# Zero Trust Defense Against Charge Manipulation Attacks in Smart EV Charging Infrastructure

Saba Marandi
*CIISE Department*
*Concordia University*
Montreal, Canada
email: saba.marandi@mail.concordia.ca

Danial Jafarigiv
*Digitization & Cybersecurity, R&D Department*
*Hydro-Québec Research Institute*
Varennes, Canada
email: jafarigiv.danial2@hydroquebec.com

Ribal Atallah
*Digitization & Cybersecurity, R&D Department*
*Hydro-Québec Research Institute*
Varennes, Canada
email: atallah.ribal@hydroquebec.com

Mohsen Ghafouri
*CIISE Department*
*Concordia University*
Montreal, Canada
email: mohsen.ghafouri@concordia.ca

Chadi Assi
textitCIISE Department
*Concordia University*
Montreal, Canada
email: chadi.assi@concordia.ca

***Abstract*—The increasing reliance on Electric Vehicle Charging Infrastructure (EVCI) within smart grids introduces new threats to grid stability and operator control. Among all, Charge Manipulation Attacks (CMAs), in which attackers use access to the charging interface or backend systems to maliciously modify the charging parameters, e.g., the amount of energy requests, charging durations, are of significant importance. These attacks generate false load profiles and disrupt grid operations. This study develops a Zero Trust-based security mechanism to safeguard charging stations and the operator's backend infrastructure, including control and management systems, from such attacks. Our solution employs a state-aware, Markov-based trust evaluation model that utilizes three behavioral indicators to continuously monitor the evolving behavior of chargers. Specifically, real-time behavioral indicators, such as request anomalies, charger usage history, and spatiotemporal consistency, are mapped into evolving trust states, which enable probabilistic prediction of charger trustworthiness. A central Zero Trust Controller issues short-lived tokens to charging stations only when the predicted state indicates acceptable trust levels. These tokens are enforced locally at each charger via a Policy Enforcement Point (PEP), which ensures strict, session-level verification. Simulation results demonstrate that the proposed architecture detects and blocks malicious charging behavior, which maintains secure and stable EVCS operations even under adversarial conditions.**

***Index Terms*—Electric Vehicle Charger Stations; Zero-Trust; OCPP; Charge Manipulation Attacks.**

## I. Introduction

As Electric Vehicles (EVs) become increasingly integrated into modern transportation systems, the supporting infrastructure, i.e., EV Charging Stations (EVCSs), has grown rapidly in scale and complexity [1]. This growth is accompanied by increasing dependence on cloud-based management platforms, charger-side control logic, and remote access protocols that facilitate automated or operator-triggered charging sessions [2], [3]. While these advancements offer operational efficiency and flexibility, they also introduce critical cybersecurity risks [4]. Threats, such as load manipulation, unauthorized firmware command injection, protocol misuse, and denial-of-service attacks, are now emerging across both physical and cyber layers of the EVCS infrastructure [5]. Traditional perimeter-based security models, which rely on fixed trust boundaries and static credential checks, are no longer sufficient to protect these dynamic and distributed systems. In broader vehicle-to-grid deployments, where EVCSs interact continuously with grid operators, aggregators, and third-party service platforms, the attack surface expands considerably [6]. These risks highlight the need for a policy-driven framework that models charger behavior, assesses trust in real time, and enforces localized protection to preserve grid stability.

To mitigate the aforementioned risks, several studies have proposed trust-based security frameworks relying on role-based or behavior-based evaluations of connected entities. For example, hierarchical and fuzzy logic-based trust models have been developed in [7] and [8] to control access to cloud-based services, energy usage records, or IoT data streams. These models typically compute trust scores using attributes, such as historical access behavior, device identity, or policy compliance. While these approaches provide a foundation for anomaly detection, they are mainly designed for user- or application-level control and often fail to capture the protocol-level behavior of EVCSs. Moreover, they suffer from static trust assumptions and lack real-time adaptability, which makes them ineffective against evolving threats, such as command manipulation anomalies. In particular, they cannot detect adversaries exploiting low-level charger protocols or adapt tactics to evade detection.

In parallel, recent research has identified three prominent categories of Charge Manipulation Attacks (CMAs) targeting EVCSs, namely sudden demand surges, coordinated switching, and market manipulation [9]. Sudden surges may result from adversaries remotely triggering large numbers of chargers or using social engineering techniques to induce synchronized charging activity [10]. Coordinated switching attacks involve repeated on-off cycling of chargers at low frequencies, which can trigger stability issues, such as inter-area oscillations [11], [12]. Furthermore, by controlling large-scale EVCS operations, adversaries can artificially inflate or deflate demand to exploit fluctuations in energy market pricing [13]. While standards like OCPP 1.6 and 2.0.1 incorporate protective features, such as randomized delays, these measures can be circumvented by attackers with elevated access privileges or through exploitation

of protocol-level controls [9]. These limitations highlight the need for adaptive, behavior-aware security strategies.

