Security Architecture for Data Protection in Intelligent Energy Prediction Systems for Non-Residential Buildings

Eleonora TUDORA, Sorin SOVIANY, Ovidiu BICA
National Institute for Research and Development in Informatics – ICI Bucharest
eleonora.tudora@ici.ro, sorin.soviany@ici.ro, ovidiu.bica@ici.ro

Abstract: This paper proposes a security architecture for an intelligent energy consumption prediction Internet-of-Thing (IoT)-based system, looking to protect sensitive data and ensure reliable operation. The proposed model is a multi-layer security architecture addressing the security requirements specific to non-residential sites (office buildings, hospitals, educational facilities), with functional components designed to mitigate the associated risks. By combining the edge-level processing, end-to-end encryption, multi-layer intrusion detection, and KPI-driven monitoring, the proposed model addresses emergent threats while ensuring a balanced performance vs. the security ratio. Furthermore, the performance indicators are defined to evaluate the effectiveness of the architecture. The paper also presents the optimization options in respect to emergent threats, as well as the potential impact of the proposed security functions on the overall performance of the designed system.
Keywords: Energy Consumption Prediction, Non-Residential Buildings, Intelligent Systems, Cyber Security.

INTRODUCTION

The global evolution towards sustainable and energy-efficient infrastructures enabled the efforts to develop intelligent monitoring and prediction systems in non-residential buildings (office complexes, hospitals, educational facilities). For many buildings the energy consumption is a significant operational cost and a key factor impacting the environmental sustainability. The integration of advanced technologies, Internet-of-Things (IoT) devices and Artificial Intelligence (AI) supports the development of advanced platforms able to continuously collect, analyze, and forecast the energy usage patterns (Poyyamozhi et al., 2024), (Barros et al., 2025). These systems enable a real-time data processing for the use-cases requiring accurate forecasts of energy

https://doi.org/10.54851/v7i2y202502

consumption, creating opportunities for cost optimization, resource allocation efficiency, and compliance with the environmental regulations (Tudora & Tîrziu, 2024).

The heterogeneity and resource constraints of the IoT devices, together with the sensitivity of the predictive models, generate major cybersecurity challenges, especially in critical infrastructures (Wyrzykowska et al., 2024). The operation of such systems involves massive data collection, storage, and processing. This is why they are targets for a variety of Cyber-attacks of increasing complexity. Threats such as unauthorized access, firmware manipulation, eavesdropping or interception, data tampering, Denial-of-Service (DoS) attacks, and adversarial manipulations can compromise both system performance and user trust. The AI-based prediction models are additionally vulnerable to data poisoning, model inversion, and adversarial manipulations; these can result in operational disruptions, privacy violations, financial losses, and even safety risks in non-residential buildings with large numbers of occupants relying on an efficient energy management (Li et al., 2020). Therefore, the deployment of reliable security solutions is justified given the roles of the collected data: both the input to Machine Learning (ML) models and the operational basis for decision-making in building management. The consequences of the data compromising or manipulation may include: the prediction accuracy degradation, undermining optimization strategies, significantly reducing the stakeholder trust. The protection of the confidentiality, integrity, and availability for energy consumption data already represents a prerequisite for the reliable operation of the intelligent energy prediction systems.

Compared to the residential sites, the non-residential buildings typically involve a higher density of IoT devices, more complex energy distribution systems, and multiple categories of users and administrators. This complexity significantly enlarges the attack surface and requires layered, adaptive security architectures able to prevent, detect, and efficiently respond to emerging threats (Korodi et al., 2024), (Amangeldy et al., 2025). On the other hand, the compliance with the strict regulatory frameworks (e.g., GDPR for personal data protection, ISO/IEC 27001 for information security) makes a security-by-design approach a mandatory requirement (Zhou et al., 2024). In this context, emerging paradigms such as Zero Trust architectures (Rose et al., 2020) and federated learning for privacy-preserving data analytics (Li et al., 2020) generate new opportunities to strengthen resilience while maintaining the interoperability and system performance.

The present research paper proposes a functional security architecture for an intelligent energy consumption prediction IoT-based system, looking to protect sensitive data and ensure reliable operation, designed specifically for non-residential buildings. This architecture will be embedded within the broader context of the ongoing „Intelligent System for Predicting Energy Consumption in Buildings (PRECONERG)" project, which will develop a secure and intelligent energy prediction solution for smart building environments. The proposed model addresses security requirements by firstly identifying vulnerabilities, risks and emerging threats, and then integrating functional components to mitigate the associated risks. The main goal is to ensure that the data-driven prediction remains trustworthy, resilient, and compliant with the regulatory requirements while minimizing the performance overhead in constrained IoT environments.

The architecture combines classical data protection mechanisms (encryption, authentication, and access control) (Mrabet et al., 2020), (Mai, 2019) with anomaly detection solutions that use ML (Mai, 2019), (El Hadj Youssef et al., 2020), (Mohanta et al., 2020). It uses a layered model including the following functional layers: IoT Devices & Sensors, Secure Gateway, Firewall & IDS/IPS (Intrusion Detection/Prevention System) with SIEM (Security Information and Event Management), Data Processing & AI Prediction, and Web/API Access. Each layer integrates specific mechanisms, including end-to-end encryption, Multi-Factor Authentication (MFA), Role-Based Access Control (RBAC), and ML-based IDS, adapted to the constraints of the IoT environments (Trabelsi et al., 2023), (Isong

et al., 2023). To evaluate the effectiveness of the architecture, performance indicators such as encryption overhead, anomaly detection accuracy, authentication success rates, and compliance levels are defined (Arshad et al., 2020), (Isong et al., 2023). The paper presents optimization options in respect to emergent threats, as well as the potential impact of the proposed security functions on the overall performance of the designed system. The optimization options look to mitigate emerging threats; these options include lightweight cryptographic primitives for the energy-constrained devices (El Hadj Youssef et al., 2020), adaptive MFA to balance usability and security, hybrid IDS for zero-day attacks detection (Isong et al., 2023), and automated incident response through SIEM integration. Overall, the architecture not only addresses confidentiality, integrity, and availability requirements for the energy consumption data but also ensures robustness against the evolving cyber risks (Mrabet et al., 2020), (Zhou et al., 2024).

