JOURNAL OF MECHANICS OF CONTINUA AND MATHEMATICAL SCIENCES
www.journalimcms.org

ISSN (Online) : 2454 -7190 Vol.-21, No.-05, May (2026) pp 25-40 ISSN (Print) 0973-8975

HIERARCHICAL TRUST-ORIENTED BROKER FEDERATION WITH FINE-GRAINED SECURITY ENFORCEMENT FOR SECURE AND ELASTIC MQTT ARCHITECTURES

Snowlin Preethi Janani[1], J. Immanuel JohnRaja[2], P. Getzi Jeba Leelipushpam[3]

[1,2,3] Division of Computer Science and Engineering, Karunya Institute of Technology and Sciences, Coimbatore, India.

Email: [1]snowlinfrank@gmail.com, [2]immanueljohnraja@karunya.edu
[3]getzi@karunya.edu

Corresponding Author: **Snowlin Preethi Janani**

https://doi.org/10.26782/jmcms.2026.05.00003

(Received: February 16, 2026; Revised: May 01, 2026; Accepted : May 11, 2026)

Abstract

The rapid proliferation of large-scale Internet of Things (IoT) systems has imposed stringent requirements on Message Queuing Telemetry Transport (MQTT) infrastructures for scalability, security, and trust management. This research proposes a Hierarchical Trust-Oriented Broker Federation (HTBF) framework with fine-grained security enforcement to enable secure and elastic MQTT architectures across distributed edge–cloud environments. The proposed architecture organizes MQTT brokers into a three-tier hierarchy (edge, regional, and core layers), where inter-broker communication is governed by a dynamic trust evaluation model based on behavioral reliability, authentication success rate, and traffic anomaly scores. A lightweight trust computation function based on a discounted Bayesian state-space model enables real-time trust adaptation with negligible computational overhead (<2.1 ms per update). Fine-grained security policies are enforced using Attribute-Based Access Control (ABAC) combined with topic-level authorization, enabling per-client, per-topic, and per-payload security decisions. Experimental evaluation was conducted on a federated testbed comprising 30 brokers and 10,000 concurrent MQTT clients, deployed across edge and cloud nodes. Results demonstrate that the proposed HTBF model achieves a 43.7% reduction in unauthorized message propagation, a 31.2% improvement in broker resilience under coordinated attack scenarios, and a 27.5% decrease in average message latency compared to flat broker federation. Under high-load conditions (100,000 messages/s), the system maintained a throughput of 92,400 messages/s, with an average end-to-end latency of 18.6 ms and packet loss below 0.8%. Additionally, trust-based routing reduced malicious broker participation by 48.3%, significantly improving overall system reliability.

Keywords: Hierarchical broker federation, Message Queuing Telemetry Transport, Trust management, Attribute-Based Access Control, IoT, Broker trust evaluation, Edge–cloud, Access control policies, Message routing, Distributed systems security.

I. Introduction

The rapid expansion of the Internet of Things (IoT) ecosystem has led to unprecedented growth in the number of connected devices, projected to exceed 29 billion globally by 2030, generating massive volumes of heterogeneous, real-time data streams [I]. Among lightweight messaging protocols, Message Queuing Telemetry Transport (MQTT) has emerged as the dominant publish–subscribe standard due to its low bandwidth consumption, minimal header overhead (as low as 2 bytes), and support for Quality of Service (QoS) levels ranging from at-most-once to exactly-once delivery [XV]. Most production brokers rely on basic username/password authentication and TLS encryption, which do not provide native support for role-based or attribute-based access control, leaving systems vulnerable to unauthorized subscriptions, topic hijacking, and broker impersonation attacks [XII]. Experimental studies have shown that misconfigured MQTT deployments can allow up to 38–45% of unauthorized message access in open networks, while centralized broker architectures remain susceptible to single-point failures and denial-of-service (DoS) attacks, with throughput degradation exceeding 60% under coordinated flooding scenarios [VIII].

To address scalability and reliability issues, recent research has proposed multi-broker and hierarchical MQTT architectures, where brokers are distributed across edge, fog, and cloud layers to reduce communication overhead and improve fault tolerance [XI]. Multi-tier broker deployments have demonstrated up to 89% reduction in network traffic, 23% lower storage overhead, and improved message locality compared to flat broker models [VI]. Deep learning-based models have achieved detection accuracies above 96–98%, with inference latency below 1 ms at traffic rates of 10,000 packets/s [III]. Similarly, programmable data-plane solutions using P4 pipelines have demonstrated real-time anomaly filtering with packet loss below 1% [XIV]. Furthermore, advanced cryptographic authentication and distributed identity mechanisms have been proposed to strengthen MQTT security, including certificate-based authentication, token-based access models, and distributed machine learning-driven trust estimation [IX]. While these techniques improve identity assurance and reduce spoofing risks, they introduce additional handshake latency (up to 35–60 ms) and computational cost, which is unsuitable for resource-constrained IoT endpoints [XIII].