As a result, recent efforts have increasingly turned toward Zero Trust Architecture (ZTA) as a resilient framework for securing EVCS operations. Given the dynamic and distributed nature of these environments—where charging sessions can be initiated remotely via backend systems, firmware triggers, or control center APIs—continuous verification and real-time enforcement are essential. Unlike traditional models that rely on pre-established trust boundaries, ZTA mandates dynamic evaluation of every access or control request based on contextual indicators, device integrity, communication behavior, and operational risk. For example, [14] presents a ZTA-inspired fault detection system using distributed attestation for intelligent vehicles; however, their work is limited to in-vehicle contexts and does not extend to charger operations. Similarly, [6] proposes a ZTA-based access control model within vehicle-to-grid systems, but this approach is constrained to cloud-side data handling and lacks real-time integration with charging station protocols. The high-level framework in [15] outlines identity and policy management components for ZTA in EVCS, but does not address protocol-specific enforcement (e.g., OCPP validation), dynamic charger-side anomaly detection, or simulation under realistic attack scenarios. Therefore, a significant gap remains in deploying Zero Trust strategies directly within EVCS control logic to counter the manipulation of charging behaviors that may jeopardize grid stability.

In addition to this architectural gap, there are methodological shortcomings that limit current approaches. Although ZTA is promising for cyber-physical system security, current applications to EVCI remain limited to cloud-side control or architectural proposals, often overlooking protocol-level vulnerabilities within EVCS deployments. In particular, the dynamic nature of charger-initiated requests, behavioral deviations at the session level, and misuse of standard protocols like OCPP 2.0.1 are rarely addressed with real-time and localized enforcement mechanisms. Moreover, existing trust evaluation methods often rely on static or short-term assessments, failing to capture how charger behavior evolves over time. To overcome these limitations, this work integrates a state-aware, Markov-based trust evaluation model into the ZTA framework, which enables predictive, dynamic assessment of charger trustworthiness based on evolving behavioral indicators. In summary, this paper makes the following key contributions:

- A novel Zero Trust-based security framework for EVCI is proposed, in which charging station behavior is continuously assessed, and energy delivery sessions are authorized using dynamically issued trust tokens based on real-time behavioral indicators.
- A Markov-based trust assessment model that is state-aware is incorporated, which captures the evolving nature of charger behavior under both normal and adversarial conditions. Key behavioral indicators, such as session frequency, recurring user associations, and statistical deviations from normal patterns, are modeled as evolving trust states with defined transition probabilities. Access control decisions are made by lightweight Policy Decision Points (PDPs), which evaluate trust conditions and issue short-lived tokens, while Policy Enforcement Points (PEPs) embedded within each EVCS locally enforce these decisions.
- The proposed Zero Trust framework is seamlessly integrated with OCPP 2.0.1 protocol exchanges, which enables real-time inspection and enforcement of session-level security policies within standard EVCS communication workflows. Its effectiveness is demonstrated through the simulation of various charger behaviors under benign and adversarial conditions, demonstrating the capability of the framework to detect and prevent unauthorized charging operations.

The remainder of this paper is organized as follows. In Section II, we present the proposed Zero Trust enforcement framework for secure EV charging infrastructure, including the system overview, threat model, trust evaluation methodology, and policy enforcement mechanisms. In Section III, we describe the simulation setup and present comprehensive results demonstrating the effectiveness of our approach under various attack scenarios. Finally, Section IV concludes the paper and discusses future research directions.

## II. Proposed Zero Trust Enforcement Framework for Secure EV Charging Infrastructure

To bring the Zero Trust approach into practice for EV charging infrastructure, this section outlines the components and mechanisms of our proposed framework. At the heart of the design is the idea that no charging request should be trusted by default, even if it comes from a known source. Instead, each interaction is evaluated in real time based on behavioral context, device integrity, and protocol activity. To enhance predictive detection of malicious behavior, a state-aware trust evaluation model based on evolving charger behavior is incorporated into the framework. The system introduces lightweight enforcement modules at the station level, supported by a central zero trust controller that issues short-term trust tokens. These tokens are granted only when a request aligns with expected patterns and policy rules. Our framework is designed with real-world constraints in mind, including the evolving nature of attacks. The following subsections describe the system setup, how trust is calculated and enforced, and how the framework integrates with existing protocols like OCPP to secure EVCS operations.

### A. System Overview & Threat Model

The proposed Zero Trust-based EVCI consists of interconnected components designed to ensure secure communication and resilient energy delivery. As illustrated in Figure 1, the architecture includes EVs, EVCSs, a centralized EV charging management system, i.e, EVCMS, Distribution System Operators (DSOs), aggregators, and energy market entities. Charging sessions are typically triggered via mobile applications or service APIs; however, under the ZTA, these requests are not directly processed by the EVCMS. Instead, they are routed through a centralized control plane that applies continuous trust evaluation. Within this control plane, behavioral indicators are assessed by a Trust Evaluation Engine in conjunction with a Policy Engine. The PDP issues short-lived trust tokens only when the contextual integrity of the request satisfies predefined policies. These tokens are verified by local PEPs

