: Can an edge gateway use a compact, explainable trust value that combines identity, behavior, protocol compliance, asset criticality, and telemetry freshness to select proportional access responses for heterogeneous IoT devices without requiring continuous transmission of raw telemetry to a central service?

II. REFERENCE LITERATURE AND DESIGN MOTIVATION

A. Findings from the supplied papers

The first supplied paper studies IoT-enabled sensors in automation systems and points to a mismatch between conventional industrial defense-in-depth and heterogeneous IoT connectivity. It discusses uncontrolled or unexpected communication paths, device authentication, network segmentation, secure boot, code signing, role-based access control, public-key infrastructure, API hardening, and privacy/confidentiality. [1]

The second supplied paper provides a broad review of IoT applications and security challenges. Its treatment of integrity, encryption, privacy, update management, and multi-layer protection supports the need for controls that extend across endpoints, networks, cloud services, and applications rather than being concentrated at one layer. [2]

The third supplied paper surveys IoT security issues across past, present, and future research. It identifies confidentiality, integrity, authenticity, data security/privacy, authorization, access control, availability, non-repudiation, and interoperability as relevant dimensions and highlights uneven research attention across these dimensions. [3]

B. Additional research context

NISTIR 8259A defines a core baseline of technical cybersecurity capabilities for IoT devices, including device identification, configuration, data protection, logical access to interfaces, software update, cybersecurity state awareness, and device security. [4] NIST SP 800-207 frames zero trust around the principle that no implicit trust should be granted merely because a device or user is located inside a particular network. [5] Recent IoT security research also explores device profiling and anomaly detection at the edge, demonstrating the value of behavioral information for detecting deviations from normal device activity. [6]–[8]

III. PROBLEM STATEMENT AND RESEARCH GAP

Existing IoT security architectures commonly combine authentication, firewalls, segmentation, encryption, and intrusion detection. However, these mechanisms may operate as separate controls with separate decision logic. Such separation creates four practical problems.

First, static authorization does not necessarily reflect current device behavior. Second, anomaly detection may produce an alert without a corresponding enforcement action. Third, centralized analysis can increase telemetry volume and response latency for constrained or remote deployments. Fourth, a binary allow/deny policy may be too coarse for situations where a device is suspicious but not yet conclusively malicious.

The research gap addressed here is the integration of lightweight behavioral evidence with identity and environmental context into an explainable, continuously refreshed access decision. The proposed system does not claim that trust scoring itself is unprecedented; rather, its contribution is a specific architecture and decision workflow that joins five evidence dimensions with proportional enforcement and local feedback.

Gap in baseline designs	TrustPulse-IoT response
Identity treated as final trust	Identity confidence is one component of a multi-factor trust vector.
Alerts disconnected from policy	Trust state directly selects an enforcement action.
Static device permissions	Permissions are re-evaluated as context changes.
Centralized raw telemetry	Edge gateway summarizes behavior locally before forwarding evidence.



Binary allow/deny response


Four response levels provide proportional containment.

IV. TRUSTPULSE-IOT ARCHITECTURE

TrustPulse-IoT is organized as five cooperating planes: device identity, observation, trust computation, policy enforcement, and evidence/learning. The architecture is intentionally modular so that cryptographic identity, a conventional firewall, or an existing IDS can be retained rather than replaced.

The device identity plane validates the device credential and records device class, ownership, expected role, and lifecycle state. The observation plane collects compact features such as connection count, destination diversity, protocol usage, message-rate deviation, failed authentication events, configuration changes, and telemetry freshness. The trust engine converts these observations into a normalized score. The policy controller maps the score and contextual constraints to an action. The evidence plane stores decisions and selected features for audit and future profile updates.

TRUSTPULSE-IoT: Context-Aware Adaptive Security Plane



Continuous loop: Observe → Score → Decide → Enforce → Learn

Fig. 1. TrustPulse-IoT architecture. The diagram is an original figure created for this paper.

A. Device and Context Model

For device d at time t , the system maintains a state vector $x_d(t) = [I_d, B_d(t), P_d(t), C_d(t), F_d(t)]$

where I is identity confidence, B is behavioral stability, P is protocol compliance, C is contextual/asset risk, and F is telemetry freshness. All terms are normalized to $[0,1]$. A lower contextual-risk value represents a safer situation in the formulation below.