The remainder of the paper is structured as follows: Section II reviews related works; Section III specifies the proposed security architecture; Section IV concludes the paper and provides future works/directions.

RELATED WORKS

The recent developments of the intelligent Internet-of-Thing (IoT)-based systems in use-cases requiring the energy consumption monitoring and prediction for non-residential buildings deal with significant challenges in terms of data protection and system resilience. These systems are based on huge amounts of data collection, storage, and processing. This is why they become targets for various cyberattacks with an increasing complexity and impact. As buildings grow more complex and data-driven, the intelligent energy management and prediction systems have attracted increasing attention in both research and practice. The convergence of the IoT devices, communication networks, edge–cloud computing, and predictive analytics enables both the real-time monitoring of energy consumption and the projection of long-term demand, contributing to sustainability objectives as well as operational efficiency. The practical deployment of such systems still presents significant challenges. The security, interoperability, and governance issues are particularly evident in the non-residential settings, where the regulatory and operational requirements are usually more stringent. To address these concerns, the recent research has explored a wide range of security-oriented architectures, including layered IoT frameworks, standards-based solutions (e.g., IEC 62443, BACnet/SC, OPC UA), and, more recently, emerging paradigms such as the Zero-Trust and edge-to-cloud security models. The following review examines these contributions, outlining their approaches, advantages, and limitations in the context of a secure energy consumption prediction for smart buildings.

(Poyyamozhi et al. 2024) proposed a layered IoT-based architecture for building security and management, organized into sensing, network, data processing, and application layers. Their framework enables real-time monitoring, predictive analytics, and intelligent control of energy consumption. One of the study's main contributions is its focus on the interoperability and security challenges. The heterogeneous device protocols and the integration of the IoT with legacy systems often lead to vulnerabilities and complicate data exchange. The authors also highlight some practical barriers such as high implementation costs, privacy concerns, and weak coordination among stakeholders. To overcome these obstacles, they recommend standardized interoperability frameworks, stronger cybersecurity mechanisms, and governance models that can support a large-scale deployment of intelligent energy prediction systems in the non-residential buildings.

Another line of research emphasizes the edge-to-cloud security models, which have gained traction as efficient and secure solutions for the energy consumption prediction. In such architectures, the edge devices are primarily tasked with local data aggregation, preprocessing, and, in certain instances, inference, whereas the cloud layer predominantly manages the large-scale model training, persistent data storage, and

global orchestration. Security is predominantly achieved through the establishment of encrypted communication channels (TLS-Transport Layer Security/mTLS), the enforcement of the role-based access control mechanisms for the IoT clients, and the application of anonymization techniques to the sensitive information at the edge. More recent studies explore the federated learning as a way to strengthen privacy by reducing the need to transfer raw data to the cloud (Barros et al., 2025).

(Amangeldy et al. 2025) introduce an intelligent and secure building automation architecture that integrates an advanced AI-driven control with the standardized industrial communication protocols. The system makes use of the digital twins of the building (DT-B) and its occupants (DT-H), enabling the predictive and adaptive management of the HVAC, lighting, and access systems. To secure the communications within the IT-managed networks, the architecture adopts the BACnet Secure Connect (BACnet/SC), which enhances the traditional BACnet protocol with encryption and authentication features, ensuring the interoperability and protecting data exchange between the automation devices. In addition, the design follows the IEC 62443 principles by segmenting the network into security zones and conduits and applying the least-privilege access policies. This results in a scalable, resilient, and secure architecture that combines the state-of-the-art AI with industry-standard cybersecurity practices, safeguarding both the operational data and critical infrastructure.

Wyrzykowska examines the role of the Intelligent Energy Management Systems (IEMS) in the Industry 5.0 era, with a strong focus on cybersecurity. The study identifies a broad spectrum of threats, ranging from malware and ransomware to Distributed Denial-of-Service (DDoS) attacks on the IoT and cyber–physical infrastructures. Such threats not only disrupt operations but may also compromise the occupant's safety in critical facilities. The author proposes a multi-layered defense strategy combining the Intrusion Detection and Prevention Systems (IDS/IPS), end-to-end encryption, and workforce training to promote a cyber-resilient organizational culture. Beyond the technical safeguards, the study stresses the need to align with the international cybersecurity standards (e.g., IEC 62443, ISO/IEC 27001) and to embed governance practices that balance the operational efficiency, sustainability targets, and regulatory compliance. The IEMS emerge not only as enablers of energy efficiency but also as critical elements in building secure, trustworthy, and future-ready infrastructures (Wyrzykowska, 2024).

The rise of the IoT-enabled smart buildings has increased the need for robust security frameworks that ensure both operational efficiency and cyber-resilience. Standards such as IEC 62443, BACnet Secure Connect (BACnet/SC), and OPC UA provide guidance for industrial and building automation, covering the authentication, encryption, and access control. However, these standards often do not fully address the constraints of the heterogeneous IoT environments or the requirements of the predictive energy management. The PRECONERG system proposes a multi-layered security architecture specifically for the energy prediction in smart buildings. By combining the edge-level processing, end-to-end encryption, multi-layer intrusion detection, and KPI (Key Performance Indicators)-driven monitoring, PRECONERG addresses emergent threats, ensuring a balanced performance vs. the security ratio. Table 1 presents a comparative analysis highlighting the distinct contributions of the PRECONERG relative to the IEC 62443, BACnet/SC, and OPC UA, emphasizing their scope, security layers, device constraints, interoperability, and optimization mechanisms.

