Intrusion Detection in Intelligent Energy Prediction Systems for Non-Residential Buildings

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

Abstract: This paper addresses the cybersecurity of intelligent energy prediction systems in non-residential buildings by proposing a formally structured Multi-Layer Correlated Intrusion Detection System (MLC-IDS) aligned with the PRECONERG security-by-design architecture. The main contribution consists of extending traditional IoT intrusion detection by integrating AI pipeline integrity protection and cross-layer alert correlation within a unified detection framework. The proposed methodology introduces resource-aware edge-level anomaly detection, protocol-aware behavioral modeling at the gateway level, hybrid signature–/Machine Learning-based network intrusion detection with adaptive fusion, and data poisoning and concept drift detection embedded directly in the prediction lifecycle. A formal alert aggregation model is defined to support the trade-off between detection accuracy and operational overhead. An experimental validation framework with Machine Learning models is proposed to demonstrate measurable improvements in detection reliability and prediction robustness under adversarial conditions.
Keywords: Intrusion Detection Systems, IoT Security, Multi-Layer Security Architecture, Intelligent Energy Consumption Prediction, Non-Residential Smart Buildings.

INTRODUCTION

The increasing deployment of IoT (Internet-of-Things)-enabled intelligent energy prediction systems in non-residential buildings supports the evolution toward sustainable infrastructure and operational efficiency. Systems such as PRECONERG (Intelligent System for Predicting Energy Consumption in Buildings) should integrate IoT devices, edge processing, and AI (Artificial Intelligence) to efficiently predict the energy consumption patterns and optimize the energy distribution.

This process deals with significant cybersecurity challenges. The issues are generated by multi-layer cyber threats affecting both communication infrastructures and AI-driven predictive pipelines. This exposure is particularly due to device heterogeneity, large-scale data collection, and the integration of AI-based forecasting models. The high density of heterogeneous IoT devices, and the

https://doi.org/10.54851/v8i1y202604

integration of AI-based forecasting in intelligent energy prediction systems for non-residential buildings, increases the cyber-attack surface. While the layered security architectures provide basic protection, the traditional intrusion detection mechanisms are not always able to efficiently address the multi-stage attacks targeting both network infrastructure and data processing pipelines. Many existing intrusion detection solutions do not properly perform in cases where data integrity directly impacts ML (Machine Learning) model accuracy, the IoT devices are resource-constrained, and multi-stage attacks may span across several layers (device, network, AI). The reliance on ML-based prediction models could expand the attack surface because of issues like data poisoning.

As stated in the previous work (Tudora et al., 2025), the intelligent energy consumption prediction systems should include layered protection mechanisms to ensure confidentiality, integrity, and availability. This paper proposes a Multi-Layer Correlated Intrusion Detection System (MLC-IDS) tailored for IoT-based intelligent energy prediction platforms such as PRECONERG. The methodology extends the previously defined security-by-design architecture (Tudora et al., 2025) by embedding intrusion detection mechanisms across 4 functional layers: lightweight edge-level detection, gateway protocol-aware anomaly detection, hybrid network IDS/IPS (Intrusion Detection/Prevention System) with SIEM (Security Information and Event Management)-based correlation, and AI-driven data integrity monitoring for protection against data poisoning and concept drift manipulation (attacks). The proposed model integrates cross-layer alert fusion, federated anomaly detection capabilities, and adaptive thresholding mechanisms to balance security (detection accuracy) with performance constraints in resource-limited IoT environments. An experimental validation framework with ML models is defined to assess the detection performance. Performance indicators are specified to evaluate the effectiveness of the proposed IDS. The present framework, together with the next few case studies with several datasets, will demonstrate that the proposed approach will enhance both cybersecurity resilience and energy consumption prediction reliability in smart non-residential buildings. The remainder of the paper is structured as follows: Section II -related works; Section III -the proposed methodology; Section IV - technical framework and experimental validation of the hybrid IDS architecture; Section V -conclusions together with future works/research and development directions.

RELATED WORKS

Recent research highlights the importance of hybrid IDS models (combining signature- and anomaly-based detection) in IoT ecosystems with promising results (Isong et al., 2024), lightweight cryptography for constrained environments (El Hadj Youssef et al., 2020), and the need for ML-aware security mechanisms (Mohanta et al., 2020). Process-aware intrusion detection for MQTT (Message Queuing Telemetry Transport) networks further shows the importance of protocol-specific monitoring (Empl et al., 2024).