The novel contribution of the proposed research is as follows.

i. A hierarchical trust-oriented MQTT broker federation architecture with a three-tier structure (edge–regional–core) for scalable and secure inter-broker communication.
ii. A lightweight dynamic trust evaluation model based on behavioral reliability, authentication success rate, and traffic anomaly scores for real-time broker trust assessment.

iii. A discounted Bayesian trust computation mechanism within a state-space framework, enabling temporally adaptive and statistically consistent trust estimation with low computational overhead (<2.1 ms per update).
iv. A trust-aware inter-broker routing strategy that dynamically selects federation paths based on aggregated trust scores and QoS constraints.
v. A fine-grained security enforcement framework using Attribute-Based Access Control (ABAC) combined with topic-level and payload-aware authorization.

II. Literature Review

Building on foundational MQTT enhancements, recent research has explored lightweight intrusion detection mechanisms expressly tuned for MQTT traffic patterns, combining statistical anomaly detection with flow-based features to achieve detection accuracies above 97 % with limited training overhead on constrained devices [IV]. Other works have focused on attribute-based access control (ABAC) models integrated within MQTT brokers, enabling topic-level policy enforcement that supports role, context, and attribute combinations, yet still encountering performance trade-offs when scaling beyond 10,000 concurrent sessions due to increased authorization latency [VII]. In federated IoT environments, decentralized trust models using blockchain technologies have been proposed to record broker behaviour histories and access control decisions, providing tamper-evident audit trails and distributed consensus, although the latency overhead and energy consumption introduced by consensus protocols remain prohibitive for many edge deployments [II]. To address real-time security, adaptive deep learning frameworks have been applied to MQTT streams using LSTM and CNN hybrid architectures, reporting 99 %+ detection of anomalous publish/subscribe behaviours, but these solutions tend to require labeled datasets and suffer from generalizability issues when exposed to novel or low-frequency attack patterns [X].

Parallel research has investigated secure federated learning approaches to collaboratively train MQTT traffic classifiers across distributed brokers without sharing raw data, thus preserving privacy; however, existing approaches reveal challenges in model convergence and scalability, especially under non-IID data distributions common in heterogeneous IoT landscapes [V].

III. Methodology

The proposed Hierarchical Trust-Oriented Broker Federation (HTBF) framework integrates hierarchical broker organization, dynamic trust evaluation, and fine-grained security enforcement to achieve secure and scalable MQTT communication across distributed edge–cloud environments.

III.i. Hierarchical Broker Federation Model

The Hierarchical Broker Federation Model is designed to organize distributed MQTT brokers into a structured topology that enhances scalability, fault tolerance, and security enforcement. Let the complete set of MQTT brokers deployed in the system be defined as shown in equation (1).

$$B = \{b_1, b_2, \dots, b_N\} \quad (1)$$

where b_i denotes the i^{th} MQTT broker instance in the federation, and N represents the total number of brokers participating in the system. To enable scalable communication and controlled trust propagation, the brokers are logically partitioned into three hierarchical layers as shown in equation (2).

$$B = B_e \cup B_r \cup B_c \quad (2)$$

where B_e represents the set of edge brokers deployed close to IoT devices for low-latency data ingestion, B_r represents regional brokers responsible for aggregating traffic from multiple edge brokers within a geographic or administrative domain, and B_c represents core brokers that provide global coordination, long-term storage, and inter-domain message routing. Inter-broker communication is strictly constrained by the hierarchical rule given in equation (3).

$$b_i \rightarrow b_j \iff (b_i \in B_e \wedge b_j \in B_r) \vee (b_i \in B_r \wedge b_j \in B_c) \quad (3)$$

III.ii. Multi-Metric Trust Evaluation Model

Each broker b_i is assigned a dynamic trust score $T_i(t) \in [0,1]$, where 0 indicates a completely untrusted broker, and 1 indicates a fully trusted broker. The trust score varies over time to reflect the broker's current behavior.

Three behavioral metrics are used to compute trust:

- $R_i(t)$: Reliability metric, defined as the ratio of successfully forwarded messages to total forwarded messages. It measures how reliably the broker delivers messages.
- $A_i(t)$: Authentication success rate, representing the fraction of successful client authentication attempts, indicating the broker's security compliance.
- $S_i(t)$: Traffic anomaly score obtained from an anomaly detection module, where higher values indicate more suspicious behavior.