Table 1: *Existing security models vs. PRECONEG comparison*

Characteristic/ Layer	IEC 62443	BACnet Secure Connect	OPC UA	PRECONERG (Proposed System)
Scope	Industrial Automation Systems	Building Automation	Industrial & Building Automation	Smart Energy Prediction in Buildings (IoT + Edge + Cloud)
Security Layers	Functional, Network, System	Device + Network	Application + Transport	IoT Device → Gateway → Perimeter → Data Processing & API
Data Protection	Encryption optional / AES	TLS / mTLS	TLS / Certificates	End-to-End TLS 1.3/ AES-256, Key rotation, HSM, MFA, OAuth, RBAC
Intrusion Detection	Not specified	Limited	Not standard	Multi-layer IDS/IPS (signature + ML), WAF, DDoS protection
Device Constraints Handling	Minimal guidance	Limited	Moderate	Lightweight edge pre-processing, telemetry minimization, secure boot, X.509 per device
Interoperability	Protocol agnostic	BACnet only	OPC UA native	Supports heterogeneous IoT protocols (MQTT/ CoAP), cloud integration
Performance vs. Security Trade-offs	Guidance at system level	Limited	Moderate	Explicit KPIs defined per layer: latency, throughput, ML accuracy, energy overhead
Monitoring & Logging	Recommended	Optional	Optional	Centralized SIEM + ELK/Splunk, real-time alerts, incident response automation
Optimization / emergent threat handling	Not detailed	Limited	Limited	ML/AI for zero-day detection, hybrid IDS, automated incident response, sandboxing, lightweight crypto

SECURITY ARCHITECTURE SPECIFICATION

Security Requirements and Challenges for Intelligent Energy Consumption Prediction Systems in Non-Residential Buildings

Vulnerabilities, Threats, and Risks

The non-residential buildings (office complexes, hospitals, shopping centers, university campuses) are increasingly adopting intelligent energy prediction systems. These systems leverage the IoT-enabled sensors, Building Automation Systems (BAS), cloud platforms, and ML algorithms to optimize the energy consumption, forecast the demand fluctuations, and enable the demand-side response strategies. While the benefits are significant - lower operational costs, compliance with sustainability goals, and improved occupant comfort - these systems also create a complex cyber-physical environment that is vulnerable to a wide range of attacks (Dumitrache & Sandu, 2020).

The typical security problems (vulnerabilities and threats) for the IoT-based intelligent energy prediction systems in smart buildings include:

- IoT specific security issues: lack of the effective security functions by-design, vulnerable firmware, attacks against the communication protocols in the IoT (MQTT-Message Queuing Telemetry Transport, CoAP- Constrained Application Protocol);
- AI/ML-specific issues: data poisoning, model inversion;
- Networks and servers-specific issues: MitM (Man-in-the-Middle) attacks, DDoS, SQL injection, digital certificates compromising.

The intelligent building infrastructures deal with 2 main categories of risks: (1) cyber risks, such as data breaches, ransomware attacks, or model poisoning, and (2) physical risks, including tampering with the IoT devices or intercepting control signals. The unauthorized access to Heating, Ventilation, and Air Conditioning (HVAC) systems could not only increase the energy costs but also endanger the occupants' health, particularly in critical facilities like hospitals. This subchapter details the security requirements and challenges associated with the deploying of the intelligent energy consumption prediction systems in the non-residential buildings.

Specific Requirements for Ensuring Security

The major security requirements for systems like PRECONERG include:

- *Data protection:* confidentiality, integrity, authentication, authorization, non-repudiation, availability;
- *Attacks detection and prevention:* firewall, IDS/IPS, secure monitoring and logging, protection against DDoS;
- *Strong authentication:* digital certificates, OAuth2.0, MFA;
- *RBAC:* a clear definition for the legitime roles, minimal privileges, audit, GDPR (General Data Protection Regulation) compliance (Savu & Mitan, 2024).

These requirements should be mapped onto several functional categories, as follows.

a) Data Protection and Encryption

The data collected from smart meters, occupancy sensors, and environmental monitoring devices are highly sensitive, as they can reveal operational patterns, business cycles, and potential security vulnerabilities. Malicious actors could exploit this information to profile organizational behavior or stage targeted intrusions.

To mitigate such risks, an end-to-end encryption must be enforced. The Transport Layer Security (TLS) 1.3 with the Perfect Forward Secrecy (PFS) ensure that compromising a key does not expose former communications. For the data at rest, the Advanced Encryption Standard (AES) with 256-bit keys (AES-256), combined with a secure key management via a Hardware Security Module (HSM), prevents an unauthorized decryption even in the event of a data breach. Additionally, the anonymization techniques, such as pseudonymization or k-anonymity, can be used when sharing datasets with third parties or for research purposes (Zhou et al., 2024).

The encryption introduces computational overhead and latency, which may be challenging for the resource-constrained IoT devices used in real-time energy monitoring.

The Attribute-Based Encryption (ABE) has been proposed as a scalable solution that supports the contextual, role-based access control while minimizing data leakage risks (Hao et al., 2019). Integrating these schemes into a large-scale BAS remains an active research challenge.

b) Secure Interoperability and Standards Compliance

Interoperability allows heterogeneous systems to work together but also increases the attack surface. Modern non-residential buildings often operate a mix of BACnet devices for HVAC, Modbus for legacy equipment, and Zigbee or MQTT for the IoT sensors. The absence of a unified security baseline makes such environments attractive targets for attackers.

Standards like IEC 62443 recommend a network segmentation through zones, isolating the critical systems from the non-critical services. For instance, separating the predictive analytics servers from the physical access control systems ensures that a breach in one subsystem does not propagate. Protocols such as the BACnet Secure Connect (BACnet/SC) and Open Platform Communications Unified Architecture (OPC UA) embed the TLS encryption, certificate-based authentication, and secure session management directly into the building automation communications.