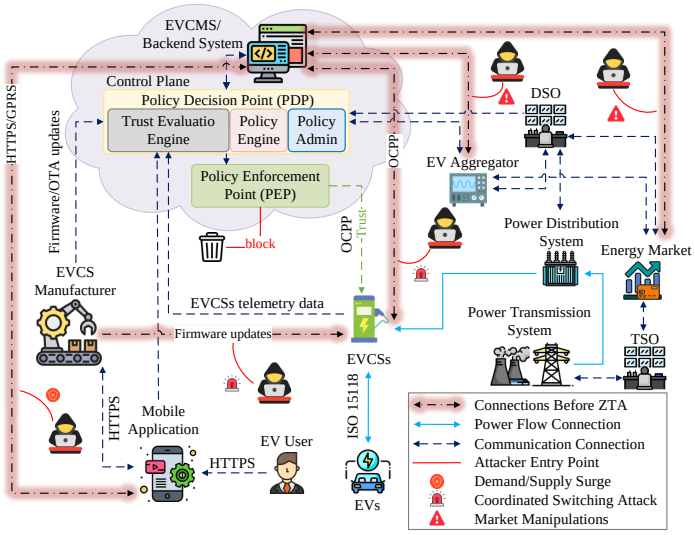


**Fig. 1:** EVCS Security Architecture with ZTA and Attack Vectors.

embedded in each EVCS, allowing or denying further action accordingly. Communications across all entities are secured using protocols, such as HTTPS and OCPP 2.0.1. Firmware updates are validated by the PDP before over-the-air (OTA) deployment to ensure source authenticity and prevent malicious manipulation. Coordination with DSOs, aggregators, and energy market systems is facilitated to preserve operational stability and compliance with grid directives. This architecture is structured to counter the range of adversarial actions characterized in the threat model, with particular emphasis on preempting unauthorized access, ensuring firmware authenticity, and detecting anomalous charging behavior.

As previously mentioned in Section I, CMAs exploit various vulnerabilities within the EV charging ecosystem. In this work, adversaries are assumed to possess moderate to advanced resources, including access to compromised user credentials, mobile applications, or backend software interfaces. As discussed in Section I, this work considers three types of CMAs: demand surge, coordinated switching, and protocol-level manipulation. Demand surge attacks are launched by abusing these assets to simultaneously initiate charging sessions at scale, potentially causing destabilizing spikes in grid demand. These attacks may utilize Man-in-the-Middle (MitM) techniques to intercept or falsify session initiation requests, especially when authentication protocols are weak or static. Coordinated switching attacks involve periodic on/off cycling of EVCSs through tampered firmware or scheduler access, which mimics natural grid oscillations that challenge detection mechanisms. Such attacks can be facilitated through unverified OTA updates or concealed using False Data Injection Attacks (FDIA) that alter telemetry data. Lastly, market manipulation attacks aim to distort price signals and trading outcomes by injecting falsified load data into aggregator platforms or reporting interfaces. These strategies exploit protocol weaknesses, spoofed measurement values, and vulnerabilities in EVCMS or aggregator APIs. Without dynamic trust validation and localized session controls, these attacks can bypass perimeter-based security measures and inflict significant disruptions at both grid and market levels.

### *B. Trust Evaluation & Policy Decision with Markov Updating*

To capture the evolving behavior of each charger, we model each trust dimension using a discrete-time Markov chain, where the trust state at any time depends only on the previous state and current behavioral observations. This probabilistic modeling approach allows the system to predict future trust levels and respond proactively to behavioral changes. Drawing inspiration from [16] and adapting the Markov-based trust score evolution technique [17], we define three distinct dimensions of charger trustworthiness: Stability Degree (SD), Intimacy Degree (ID), and Abnormality Degree (ED). Each dimension independently tracks a specific behavioral aspect of the EVCS and evolves over time through its corresponding state transition process, guided by real-time logs and observed features.

*Time Model Clarification:* In our framework, time $t$ is discretized into evaluation periods, where each period corresponds to a 3-hour operational window during peak charging hours. When we refer to time $t-1$, we mean the previous evaluation period (i.e., the preceding 3-hour window). Trust scores and behavioral indicators are updated once per evaluation period based on the sessions observed during that window. This discrete-time model aligns with realistic operational monitoring intervals in EVCS management systems.

Each of these scores is updated over time using its transition mechanism. The global trust score is then defined as:

$$T_n^{(t)} = w_s \cdot T_s^{(t)} + w_i \cdot T_i^{(t)} - w_e \cdot T_e^{(t)} \quad (1)$$

where $T_s^{(t)}$, $T_i^{(t)}$, and $T_e^{(t)}$ represent the normalized stability, intimacy, and abnormality scores at time $t$, respectively, and $w_s$, $w_i$, $w_e$ are their corresponding weights, with $w_s + w_i + w_e = 1$.

*1) Stability Degree (SD):* This score evaluates the consistency of charger communication by incorporating the number of connection interruptions and request timeouts observed during a given evaluation period. Specifically, $l_s$ represents the number of connection disruptions observed during the current evaluation period, $r_s$ represents the number of request timeout events during the current period, while $T_l$ and $T_r$ denote the maximum allowable thresholds for disruptions and timeouts, respectively (set based on operational requirements). The parameters $V_l$ and $V_r$ quantify the base penalty weight associated with each type of failure, and $\lambda_l$, $\lambda_r$ are scaling coefficients that adjust the severity based on the frequency of failures. The stability score is updated as follows:

$$T_s^{(t)} = T_s^{(t-1)} + ((V_l + \lambda_l \cdot l_s) \cdot (T_l - l_s)) + ((V_r + \lambda_r \cdot r_s) \cdot (T_r - r_s)) \quad (2)$$

As long as $l_s < T_l$ and $r_s < T_r$, the penalty remains bounded. However, when either metric approaches or exceeds its threshold, the penalty term increases rapidly, reducing the stability score and reflecting decreased reliability in communication.

