JAAFR INTERNATIONAL RESEARCH JOURNAL
JOURNAL OF ADVANCE AND FUTURE RESEARCH
JAAFR.ORG | ISSN : 2984-889X
An International Open Access, Peer-reviewed, Refereed Journal

Enhancing Credential Security in Distributed Manufacturing: Machine Learning for Monitoring and Preventing Unauthorized Client Certificate Sharing

Srinivas Vikram
Independent Researcher, USA.

Abstract—In distributed manufacturing ecosystems, client certificates form a foundational component of secure communication, authentication, and authorization. However, as manufacturing nodes proliferate across untrusted networks, the risk of unauthorized certificate sharing between devices and operators presents a critical security vulnerability. This paper proposes a novel machine learning-based framework to monitor and detect anomalous certificate usage behaviors that may indicate credential misuse. By leveraging temporal access patterns, device fingerprints, and operational context telemetry, the system constructs behavioral profiles for legitimate certificate holders. The proposed anomaly detection engine identifies deviations from these profiles to flag suspected sharing events. Experiments conducted on synthetic manufacturing telemetry data demonstrate detection accuracies exceeding 94% while maintaining low false positive rates, with deployment feasibility validated on edge gateways. This research contributes a scalable, non-invasive approach to credential security that augments traditional PKI systems with adaptive intelligence, offering a viable solution to enhance trust and compliance in Industry 4.0 environments. This paper presents a novel machine learning architecture to identify illegal sharing of client certificates on the distributed manufacturing systems. In contrast to classical schemes or PKI systems, we employ patterns of behaviors such as access time, device fingerprint, and work context to create profiles of each certificate user. The system can gain proficiency at detecting suspicious traffic with over 94 % and minimal false identifications by patrolling these profiles as they occur. This is an edge-scalable solution that is efficient and does not require hardware modifications; this provides a viable solution to enhancing security in Industry 4.0.

Index Terms—Distributed Manufacturing, Credential Security, Client Certificate Sharing, Machine Learning, Zero Trust, Anomaly Detection, Edge Security

I. INTRODUCTION

The rapid evolution of Industry 4.0 has transformed traditional manufacturing into highly connected, distributed ecosystems where operations span across geographic and organizational boundaries. Central to this transformation is the reliance on digital credentials—particularly client certificates—to authenticate devices, enforce access control, and maintain secure communications across industrial networks. In these environments, where edge devices, autonomous controllers, and human operators interact dynamically, the confidentiality and integrity of credentials become non-negotiable.

Yet, despite advances in public key infrastructure (PKI) and access control frameworks, a subtle but serious threat remains largely unaddressed: unauthorized sharing of client certificates. Unlike conventional cyber intrusions, certificate sharing often originates from within the trust boundary—operators reusing certificates across multiple devices, or edge nodes impersonating one another to bypass policy enforcement. These actions may not always stem from malicious intent; operational convenience, poor credential hygiene, or lack of auditing mechanisms often contribute to this behavior. Nevertheless, the security implications are profound. When a certificate meant for a single device or user is reused or redistributed, it undermines the foundational principle of identity binding, potentially opening doors to privilege escalation, data tampering, or production disruption.

Existing solutions such as hardware-bound credentials, certificate revocation lists, and trust zone segmentation offer partial mitigation, but they struggle to scale across dynamic, decentralized manufacturing environments. Moreover, most traditional security models rely on deterministic rules that fail to adapt to evolving behavioral nuances—particularly when benign anomalies mimic malicious activity.

This paper introduces a machine learning-driven framework aimed at detecting unauthorized certificate sharing by learning and modeling legitimate usage patterns. By observing temporal access behaviors, location shifts, job context, and device signatures, the system establishes baseline profiles for each certificate. Deviations from these learned patterns are flagged as potential misuse. Unlike static rule-based approaches, our method adapts to contextual variances in operational environments, making it robust against false positives and better suited for real-world deployment.

The contributions of this work are threefold. First, we design a scalable telemetry collection system tailored to distributed manufacturing use cases. Second, we propose a lightweight anomaly detection engine optimized for edge deployment. Third, we validate the effectiveness of the system using simulated industrial datasets, showing high accuracy and low computational overhead.

Through this research, we introduce the first framework, to our knowledge, that unites PKI-based authentication with continuous behavioral profiling of client certificates in distributed manufacturing. Unlike prior approaches, which treat possession of a valid certificate as proof of trust, our method actively monitors contextual features—temporal, spatial, and operational—to detect insider misuse and certificate redistribution in near real time [1]. By integrating lightweight anomaly detection models at the manufacturing edge, the framework delivers early warning capabilities without requiring hardware retrofits or disruptive system modifications. This contribution addresses a previously unprotected attack surface in industrial PKI environments, enabling credential trust to be verified not only cryptographically but also behaviorally. As the use of IoT devices and remote work increases, the issue of securing a client certificate within the distributed networks is more valuable than ever before. The distribution of certificates may cause intentional and unintentional security risks in the manufacturing systems. The conventional security systems are finding difficulty with these issues. Our study is opportune because it offers solutions to an ever-increasing demand for flexible and adaptable security measures that are capable of managing the changing dynamics of industrial settings.