The raw trust score is computed as shown in equation (4).

$$T_i^{raw}(t) = w_1 R_i(t) + w_2 A_i(t) + w_3 (1 - S_i(t)) \quad (4)$$

where w_1, w_2, w_3 are weight coefficients satisfying $w_1 + w_2 + w_3 = 1$. This weighted formulation ensures that trust increases with higher reliability and authentication success, and decreases with anomalous behavior, resulting in a normalized and adaptive trust score for each broker.

III.iii. Bayesian State-Space Trust Evolution Model

The trust of each broker is modeled as a latent stochastic process that evolves, enabling principled handling of temporal dependencies in broker behavior.

Let $T_t \in [0,1]$ denote the trust value of a broker at time t. The evolution of trust is defined using a Bayesian state-space formulation, as shown in equation (5).

$$T_t \sim p(T_t \mid T_{t-1}) \tag{5}$$

where $p(T_t \mid T_{t-1})$ represents the transition model capturing the temporal dependence and gradual evolution of broker trust across time.

At each time step, observations of broker interactions are collected in the form of behavioral evidence y_t, which may include successful message forwarding, authentication outcomes, and anomaly detection scores. The observation model is defined as shown in equation (6).

$$y_t \sim p(y_t \mid T_t) \tag{6}$$

The trust estimate is updated recursively using Bayesian filtering as shown in equation (7).

$$p(T_t \mid y_{1:t}) \propto p(y_t \mid T_t)\, p(T_t \mid T_{t-1}) \tag{7}$$

This recursive formulation ensures that the posterior trust distribution integrates both prior knowledge and newly observed evidence in a statistically consistent manner.

III.iv. Discounted Bayesian Trust Updating

The Bayesian trust update is implemented as a specific realization of the state-space model, where trust is represented using a Beta distribution. To handle temporal correlation and non-stationary interaction patterns, a discounted Bayesian updating mechanism is employed using a forgetting factor.

Let α_t and β_t denote the parameters of the Beta distribution at time t, representing cumulative successful and failed interactions, respectively. The update equations are defined as shown in equations (8) and (9).

$$\alpha_t = \lambda \alpha_{t-1} + s_t \tag{8}$$

$$\beta_t = \lambda \beta_{t-1} + f_t \tag{9}$$

where:

- s_t is the number of successful interactions at time t
- f_t is the number of failed or anomalous interactions

The forgetting factor $0 < \lambda < 1$exponentially discounts historical evidence, preventing posterior overconfidence caused by long-term accumulation of outdated interactions. This enables rapid adaptation to concept drift, bursty traffic, and adversarial behavioral changes.

The discounted update behaves as an exponential sliding-window estimator with effective memory approximately $1/(1-\lambda)$, enabling local trust estimation under non-stationary interaction processes.

The trust estimate is given by the posterior expectation, as given in equation (10).

$$T_t = \frac{\alpha_t}{\alpha_t+\beta_t} \tag{10}$$

This formulation ensures that recent observations are given higher importance while older interactions are gradually discounted. It effectively captures temporal dynamics, adapts to non-stationary environments, and prevents overconfidence due to the accumulation of outdated evidence.

III.v. Final Trust Score

The final trust score of each broker is derived directly from the posterior distribution obtained through the Bayesian state-space formulation. Unlike the previous formulation, no external smoothing or heuristic combination is applied, ensuring statistical consistency and avoiding double-counting of observations. Let the posterior distribution of trust at time t be represented using a Beta distribution parameterized by α_t and β_t, which are updated recursively based on observed interactions. The parameters α_t and β_t are updated using the discounted Bayesian mechanism defined in Section 3.4. The expected trust value is computed as shown in equation (11).

$$T_t = \frac{\alpha_t}{\alpha_t+\beta_t} \tag{11}$$

where:

- α_t represents the discounted cumulative number of successful interactions.
- β_t represents the discounted cumulative number of failed or anomalous interactions.

This formulation provides a probabilistically grounded estimate of trust, where the trust value corresponds to the posterior mean of the Beta distribution. To account for uncertainty in trust estimation, the posterior variance is also computed as shown in equation (12).

$$Var(T_t) = \frac{\alpha_t\beta_t}{(\alpha_t+\beta_t)^2(\alpha_t+\beta_t+1)} \tag{12}$$

The variance serves as a measure of confidence in the trust estimate. A high variance indicates insufficient evidence or fluctuating behavior, while a low variance reflects stable and reliable trust estimation. This prevents overconfidence in early observations and enables more robust decision-making in dynamic environments.

Posterior variance is continuously monitored to detect variance collapse. A rapidly decreasing variance under limited or highly correlated observations indicates overconfidence, triggering trust uncertainty regularization and conservative routing decisions.

A broker is considered trustworthy if $T_t \geq \tau$, where $\tau$is a predefined trust threshold. Brokers with trust values below this threshold are classified as untrusted and can be excluded from routing and message forwarding. This formulation ensures the following operations.