Identity confidence is derived from credential validity, device lifecycle state, and optional attestation evidence. Behavioral stability compares recent activity with a learned baseline. Protocol compliance measures whether observed services, ports, message types, and destinations are consistent with the device profile. Contextual risk represents factors such as asset criticality, maintenance windows, unusual network location, or an administrative risk flag. Telemetry freshness measures whether the evidence used for the decision is sufficiently recent.

B. Trust Score

The proposed trust score is

$$T_d(t) = w_I I_d + w_B B_d(t) + w_P P_d(t) + w_F F_d(t) + w_C (1 - C_d(t))$$

subject to $w_I + w_B + w_P + w_F + w_C = 1$.



The weights are policy parameters rather than universal constants. A safety-critical deployment may assign greater weight to identity and protocol compliance, while a consumer environment may place more emphasis on behavior. This parameterization makes the model testable instead of assuming one fixed weighting scheme.

C. Behavior Stability

A compact behavior profile can store expected ranges rather than raw payloads. For a feature vector $z_d(t)$, the deviation from the baseline can be calculated as

$$D_d(t) = \sum_k a_k |z_{\{d,k\}}(t) - \mu_{\{d,k\}}| / (\sigma_{\{d,k\}} + \epsilon),$$

where μ and σ are baseline statistics, a_k is a feature weight, and ϵ prevents division by zero. The stability component can then be defined as $B_d(t) = \exp(-D_d(t))$ after appropriate normalization. The profile can be updated only when the device is in a trusted state, reducing the risk that an attack rapidly becomes part of the normal baseline.

V. ADAPTIVE ACCESS DECISION ALGORITHM

The key design choice is to convert the trust value into proportional action rather than treating anomaly detection as a separate dashboard function. Four states are proposed:

- Permit: normal access to the device's approved services.
- Restricted: allow only a reduced service set, such as telemetry upload to the gateway.
- sQuarantine: isolate the device in a remediation segment while preserving diagnostic communication.
- Deny: block the session and create a high-priority security event.

Thresholds τ_1 , τ_2 , and τ_3 are deployment parameters with $\tau_1 > \tau_2 > \tau_3$. The controller also applies hard constraints. For example, a revoked credential or failed cryptographic verification can force denial even when the behavioral score is high.

Algorithm 1: Adaptive Device Decision

Input: device d , request r , context c

1. Validate device identity and lifecycle state.
2. Collect compact behavioral and protocol features.
3. Update $x_d(t)$ only with trusted observations.
4. Compute $T_d(t)$ using the weighted trust function.
5. Check hard constraints: revoked identity, forbidden protocol, or critical policy violation.
6. If any hard constraint fails, action $\leftarrow$ DENY.
7. Else if $T_d \geq \tau_1$, action $\leftarrow$ PERMIT.
8. Else if $T_d \geq \tau_2$, action $\leftarrow$ RESTRICTED.
9. Else if $T_d \geq \tau_3$, action $\leftarrow$ QUARANTINE.
10. Else action $\leftarrow$ DENY.
11. Record action, score, evidence identifiers, and policy version.
12. Return action and a short explanation.






Fig. 2. Adaptive decision workflow and evidence sources. Original figure.

VI. SECURITY MODEL AND THREAT ANALYSIS

The threat model assumes that an attacker may obtain a device credential, compromise a device after legitimate onboarding, abuse an exposed service, manipulate traffic, or attempt lateral movement. The gateway, trust engine, and policy store are assumed to be protected components, although their compromise is discussed as a limitation.

The architecture provides defense through diversity: cryptographic identity addresses impersonation; protocol policy limits unexpected services; behavior monitoring identifies deviations; contextual rules raise protection for critical assets; segmentation limits the blast radius; and audit evidence supports investigation. No single component is assumed to detect every attack.