A new IDS for IoT networks, SEHIDS (Self Evolving Host-based Intrusion Detection System), was introduced in (Baz, 2022); each IoT device is associated with an ANN (Artificial Neural Network) architecture in a reduced version to the almost minimal configuration. The system includes a resource-friendly mechanism aiming to optimize the model once its performance degradation occurs; 2 types of IDS were developed: a signature-based IDS with a supervised approach used to classify different attacks, and an anomaly-based IDS with an unsupervised approach-Replicator Neuron Network used to differentiate between malicious and normal traffic. A well-optimized IDS using DL (Deep Learning) and ML was proposed in (Awad et al., 2024) to increase IoT security; it should detect and mitigate 14 different attack types from the following 5 threat groups: information gathering, malware, DDoS (Distributed Denial-of-Service), MitM (Man-in-

the-Middle), and injection attacks. An evolved IDS framework for IoT is proposed in (Ghadami, 2025), using: game-theory-based Generative Adversarial Networks for dataset balancing, hybrid Arithmetic Optimization Algorithm and a Sine Cosine Algorithm for feature selection and Parallel Convolutional Neural Network (PCNN) combined with Long Short-Term Memory (LSTM) layer for accurate attack detection. An analysis of DL methods used for IDS systems in IoT environments is made in (Albulayhi et al., 2021), also with their associated datasets, aiming to determine gaps and weaknesses and find a neutral reference architecture; for the DL approaches in anomaly-based IDS systems, a detection model was proposed. (Sanjalawe et al., 2026) made an analysis of how AI supports the development of intrusion detection solutions in the IIoT (Industrial IoT) context over the last decade; the research identified practical limitations of common solutions such as transparency, data quality and development costs, together with future directions, in the short, medium and long term that should be pursued to improve IDS in IIoT environments, for example edge-based detection and the use of XAI (eXplainable AI). In (Sadhwani et al., 2024), the authors classified attacks using Edge-IIoT set; for the SmartSentry model proposed in this research, several ML models were used: RF (Random Forest), DT (Decision Tree), SVM (Support Vector Machine), KNN (k-Nearest Neighbor) and ETC (Extra Tree Classifier), and also a single DL model.

IoT-based smart building systems have been widely analyzed from security and performance perspectives (Poyyamozhi et al., 2024). Multi-layer IoT security architectures typically combine encryption, access control, and perimeter IDS mechanisms (Mrabet et al., 2020).

Industrial security standards such as IEC 62443 recommend segmentation and zone-based protection, yet they do not fully address AI-driven prediction pipelines (Korodi et al., 2024). Similarly, Zero Trust models emphasize continuous verification (Rose et al., 2020) but lack integration with data-driven anomaly detection for ML systems.

In ML-centric systems, data poisoning and adversarial attacks remain major vulnerabilities (Mohanta et al., 2020). Federated learning has been proposed as a privacy-preserving alternative (Li et al., 2020), yet its integration with intrusion detection remains an open research challenge.

The proposed MLC-IDS model advances beyond these approaches through the following elements: a multi-layer design approach for intrusion detection, where IDS mechanisms are embedded at every architectural layer; a cross-layer correlation logic; the protection of the ML training and inference pipeline.

INTRUSION DETECTION METHODOLOGY

The multi-layer security architecture for intelligent energy prediction in non-residential buildings (Tudora et al., 2025), requires an adaptive IDS aligned with the heterogeneous and resource-constrained IoT environment. Unlike traditional enterprise IDS solutions, intelligent energy prediction systems introduce 3 specific constraints: high density of heterogeneous IoT devices, real-time data ingestion for ML-based prediction, and direct coupling between data integrity and prediction accuracy. Therefore, intrusion detection must extend beyond classical network anomaly detection to integrate protection of the AI-driven prediction pipeline.