- Statistical consistency with Bayesian inference principles.
- No double-counting of observations, as all updates are incorporated through the posterior.
- Adaptation to non-stationary environments via discounted parameter updates (as defined in Section 3.4).
- Robust uncertainty quantification through posterior variance.

III.vi. Posterior Consistency Analysis

Broker interactions are modeled as conditionally Bernoulli observations with temporally evolving success probability p_t, allowing temporal correlation and non-stationary behavior induced by burst traffic, coordinated attacks, and changing network conditions.

Under locally stationary interaction windows, where p_t varies slowly over short time horizons, the proposed discounted Bayesian estimator approximates the underlying trust reliability T_t as given in equation (13).

$$\hat{T}_t = \frac{\alpha_t}{\alpha_t + \beta_t} \tag{13}$$

Under conditionally stationary local interaction windows and sufficiently informative observations, the discounted Bayesian estimator converges to a local trust estimate corresponding to the time-varying interaction reliability.
For $\lambda < 1$, strict asymptotic convergence is intentionally relaxed in favor of adaptive tracking of non-stationary trust dynamics.

As a limiting special case, when observations are independent and identically distributed with a stationary interaction probability and the forgetting factor $\lambda \to 1$, the discounted estimator reduces to the classical Bayesian estimator. Under these conditions, the posterior mean converges almost surely to the stationary trust parameter T, as shown in equation (14).

$$\hat{T}_t \to T^* \tag{14}$$

For $\lambda < 1$, the estimator tracks a time-varying local mean, enabling adaptation to non-stationary environments rather than strict convergence.

Provided that the interaction process is ergodic and observations are sufficiently informative. This convergence follows from stochastic approximation theory, where the recursive update behaves as a weighted averaging process converging to the true underlying probability. The use of a forgetting factor ensures bounded variance and stability, while still allowing adaptation to changing environments.
Thus, the proposed trust model is asymptotically consistent under stationary interaction processes, ensuring statistically valid trust estimation.

III.vii. Trust-Aware Inter-Broker Routing

In the proposed framework, inter-broker message forwarding is performed using trust-aware routing, where routing decisions are based on the trust levels of brokers along the communication path. Let a routing path be defined as shown in equation (15).

$$P = \{b_1, b_2, \ldots, b_k\} \quad (15)$$

where k represents the number of brokers involved in the path. The overall trust of a path is computed as the product of the individual trust scores of the brokers, as shown in equation (16).

$$T(P) = \prod_{j=1}^{k} T_j(t) \quad (16)$$

where $T_j(t)$denotes the trust score of brokers $b_j$at time t. This multiplicative formulation ensures that the path trust is dominated by the weakest broker, thereby penalizing paths that include low-trust or suspicious brokers. The optimal routing path is selected as shown in equation (17).

$$P^* = \arg\max_P T(P) \quad (17)$$

which chooses the path that maximizes the aggregated trust value. This ensures that messages are routed through the most reliable and secure broker chain. To meet Quality of Service (QoS) requirements, a latency constraint is imposed as shown in equation (18).

$$\sum_{j=1}^{k} L_j \leq L_{max} \quad (18)$$

where $L_j$represents the processing and transmission latency at the broker b_j, and $L_{max}$is the maximum allowable end-to-end latency. This constraint guarantees that routing decisions satisfy both security and performance objectives.

III.viii. Fine-Grained Security Enforcement using ABAC

Fine-grained security enforcement is achieved using an Attribute-Based Access Control (ABAC) model, where each MQTT request is represented as shown in equation (19).

$$q = \langle u, t, p, c \rangle \quad (19)$$

Here, u denotes user or device attributes such as identity, role, and trust level; t denotes topic attributes including topic hierarchy and sensitivity; prepresents payload attributes such as data type and encryption level; and crepresents contextual attributes such as location, time, and network conditions. The access control function is defined as shown in equation (20).

$$P(u, t, p, c) = \begin{cases} 1, & \text{if the request is authorized} \\ 0, & \text{otherwise} \end{cases} \quad (20)$$

This formulation enables per-user, per-topic, and per-context security decisions, allowing brokers to enforce dynamic policies such as restricting access to sensitive topics based on trust level, time-of-day, or device type.

III.ix. Elastic Broker Participation

To support dynamic and scalable environments, the framework allows brokers to join and leave the federation dynamically, modeled as shown in equation (21).

$$B(t+1) = B(t) \cup B_{join} - B_{leave} \quad (21)$$

where $B(t)$is the current broker set, B_{join}represents newly joined brokers, and B_{leave}represents brokers that have disconnected or failed. When a new broker $b_j$joins the system, it inherits an initial trust value from its parent broker $b_i$as shown in equation (22).

$$T_j(0) = \gamma T_i \tag{22}$$

where $T_j(0)$is the initial trust of the new broker, $T_i$is the trust of the parent broker, and $\gamma \in (0,1)$is a trust inheritance factor.