The Legacy devices remain a key challenge. The older systems often lack sufficient processing power to perform cryptographic operations, leaving them vulnerable to replay attacks or unauthorized control. Security gateways that encapsulate the legacy traffic within the secure tunnels show promise, but their large-scale deployment is limited by their cost and integration complexity. Consequently, any compliance with standards is often partial, resulting in a heterogeneous landscape of secure and insecure communication channels (Korodi et al., 2024).

c) Identity, Authentication and Authorization

The access control mechanism in smart buildings involves human users, automated services, and IoT devices. A facility manager, a cloud-based energy analytics service, and a predictive HVAC controller all require different access rights. Without a robust Identity and Access Management (IAM), unauthorized privilege escalation could have serious consequences.

OAuth 2.0 provides a practical solution by issuing scoped tokens that restrict the privileges for each actor (Anon OAuth 2.0, 2025). An energy provider may receive a token allowing read-only access to consumption data without control over devices. OpenID Connect (OIDC) complements OAuth by enabling federated authentication, facilitating the integration with the enterprise identity providers, and reducing the password fatigue via Single Sign-On (SSO) (Anon OpenID, 2025).

Zero Trust architectures further strengthen the IAM by continuously verifying user's and device identities, regardless of the network location (Rose et al., 2020). A building management system may require a re-authentication if an access request originates from an unusual location or device. The IETF (2025) recommends avoiding implicit grants, using proof-of-possession tokens, and enforcing short-lived refresh tokens to mitigate token replay attacks.

A major challenge is to balance security with usability. While frequent re-authentication increases security, it may frustrate users and administrators. Context-aware authentication, which incorporates real-time risk scores based on the device trustworthiness, location, and behavioral patterns, offers a promising compromise.

d) Machine Learning and Data Governance Security

ML is central to energy prediction but is also a target for adversarial attacks. Data poisoning, for example, can bias forecasts, leading to inefficient energy allocation or financial losses. Injecting false peak consumption data could mislead demand-response strategies, increasing costs for building operators.

Secure MLOps pipelines are essential to counter such threats. These pipelines should integrate:

- OAuth-based access control for training datasets and model APIs.
- Digital signatures for model verification, ensuring only authentic models are deployed.
- Differential privacy to prevent inference attacks on occupant data.

- Federated Learning (FL) to allow local training on edge devices, sharing only encrypted model updates centrally (Li et al., 2020).

Foundation models for time-series analytics can improve energy prediction accuracy (Lin et al., 2024), but their large-scale deployment raises security concerns. Auditing, explaining, and securing these models against adversarial manipulation remains a challenge. Ensuring data governance through regulatory compliance (e.g., GDPR, CCPA) requires anonymization and strict access logging.

e) Cyber-Physical and Operational Security

Unlike the conventional IT systems, BAS devices are physically accessible. Attackers could tamper with occupancy sensors, disconnect the IoT devices, or install malicious firmware.

Countermeasures include:

- Secure boot processes that verify firmware integrity before execution.
- Tamper detection hardware that triggers alerts if a device casing is opened.
- Lightweight OAuth profiles for constrained devices to ensure authentication on low-power sensors.
- Protocol-aware Intrusion Detection Systems (IDS) trained on BACnet or MQTT traffic patterns to detect anomalies, such as the unusual broadcasts or flooding attacks (Empl, Böhm, & Pernul, 2024).

The operational resilience is also very important for such use-cases. The redundancy in critical sensors and systems can prevent a single compromised device from disrupting the entire predictive control loop.

f) Holistic and Multilayered Security Approach

Security for the intelligent energy prediction systems requires a defense-in-depth strategy, integrating protections across all layers:

- Device layer: hardware-backed encryption, secure boot, tamper resistance.
- Network layer: TLS/mTLS, BACnet/SC, IEC 62443 segmentation, anomaly detection.
- Application layer: OAuth 2.0, OIDC, context-aware authorization (e.g., XACML policies).
- Analytics layer: differentially private datasets, federated learning, signed ML models.
- Governance layer: continuous compliance monitoring, penetration testing, adherence to ISO/IEC 27001 and IEC 62443.

This approach emphasizes that no single mechanism suffices. The resilience emerges from the overlapping protections, and the governance ensures both technical robustness and regulatory compliance, enhancing stakeholder trust in the long-term deployments.

Architectural Model for Data Security

Security Architecture

The PRECONERG System is designed as an intelligent IoT- and AI-driven platform for the energy consumption prediction in non-residential buildings. It should incorporate a multi-layered security architecture to protect data confidentiality, integrity, and availability across the entire lifecycle of the energy prediction. Given the exposure of the IoT devices to cyber threats, security must be embedded as a core design principle rather than an add-on.

The proposed security architecture follows the previously specified principles and essentially it is based on 2 major conceptual layers:

- *Data Protection Layer:* end-to-end encryption (TLS 1.3/AES-256); cryptographic keys management (HSM, periodical keys rotation); strong authentication (MFA, digital certificates, OAuth); RBAC.
- *Cyber Attack Detection and Prevention Layer:* multi-layer firewall (perimeter, application, device level or WAF- Web Application Firewall), IDS/IPS (signatures and ML-based for anomaly detection), monitoring and logging (ELK, Splunk, SIEM), automated incident response and DDoS protection.

These layers are very generic; for a proper correlation with more specific security functionalities, the multi-layer architectural model for PRECONERG should contains 4 functional layers (Figure 1). For each layer one can specify: a) its basic features (specific mechanisms and controls for data protection, intrusion detection and secure system working); b) design optimization/trade-off options concerning security vs. performance (in respect

to the IoT constraints); c) Key Performance Indicators (KPI) that should be used to evaluate or at least to predict the performance in a real use-case; d) the potential impact regarding the system functionalities (using as criteria the performance, reliability, ML accuracy, basic operations).