Multi-Layer Correlated IDS Architecture

Conceptual Framework: Multi-Layer IDS for IoT-Based Intelligent Energy Prediction Systems. Design Principles for Multi-Layer Intrusion Detection in PRECONERG

The methodology is based on the previously-defined security architecture for intelligent energy prediction systems (Tudora et al. 2025). Its specification is aligned with the intelligent energy prediction system design. The present work makes a design specification for one of the basic functionalities of the defined architecture, intrusion detection. It introduces a distributed

and correlated detection model with 4 layers: 1) Edge/Device-Level (Lightweight) IDS); 2) Gateway-Level IDS (Protocol-Aware Detection); 3) Perimeter/Network IDS-IPS (Hybrid Detection + SIEM Correlation); 4) Application & AI-Level IDS (Data & Model Integrity Monitoring). The main goals are to reduce false positives and detect multi-stage attacks.

The proposed approach follows a defense-in-depth and cross-layer correlation strategy. Each layer performs several operations: local detection of layer-specific anomalies, the relevant metadata forwarding to the upper level, and supporting correlation within SIEM, enabling a FPR (False Positive Rate) reduction and improved detection of complex multi-stage attacks. The conceptual architecture of MLC-IDS is depicted in Figure 1, while the experimental workflow for the future MLC-IDS evaluation is represented in Figure 2.

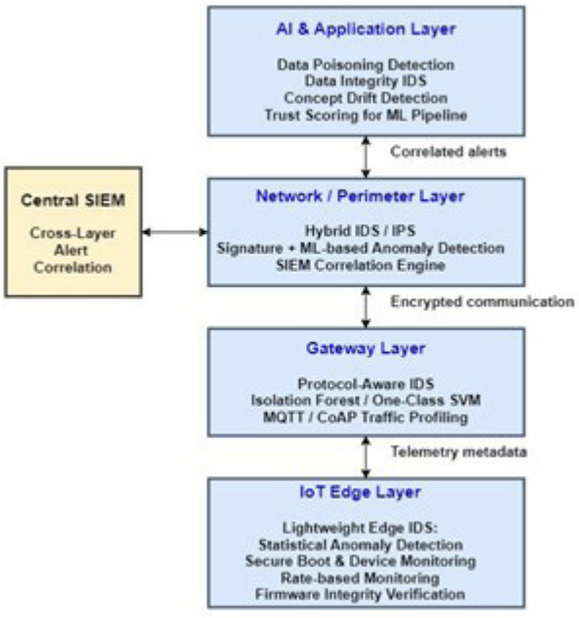

Figure 1. *Conceptual architecture of the Multi-Layer Correlated IDS for the PRECONERG platform*

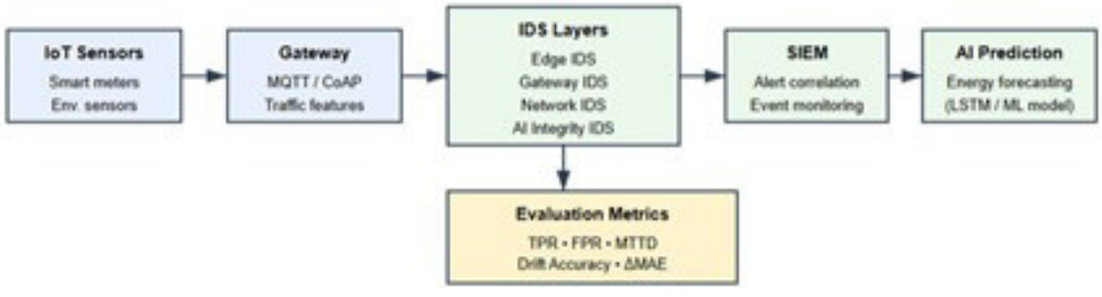

Figure 2. *Experimental workflow for future evaluation of the Multi-Layer Correlated IDS*

Threat Model Specific to IoT-Based Energy Prediction Systems. Attack Taxonomy

The proposed methodology is based on a threat model tailored to the operational environment of the predictive system. This environment integrates IoT sensing devices, network infrastructure, and ML–based energy prediction pipelines. The specific attacks are grouped according to the following taxonomy:

- *Attacks at the IoT layer.* Potential threats against IoT devices include:
  - firmware tampering, which may change the device behavior or introduce backdoors;
  - device spoofing, when a malicious node impersonates a legitimate one;
  - MQTT flooding, which overwhelms the message broker or communication channels;
  - replay attacks, with the retransmission of previously captured valid messages;
  - battery-drain attacks, trying to exhaust the limited energy resources of IoT devices.
- *Attacks at the Network layer.* At the network level, the types of attacks include:
  - MitM attacks, enabling data interception or modification;
  - DDoS attacks against the gateway, which can disrupt data collection and system availability;
  - lateral scanning across VLANs (Virtual Local Area Network), used to identify additional attack surfaces within the network.
- *Attacks against ML pipeline.* The ML pipeline used for energy prediction may be exposed to adversarial actions, such as:
  - data poisoning, where malicious samples are injected into the training dataset, altering the dataset used to design and develop ML-based predictive models;
  - outlier injection, to distort statistical patterns in the data;
  - artificially induced model drift, which gradually degrades the prediction performance;
  - model extraction, where an attacker attempts to reconstruct the trained model through repeated queries.

Layer 1 – Lightweight Edge-level IDS (Resource-Constrained IoT Adaptation)

The 1st layer of MLC-IDS should be deployed directly on IoT edge devices. These devices include smart meters, environmental sensors, and occupancy monitoring nodes that collect operational data to be used by the intelligent energy prediction pipeline. Such devices typically work under strict constraints (resources, energy), therefore complex ML algorithms are unsuitable for deployment at this layer. Consequently, a lightweight statistical anomaly detection model should be considered to identify abnormal device behaviors with minimal computational overhead.

a) Design Principles

Given that the IoT devices are usually resource-limited (CPU, storage), the edge-level IDS should be designed to meet the following constraints: not to run very complex ML models; to use static rule-based detection together with lightweight statistical anomaly detection; to minimize energy consumption and computational overhead.

b) Techniques to be considered

The following techniques should be considered (Al Maawi et al., 2025), (Dikii et al., 2021), (Katole et al, 2025): rate-based anomaly detection (dynamic thresholding); detection of abnormal publish frequency in MQTT traffic; firmware hash verification (secure boot monitoring); device identity consistency check (certificate mismatch detection).

c) Modeling elements: Statistical Anomaly Detection Model for Edge-level IDS

The edge-level detection mechanism relies on a normalized deviation model that evaluates whether the current observation produced by a device significantly deviates from its historical behavioral profile.

The statistical model used for the anomaly detection at the IoT device level is defined by the following Anomaly Scoring function, which must be estimated/computed at time t for each device d:

$$AS_d(t) = \frac{|x_d(t) - \mu_d|}{\sigma_d} \quad (1)$$

where $AS_d(t)$ is the anomaly score of device d at time t, $x_d(t)$ is the instantaneous (measured) value. The historical profile of the device working is given by the statistical amounts μ_d and σ_d. The anomaly score represents the standardized deviation between the current measurement and the historical baseline of the device. The

local alert should be issued if the (detection) thresholding condition $AS_d(t) > \theta_d$ is met for the specified threshold value. The significance of the amounts/parameters in equation (1) is as follows:

d: device identifier. This is the unique IoT device identifier generating the monitored data stream within the system (sensor, smart meter, occupancy sensor). Each device has its own statistical behavior profile. One can have, for example, a temperature sensor within the HVAC (Heating, Ventilation and Air-Conditioning) system, a smart meter for electricity consumption, indoor occupancy sensor.

t: time instance. This is the timestamp (time index) corresponding to the current observation (measurement) moment. It could be given in seconds, minutes, or as aggregation intervals (e.g., 5 minutes). In energy consumption prediction systems, it is common to use $t \in \{1,2,\ldots,T\}$, where T denotes the specified time period over which the total number of observations is provided.

x_d (t): the instantaneous (observed) value. It denotes the current observed value measured by device d at time t. This may represent: instantaneous energy consumption, temperature, number/frequency of MQTT messages, publishing frequency, and payload size. For example: smart meter -val kWh, temperature sensor -val °C, MQTT publisher –val messages/sec.

μ_d: historical mean. This is the historical mean value of the monitored parameter for device d, computed from historical observations during normal operation. It is given by

$$\mu_d = \frac{1}{N} \cdot \sum_{i=1}^{N} x_d(i) \quad (2)$$

where N is the number of measurements. This amount provides a statistical representation of the device's normal behavior. It enables the definition of the device's normal operating profile and is computed using accurate historical data.