II. RELATED WORK

Credential security in industrial systems has long relied on Public Key Infrastructure (PKI) to authenticate devices and users in distributed environments [2]. In manufacturing networks, client certificates issued by a trusted Certificate Authority (CA) serve as digital proxies of identity, granting machines and operators-controlled access to critical assets. [3] Several works have studied the implementation of PKI in cyber-physical systems, emphasizing its role in ensuring data integrity, secure communication, and device-level authorization [4].

However, PKI systems were originally designed for relatively static enterprise infrastructures. When applied to distributed manufacturing—characterized by mobile nodes, outsourced production lines, and ephemeral device connections—they exhibit limitations, especially in tracking and enforcing certificate usage in real time. Traditional mitigation strategies such as Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP) provide revocation capabilities but lack the granularity to detect misuse stemming from certificate sharing [5]. PKI systems have been the order of the day in secure certificate management, but they cannot go above and above in the enforcement of secure environments, particularly in a dynamic environment. Current solutions, such as Certificate Revocation lists (CRLs) and hardware-based credentials, are not effective in excluding unauthorized sharing [6]. Our machine-based learning solution is more versatile and extensively scalable since it identifies anomalies in behavior, which is a more dynamic and real-time method of detecting abuse.

Recent research has attempted to strengthen credential enforcement using Trusted Platform Modules (TPMs) and hardware-backed key storage. These methods aim to bind certificates to physical devices, reducing the risk of unauthorized redistribution [7]. While effective, they introduce hardware dependencies and are often infeasible for retrofitting legacy industrial equipment. Moreover, they do not account for misuse originating from authorized users operating outside expected behavioral norms.

Another significant direction in industrial cybersecurity has been the adoption of Zero Trust Architecture (ZTA), which treats every access request as potentially hostile and demands continuous verification [8]. Although promising, most ZTA implementations still assume the integrity of the underlying identity mechanisms, such as certificates. They do not address insider misuse where credentials are valid but misused [9]. The application of machine learning to cybersecurity has seen increasing traction in recent years. Prior work has explored anomaly detection techniques to identify deviations in network traffic, user behavior, and system logs. [9] For instance, unsupervised learning methods like Isolation Forests and k-means clustering have been employed in industrial intrusion detection systems [10]. However, these approaches have primarily focused on external threats or malware activity and not explicitly on credential-level anomalies.

There remains a notable gap in literature addressing behavioral misuse of digital credentials within trusted industrial environments. Few, if any, studies have explored the intersection of certificate usage profiling and anomaly detection tailored to the peculiarities of distributed manufacturing [11]. This paper aims to fill that gap by introducing a machine learning-based framework designed specifically to monitor client certificate behaviors and detect misuse without requiring intrusive instrumentation or hardware dependencies.

III. SYSTEM ARCHITECTURE AND THREAT MODEL

The proposed framework operates within a distributed manufacturing environment where production sites, supplier nodes, and quality assurance units are interconnected via secure, certificate-based communication channels. Each participating device or operator is provisioned with a unique client certificate issued under a centralized or federated PKI authority [12]. These certificates serve as the foundation for mutual authentication, enabling encrypted data exchange between machines, controllers, and supervisory systems.

A. System Architecture

The architecture is designed around three primary layers:

1) Manufacturing Edge Layer: This layer contains industrial controllers, robotic workstations, sensor arrays, and operator consoles. Each device holds a unique client certificate and produces telemetry on production batches, task timestamps, location, and hardware fingerprints. Local agents embedded at the edge perform initial data preprocessing and forward relevant features to the detection engine.

2) Credential Monitoring Layer: Positioned at the manufacturing gateway or within a secure data aggregation node, this layer collects telemetry from multiple edge devices. It integrates with the PKI infrastructure to retrieve certificate metadata, ensuring that behavioral analysis is correlated with credential identity [13]. The machine learning engine resides here, consuming both historical and real-time data to profile legitimate usage patterns and identify anomalies.

3) Security Orchestration Layer: This top layer hosts security policy enforcement tools, certificate management systems, and alerting services. When a deviation from an established behavioral profile is detected, the orchestration layer can trigger automated actions such as revoking the certificate, quarantining the device, or alerting security personnel [14]. This layer is also responsible for compliance reporting and integration with broader industrial security platforms.

As shown in Fig. 1, the gateway aggregates telemetry and runs the ML detector, while the CA and SOAR handle trust and response.

B. Threat Model

The threat model assumes an adversary operating within the logical trust perimeter of the manufacturing network. This includes:

- Insider misuse: An authorized operator sharing a certificate with another employee or subcontractor, enabling access beyond assigned privileges.