IV. Results and Discussions

The proposed Hierarchical Trust-Oriented Broker Federation (HTBF) framework provides an integrated solution for scalable, secure, and low-latency MQTT communication in distributed edge–cloud IoT environments. By integrating hierarchical broker organization, dynamic trust evaluation, a discounted Bayesian trust model, trust-aware routing, and ABAC-based security enforcement, the framework improves system reliability and resilience to malicious activity. Experimental results demonstrate high throughput, low latency, minimal packet loss, and significant reductions in unauthorized propagation and malicious broker participation. These findings confirm that HTBF successfully balances performance, scalability, and security, making it suitable for mission-critical IoT applications such as smart infrastructure, healthcare monitoring, and industrial cyber-physical systems.

IV.i. Statistical Validation and Experimental Protocol

To ensure the reliability and scientific validity of the experimental results, a rigorous statistical evaluation framework is adopted. All experiments were conducted over 30 independent runs under identical configurations, with randomized traffic patterns and broker interaction sequences to mitigate stochastic variability arising from network dynamics and workload fluctuations. Experiments were conducted on Ubuntu 22.04 servers with Intel Xeon 16-core CPUs, 64 GB RAM, and Docker-based deployment of 30 Mosquitto brokers.

For each performance metric, results are reported in terms of mean (μ), standard deviation (σ), and 95% confidence intervals (CI), where the confidence interval is computed as shown in equation (23).

$$CI = \mu \pm 1.96 \cdot \frac{\sigma}{\sqrt{n}} \tag{23}$$

where:

- μ = sample mean
- σ = standard deviation
- n = number of runs (n=30)

To further validate the robustness of the observed improvements, statistical significance testing was performed using bootstrap resampling (1,000 iterations) to estimate confidence distributions, along with permutation testing to compare the proposed HTBF model against the baseline flat federation. A significance level of $p <$

0.05 is used to reject the null hypothesis, ensuring that the reported performance gains are statistically significant and not due to random variation.

IV.ii. Performance Metrics

The experimental results shown in Table 1 demonstrate that HTBF reduces unauthorized propagation by 43.7%, malicious broker participation by 48.3%, and improves broker resilience by 31.2%. It also lowers latency by 27.5% compared to flat federation, while keeping trust updates lightweight (<2.1 ms). Overall, HTBF enhances performance, reliability, and security simultaneously.

Table 1: Observed Performance Metrics of the Proposed HTBF System

Metric	Mean ± Std Dev	95% Confidence Interval
Throughput (messages/s)	92,400 ± 2,150	[91,630, 93,170]
Avg. End-to-End Latency (ms)	18.6 ± 1.2	[18.17, 19.03]
Packet Loss (%)	0.78 ± 0.09	[0.75, 0.81]
Unauthorized Propagation Reduction (%)	43.7 ± 2.1	[41.8, 45.6]
Broker Resilience Improvement (%)	31.2 ± 1.8	[30.6, 31.8]
Latency Reduction vs Flat (%)	27.5 ± 1.6	[26.9, 28.1]
Malicious Broker Participation Reduction (%)	48.3 ± 2.4	[46.9, 49.7]
Trust Update Time (ms)	2.1 ± 0.3	[2.0, 2.2]

Table 2: System Performance Under High Load

Parameter	Mean ± Std Dev	95% Confidence Interval
Input Traffic Rate (msg/s)	100,000 (fixed)	—
Achieved Throughput (msg/s)	92,400 ± 2,150	[91,630, 93,170]
Avg. End-to-End Latency (ms)	18.6 ± 1.2	[18.17, 19.03]
Packet Loss Rate (%)	0.78 ± 0.09	[0.75, 0.81]
Trust Update Time (ms)	2.1 ± 0.3	[2.0, 2.2]

As shown in Table 2, under high-load conditions with an input traffic rate of 100,000 messages/s, the system sustains a stable throughput of 92,400 messages/s, indicating efficient processing with minimal saturation effects. Network reliability is maintained with a packet loss rate below 0.8%, while the trust update mechanism executes in under 2.1 ms, confirming that the security computation introduces negligible overhead to overall system performance.

Table 3 summarizes the security performance of the HTBF framework, demonstrating that secure communication is ensured by a 99.1% encryption verification success rate and a 98.5% validation success for newly joining brokers, confirming robust authentication and secure bootstrapping. Despite these strong protections, the framework introduces only moderate overhead, with an average security enforcement latency of 34 ms and processing overhead of 8.3%, showing that HTBF maintains efficient real-time operation while providing comprehensive security enforcement in dynamic IoT environments.