Layer 1: IoT Devices, Sensors and Edge Security

The Physical layer includes environmental sensors, heterogeneous smart meters, and Raspberry Pi edge devices. These devices are deployed through the building to collect energy-related data. Usually, the IoT devices are featured by limited (low) CPU/RAM and battery power. They produce time-series energy consumption data. The typical communication protocols are MQTT/CoAP over IP. The vulnerabilities of the IoT devices are caused by their limited computational and energy resources. To mitigate the inherent vulnerabilities of the IoT devices, a lightweight pre-processing is performed at the edge, including data aggregation, filtering, and rate limiting, thereby reducing any exposure to injection or flooding attacks. Device hardening techniques (minimal OS footprint) are combined with secure boot to ensure firmware integrity. The goal of the security measures in this layer is to protect data at source and to preserve the device identity.

Additional optimization options could be considered for this layer after the full development process will be completed. These options for data protection may include:

- end-to-end cryptographic process at device, looking to encrypt the payloads at the sensor/gateway boundary. TLS 1.3 will be the 1st option. Otherwise, TLS 1.2 with AES/GCM Mode (Galois/Counter Mode) could be considered. The keys per device should be generated with strong RNG (Random Number Generator) and should also be periodically rotated;
- device identity with X.509: per-device certificates, storing the private keys in secure locations when available; mutual TLS if supported by devices;
- local access control: the minimization of the local services; host-level firewall on the devices that allow it;
- telemetry minimization: only the most significant features should be sent in cases with bandwidth/energy constraints.

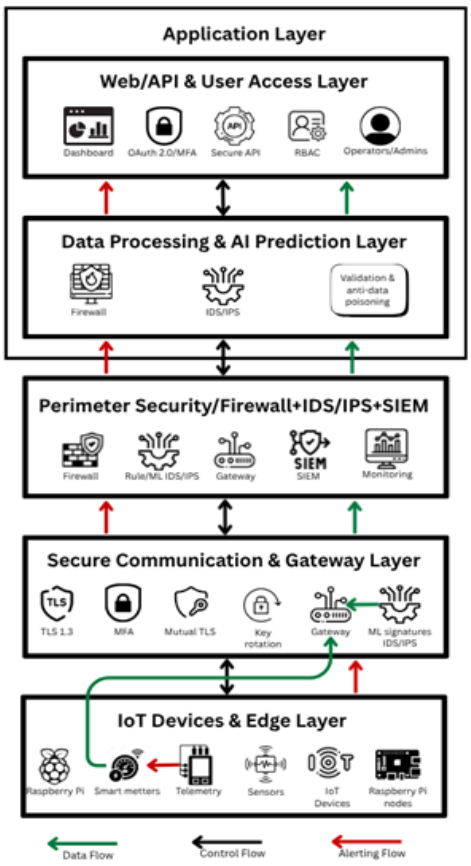

Figure 1: *The multi-layer security architecture for PRECONERG*

KPIs for this layer include: the number (or percentage) of devices with still valid certificates; key rotation interval; end-to-end encrypted payload ratio; mean battery impact (mWh/day) assumed for the cryptographic process.

The impact on the system concerns: performance vs. energy, reliability (some misconfigurations can cause failures), ML accuracy for the predictions (as much as the security on data sources could prevent data tampering/poisoning and allows to generate and run more stable models with proper training and reliable predictions about the energy consumption).

Layer 2: Secure Communication/Gateway

All data flows between the IoT devices, edge nodes, and cloud components are protected using end-to-end encryption (TLS 1.3). The Certificate-based authentication and mutual TLS (mTLS) guarantee that only trusted devices can communicate with the backend, while the periodic key rotation strengthens the resistance against a long-term credential compromise. These mechanisms prevent eavesdropping, tampering, and replay attacks, which are prevalent in the IoT-enabled systems.

Additional optimization options for this layer after the full development process may include:

- strong authentication: MFA for human users, certificates for devices; app-to-app authorization with OAuth 2.0;
- key management: integration with HSM-backed key vault for generation, storage, rotation.

KPIs for this layer include: successful vs. failed authentications, handshake latency, certificate renewal success rate.

The impact on the system concerns: latency, resilience (the critical role of the gateway).

Layer 3: Perimeter Security with Firewall and IDS/IPD, SIEM

At the network boundary, the architecture integrates a firewall and intrusion detection/prevention system (IDS/IPS). The multi-layer firewall includes: perimeter, WAF for API/web, host-level. The firewall enforces IoT-specific rules, VLAN (Virtual Local-Area Network)-based segmentation, and DDoS mitigation, while the IDS/IPS layer leverages both signature-based detection and ML-based anomaly detection to identify zero-day attacks. The centralized SIEM is useful for log aggregation, correlation, alerting. Together, these controls ensure that malicious traffic is filtered before reaching critical components of the prediction pipeline.

Additional optimization options for this layer after the full development process may include:

- hybrid detection (IDS): signature-based IDS for fast, low-cost blocks but with periodical updating of the detection rules; ML-based anomaly detection adapted for IoT traffic patterns and constraints, including clustering (unsupervised learning) to perform the outliers' detection in device behaviors;
- automated response;
- protection against DDoS: local rate-limiters, black/white lists.

KPIs for this layer include: detection rate -TPR (True Positive Rate), false alarm rate-FPR (False Positive Rate), MTTD (Mean-Time-To-Detect), MTTR (Mean-Time-To-Response), number of quarantined devices, WAF blocks.

The impact on the system concerns: performance (WAF, anomaly scoring with impact on CPU, added latency), availability (as much as the auto-quarantine can disrupt the normal operations if the thresholding is improperly fixed), ML accuracy (the early detection of poisoned data/outliers preserve the quality of the training/testing/validation datasets), operations.

Layer 4: Security for Data Processing and Storage (Application-level), Web/API Access

Within the cloud environment, the proposed security model for PRECONERG applies encrypted storage and integrity verification for all ingested data. The access control policies are enforced through the role-based authentication and fine-grained authorization, ensuring that only authorized users and services can access the prediction models or sensitive energy datasets. The system also integrates continuous monitoring and logging, enabling anomaly detection and forensic analysis. By combining monitoring with resilience-by-design principles, one can maintain a robust protection against both inside and external threats. The full design and development process will include user/admin portals and APIs for dashboards, configuration.