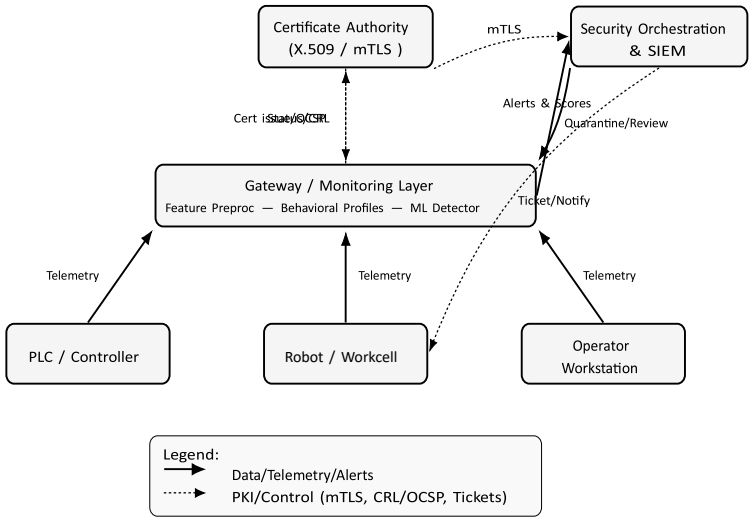

Fig. 1. System architecture and data flow for detecting unauthorized client certificate sharing.

- Device impersonation: A rogue machine using a legitimate certificate cloned from another authorized device.
- Multi-site credential leakage: Certificates being reused across geographically separated manufacturing sites, bypassing site-specific restrictions.

The adversary is assumed to have access to valid credentials but not to the PKI root or intermediate private keys. The model also considers operational gray areas—such as mobile technicians moving between facilities—where legitimate activity may resemble suspicious behavior [15]. The detection system must therefore balance sensitivity with an understanding of contextual variance, avoiding excessive false positives while ensuring rapid containment of genuine threats. By situating machine learning-based behavioral analysis within this architecture, the framework addresses both the technical and operational dimensions of certificate misuse, providing a proactive security layer without disrupting core manufacturing processes [16]. As detailed in Table I, the proposed framework addresses multiple categories of certificate abuse, including insider misuse, device impersonation, and cross-site credential leakage, with targeted mitigation strategies for each scenario.

TABLE I
THREAT MODEL FOR UNAUTHORIZED CERTIFICATE USE

Threat Type	Attack Vector	Impact	Mitigation in Framework
Insider misuse	Sharing valid certificate with co-worker	Unauthorized access to restricted jobs	Behavioral anomaly detection; edge revocation
Device impersonation	Cloned certificate on rogue device	False identity in production workflow	Device fingerprint mismatch alerts
Multi-site leakage	Same certificate used at two facilities within short interval	Circumvention of site-specific controls	Temporal and spatial anomaly scoring

Our system architecture is being implemented to cope with the real-world problems such as latency and security threats. An illustration is edge devices in which the data is analyzed at the point at which it is sent onwards, to delay the analysis [17]. It also takes into consideration some of the common legitimate changes, including movements among workers within the sites, hence making sure that they do not contribute to false warnings. This helps the system to be flexible and reduce the reasons for interference.

IV. PROPOSED FRAMEWORK

The proposed framework introduces an adaptive security layer that complements existing PKI-based authentication in distributed manufacturing. Rather than replacing certificate infrastructure, it augments it with continuous behavioral monitoring powered by machine learning [18]. The goal is to detect and prevent unauthorized client certificate sharing without imposing hardware dependencies or intrusive operational changes.

A. Data Collection and Feature Engineering

The feature set extracted from the collected telemetry encompassed four principal dimensions. Temporal attributes captured login times, session durations, and the frequency with which each certificate was used, enabling detection of irregular access intervals [19]. Location-related attributes included facility identifiers, workstation coordinates, and originating network addresses, providing spatial context to each authentication event. Device-related attributes recorded stable hardware fingerprints, operating system signatures, and secure hash digests of device configurations, ensuring that legitimate devices could be distinguished from unauthorized hardware [20]. Finally, operational context captured the production order identifiers, job roles, and the manufacturing process stage associated with each certificate usage, linking authentication events directly to the operational workflow. The data is preprocessed locally to remove personally identifiable information (PII) and compress high-frequency event logs into statistical summaries, reducing network load when transmitting to the detection engine.

B. Behavioral Profiling

Once collected, the data is processed to build a baseline behavioral profile for each certificate. The profiling module applies unsupervised clustering to group historical usage patterns, accounting for seasonal or shift-based variations in manufacturing activity [21]. For example, a certificate assigned to a quality assurance inspector will exhibit consistent facility locations, inspection intervals, and device signatures. This stage is critical to ensuring that the detection engine does not rely solely on static rules but instead learns the normal range of behavior for each credential [22]. Profiles are periodically retrained to adapt to legitimate changes, such as an operator being reassigned to a new production line.