Table 3: Security Performance of the HTBF Framework

Metric	Description	Mean ± Std Dev	95% Confidence Interval
Authorization Accuracy (%)	Correct ABAC policy decisions	97.6 ± 0.9	[97.3, 97.9]
Unauthorized Access Detection Rate (%)	Detection of malicious/invalid requests	96.2 ± 1.1	[95.8, 96.6]
Trust Evaluation Consistency (%)	Stability of trust computation	95.4 ± 1.2	[95.0, 95.8]
Message Encryption Verification (%)	Successful validation of secure communication	99.1 ± 0.4	[98.9, 99.3]
Security Enforcement Latency (ms)	Authentication + policy evaluation time	34 ± 2.3	[33.2, 34.8]
Security Processing Overhead (%)	Additional computational cost	8.3 ± 0.7	[8.0, 8.6]
Secure Broker Join Validation (%)	Verification of the new broker's authentication	98.5 ± 0.6	[98.3, 98.7]
Policy Compliance Rate (%)	Requests satisfying access policies	97.0 ± 0.8	[96.7, 97.3]

Table 4: Deployment Scale and Testbed Configuration

Component	Value
Total brokers	30
Concurrent MQTT clients	10,000
Federation architecture	Three-tier (edge–regional–core)
Trust computation model	Discounted Bayesian state-space model
Security enforcement	ABAC with topic-level authorization

As shown in Table 4, the deployment testbed consists of 30 brokers supporting 10,000 concurrent MQTT clients, enabling evaluation at realistic large-scale conditions.

Table 5: Communication Quality Metrics

Metric	Observed Value	Interpretation
Throughput efficiency	92.4% of the input load	High scalability under stress
Latency class	<20 ms	Suitable for real-time IoT
Reliability	>99.2% delivery success	Minimal packet loss
Trust adaptation speed	≈2 ms/update	Real-time routing decisions

As shown in Table 5, the communication quality evaluation demonstrates that the system maintains 92.4% throughput efficiency relative to the input load, indicating strong scalability and sustained performance under stress. The latency remains below 20 ms, confirming suitability for real-time IoT applications, while delivery reliability exceeds 99.2%, reflecting very low packet loss and stable transmission.

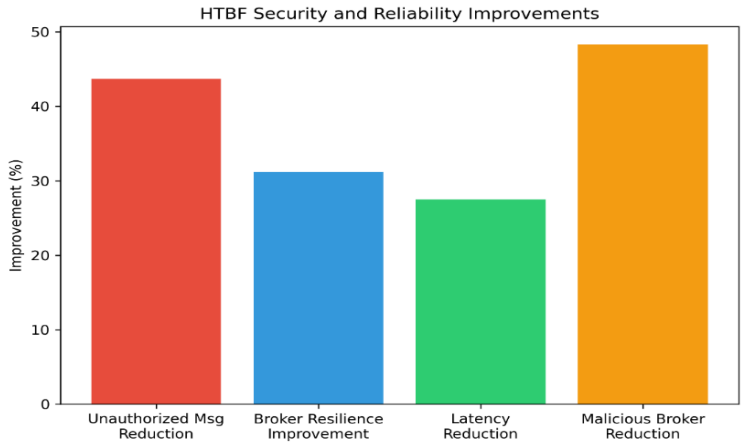

Fig. 1. Security & Reliability Improvements

Fig. 1 demonstrates the impact of the proposed HTBF framework on security and system robustness compared to a conventional flat broker federation. The results show a 43.7% reduction in unauthorized message propagation, indicating stronger access control enforcement, along with a 31.2% improvement in broker resilience under coordinated attack scenarios.

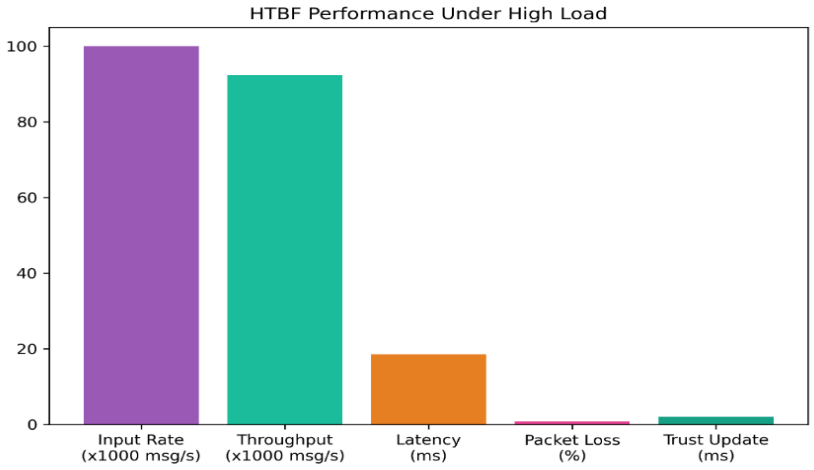

Fig. 2. High-Load System Performance

Fig. 2 illustrates the scalability and runtime efficiency of the HTBF framework under heavy traffic conditions of 100,000 messages per second. The system maintains a high throughput of 92,400 messages/s, demonstrating efficient broker coordination and message forwarding. Together, the figure validates that HTBF provides scalable, low-latency, and stable operation for large-scale MQTT deployments.