Additional optimization options for this layer after the full development process may include:

- secure ingestion: accept only traffic from trusted gateways, enforce size and frequency limits, integrity verification (MAC-Message Authentication Code/signatures);

- confidentiality and integrity; encryption for data in transit (TLS);
- access control: datasets and models protection with RBAC;
- anti-data poisoning: data validation, isolate low-trust sources;
- federated authorization with OAuth 2.0; MFA for the most privileged roles; mutual TLS; HTTPS;
- RBAC enforcing least privilege; role hierarchy and context-aware rules; logging of all access and admin actions;
- WAF.

KPIs for this layer include: model accuracy, authentication success/failure rates, RBAC policy violations that were prevented, WAF blocks.

The impact on the system concerns: latency (caused by integrity checks and encryption), throughput, model quality (improved due to the cleaning process).

The security architectural model for PRECONERG shows the layered flow from the IoT devices to the application level, secure Web/API Access with RBAC and OAuth. The protection layers (TLS, Firewall+IDS/IPS, SIEM) and the processing layer can be seen.

The multi-layer security design does not only prevent cyber threats but also improves the reliability and trustworthiness of the energy consumption forecasts. By guaranteeing data authenticity and secure model execution, the proposed security model ensures that predictive outputs are free from adversarial manipulation—an essential requirement in the energy management applications where inaccurate forecasts can lead to economic and operational inefficiencies.

Cross-layer Functions

The architectural specifications include as cross-cutting functionalities the monitoring, secure logging and compliance spanning across all layers. The centralized log collection ensures the tamper-evident event tracking, while the SIEM correlation enhances the incident readiness. The logs are securely retained for audits and forensic investigations, with compliance in respect to the data protection regulations.

Functional Components and Corresponding Performance Indicators

The design and development process will consider several performance indicators (KPI):

- for the IoT devices and sensors: battery impact (the increase of daily energy consumption of the devices due to the cryptographic tasks), certificate validity (percentage of the devices with valid and up-to-date certificates);
- for secure gateway (TLS 1.3, MFA, certificates): TLS handshake latency, authentication success/failure rate (percentage of valid/rejected connections), key rotation efficiency (percentage of keys/certificates renewed before their expiring time);
- for encryption: encryption overhead (average processing time/CPU, storage/memory usage per device for TLS, bandwidth, latency), percentage of the end-to-end encrypted data;
- for data integrity: percentage of the sensor messages successfully verified with MAC/signature;
- for authentication: rate of failed detected authentications, TTR (Time-To-Response);
- for the RBAC component: access control granularity, number of fraudulent access attempts prevented or blocked;
- for IDS/IPS: anomaly detection rate -TPR (True Positive Rate), false alarm rate- FPR (False Positive Rate), accuracy, precision;
- for SIEM/Logging, firewall: mean time to detect the incident (MTTD), mean time to response (MTTR), quarantine effectiveness (counting for the number of compromised devices successfully isolated);
- for data processing & the AI-based prediction: data ingestion latency (average delay from sensor data capture to processing), encryption coverage (percentage of encrypted data in site and in transfer), data poisoning resilience (percentage of anomalous samples successfully flagged and removed), model drift alerts (counting the significant outliers signaled in the prediction accuracy);

- for Web/API Access (RBAC, OAuth 2.0, MFA): authorization latency (time to generate an access token or to validate credentials), prevented access control violations (counting for the blocked unauthorized access attempts), MFA adoption rate (percentage of privileged accounts having MFA-based protection), audit log completeness (percentage of critical processes recorded in resilient logs);
- for monitoring and logging (cross-cutting): log coverage (percentage of components generating standardized logs), log ingestion latency (delay between the event occurrence and SIEM availability), compliance rate (percentage of security reports/logs - compliant with GDPR/ISO), incident investigation lead time (time required to reconstruct the attack chain from logs).

Optimization Solutions Against Emergent Threats

Additional solutions to enhance the system protection against emergent threats could include:

- ML/AI integration to develop high-performance mechanisms for zero-day attacks detection;
- the continuous updating of the signatures list for firewall/IDS;
- hybrid IDS: signature-based detection for known malicious cases and ML-based anomaly detection for unknown/zero-day cases, IDS thresholds adjustment with contextual elements to reduce FPR;
- federated learning to train the anomaly detection models across gateways avoiding the centralization of the sensitive IoT data;
- improving the ML model's robustness through differential privacy or adversarial training;
- isolating the low-confidence data sources: segment training data by trust level;
- the IoT network segmentation to reduce the attack surface;
- lightweight cryptographic primitives to reduce the device resources loading and HSMs for accelerated cryptographic processes;
- session resumption and persistent TLS connections to bypass/reduce the full handshake process;
- automate certificate lifecycle management;
- firmware with secure boot and signed updates to prevent new malware targeting the IoT devices;
- automated incident response;
- adaptive MFA (for high-risk access only) to ensure a proper trade-off usability vs. security;
- integrate an anomaly detection engine within the data processing chain to filter the poisoned data, automatic retrain of the model with clean datasets to properly adapt against data poisoning cases;
- fine-trained RBAC with just-in-time access for the most privileged processes, OAuth 2.0 with short-live tokens and Proof Key for Code Exchange (PKCE);
- tamper-evident, encrypted logging with secure time-stamping;
- cloud-based retention and tiered storage for scalability
- applying ML on SIEM data to prioritize the critical alerts;
- sandboxing or using containerization and secure enclaves for processing sensitive data.