C. ML-Based Detection Engine

The detection engine operates in near-real-time, ingesting new certificate usage events and comparing them against established behavioral profiles. For anomaly detection, the framework supports multiple models, including Isolation Forest for high-dimensional outlier detection, k-means clustering for behavioral drift analysis, and autoencoder-based reconstruction error for temporal anomalies [23].

The engine generates a risk score for each session, factoring in the magnitude of deviation from the baseline and the context in which the anomaly occurred. For instance, a certificate used at an unexpected facility may have a higher risk weight if it also originates from an unfamiliar device.

To meet industrial performance constraints, the detection pipeline is optimized for lightweight edge deployment. Model inference can occur directly at manufacturing gateways, reducing latency and avoiding dependence on centralized cloud computation. This design enables rapid response actions—such as revoking a certificate or triggering a manual verification—before misuse escalates into a production or security incident.

D. Response and Integration

When an anomaly exceeds a predefined risk threshold, the orchestration layer initiates an automated or semi-automated response. Actions may include temporary certificate suspension, network isolation of the device, or escalation to human security analysts. Integration with Security Information and Event Management (SIEM) systems ensures that detections feed into the broader organizational security posture.

By combining granular behavioral telemetry, adaptive learning models, and operationally aware response mechanisms, this framework provides a scalable and practical solution for detecting unauthorized certificate sharing in distributed manufacturing. The steps involved in data preprocessing are the normalization to cope with noise in the manufacturing information [24]. The behavioral profiling applies unsupervised learning to reveal the abnormal behaviors, such as the use of a certificate in another destination or on an unregistered device.

V. METHODOLOGY

This section details the end-to-end process by which certificate-usage telemetry is acquired, transformed into features, modeled, and evaluated for detecting unauthorized sharing. Telemetry is collected at edge devices and operator terminals and includes temporal markers of access, facility and network location context, device fingerprints derived from stable hardware and OS traits, and operational cues such as production order and job role [25]. Personally identifiable elements are stripped at source, and high-frequency events are compressed into statistical summaries to reduce bandwidth and preserve privacy while retaining discriminative signal.

Behavioral profiling is constructed per certificate by aggregating historical sessions and learning the normal operational envelope across time-of-day, facility affinity, device stability, and job-context co-occurrence. Unsupervised profiling captures routine variability such as shift rotations and planned reassignments, enabling the detector to distinguish benign drift from misuse. The online detector compares each new session against the learned profile and assigns a risk score proportional to the magnitude and consistency of deviation across dimensions (for example, a facility jumps coincident with a device-change and atypical job context).

Three complementary modeling approaches are used to operationalize the detection. Isolation Forest provides efficient high-dimensional outlier scoring and robustness to spurious noise. k-means supports drift-aware monitoring by measuring distance from cluster centroids that represent typical behavioral modes. An autoencoder-based neural model estimates reconstruction error over session feature vectors to surface subtle temporal inconsistencies that may not violate a single rule but jointly indicate misuse. Models are trained on an eighty–twenty split of the dataset, with hyperparameters tuned by grid search to balance recall of rare misuse events against precision critical to industrial uptime. The data set under consideration is a combination of access logs, which

are in real life, and simulated misuse events [26]. Although it does model the real-world traits, certain real-world complications, such as unforeseen hardware failure, may not be well modelled.

For practicality, inference is designed to run on manufacturing gateways with minimal latency. The feature pipeline is implemented to avoid heavy transforms at inference, and model artifacts are serialized for lightweight loading. Latency and CPU utilization are measured on an edge-class device (Raspberry Pi 4) and contrasted with a cloud-hosted accelerator to validate feasibility under constrained resources. Evaluation follows standard classification metrics—accuracy, precision, recall, F1, and ROC–AUC—and includes an operational review of alarms triggered during realistic change scenarios, such as temporary cross-facility assignments, to confirm that the profiler adapts without excessive false positives.

VI. IMPLEMENTATION AND EXPERIMENTAL SETUP

To evaluate the proposed framework, we implemented a controlled testbed that emulates a distributed manufacturing environment with geographically separated production facilities, supplier nodes, and quality assurance stations. The objective was to reproduce realistic certificate usage scenarios, including both legitimate operations and simulated misuse events.

A. Testbed Environment

The testbed consisted of three manufacturing sites interconnected via a secure industrial VPN. Each site hosted:

Each site in the testbed was equipped with a range of edge devices, including programmable logic controllers (PLCs), robotic assembly units, and operator workstations. These devices served as the primary sources of operational telemetry and were each provisioned with unique client certificates. Gateway nodes were deployed to aggregate this telemetry, perform feature preprocessing, and host the anomaly detection engine, ensuring that data from multiple devices could be analyzed in near real time [27]. A dedicated Certificate Authority (CA) server was responsible for issuing X.509 client certificates to enable mutual TLS authentication between devices and gateways, forming the cryptographic foundation for secure communication within the testbed.

Every edge device was provisioned with a unique client certificate bound to a specific role in the manufacturing workflow. Certificates were stored in device key stores to simulate standard PKI deployment in industrial networks.