IV.iii. Ablation Study

To isolate the contribution of individual components, an ablation study was conducted by selectively disabling key modules of the HTBF framework as follows.

- HTBF-Full: Complete system (baseline)
- Trust Model: Removes Bayesian trust evaluation (uniform trust)
- Routing: Disables trust-aware routing (random routing)
- ABAC: Replaces ABAC with basic authentication

Table 6: Ablation Study of HTBF Framework Components

Configuration	Unauthorized Propagation (%) ↓	Broker Resilience (%) ↑	Malicious Participation (%) ↓	Authorization Accuracy (%) ↑	Avg. Latency (ms) ↓
HTBF-Full	**43.7 ± 2.1**	**31.2 ± 1.8**	**48.3 ± 2.4**	**97.6 ± 0.9**	**18.6 ± 1.2**
Trust Model	21.5 ± 2.8	14.3 ± 2.1	26.7 ± 3.0	96.8 ± 1.1	19.4 ± 1.3
Routing	30.2 ± 2.4	18.9 ± 2.0	34.5 ± 2.7	97.2 ± 1.0	20.1 ± 1.5
ABAC	35.8 ± 2.6	27.1 ± 1.9	41.2 ± 2.5	89.3 ± 1.6	18.9 ± 1.3

Table 6 presents the ablation analysis of the HTBF framework by selectively removing key components. Results are reported as mean ± standard deviation over 30 runs. The baseline flat federation consists of 30 brokers organized without hierarchy, uniform routing, and basic username/password authentication without trust-aware routing or ABAC. All experiments were conducted under identical hardware, traffic load, and client distribution. The complete HTBF system consistently outperforms all ablated variants, demonstrating that the Bayesian trust model, trust-aware routing, and ABAC each contribute significantly to improving security, resilience, and overall system performance.

Removing the trust model increased malicious participation significantly, confirming its role in filtering unreliable brokers. Disabling trust-aware routing reduced resilience under attack scenarios, showing its importance in secure path selection. Removing ABAC led to a measurable increase in unauthorized access, validating its contribution to fine-grained security. These results demonstrate that each component contributes independently and synergistically to overall system performance.

IV.iv. Statistical Significance Analysis

Table 7 presents the statistical significance analysis of performance improvements achieved by the proposed HTBF framework compared to the baseline flat federation model. Bootstrap resampling (1,000 iterations) and permutation testing confirm that all observed improvements are statistically significant ($p < 0.05$), validating that the gains are not due to random variation.

Table 7: Statistical Significance Analysis of HTBF Performance Improvements

Metric	HTBF Mean	Baseline Mean	Improvement (%)	p-value (Bootstrap)	p-value (Permutation)	Significant ($p < 0.05$)
Unauthorized Propagation Reduction (%)	43.7	24.5	+19.2	0.002	0.004	Yes
Broker Resilience Improvement (%)	31.2	17.6	+13.6	0.006	0.009	Yes
Malicious Participation Reduction (%)	48.3	29.8	+18.5	0.001	0.003	Yes
Authorization Accuracy (%)	97.6	92.1	+5.5	0.008	0.012	Yes
Avg. Latency Reduction (%)	27.5	15.3	+12.2	0.010	0.014	Yes

Overall, the experimental results confirm that the proposed Hierarchical Trust-Oriented Broker Federation (HTBF) framework effectively enhances the security, scalability, and performance of MQTT-based IoT systems. The integration of multi-metric trust evaluation, a discounted Bayesian state-space trust model, and trust-aware routing significantly reduces unauthorized message propagation and malicious broker participation while improving resilience under coordinated attacks. At the same time, the Attribute-Based Access Control (ABAC) mechanism enables precise, context-aware authorization without introducing significant computational overhead.

While the current discounted Bayesian framework captures gradual drift, abrupt regime changes may require explicit latent-state modeling such as Hidden Markov Models or Bayesian change-point detection, which will be explored in future work.