The Impact of Security Functions on the General Functions of the Intelligent Energy Consumption Prediction Systems in Non-Residential Buildings

As concerning the impact of the specified security architecture on the general operations of the Intelligent Energy Consumption Prediction Systems in Non-Residential Buildings (PRECONERG), there are 2 sides:

- the positive impact (benefits): increased trustiness, GDPR compliance (by ensuring the privacy protection), accurate predictions with low chances for being compromised;
- the potential negative impact (drawbacks): increased latency (for encryption, IDS/IPS), increased energy consumption for the IoT devices (generated by TLS and MFA), need for additional Hardware/Software resources with increased expenses;
- the specific advantage for the PRECONERG: The integrated security could actually improve the

performance of the prediction functionality by providing the data integrity. The accurate data allows to make more precise predictions about the energy consumption.

The impact of the proposed security architectural model on the system to be developed can be resumed as follows:

- performance: overhead introduced by cryptography (TLS), WAF, anomaly detection;
- availability: resilient certificate lifecycle;
- scalability: RBAC centralization, token-based authentication, horizontally-scaled IDS/IPS supporting the growth with the number of devices;
- data quality and ML accuracy: integrity-checked and properly-monitored processing may reduce the opportunities for data poisoning, ensuring the model stability for prediction performance.

The proposed security architecture ensures a measurable protection through the systematically monitoring of the KPIs across all layers. The optimization strategies, including the lightweight cryptography, adaptive authentication, ML-based intrusion detection and automated response, addresses the extending landscape of the cyber threats, with a proper trade-off security vs. performance and usability. The layered architecture introduces computational and operational overhead, caused by TLS handshakes, anomaly scoring, WAF filtering. This drawback is outweighed by the main benefits: the protection of the sensitive energy consumption data, resilience against cyberattacks, improved robustness of the AI predictions. The security integration per layer prevents or reduces the opportunities for data tampering, ensuring the enhanced trustiness of forecasts and supports, the compliance with the international standards for the critical infrastructures protection.

CONCLUSIONS AND FUTURE DIRECTIONS

This paper addresses the need for resilient security architectures in intelligent energy consumption prediction systems for non-residential buildings. The proposed multi-layer security architecture is designed to ensure the protection of data confidentiality, integrity, and availability for the non-residential cases. It looks to properly manage the vulnerabilities specific to the complex IoT-based infrastructures with a risk mitigation approach. The proposed model provides a layered structured and adaptive security framework that integrates complementary mechanisms (encryption, authentication, access control and ML-based intrusion detection) applied consistently across all layers, from the IoT devices to the cloud-based AI prediction. It follows a security-by-design approach able to improve both the resilience and trustworthiness of the intelligent energy prediction systems. This comprehensive defense-in-depth approach is supported by the mapping of the security requirements onto functional components across multiple layers.

Compared to the existing industrial and building automation security standards (e.g., IEC 62443, BACnet/SC, OPC UA), the PRECONERG security architecture follows a more granular and KPI-driven approach that explicitly correlates the security layers with the measurable performance indicators. The security model supports the end-to-end protection and resilience through defense-in-depth strategies, including the MFA, RBAC, federated authorization (OAuth 2.0), SIEM-based monitoring, and hybrid IDS/IPS detection combining the signature-based and ML-driven models. The adoption of the mechanisms such as MFA, RBAC, and ML-based IDS strengthens both the confidentiality and integrity of the sensitive energy data, while also supporting the availability and operational reliability. This approach advances beyond the conventional existing static models by incorporating the adaptability and data integrity protection directly into the predictive energy management framework.

The main contributions are: 1) the formal specification of a security-by-design layered security functional architecture for energy consumption prediction systems in the complex IoT environments; 2) the integration of performance indicators (latency, energy overhead, detection accuracy, encryption efficiency/

overhead, authentication success rates) for the evaluation of the trade-off between security and the system operational performance. The KPI specification provides a measurable framework for the effectiveness of the security architecture; 3) the inclusion of the ML-driven mechanisms for the anomaly detection and incident response automation within the architectural model. These contributions provide an advance beyond the current state through a quantifiable, modular, and scalable security framework for critical non-residential infrastructures. The specified optimization strategies (lightweight cryptographic primitives, adaptive MFA, hybrid IDS) illustrate its adaptability in respect to the emerging threats and resource-constrained IoT environments. The optimal trade-off security vs. performance still remains a major challenge for the development of the smart systems in non-residential buildings with complex IoT infrastructures.

Future work will focus on: extending the architecture towards large-scale pilot implementations; the simulation-based performance evaluation under various attack scenarios to prove the resilience and operational efficiency of the model; the experimental validation and pilot testing of the proposed architecture within real non-residential buildings (measuring security KPIs such as encryption latency, IDS accuracy, and energy impact); refining optimization strategies to ensure scalability and resilience. The main experimental directions for the future specific case studies will include: the architecture testing in real-world pilot scenarios, measuring its performance under simulated attack conditions, the validation of the energy cost for cryptographic mechanisms on the IoT devices. Further research will also investigate: the integration of the emerging paradigms such as federated learning for privacy-preserving analytics with secure distributed model training on edge devices, the adoption of the Zero Trust principles for adaptive identity management (to strengthen the dynamic identity verification and continuous authentication); the AI-driven autonomous response mechanisms to counter the increasingly sophisticated Cyber-attacks; the development of a testing environment allowing to assess the scalability, interoperability, and cost–benefit trade-offs of the lightweight cryptographic mechanisms in the IoT environments.

The security architecture for the PRECONERG will provide a robust, adaptive, and future-ready security support for the intelligent energy consumption prediction systems, aiming to balance their operational efficiency, resilience, sustainability objectives, and trustworthiness in the non-residential smart buildings.

ACKNOWLEDGEMENTS

This work was supported from the project “Intelligent system for predicting energy consumption in buildings (PRECONERG)”, funded by the Romanian Core Program of the Ministry of Research, Innovation and Digitalization (2023-2026).

REFERENCE LIST

Amangeldy, B., Tasmurzayev, N., Imankulov, T., Baigarayeva, Z., Izmailov, N., Riza, T., Abdukarimov, A., Mukazhan, M. & Zhumagulov, B. (2025) AI-Powered Building Ecosystems: A Narrative Mapping Review on the Integration of Digital Twins and LLMs for Proactive Comfort, *IEQ, and Energy Management. Sensors.* 25(17), 5265. doi: org/10.3390/s25175265.