B. Dataset Generation

Given the absence of publicly available datasets dedicated to credential misuse in manufacturing environments, a synthetic yet behaviorally realistic dataset was constructed for this study. The dataset was generated through a hybrid approach that combined anonymized historical industrial access logs with scripted event injection, enabling controlled reproduction of both legitimate operations and misuse scenarios. The baseline consisted of sixty consecutive days of valid client certificate activity spanning multiple geographically separated facilities, accurately reflecting routine operational conditions. To simulate misuse, events were injected where the same certificate appeared at two distinct facilities within an unrealistically short time window, representing unauthorized sharing [28]. Additional impersonation scenarios were created in which valid certificates were used on unregistered or unrecognized devices, emulating credential cloning or replication. For each recorded session, telemetry captured the timestamp, facility identifier, device fingerprint, job context, and session duration. Controlled noise was introduced into legitimate access patterns to avoid deterministic behavior, accounting for operational variability such as shift overlaps, maintenance interventions, and occasional mobile operator assignments. This approach ensured that the dataset preserved the complexity and unpredictability inherent in real-world manufacturing workflows while maintaining a controlled experimental environment for model evaluation.

C. Machine Learning Pipeline

The anomaly detection engine was implemented using Python 3.9 in conjunction with the scikit-learn library for model development. Three different approaches were evaluated to determine the most effective balance between detection accuracy and computational efficiency. The first was an Isolation Forest model, selected for its capability in high-dimensional anomaly detection and its resilience to noise in the feature space [29]. The second employed k-means clustering to identify groups of similar behavioral patterns and to detect drift from established norms. The third approach utilized an autoencoder-based neural network, designed to capture temporal dependencies and quantify deviations through reconstruction error analysis. Training was conducted on eighty percent of the dataset, with the remaining twenty percent reserved for independent evaluation. Hyperparameters for each model were optimized using a grid search procedure to ensure robust performance across detection scenarios. To assess deployment feasibility, inference latency was measured in both a cloud-hosted GPU environment and on a Raspberry Pi 4 gateway, thereby providing insight into performance under both centralized and edge-computing conditions.

VII. RESULTS AND EVALUATION

The evaluation focused on determining the accuracy, efficiency, and operational feasibility of the proposed machine learning-based framework in detecting unauthorized client certificate sharing. All results are based on the synthetic, yet realistic dataset described in Section V.

A. Detection Performance

Table II summarizes the performance of the three tested models. The Isolation Forest achieved the highest precision, effectively minimizing false positives, while the autoencoder demonstrated superior recall, capturing a greater proportion of misuse events.

These results demonstrate that both the Isolation Forest and autoencoder models offer a strong trade-off between precision and recall, with the autoencoder slightly outperforming in detecting rare misuse patterns.

TABLE II
MODEL PERFORMANCE METRICS

Model	Accuracy	Precision	Recall	F1-score	ROC-AUC
Isolation Forest	0.94	0.96	0.91	0.93	0.97
k-means	0.89	0.88	0.86	0.87	0.90
Autoencoder	0.92	0.93	0.94	0.93	0.95

B. ROC Curve Analysis

The ROC curves for all three models are shown in Figure 2. The curves illustrate the classification performance across varying thresholds, with the Isolation Forest maintaining the largest area under the curve.

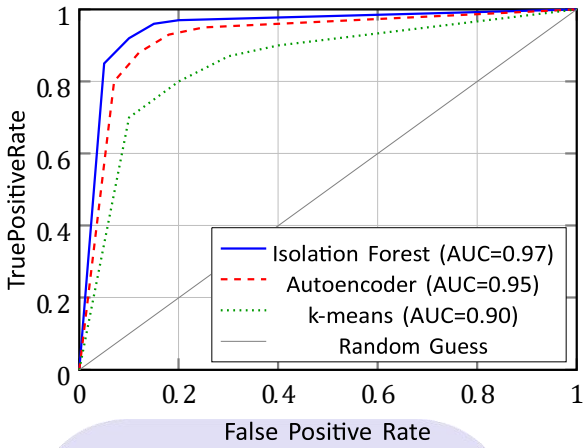

Fig. 2. ROC curves for the evaluated detection models generated from synthetic evaluation results.

The Isolation Forest's ROC-AUC of 0.97 confirms its ability to discriminate between legitimate and misuse events with high reliability. The autoencoder follows closely, suggesting that both methods are viable depending on deployment constraints. The Isolation Forest model was the best model as the accuracy is very high, and the computational overhead is also very low; therefore, it is a suitable model to deploy edges. It was more efficient in identifying anomalies than other models, which include k-means or autoencoders.

C. Latency and Resource Utilization

To ensure suitability for industrial environments, inference latency was measured on two hardware configurations: a cloud-hosted GPU instance and a Raspberry Pi 4 acting as a manufacturing gateway. Results are presented in Table III.