σ_d: standard deviation. It is the standard deviation of the historical observations, representing the expected variability of the device behavior. It is given by

$$\sigma_d = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_d(i) - \mu_d)^2} \quad (3)$$

This indicator measures data variability and enables the differentiation between minor normal variations and anomalies. A low value indicates stable behavior, whereas a high value shows increased variability.

$AS_d(t)$: anomaly score (Anomaly Scoring function). It is the normalized deviation between the current observation and the historical baseline. It is estimated using equation (1). The following significance levels can be assumed or assigned:

$AS_d(t) \approx 0$: normal behaviour without any variations;

$AS_d(t) \approx 1$: minor variations not exceeding the normal behavior;

$AS_d(t) > 2$: potential anomaly (to be further investigated);

$AS_d(t) > 3$: major (severe) anomaly.

This scoring function is quite similar to a Z-score. Actually, this modeling corresponds to a Z-score–based statistical anomaly detection method, which allows the fast identification of abnormal cases without requiring complex training.

θ_d: detection threshold. The system continuously monitors the anomaly score for each device. A local alert should be triggered if the scoring function value (anomaly score) exceeds a predefined detection threshold: $AS_d(t) > \theta_d$. θ_d is a device-specific detection threshold. It could be determined during the system calibration phase based on the expected variability of the monitored parameter. Different types of sensors exhibit different levels of variability and are associated with distinct thresholds for anomaly detection. Environmental sensors typically exhibit small deviations around their mean values, while energy consumption measurements may present larger variations due to operational dynamics. The generated alert is transmitted to the next layer of the IDS architecture.

The interpretation of these modeling elements, within the context of an intelligent energy consumption prediction system for buildings, addresses the detectable cases as follows:

- MQTT flooding, as concerning the message frequency (rate), with $x_d(t) >> \mu_d$, that leads to a high Anomaly Score;
- Device spoofing, in which the message arrivals follow a different or abnormal pattern, also with unusual frequency. This increases $AS_d(t)$;
- Device malfunction, for example a case in which the measured value provided by a temperature sensor is significantly out of the established range. This also increases $AS_d(t)$.

The modeling is suitable for IoT edge detection (Layer 1-Edge IDS) due to its specific advantages: low complexity, low storage, fast computation and reduced energy consumption for the devices. Actually, this statistical model enables the early detection of abnormal behavioral cases associated with several classes of IoT-level attacks. A sudden increase in message transmission frequency may indicate a flooding attack, while unusual sensor readings may indicate firmware manipulation or device malfunction. Since only the anomaly score and minimal metadata are transmitted to higher layers, the communication overhead remains low, ensuring the energy efficiency of the IoT infrastructure. In addition to computational efficiency, the model supports adaptive threshold adjustment based on contextual information such as daily operational patterns or seasonal energy consumption variations. This is useful to reduce false positives while preserving the target sensitivity to abnormal cases.

The proposed edge-level anomaly detection model provides a computationally efficient mechanism to identify suspicious device behavior at the earliest stage of data processing in the system architecture. By detecting anomalies directly at the data source, the model aims to reduce attack propagation across the network and supports the defense-in-depth strategy to be applied for the PRECONERG multi-layer security architecture.

d) Output

The device does not transmit all raw data. Instead, the transmitted content includes the anomaly flag, score, timestamp, and device identifier: $Alert_d$=(device_id,AS_d (t),timestamp). The advantages are: reduced traffic, data protection, scalability. This approach follows the telemetry minimization principle specified for Layer 1 in the previous work (Tudora et al., 2025).

Optional extensions: The proposed model for Layer 1-IDS could be further extended with adaptive thresholding or multi-feature detection, respectively. These will be approached in future works.

Layer 2 – Gateway-Level Protocol-Aware IDS

The 2nd layer of MLC-IDS should be deployed at the gateway level. Unlike the edge devices, the gateways have higher computational capacity and can therefore execute ML–based anomaly detection algorithms. This layer performs protocol-aware behavioral monitoring of IoT communication traffic to detect abnormal patterns associated with network-level attacks.