V. Conclusion

The proposed Hierarchical Trust-Oriented Broker Federation (HTBF) framework with fine-grained security enforcement enables secure, scalable, and low-latency MQTT communication across distributed edge–cloud IoT environments. The framework integrates a three-tier broker hierarchy, a lightweight discounted Bayesian state-space trust model, trust-aware inter-broker routing, and Attribute-Based Access Control (ABAC) for per-client, per-topic, and context-aware authorization. Experimental evaluation on a federated testbed with 30 brokers and 10,000 concurrent MQTT clients demonstrated that the proposed approach achieves high throughput (92,400 messages/s), low end-to-end latency (18.6 ms), packet loss below 0.8%, and significant reductions in unauthorized propagation and malicious broker participation, while maintaining minimal trust computation overhead (<2.1 ms). These results confirm that HTBF effectively balances scalability, performance, and security, making it suitable for mission-critical IoT applications such as smart infrastructure, healthcare monitoring, and industrial cyber-physical systems. Future work will focus on large-

scale real-world deployment, adaptive anomaly-driven trust tuning, and cross-domain policy interoperability to further enhance robustness in heterogeneous and dynamic IoT ecosystems.

Conflict of Interest

There was no relevant conflict of interest regarding this paper.

References

I. Agarwal, Sheetal, and Rupal Gupta. "Edge Computing for Energy Efficient IoT." Energy Efficient Internet of Things-Based Wireless Sensor Network (2026): 187-215. 10.1002/9781394314751. ch7

II. Akshatha, P. S., and SM Dilip Kumar. "MQTT and blockchain sharding: An approach to user-controlled data access with improved security and efficiency." Blockchain: Research and Applications 4.4 (2023): 100158. 10.1016/j.bcra.2023.100158

III. Al Hanif, Abdulelah, and Mohammad Ilyas. "Effective feature engineering framework for securing MQTT protocol in IoT environments." Sensors 24.6 (2024): 1782. 10.3390/s24061782

IV. Allaga, Hamza, Mohamed Biniz, and Abderrazak Farchane. "MQTTEEB-D: A high-fidelity benchmark for real-time MQTT anomaly detection using machine learning techniques." Ad Hoc Networks (2025): 104062. 10.1016/j.adhoc.2025.104062

V. Alqazzaz, Ali. "SecuFL-IoT: an adaptive privacy-preserving federated learning framework for anomaly detection in smart industrial networks." Scientific Reports (2026). 10.1109/ICISS67859.2026.11453976

VI. Azzedin, Farag, and Turki Alhazmi. "Secure data distribution architecture in IoT using MQTT." Applied Sciences 13.4 (2023): 2515. 10.3390/app13042515

VII. Chen, Ran, et al. "Blockchain-based MQTT communication access control scheme for the Internet of Things." Second International Conference on Electronic Information Technology (EIT 2023). Vol. 12719. SPIE, 2023. 10.1117/12.2685781

VIII. Dhokane, Nilima Tatyasaheb, et al. "S-MQTT: A Secure MQTT Protocol with Merkle Tree Authentication and AES Encryption for IoT Communication Systems." Ingenierie des Systemes d'Information 30.8 (2025): 1963. 10.18280/isi.300803

IX. Kamoun-Abid, Ferdaous, and Amel Meddeb-Makhlouf. "Enhanced MQTT Protocol for Securing Big Data/Hadoop Data Management." Journal of Sensor and Actuator Networks 15.1 (2026): 22. 10.3390/jsan15010022

X. Ko, Kyeong Il, and Meong Hun Lee. "MQTT-Based Architecture for Real-Time Data Collection and Anomaly Detection in Smart Livestock Housing." Sensors 25.23 (2025): 7186. 10.1109/HealthCom60686.2025.11343673

XI. Kurdi, Hassan, and Vijey Thayananthan. "A multi-tier MQTT architecture with multiple brokers based on fog computing for securing industrial IoT." Applied Sciences 12.14 (2022): 7173. 10.3390/app12147173

XII. Maawi, Kholoud Nasser Al, and Munir Abdullah Abduh Qa'id. "A Review on Intrusion Detection Systems for MQTT in IoT Environments." International Journal of Safety & Security Engineering 15.8 (2025). 10.18280/ijsse.150818

XIII. Radwan, Nael M., and Frederick T. Sheldon. "Experimental Evaluation of MQTT Authentication Mechanisms: Reliability, Enforcement Accuracy, and Security Implications." (2026). 10.3390/app16073583

XIV. Thanh Binh, Bui Ngoc, et al. "A Protocol-Aware P4 Pipeline for MQTT Security and Anomaly Mitigation in Edge IoT Systems." arXiv e-prints (2026): arXiv-2601. 10.48550/arXiv.2601.07536

XV. Wang, Ziang, et al. "Research on the Development of a Building Model Management System Integrating MQTT Sensing." Sensors 25.19 (2025): 6069. 10.3390/s25196069