Anon OAuth 2.0. (2025) *Oauth 2.0.* https://oauth.net/2/ [Accessed 27th August 2025].

Anon OpenID (2025) OpenID. https://openid.net/connect/ [Accessed 27th August 2025].

Arshad, J., Azad, M. A., Amad, R., Salah, K., Alazab, M. & Iqbal, R. (2020) A Review of Performance, Energy and Privacy of Intrusion Detection Systems for IoT. *Electronics.* 9(4), 629. doi:org/10.3390/electronics9040629.

Barros, E.B.C., Souza, W.O., Costa, D.G., Rocha Filho, G.P., Figueiredo, G.B. & Peixoto, M.L.M. (2025) Energy Management in Smart Grids: An Edge-Cloud Continuum Approach with Deep Q-learning. *Future Generation Computer Systems.* 165, 107599. doi: org/10.1016/j.future.2024.107599.

Dumitrache, M. & Sandu, I.-E. (2020) Network security and communication systems in Smart environments. *Romanian Journal of Information Technology and Automatic Control.* 30(1), 61-70. doi: org/10.33436/v30i1y202005.

El Hadj Youssef, W., Abdelli, A., Dridi, F. & Machhout, M. (2020) Hardware Implementation of Secure Lightweight Cryptographic Designs for IoT Applications. *Security and Communication Networks*. 8860598. doi: org/10.1155/2020/8860598.

Empl, P., Böhm, F. & Pernul, G. (2024) Process-Aware Intrusion Detection in MQTT Networks. In *Proceedings of the Fourteenth ACM Conference on Data and Application Security and Privacy (CODASPY '24), 19–21 June 2024, Porto, Portugal. ACM.* doi: org/10.1145/3626232.3653271.

Hao, J. & Wang, H. (2019) Fine-grained data access control with attribute-hiding policy for cloud-based IoT. *Computer Networks.* 152, 1–11. doi: org/10.1016/j.comnet.2019.02.008.

IETF. (2025) RFC 9700: Best Current Practice for OAuth 2.0 Security. *Internet Engineering Task Force.* https://datatracker.ietf.org/doc/rfc9700/ [Accessed 27th August 2025].

Isong, B., Kgote, O. & Abu-Mahfouz, A. (2024) Insights into Modern Intrusion Detection Strategies for Internet of Things Ecosystems. *Electronics.* 13(12), 2370. doi: org/10.3390/electronics13122370.

Korodi, A., Nițulescu, I.-V., Fülöp, A.-A., Vesa, V.-C., Demian, P., Braneci, R.-A. & Popescu, D. (2024) Integration of Legacy Industrial Equipment in a Building-Management System Industry 5.0 Scenario. *Electronics.* 13(16), 3229. doi: org/10.3390/electronics13163229.

Li, T., Sahu, A. K., Talwalkar, A. & Smith, V. (2020) Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine.* 37(3), 50–60. doi: org/10.1109/MSP.2020.2975749.

Lin, X., Prabowo, A., Razzak, I., Xue, H., Amos, M., Behrens, S. & Salim, F.D. (2024). A Gap in Time: The Challenge of Processing Heterogeneous IoT Data in Digitalized Buildings. *IEEE Pervasive Computing.* 24, 19-31. https://arxiv.org/abs/2405.14267.

Mai, T.D. (2019) Research on Internet of Things security architecture based on fog computing. *International Journal of Distributed Sensor Networks.* 15(11). doi: org/10.1177/1550147719888166.

Mohanta, B.K., Debasish, J., Satapathy, U. & Patnaik, S.K. (2020) Survey on IoT security: Challenges and solution using machine learning, artificial intelligence and blockchain technology. *Internet of Things.* 11, 100227, doi: org/10.1016/j.iot.2020.100227.

Mrabet, H., Belguith, S., Alhomoud, A. & Jemai, A. (2020) A Survey of IoT Security Based on a Layered Architecture of Sensing and Data Analysis. *Sensors*. 20(13), 3625. doi: org/10.3390/s20133625.

Poyyamozhi, M.; Murugesan, B.; Rajamanickam, N.; Shorfuzzaman, M. & Aboelmagd, Y. (2024) IoT—A Promising Solution to Energy Management in Smart Buildings: A Systematic Review, Applications, Barriers, and Future Scope. *Buildings.* 14, 3446. doi: org/10.3390/buildings14113446.

Rose, S., Borchert, O., Mitchell, S. & Connelly, S. (2020) Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology. doi: org/10.6028/NIST.SP.800-207.

Savu, D. & Mitan, E. (2024) A Proactive Approach to Mitigate Cyber Risks in IoT Systems. *Romanian Cyber Security Journal.* 6(2), 107-121. doi: org/10.54851/v6i2y202410.

Trabelsi, R., Fersi, G. & Jmaiel, M. (2023) Access control in Internet of Things: A survey. *Computers & Security.* 135, 103472, doi: org/10.1016/j.cose.2023.103472.

Tudora, E. & Tîrziu, E. (2024) Advanced approaches in building energy consumption prediction. *Romanian Journal of Information Technology and Automatic Control.* 34(2), 21-34. doi: org/10.33436/v34i2y202402.

Wyrzykowska, B.; Szczepaniuk, H., Szczepaniuk, E.K., Rytko, A. & Kacprzak, M. (2024) Intelligent Energy Management Systems in Industry 5.0: Cybersecurity Applications in Examples. *Energies.* 17, 5871. doi: org/10.3390/en17235871.

Zhou, J.; Fu, W., Hu, W., Sun, Z., He, T. & Zhang, Z. (2024) Challenges and Advances in Analyzing TLS 1.3-Encrypted Traffic: A Comprehensive Survey. Electronics. 13, 4000. https://doi.org/10.3390/electronics13204000.