More complex ML techniques should be considered for this layer functionality.

a) Anomaly Detection in IoT Traffic: basic requirement

To achieve IoT traffic anomaly detection as a Layer 2 IDS functionality, it will be necessary to build a normal behavior model for: message frequency, topic distribution (MQTT), payload size distribution, and TLS (Transport Layer Security) session patterns.

b) Recommended algorithms (lightweight ML)

The algorithms recommended to implement the functionalities of this layer are: Isolation Forest (for low computational cost and suitable for real-time anomaly detection), One-Class SVM, DBSCAN (Density-Based Spatial Clustering of Applications with Noise), and RF.

c) Features. Minimal Feature Set

The gateway aggregates traffic flows originating from multiple IoT devices and extracts a set of behavioral features describing the communication profile of each device. These features describe the normal operational patterns of IoT communication protocols such as MQTT or CoAP (Constrained Application Protocol). The feature vector associated with each observation interval is defined as: $F=\{f_1,f_2,f_3,f_4,f_5\}$ in which: f_1 denotes

message transmission frequency, f_2 is inter-arrival time between messages, f_3 is payload size distribution, f_4 represents topic entropy in MQTT communication and f_5 denotes TLS session duration. These amounts provide a compact behavioral representation for IoT traffic patterns.

The minimal feature set could be *F={freq,interarrival_time,payload_size,topic_entropy}*. The last amount evaluates the unpredictability, diversity, or information content within MQTT topic structures and the messages published to them. Topics are strings allowing to filter messages for a connected client in MQTT (HiveMQ, 2026).

d) Anomaly Detection in the extracted feature space: Behavioral Detection at Gateway-level

For the anomaly detection in the extracted feature space, the gateway-level IDS could use unsupervised ML techniques. One of the primary algorithms to be used in this layer is the Isolation Forest, which identifies anomalies by isolating observations through recursive partitioning of the feature space. The anomaly score produced by the Isolation Forest model is defined as (Matlab Help Center, n.d.):

$$AS(x,n) = 2^{-\frac{E(H(x))}{c(n)}} \quad (4)$$

in which: AS(x,n) is the anomaly score of observation x with n samples, E(H(x)) is the average path length over all isolation trees in the isolation forest (number of splits required to isolate the point across the trees) and c(n) denotes a normalization factor depending on dataset size, representing the average path length of unsuccessful searches in a binary search tree of n observations.

Observations with short path lengths are more easily isolated and therefore more likely to represent anomalies. The abnormal cases correspond to datapoints that are easy to isolate, resulting in a short path length. Their corresponding scores are close to 1 (Matlab Help Center, n.d.).

In addition to Isolation Forest, the gateway layer may use One-Class SVM (OC-SVM) to learn the boundary of normal network behavior. OC-SVM model builds a decision function that separates normal observations from anomalous ones in a high-dimensional feature space. The decision function is given by (Hejazi & Singh, 2013):

$$f(x) = w \cdot \phi(x) - \rho \quad (5)$$

where: x is the input feature vector, $\phi(x)$ denotes the nonlinear mapping (kernel) to feature space, w is the model weight vector in the feature space and ρ is a decision boundary parameter (offset or bias term of hyperplane). An observation is considered anomalous when (Hejazi & Singh, 2013) f(x)<0.

e) Specific attacks Detection

The gateway-level IDS is capable to detect several classes of specific attacks targeting IoT communication infrastructures, including: MQTT flooding attacks recognized through sudden increases in message frequency, replay attacks identified through repeated payload signatures, device spoofing attempts detected through inconsistencies between device certificates and network identifiers, abnormal inter-device communication patterns indicating lateral movement- class of techniques that the intruders can use to move from a compromised IoT device to other sensitive resources (Baker, 2025).

By combining protocol-aware feature extraction with lightweight unsupervised learning models, this layer provides accurate detection of communication anomalies while maintaining moderate computational overhead. Detected alerts are forwarded to the SIEM correlation engine located in the upper IDS layer.