*2) Intimacy Degree (ID):* This metric quantifies the historical interaction strength between the charger and the cloud-based management platform. The variables $R_t$ and $R'_t$ represent the percentile rank of interaction frequency in the current evaluation period $t$ and the previous evaluation period $t-1$, respectively (ranging from 0 to 100, computed based on the charger's relative frequency compared to all chargers in the network). The term $q_i$ serves as a fluctuation control parameter that moderates

the score update (preventing excessive sensitivity to minor rank changes). The intimacy score is computed as:

$$T_i^{(t)} = T_i^{(t-1)} + \left(q_i + \eta_1 R_t + \eta_2 R_t'\right)\left(R_t' - R_t\right) \tag{3}$$

where $\eta_1$ and $\eta_2$ are scaling coefficients that tune the sensitivity of the score to temporal shifts in interaction rank. A decrease in activity ranking over time reduces the intimacy score, reflecting diminished trust, whereas stable or improved interaction frequency contributes to trust accumulation.

*3) Abnormality Degree (ED):* The abnormality score quantifies deviations from expected behavioral norms observed in charger activity logs. It reflects the occurrence of suspicious events during the current evaluation period, such as repeated protocol misuse, attempts to access unauthorized commands, or operating from unexpected geographic regions. Since EV chargers are typically stationary and operate within fixed service zones, any sudden change in their geolocation compared to historical records is treated as a potential anomaly. Let $T_e^{*(t)}$ denote the abnormality degree at the end of time period $t-1$ (i.e., the previous evaluation window), and suppose $N$ abnormal behaviors are detected during the current period $t$. Each anomaly is assigned a severity score $Q(\text{Exception}_i)$ (a predefined weight based on the type of anomaly, e.g., unauthorized command = 5, location anomaly = 3), where $i \in \{1,\ldots,N\}$ indexes the observed exceptions. The abnormality degree is then updated according to:

$$T_e^{(t)} = \begin{cases} T_e^{*(t)} + \sum_{i=1}^{N} Q(\text{Exception}_i), & N \neq 0 \\ T_e^{*(t)} - p_e, & N = 0 \end{cases} \tag{4}$$

here, $p_e$ is a recovery coefficient (a positive constant) that gradually decreases the abnormality score in the absence of violations, promoting long-term behavioral rehabilitation. The function $Q(\cdot)$ captures the relative severity of each abnormality type—such as excessive charging power, unauthorized command usage, unverified firmware updates, or location inconsistency—and allows fine-grained adjustment of the trust score based on context. This formulation enables the framework to adaptively penalize malicious behavior while allowing trust to be restored when normal patterns are maintained.

*4) Markov-Based State Update Process:* To capture the temporal evolution of trust across various dimensions, each dimension—stability ($T_s$), intimacy ($T_i$), and abnormality ($T_e$)—is modeled using a discrete-time Markov chain. The state space for each dimension $d \in \{s,i,e\}$ is defined as $\mathscr{S}_d = \{S_1, S_2, S_3, S_4, S_5\}$, corresponding to semantic levels of *Very Low Trust*, *Low Trust*, *Moderate Trust*, *High Trust*, and *Very High Trust*, respectively. This classification enables fine-grained tracking of trust dynamics in the presence of varying behavioral patterns. State transitions are governed by a stationary Markov transition matrix $P_d \in \mathbb{R}^{5\times5}$, where each entry $P_d(i,j)$ specifies the probability of transitioning from state $S_i$ at time period $t-1$ to state $S_j$ at time period $t$. The transition process is influenced by behavioral evidence collected during the current evaluation window, such as interaction frequency, connection reliability, or protocol misuse anomalies.

The belief distribution over the 5 trust states at time $t$ is denoted by $\pi_d^{(t)} \in \mathbb{R}^5$, where each element represents the probability that the charger is in a particular trust state. Given the belief from the previous period $\pi_d^{(t-1)}$ and an observation $z_d^{(t)}$ (the computed trust score for dimension $d$ in the current period) associated with the current trust scores, the belief is updated using a Bayesian filtering approach:

$$\pi_d^{(t)} = \frac{P_d^\top \cdot \left(\pi_d^{(t-1)} \odot \mathscr{O}_d(z_d^{(t)})\right)}{\left\| P_d^\top \cdot \left(\pi_d^{(t-1)} \odot \mathscr{O}_d(z_d^{(t)})\right)\right\|_1} \tag{5}$$

where $\odot$ denotes the element-wise product, and $\mathscr{O}_d(z_d^{(t)}) \in \mathbb{R}^5$ represents the observation likelihood vector, which quantifies the alignment between the current observation $z_d^{(t)}$ and each trust state in $\mathscr{S}_d$. These likelihood vectors are derived from analyzing historical behavioral logs and discretizing trust-related features (e.g., stability, interaction frequency, anomaly levels) into state-dependent distributions. In scenarios where labeled data is limited, domain-specific thresholds are used to construct $\mathscr{O}_d(z_d^{(t)})$ based on expert-defined boundaries. Following the belief update, the scalar trust score for dimension $d$ is computed as the expectation of the updated belief distribution:

$$T_d^{(t)} = \pi_d^{(t)} \cdot \mathbf{v}_d \tag{6}$$

where $\mathbf{v}_d = [0.0, 0.25, 0.5, 0.75, 1.0]^\top$ represents numerical values assigned to the five trust states, from *Very Low* to *Very High*. Higher values of $T_d^{(t)}$ indicate stronger trustworthiness in the given dimension.