Threat	Relevant evidence	Primary response
Stolen credential	Identity state + behavior deviation	Restricted / deny
Compromised legitimate device	Behavior + protocol deviation	Quarantine
Lateral movement	Destination diversity + segmentation rule	Deny / quarantine
Unauthorized service	Protocol compliance failure	Deny
Telemetry manipulation	Freshness + cross-checks	Re-evaluate / quarantine
Rapid baseline poisoning	Trusted-state-only profile update	Freeze profile update

VII. PROTOTYPE AND EXPERIMENTAL METHOD

A reproducible prototype can be implemented using a small edge gateway, several heterogeneous IoT nodes, a network switch or virtual network, an MQTT or equivalent application protocol, and a logging service. The prototype should contain a baseline configuration in which authentication, ordinary firewall rules, and conventional monitoring are enabled, followed by the TrustPulse-IoT configuration.

The evaluation should use both benign workload variation and attack scenarios. Benign variation is important because a security model that treats every workload change as malicious will generate excessive false positives. The device profiles should therefore include multiple normal operating periods, maintenance activity, temporary connectivity loss, and changes in reporting frequency.



The experiment should be divided into training/profile-building, calibration, and held-out testing periods. Thresholds and weights must be selected without using the final test set. This separation is necessary to avoid optimistic evaluation.

Test Scenarios

- T1: valid device with normal periodic telemetry.
- T2: valid device with an unexpected destination.
- T3: valid credentials reused from an unauthorized location.
- T4: abnormal message burst or connection-frequency change.
- T5: attempt to use a service outside the device profile.
- T6: device moved to a different network segment.
- T7: configuration/software change followed by reconnection.
- T8: maintenance-window behavior to test false-positive control.

VIII. EVALUATION METRICS AND RESEARCH HYPOTHESES

The primary metrics should be measured rather than assumed. Security effectiveness can be represented by precision, recall, F1-score, false-positive rate, and mean detection time. Enforcement effectiveness should include unauthorized-session blocking rate and time-to-quarantine. Operational cost should include gateway CPU utilization, memory usage, decision latency, and additional network traffic. Audit quality can be assessed by the percentage of decisions that have complete evidence and policy-version records.

Metric	Definition	Expected interpretation
Precision	Correct security decisions / all positive decisions	Lower alert noise
Recall	Detected attacks / all attacks	Coverage of abnormal behavior
F1	Harmonic mean of precision and recall	Balanced detector quality
Decision latency	Time from observation to enforcement	Real-time suitability
Gateway overhead	CPU/memory added by TrustPulse-IoT	Deployment cost
Containment time	Detection-to-quarantine interval	Response effectiveness

Research Hypotheses

- H1: Multi-factor trust decisions will reduce unauthorized access compared with identity-only authorization under compromised-device scenarios.
- H2: Local behavioral profiling will reduce response latency compared with a cloud-only decision path.
- H3: Proportional enforcement will reduce unnecessary full isolation during benign workload changes compared with binary deny policies.
- H4: Freezing profile updates during low-trust states will reduce the probability of rapid behavioral-baseline poisoning.

IX. COMPARISON WITH BASELINE SECURITY APPROACHES

Capability	Static ACL	Perimeter/Firewall	IDS only	TrustPulse-IoT
Device identity	Optional	Indirect	Usually not decision owner	Core input
Behavior context	No	Limited	Yes	Yes
Adaptive authorization	No	Limited	No	Yes
Proportional response	Rare	Limited	Alert dependent	Four states
Edge-local decision	Yes	Yes	Possible	Designed for it
Audit explanation	Rule based	Rule based	Alert based	Score + evidence + policy
Profile poisoning control	N/A	N/A	Varies	Trusted-state update rule



This comparison is architectural rather than an experimental benchmark. The objective is to identify which functions are integrated by the proposed design. Existing IDS, firewall, and identity systems can remain as implementation components inside TrustPulse-IoT.

X. DISCUSSION

The main benefit of the proposed architecture is decision continuity. Instead of generating an anomaly alert and leaving an administrator to decide what to do, the anomaly evidence can immediately influence the device's permitted service set. This can shorten the interval between detection and containment.

The design also supports gradual deployment. An organization can initially run the trust engine in monitor-only mode, then use it to recommend policies, and finally enable automated restricted/quarantine actions. This staged adoption is important because incorrect thresholds can disrupt legitimate industrial or household workflows.