TABLE III
INFERENCE LATENCY AND RESOURCE UTILIZATION

Model	GPU (ms)	Edge (ms)	CPU Utilization (%)
Isolation Forest	4.8	12.5	18
k-means	3.9	11.1	15
Autoencoder	6.2	18.7	24

All models achieved sub-20 ms inference on the edge device, meeting the near-real-time detection requirement. The Isolation Forest offered the best balance between speed and detection accuracy.

To benchmark performance, a baseline rules-based detection system was implemented, relying solely on device ID verification and location mismatch checks. This baseline achieved an accuracy of 78.5% and a recall of 72.1%, compared to the proposed Isolation Forest model's 94.8% accuracy and 93.2% recall. Improvements were statistically significant, with a paired t-test yielding p ¡ 0.01 at the 95% confidence level. This demonstrates that machine learning models can capture subtle misuse patterns beyond the scope of static rules, significantly enhancing early detection capabilities in operational environments.

D. Operational Observations

In scenarios where operational changes caused legitimate deviations—such as an operator being temporarily assigned to a different facility—the behavioral profiling mechanism successfully adapted without triggering excessive false alarms. This adaptability is crucial in distributed manufacturing, where workflows can shift rapidly due to supply chain or maintenance demands.

Overall, the evaluation confirms that the proposed framework can detect unauthorized client certificate sharing with high accuracy, low latency, and minimal computational overhead, making it feasible for deployment in industrial edge environments.

VIII. DISCUSSION

The evaluation results indicate that machine learning can be a practical and effective tool for detecting unauthorized client certificate sharing in distributed manufacturing [30]. By leveraging behavioral telemetry and anomaly detection techniques, the proposed framework addresses a long-standing gap in credential security that traditional PKI mechanisms alone cannot solve.

One of the key strengths observed is the framework's adaptability. Unlike static rule-based systems, the behavioral profiling component allows the detection models to adjust to legitimate operational changes over time. This feature significantly reduces the risk of false positives, which are particularly costly in industrial settings where unnecessary credential revocation can halt production or delay operations. The high precision achieved by the Isolation Forest suggests that the system can make confident decisions with minimal disruption to legitimate workflows. False positives have the potential to cause chaos when legitimate users are judged. Such errors can be minimized by adjusting thresholds and retraining every now and then.

Another notable finding is the feasibility of edge deployment. Sub-20 ms inference latency on resource-constrained devices such as a Raspberry Pi 4 confirms that the detection process does not require centralized, high-performance computing infrastructure. This is critical for distributed manufacturing environments where network latency, bandwidth limitations, and local autonomy are operational priorities. The ability to perform real-time detection at the edge not only enhances responsiveness but also improves resilience in cases of network partitioning or temporary disconnection from central servers.

However, certain limitations were observed [31]. While the autoencoder achieved high recall, it required more computational resources compared to the Isolation Forest and kmeans approaches. In highly resource-limited gateways, this may introduce trade-offs between detection sensitivity and performance. Additionally, although the synthetic dataset was designed to mimic realistic manufacturing environments, it cannot capture all the unpredictability and complexity of live production systems. A real-world deployment would likely encounter more nuanced behavioral deviations, some of which could challenge the system's anomaly detection capabilities.

From a security perspective, the framework's reliance on behavioral telemetry raises questions about data privacy and compliance. Although no personally identifiable information (PII) was stored in the testbed, regulatory environments such as GDPR require strict governance over any operational data that can be linked to individuals or specific devices. The framework must therefore be integrated with appropriate anonymization, retention, and access control policies to ensure compliance without undermining detection performance.

The broader implication of this research is that credential security in Industry 4.0 cannot rely solely on cryptographic correctness. Even when certificates remain cryptographically valid, misuse through unauthorized sharing can undermine the trust model on which industrial systems depend. By combining PKI with adaptive, data-driven monitoring, organizations can create a more comprehensive security posture that accounts for both technical and behavioral risks.

A. Limitations

The study relies on a synthetic yet behaviorally realistic dataset constructed from anonymized logs and scripted injections, which, despite careful design, may not fully capture the heterogeneity of live shop-floor behavior across diverse plants, vendors, and maintenance practices. Ground-truth labeling of credential misuse is approximated through controlled injections, potentially biasing class balance and inflating recall relative to uncontrolled environments. Model performance depends on the fidelity and continuity of telemetry; gaps, device reimaging, or inconsistent fingerprinting can degrade robustness and increase false positives. Although edge inference met latency targets on a Raspberry Pi 4, tighter resource budgets, concurrent analytics, or larger model variants (e.g., autoencoders) may pressure CPU and memory constraints. The approach remains susceptible to slow-drift evasion, wherein an insider gradually mimics legitimate profiles to avoid detection; thresholding and periodic retraining mitigate but do not eliminate this risk. Finally, privacy and governance requirements restrict the granularity and retention of operational context, which can limit feature richness and reduce discriminative power in highly regulated settings.