This Markov-based process enables robust, probabilistic modeling of charger behavior under partial observability, which allows the system to smoothly adapt to both short-term anomalies and long-term behavioral shifts.

*5) Trust-Based Authorization:* A charging request is approved only if the aggregated trust score $T_n^{(t)}$ exceeds a predefined threshold $\tau$, which is configured by the system administrator based on organizational security policies and risk tolerance levels. This threshold represents the minimum trustworthiness required for a charger to obtain session approval. If the trust score falls below $\tau$, the PDP withholds the access token, thereby blocking the EVCS from executing further protocol operations.

$$\kappa_n^{(t)} = \begin{cases} \text{JWT issued} & \text{if } T_n^{(t)} \geq \tau \\ \text{Request denied} & \text{otherwise} \end{cases} \tag{7}$$

where $\kappa_n^{(t)}$ represents the authorization decision for charger $n$ at time period $t$, and JWT (JSON Web Token) is the short-lived cryptographic token issued by the PDP.

### C. Policy Enforcement & Protocol Integration

Once the PDP authorizes a session by issuing a token, enforcement is carried out at the edge via lightweight PEPs embedded in each EVCS. Each incoming request is intercepted and validated by the PEP, which parses the OCPP 2.0.1 message and inspects the embedded JWT token for scope, expiration, and signature integrity. A control request $r_i^{(t)}$ from EVCS $i$ at time $t$ is authorized only if:

$$r_i^{(t)} \in \mathscr{A}(\mathscr{P}, \kappa) \tag{8}$$

where $\mathscr{A}$ represents the set of allowable actions under system policy $\mathscr{P}$ and token $\kappa$. If a mismatch occurs or the token is invalid, the request is blocked and flagged. All EVCS-bound messages, including OCPP control commands, firmware updates, and metering requests, are routed through the PEP for real-time, context-aware enforcement. The PEP ensures that each request is authenticated, authorized, and compliant with session rules. For instance, an `UpdateFirmware` command is accepted only if authorized by the PDP and digitally signed:

$$\text{Verify}(\sigma_{\text{firm}}, K_{\text{pub}}^{\text{PDP}}) = \texttt{True} \tag{9}$$

where $\sigma_{\text{firm}}$ is the firmware's digital signature and $K_{\text{pub}}^{\text{PDP}}$ is the PDP's public key. If verification fails, the command is dropped, and an alert is triggered.

Token-based enforcement operates in coordination with the trust authorization mechanism discussed in Section II-B5. Only EVCSs with an aggregated trust score exceeding the threshold $\tau$ are granted access tokens. This ensures that subsequent protocol interactions, including firmware updates, control messages, and session operations, originate exclusively from authenticated and trustworthy sources.

## III. Simulation Results

To validate the effectiveness of the proposed ZTA-based defense framework, a comprehensive simulation environment was developed to emulate realistic EV charging station operations and behavioral patterns. The simulation was run in a virtual machine hosting all essential software modules, including charger emulation, trust computation, and policy enforcement mechanisms. Each EVCS was represented as a logical instance, emulating real-world interactions through Python scripts that implement the OCPP 2.0.1 protocol at the application layer. This abstraction ensured a realistic cyber-physical representation without requiring dedicated hardware. A total of five EVCSs, labeled Chargers A through E, were selected from a real-world dataset, each supporting two independent power outlets. Charging session logs from these stations were used to extract behavioral data during peak demand periods across multiple operational days. To capture the temporal dynamics of charging activity, the simulation was divided into ten evaluation windows, each representing a 3-hour peak period sampled from the dataset. The behavioral indicators (SD, ID, ED) were updated once per simulation window, resulting in the 10 evaluation periods shown in the results. This segmentation reflects realistic session density and aligns with the operational capabilities of dual-outlet chargers, since the original dataset exhibits 2-3 charging sessions per 3-hour peak window per dual-outlet charger, with activity representing high-utilization periods. Accordingly, the number of sessions in each time window was extracted directly from the empirical dataset to preserve authentic load patterns.

To simulate adversarial scenarios, we used Python scripts to inject charge manipulation attacks into 30% of the sessions in the dataset. These attacks changed key session attributes like requested energy and charging duration by modifying their values within a realistic range, based on typical usage patterns. All modified sessions followed the format and rules of the OCPP 2.0.1 protocol, making them appear legitimate. The attacks were applied randomly across different chargers and time windows to reflect stealthy, coordinated behavior that evades basic rule-based detection.

It is important to note that the baseline "traditional" approach in our evaluation represents a purely trust-agnostic access model where all syntactically valid OCPP messages are accepted without behavioral assessment. While real-world operators typically employ some security measures, such as rate limiting or basic anomaly detection, our baseline was intentionally simplified to isolate and highlight the specific contribution of the proposed trust-based framework. This choice allows us to demonstrate the fundamental value of continuous behavioral assessment and dynamic trust evaluation. However, we acknowledge that a more realistic baseline would incorporate common security controls. Future work will include comparative evaluation against baselines that integrate rate limiting, session constraints, and threshold-based anomaly detection to provide a more comprehensive assessment of the added value of our ZTA approach. Additionally, the impact of different parameter choices on system performance requires systematic investigation through sensitivity analysis, as discussed below.