There are also trade-offs. Behavioral profiling introduces a learning problem and may be affected by device firmware changes, seasonal workloads, maintenance activity, or sparse observations. Trust scoring can also create a false sense of precision if weights are not calibrated. For these reasons, the score should be treated as decision evidence, not as a universal probability of compromise. Critical cryptographic failures should remain hard policy violations rather than being overridden by a favorable behavioral score.

XI. LIMITATIONS AND ETHICAL RESEARCH PRACTICE

This paper presents an architecture and evaluation protocol, not completed experimental validation. Therefore, no numerical improvement, accuracy, or latency result is claimed. A future implementation should report the dataset, device hardware, network conditions, attack generation method, parameter-selection procedure, random seeds where applicable, and complete evaluation results.

The paper also does not reproduce source-paper figures or long passages. The supplied papers are used to identify the problem space and research motivation, while the architecture, equations, workflow, and evaluation plan are newly formulated. All external works are cited so that ideas can be traced to their original sources. Similarity checking should still be performed on the final manuscript before academic submission.

XII. CONCLUSION

This paper proposed TrustPulse-IoT, a context-aware adaptive access-control architecture for heterogeneous IoT networks. The design builds on security concerns documented in the supplied literature and extends them into a coordinated decision loop that combines identity confidence, behavioral stability, protocol compliance, asset/context risk, and telemetry freshness. The resulting trust value controls proportional responses ranging from normal access to restriction, quarantine, and denial.

The contribution is deliberately presented as an implementable research hypothesis rather than a fabricated performance claim. A prototype and controlled evaluation can test whether continuous edge-based trust decisions improve unauthorized-access prevention and response time while keeping computational overhead and false positives within acceptable limits. Future work should validate the model across consumer, smart-building, and industrial IoT environments and investigate automated calibration, privacy-preserving telemetry, and robust defenses against behavior-profile poisoning.

REFERENCES

- [1] T. Sauter and A. Treytl, "IoT-enabled sensors in automation systems and their security challenges," *IEEE Sensors Lett.*, vol. 7, no. 12, Art. no. 7500904, Dec. 2023, doi: 10.1109/LSSENS.2023.3332404.
- [2] M. K. Saini and R. K. Saini, "Internet of Things (IoT) applications and security challenges: A review," *Int. J. Eng. Res. Technol. (IJERT)*, Special Issue, 2019.
- [3] Imran, S. M. Ali, M. M. Alam, and M. M. Su'ud, "A survey of IoT security issues—From past to future trends," *J.*



Comput. Sci., vol. 17, no. 11, pp. 1031–1045, 2021, doi: 10.3844/jcssp.2021.1031.1045.

[4] M. Fagan, K. Megas, K. Scarfone, and M. Smith, “IoT device cybersecurity capability core baseline,” NISTIR 8259A, National Institute of Standards and Technology, 2020, doi: 10.6028/NIST.IR.8259A.

[5] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero trust architecture,” NIST SP 800-207, National Institute of Standards and Technology, 2020, doi: 10.6028/NIST.SP.800-207.

[6] “Deep reinforcement learning aided device profiling for anomaly detection in IIoT over zero trust security network,” in Proc. IEEE Asia-Pacific Conf. Comput. Sci. Data Eng. (CSDE), 2023, doi: 10.1109/CSDE59766.2023.10487721.

[7] “Anomaly detectors for self-aware edge and IoT devices,” in Proc. IEEE Int. Conf. Software Quality, Reliability, and Security (QRS), 2023, doi: 10.1109/QRS60937.2023.00013.

[8] “Multidimensional behavioral profiling of Internet-of-Things in edge networks,” in Proc. IEEE/ACM Int. Symp. Quality of Service (IWQoS), 2019, doi: 10.1145/3326285.3329072.

[9] “IoT intrusion detection: A review of ML and DL-based approaches,” in Proc. 4th Int. Conf. Innovative Research in Applied Science, Engineering and Technology (IRASET), 2024, doi: 10.1109/IRASET60544.2024.10548107.

[10] M. J. Fagan et al., “Foundational cybersecurity activities for IoT product manufacturers,” NISTIR 8259 Rev. 1, National Institute of Standards and Technology, 2026, doi: 10.6028/NIST.IR.8259r1.