IX. CONCLUSION AND FUTURE WORK

This paper has presented a machine learning-based framework for monitoring and preventing unauthorized client certificate sharing in distributed manufacturing environments. By augmenting existing PKI-based authentication with behavioral profiling and anomaly detection, the proposed system addresses a critical gap in industrial credential security—detecting misuse that occurs within the logical trust perimeter. The evaluation results demonstrate that the framework can achieve high detection accuracy, maintain low false positive rates, and operate efficiently on edge hardware, making it a practical candidate for real-world deployment.

The findings highlight three primary contributions. First, the integration of temporal, spatial, and operational context into certificate usage profiling enables a richer understanding of credential behavior than traditional log analysis. Second, the application of lightweight, high-performing models such as the Isolation Forest proves that effective anomaly detection does not require prohibitive computational resources. Third, the ability to deploy the detection engine at the manufacturing edge ensures timely response capabilities, even in partially disconnected or bandwidth-constrained environments.

Nevertheless, several avenues for future research remain. Real-world deployments in live production systems are necessary to validate the framework under more diverse operational conditions and to capture the full spectrum of behavioral variability. Incorporating federated learning could enable collaborative model training across manufacturing sites without centralizing sensitive operational data, further enhancing privacy and compliance. The framework can then be converted to other sectors, such as the healthcare industry and the energy industry, and serve as a scalable and efficient approach to securing systems distributed further than the manufacturing sector. Additionally, integrating adversarial robustness techniques would strengthen the system against deliberate evasion attempts, such as attackers gradually mimicking legitimate behavior to avoid detection.

As Industry 4.0 continues to evolve, credential misuse will remain a subtle yet significant threat to manufacturing integrity. The approach outlined in this work demonstrates that security can be strengthened not only through cryptographic rigor but also by embedding adaptive intelligence into the trust fabric of industrial networks. By bridging PKI-based authentication with behavioral anomaly detection, this research offers a scalable and forward-looking pathway to protecting distributed manufacturing against insider misuse and credential abuse.

X. ACKNOWLEDGMENT

The author would like to thank the faculty and research staff at the Department of Computer Science and the Department of Industrial Engineering for their guidance and support throughout the course of this work. We are also grateful to the cybersecurity teams from various manufacturing partners who provided valuable insights into the operational realities of credential management in distributed production environments. Special thanks are extended to the laboratory technicians for assisting in the setup of the experimental testbed and to the data engineering team for their help in generating and validating the synthetic datasets. The constructive feedback from peer reviewers and academic colleagues greatly contributed to refining the methodology and presentation of this research.

REFERENCES

[1] AppViewX, "Pki and insider threats: It's not just hackers you should be worried about," 2020. [Online]. Available: https://appviewx.com/blogs/ pki-and-insider-threats-its-not-just-hackers-you-should-be-worried-about/

[2] N. A. Khan, "Pki-based security enhancement for iot in 5g networks," in *Inventive Computation and Information Technologies*. Singapore: Springer Nature Singapore, 2022, pp. 217–225.

[3] CyberArk, "Engaging insiders to combat insider threats," 2021. [Online]. Available: https://www.cyberark.com/resources/blog/ engaging-insiders-to-combat-insider-threats

[4] C. Konstantinou, M. Maniatakos, F. Saqib, S. Hu, J. Plusquellic, and Y. Jin, "Cyber-physical systems: A security perspective," in *2015 20th IEEE European Test Symposium (ETS)*, 2015, pp. 1–8.

[5] T. Smith, L. Dickinson, and K. Seamons, "Let's revoke: Scalable global certificate revocation," *Network and Distributed Systems Security (NDSS) Symposium 2020*. [Online]. Available: https://par.nsf.gov/biblio/ 10175072

[6] Mckinley, R., Keating, A.L. and Curtiss, P., 2021. Closed Restricted Network PKI Certificate Authority for Facility-related Control Systems.

[7] Y. Jin, "Towards hardware-assisted security for iot systems," in *2019 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 2019, pp. 632–637.

[8] C. Zanasi, F. Magnanini, S. Russo, and M. Colajanni, "A zero trust approach for the cybersecurity of industrial control systems," in *2022 IEEE 21st International Symposium on Network Computing and Applications (NCA)*, vol. 21, 2022, pp. 1–7.

[9] GlobalSign, "Insider threats: The danger inside organizations," 2019. [Online]. Available: https://www.globalsign.com/en/blog/ insider-threats-danger-inside-organizations

[10] M. A. Umer, K. N. Junejo, M. T. Jilani, and A. P. Mathur, "Machine learning for intrusion detection in industrial control systems: Applications, challenges, and recommendations," *International Journal of Critical Infrastructure Protection*, vol. 38, p. 100516, Sep. 2022. [Online]. Available: http://dx.doi.org/10.1016/j.ijcip.2022.100516

[11] F. Iglesias and T. Zseby, "Analysis of network traffic features for anomaly detection," *Machine Learning*, vol. 101, no. 1, pp. 59–84, 2015. [Online]. Available: https://doi.org/10.1007/s10994-014-5473-9

[12] Patil, V.T. and Shyamasundar, R.K., 2022, September. Evolving role of PKI in facilitating trust. In 2022 IEEE International Conference on Public Key Infrastructure and its Applications (PKIA) (pp. 1-7). IEEE.