TABLE I: TRUST LEVEL HIERARCHY.

| Trust Level | Controlling Measure | Permission | Trust Score Range |
|---|---|---|---|
| I | N/A | Monitored access | [90,100] |
| II | Downgrading permissions | Restricted access; Monitored access | [70,90) |
| III | Re-authenticating | Access temporarily denied | [50,70) |
| IV | Blocking connections | Access denied | [0,50) |

Three behavioral indicators were used to continuously monitor and rate each EVCS, computed from historical logs and simulated session events acting as real-time inputs. Stability measured the temporal consistency in session frequency and duration; intimacy reflected recurring associations with known users or vehicles; and abnormality captured statistical deviations from baseline behaviors, such as irregular inter-arrival times or anomalous energy requests. Each indicator evolved according to a discrete-time Markov process with 5 trust states, and transition probabilities were learned from empirical distributions. The indicators were normalized to the [0,100] scale, enabling fair comparison. The overall trust score for each charger was calculated as a weighted average of the expected values of the three indicators, with empirically determined weights: 0.4 (stability), 0.3 (intimacy), and 0.3 (abnormality). These weights emphasize behavioral consistency, which is critical for mitigating CMAs.

*Parameter Sensitivity:* The choice of weights ($w_s = 0.4$, $w_i = 0.3$, $w_e = 0.3$) and the trust threshold ($\tau = 70$) significantly influence system performance. While our current parameterization is based on empirical tuning using the available dataset, we acknowledge that optimal values may vary across different operational contexts, charger types, and threat landscapes. Preliminary experiments (not shown here due to space constraints) indicate that increasing $w_s$ improves detection of communication-based attacks but may reduce sensitivity to user behavior anomalies. Similarly, lowering $\tau$ increases security at the cost of more frequent false rejections of legitimate sessions.

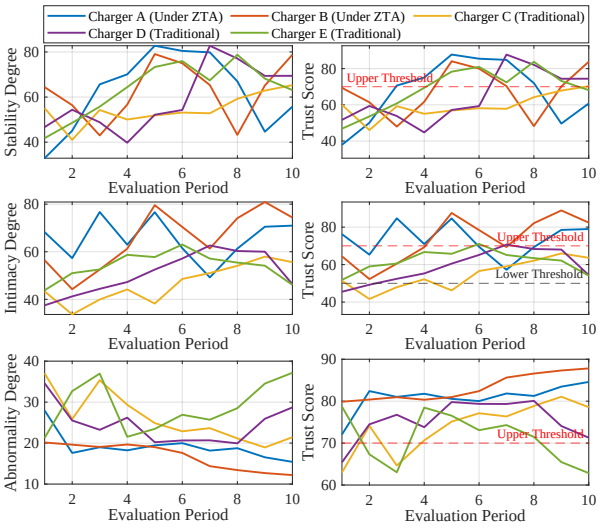


**Fig. 2:** Trust evolution under ZTA versus traditional access after attack.

A comprehensive sensitivity analysis examining the impact of weight variations, recovery coefficients ($p_e$), and threshold settings on detection rates, false positives, and overall system responsiveness will be conducted in future work. This analysis is essential for understanding the framework's portability and scalability across diverse operational environments and operator requirements.

Based on the resulting trust score, each charger was assigned a trust level in accordance with the hierarchy shown in Table I, guiding the system's dynamic access control logic. Chargers A and B were subject to ZTA enforcement, where sessions were only approved if the trust score remained above the threshold corresponding to Level II (i.e., $\geq 70$). Otherwise, control measures, such as re-authentication (Level III) or access denial (Level IV), were triggered based on the degree of trust degradation, as defined in Table I. In contrast, Chargers C, D, and E operated under a traditional access model without dynamic trust evaluation and approved all syntactically valid session requests regardless of behavioral anomalies. To assess the system's responsiveness to cyber-physical threats, all five chargers were subjected to CMAs, wherein a subset of session requests was synthetically modified to exhibit abnormal behavioral characteristics, such as irregular inter-arrival patterns, abrupt load changes, and inconsistent user associations. This attack strategy aimed to mimic real-world stealthy manipulation while degrading the behavioral indicators over time. The divergence in session approval logic between ZTA-enforced and trustless policies enabled a clear comparative analysis of how dynamic trust evaluation can mitigate attacks by proactively denying low-trust sessions while maintaining uninterrupted service for compliant chargers.