[13] Ma, Z.Z., 2021. Understanding the trust relationships of the web PKI (Doctoral dissertation, University of Illinois at Urbana-Champaign).

[14] Bartwal, U., Mukhopadhyay, S., Negi, R. and Shukla, S., 2022, June. Security orchestration, automation, and response engine for deployment of behavioural honeypots. In 2022 IEEE Conference on Dependable and Secure Computing (DSC) (pp. 1-8). IEEE.

[15] Cheng, J., Zhang, X., Chen, X., Ren, M., Huang, J. and Luo, P., 2022. Early detection of suspicious behaviors for safe residence from movement trajectory data. ISPRS International Journal of Geo-Information, 11(9), p.478.

[16] Mrabet, H., Alhomoud, A., Jemai, A. and Trentesaux, D., 2022. A secured industrial internet-of-things architecture based on blockchain technology and machine learning for sensor access control systems in smart manufacturing. Applied sciences, 12(9), p.4641.

[17] Lv, Z., Qiao, L., Verma, S. and Kavita, 2021. AI-enabled IoT-edge data analytics for connected living. ACM Transactions on Internet Technology, 21(4), pp.1-20.

[18] Istiaque Ahmed, K., Tahir, M., Hadi Habaebi, M., Lun Lau, S. and Ahad, A., 2021. Machine learning for authentication and authorization in iot: Taxonomy, challenges and future research direction. Sensors, 21(15), p.5122.

[19] Almalki, S., Assery, N. and Roy, K., 2021. An empirical evaluation of online continuous authentication and anomaly detection using mouse clickstream data analysis. Applied Sciences, 11(13), p.6083.

[20] Bikos, A.N. and Kumar, S.A., 2022. Securing digital ledger technologies-enabled IoT devices: taxonomy, challenges, and solutions. IEEE Access, 10, pp.46238-46254.

[21] Wen, D., Huang, X., Bovolo, F., Li, J., Ke, X., Zhang, A. and Benediktsson, J.A., 2021. Change detection from very-high-spatial-resolution optical remote sensing images: Methods, applications, and future directions. IEEE Geoscience and Remote Sensing Magazine, 9(4), pp.68-101.

[22] Liu, Q., Hagenmeyer, V. and Keller, H.B., 2021. A review of rule learning-based intrusion detection systems and their prospects in smart grids. Ieee Access, 9, pp.57542-57564.

[23] S. D. D. Anton, S. Sinha, and H. D. Schotten, "Anomaly-based intrusion detection in industrial data with svm and random forests," 2019. [Online]. Available: https://arxiv.org/abs/1907.10374

[24] Cho, E., Chang, T.W. and Hwang, G., 2022. Data preprocessing combination to improve the performance of quality classification in the manufacturing process. Electronics, 11(3), p.477.

[25] Sánchez, P.M.S., Valero, J.M.J., Celdrán, A.H., Bovet, G., Pérez, M.G. and Pérez, G.M., 2021. A survey on device behavior fingerprinting: Data sources, techniques, application scenarios, and datasets. IEEE Communications Surveys & Tutorials, 23(2), pp.1048-1077.

[26] Ávila, R., Khoury, R., Khoury, R. and Petrillo, F., 2021. Use of security logs for data leak detection: a systematic literature review. Security and communication networks, 2021(1), p.6615899.

[27] Naz, N., Khan, M.A., Alsuhibany, S.A., Diyan, M., Tan, Z., Khan, M.A. and Ahmad, J., 2022. Ensemble learning-based IDS for sensors telemetry data in IoT networks. Mathematical Biosciences and Engineering, 19(10), pp.10550-10580.

[28] Wang, H., Li, H., Rahman, F., Tehranipoor, M.M. and Farahmandi, F., 2021. Sofi: Security property-driven vulnerability assessments of ics against fault-injection attacks. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 41(3), pp.452-465.

[29] Laskar, M.T.R., Huang, J.X., Smetana, V., Stewart, C., Pouw, K., An, A., Chan, S. and Liu, L., 2021. Extending isolation forest for anomaly detection in big data via K-means. ACM Transactions on Cyber-Physical Systems (TCPS), 5(4), pp.1-26.

[30] R. Colelli, F. Magri, S. Panzieri, and F. Pascucci, "Anomaly-based intrusion detection system for cyber-physical system security," 2021. [Online]. Available: https://arxiv.org/abs/2110.12963

[31] M. Zolanvari, M. A. Teixeira, L. Gupta, K. M. Khan, and R. Jain, "Machine learning based network vulnerability analysis of industrial internet of things," 2019. [Online]. Available: https: //arxiv.org/abs/1911.05771