In addition to trust score computation, the simulation incorporated an access control mechanism to monitor the number of approved and denied sessions per charger across ten time windows. For Chargers A and B, which were governed by the ZTA policy, access decisions were strictly enforced based on trust levels: sessions were approved only when the trust score remained within Levels I or II (i.e.,$\geq 70$), and were denied or required re-authentication for Levels III and IV. As a result, these chargers exhibited a selective but secure

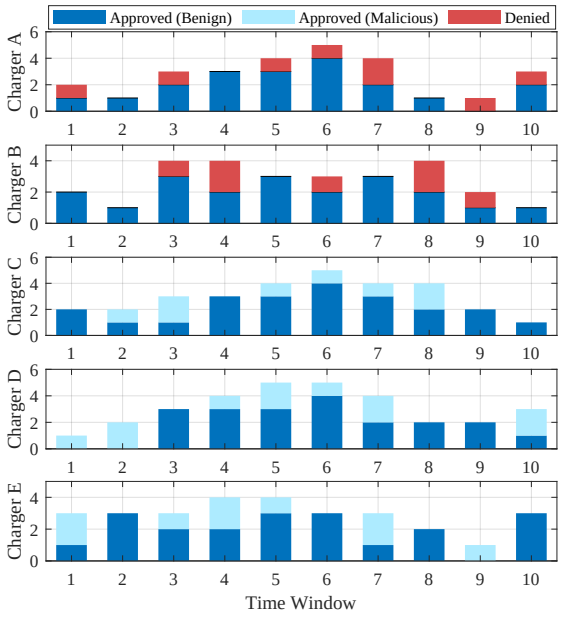


**Fig. 3:** Session-level access control decisions across EVCSs.

approval pattern, while most sessions were permitted, access was consistently denied during intervals where trust degradation was detected due to ongoing behavioral anomalies introduced by the attack. In contrast, Chargers C, D, and E, which operate without trust-based access control, approved all session requests regardless of their underlying behavioral indicators, which allowed malicious activity to persist undetected. All behavioral indicators were evaluated in their normalized form to ensure consistent trust score computation across chargers. Figure 2 presents six subplots depicting the evolution of behavioral degrees and their corresponding trust scores across chargers over time. Each indicator's deviation is shown on the left, while the right plots reflect the computed trust scores with respect to defined trust level thresholds. Due to the coordinated attack, all chargers experienced noticeable trust score variations, oscillating between Level I and Level III. Despite the uniform attack conditions, Chargers A and B leveraged ZTA to dynamically regulate access, while Chargers C–E lacked such enforcement. Figure 3 visualizes the resulting session-level decisions. Chargers A and B show a proportion of sessions denied due to sub-threshold trust scores, which demonstrates ZTA's enforcement. Chargers C–E, despite experiencing similar trust score declines, continued to permit all sessions due to the absence of a behavioral assessment. These findings validate the ZTA framework's capability to enforce proactive access decisions based on real-time trust evaluation, thereby limiting the system's exposure to charge manipulation attacks. Table II quantifies these outcomes. With 30% of sessions containing injected attacks, ZTA-enabled chargers successfully detected and blocked 87.5-100% of malicious sessions, while traditional chargers approved all requests regardless of malicious intent.

### *A. Computational Overhead and Scalability Considerations*

While our current evaluation focuses on security effectiveness, computational latency is a critical factor for real-world deployment. The proposed framework introduces overhead at two levels: (1) trust score computation at the central PDP,

TABLE II: ACCESS CONTROL EFFECTIVENESS UNDER CMA.

| Charger Type | Approved | Denied | Denial Rate |
|---|---|---|---|
| A (ZTA-enabled) | 19 | 8 | 100% |
| B (ZTA-enabled) | 20 | 7 | 87.5% |
| C (Traditional) | 30 | 0 | 0.0% |
| D (Traditional) | 31 | 0 | 0.0% |
| E (Traditional) | 29 | 0 | 0.0% |

and (2) token verification at each local PEP. In our simulation environment, trust score updates (performed once per 3-hour evaluation window) required an average of 120-150 milliseconds per charger on a standard virtual machine (Intel Xeon E5-2680 v4, 8GB RAM). Token verification at the PEP level, which occurs per session request, adds approximately 2-5 milliseconds of latency due to JWT signature validation. Given that typical charging session initiation times are on the order of seconds, this overhead is unlikely to impact user experience in practice. However, formal latency measurements on edge devices and under concurrent load conditions are needed to fully validate deployment feasibility.

Regarding scalability, our evaluation is limited to five chargers from a single dataset, which represents a small-scale deployment scenario. The generalizability of our findings to larger networks with diverse charger types (AC Level 2 versus DC fast charging), different operators, and varying traffic patterns remains an open question. Different charging environments may exhibit distinct behavioral patterns, requiring adaptation of trust indicator definitions, weight assignments, and transition probability matrices. Future work will address these limitations through: (1) evaluation on larger-scale datasets spanning multiple operators and geographic regions, (2) investigation of AC versus DC charger behavioral differences, (3) real-time edge deployment and latency profiling, and (4) adaptive parameter tuning mechanisms that can adjust to local operational contexts without manual reconfiguration.

## IV. Conclusion

This study introduced a Zero Trust Architecture (ZTA)-based predictive trust evaluation framework tailored for Electric Vehicle Charging Infrastructure (EVCI). The proposed system leverages a Markov-based behavioral model to assess charger trustworthiness across three key dimensions: stability, intimacy, and abnormality. Through a simulation involving five EVCSs, all subjected to coordinated behavioral anomalies, we demonstrated the framework's capability to track behavioral trends and predict declining trust levels dynamically. While all chargers experienced trust-score degradation under attack, only those with ZTA enforcement could prevent access based on real-time trust evaluation, thereby illustrating the framework's ability to contain malicious sessions. This predictive, context-aware approach enables fine-grained access control and enhances the resilience of EVCS networks against application-layer threats such as Charge Manipulation Attacks (CMAs). As part of future work, we aim to integrate additional real-world datasets across varied environments and operational contexts and to implement a live, edge-based prototype for real-time trust computation and session authorization. This will further validate the scalability and adaptability of the proposed framework in practical smart grid deployments.

## Acknowledgment

This work is supported by Fonds de recherche du Québec – Nature et technologies (FRQNT) and the Concordia University/Hydro-Québec/NSERC Research Collaboration Project.

## References

[1] K. KSarieddine, M. A. Sayed, C. Assi, R. Atallah, S. Torabi, J. Khoury, M. S. Pour, and E. Bou-Harb, "Ev charging infrastructure discovery to contextualize its deployment security," *IEEE Transactions on Network and Service Management*, vol. 21, no. 1, pp. 1287–1301, 2024.

[2] A. Vishnoi and V. Verma, "The analysis on impact of cyber security threats on smart grids," in *Security and Risk Analysis for Intelligent Edge Computing*. Springer, 2023, pp. 111–118.

[3] K. Sarieddine, M. A. Sayed, S. Torabi, R. Attallah, D. Jafarigiv, C. Assi, and M. Debbabi, "Uncovering covert attacks on ev charging infrastructure: How ocpp backend vulnerabilities could compromise your system," in *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security*, 2024, pp. 977–989.

[4] S. Marandi, A. Moradzadeh, H. Moayyed, C. Assi, M. Ghafouri, and Z. Vale, "Anomaly detection in load forecasting for electric vehicles using image processing techniques," in *2024 IEEE International Conference on Environment and Electrical Engineering and 2024 IEEE Industrial and Commercial Power Systems Europe (EEEIC / ICPS Europe)*, 2024, pp. 1–6.

[5] J. Johnson, B. Anderson, B. Wright, J. Quiroz, T. Berg, R. Graves, J. Daley, K. Phan, M. Kunz, R. Pratt *et al.*, "Cybersecurity for electric vehicle charging infrastructure," Sandia Natl. Lab., Albuquerque, NM (US), Tech. Rep., 2022.

[6] N. Zhou, S. Ji, and Q. Mao, "Research on trust evaluation model for vehicle-to-grid interaction based on zero trust architecture," in *2024 5th International Conference on Information Science, Parallel and Distributed Systems (ISPDS)*. IEEE, 2024, pp. 189–197.

[7] A. Kesarwani and P. M. Khilar, "Development of trust-based access control models using fuzzy logic in cloud computing," *Journal of King Saud University-Computer and Information Sciences*, vol. 34, no. 5, pp. 1958–1967, 2022.

[8] S. Ashtari, M. Danesh, and h. shirgahi, "A novel user profile-based fuzzy approach for evaluating trust in semantic web," *IIUM Engineering Journal*, vol. 20, no. 1, pp. 158–176, 2019.

[9] H. Jahangir, S. Lakshminarayana, and H. V. Poor, "Charge manipulation attacks against smart electric vehicle charging stations and deep learning-based detection mechanisms," *IEEE Transactions on Smart Grid*, 2024.

[10] B. Wang, P. Dehghanian, S. Wang, and M. Mitolo, "Electrical safety considerations in large-scale electric vehicle charging stations," *IEEE Transactions on Industry Applications*, vol. 55, no. 6, pp. 6603–6612, 2019.

[11] M. E. Kabir, M. Ghafouri, B. Moussa, and C. Assi, "A two-stage protection method for detection and mitigation of coordinated evse switching attacks," *IEEE Transactions on Smart Grid*, vol. 12, no. 5, pp. 4377–4388, 2021.

[12] M. A. Sayed, R. Atallah, C. Assi, and M. Debbabi, "Electric vehicle attack impact on power grid operation," *International Journal of Electrical Power & Energy Systems*, vol. 137, p. 107784, 2022.

[13] S. Acharya, R. Mieth, R. Karri, and Y. Dvorkin, "False data injection attacks on data markets for electric vehicle charging stations," *Advances in Applied Energy*, vol. 7, p. 100098, 2022.

[14] Y. Wang, Y. Na, S. Xu, and D. Huang, "Distributed fault detection scheme for intelligent connected vehicles in a zero-trust environment," in *2023 2nd Conference on Fully Actuated System Theory and Applications (CFASTA)*. IEEE, 2023, pp. 526–531.

[15] T. E. Carroll, L. H. Chang, and C. L. Wright-Hamor, "The design and evaluation of zero trust architecture for electric vehicle charging infrastructure: Evs@ scale series on ev charging station cybersecurity," Pacific Northwest National Laboratory (PNNL), Richland, WA (United States), Tech. Rep., 2024.

[16] P. Li, W. Ou, H. Liang, W. Han, Q. Zhang, and G. Zeng, "A zero trust and blockchain-based defense model for smart electric vehicle chargers," *Journal of Network and Computer Applications*, vol. 213, p. 103599, 2023.

[17] M. Bampatsikos, I. Politis, T. Ioannidis, and C. Xenakis, "Trust score prediction and management in iot ecosystems using markov chains and madm techniques," *IEEE Transactions on Consumer Electronics*, 2025.