Layer 3 – Hybrid Network IDS/IPS with SIEM Correlation

This layer was described in a generic manner in the previous paper (Tudora et al., 2025). In the present work, the correlation methodology is formalized.

a) Hybrid Detection Model

The hybrid detection model (scoring) is formally defined according to

$$SD_{hybrid} = \alpha \cdot SD_{sign} + \beta \cdot SD_{ML} \quad (6)$$

where: SD_{sign} is the score for signature-based detection, SD_{ML} is the score for anomaly-based

detection, SD_{hybrid} denotes the overall score for hybrid detection, α, β are adaptive weights.

b) Multi-layer Correlation for IoT Traffic Anomaly Detection

An attack is confirmed if the overall alert score exceeds a given threshold: $Alert_{total}>\delta$, where

$$Alert_{total} = Alert_{edge} + Alert_{gateway} + Alert_{network} \quad (7)$$

This correlation reduces FPR and enables the detection of slow, distributed attacks.

c) Adaptive Thresholding

The thresholds should be adjusted based on time of day (day/night), season and/or building type (hospital vs. campus).

Layer 4 – AI & Data Integrity IDS and Anti-Data Poisoning

This is the new contribution compared to the previous paper (Tudora et al., 2025). The 4th layer of MLC-IDS architecture addresses a critical security challenge specific to AI-driven systems: the protection of the ML pipeline against adversarial manipulation. In intelligent energy prediction systems, the integrity of the training data directly influences the accuracy of the predictive models. Consequently, malicious data injection or distribution manipulation may significantly degrade the prediction accuracy. To mitigate these risks, the AI-level IDS integrates data integrity monitoring mechanisms able to detect data poisoning attacks and concept drift manipulation.

a) Data Poisoning Detection

The design and development process should consider: outlier detection on the training batch, concept drift detection (ADWIN - ADaptive WINdowing, Page-Hinkley), and validation against a trusted subset of sensors. This will be approached in future work.

b) Data Trust Scoring

This component of the model evaluates the trustworthiness of each incoming data stream using a data trust scoring function defined as:

$$TS_i = \sum_{j=1}^{3} w_j \cdot C_i \quad (8)$$

TS_i is the trust score associated with data stream i. Each flow (i) will be assigned a trust score TS measuring the conformance with the following criteria: C1 Integrity (I), C2 Behavioral Consistency (BC), C3 Historical Reliability (HR). Each criterion is weighted according to its importance:

$$TS_i = w_1 \cdot I_i + w_2 \cdot BC_i + w_3 \cdot HR_i \quad (9)$$

where: I_i is a data integrity verification score, BC_i is a behavioral consistency indicator relative to expected patterns, HR_i denotes historical reliability of the data source, w_1, w_2, w_3 are weighting coefficients such that $\Sigma^3_{j=1} w_j=1$. Data streams whose trust score falls below a predefined threshold (therefore having a low overall trust score) should be flagged as suspicious and excluded from the training dataset. Such data should be isolated or not included in retraining.

c) Model Drift Monitoring

In addition to trust scoring, the system monitors the performance stability of the prediction model in order to detect adversarial drift conditions. Concept drift is evaluated by measuring deviations between the current prediction error and the baseline error recorded during normal operation. The drift indicator is defined as:

$$Drift = |MAE_{crt} - MAE_{base}| \quad (10)$$

where: MAE_{crt} is the mean absolute error of the prediction model in the current observation window, MAE_{base} is the reference error obtained during validated training conditions. When the drift indicator exceeds a predefined threshold, the system triggers an integrity alert for a possible adversarial attack and initiates additional validation procedures. These may include: outlier detection on the current training batch, comparison with trusted sensor subsets, and temporary suspension of model retraining.

This multi-layer integrity monitoring mechanism ensures that malicious data cannot silently influence the training process. By integrating AI-level intrusion detection within

the IDS architecture, the system protects not only communication infrastructures but also the reliability of the predictive models. Consequently, the AI-level IDS becomes a critical extension of traditional IDS frameworks by incorporating ML integrity protection as an important component of the cybersecurity architecture.

Cross-Layer Correlation & Federated IDS

An important limitation of traditional IDS deployed in IoT environments is the high FPR generated by independent detection mechanisms operating at different architectural layers. In intelligent energy prediction systems, the attacks may evolve gradually across multiple layers of the infrastructure, including IoT devices, network communication channels, and ML pipelines. To address this challenge, the proposed MLC-IDS should integrate a cross-layer alert correlation mechanism that aggregates detection signals generated by all IDS layers. The correlation engine should be implemented within the centralized monitoring infrastructure (SIEM component) and will operate by combining alert scores originating from the edge, gateway, network, and AI layers.

The given set of IDS layers for the specified model is $L=\{L_1,L_2,L_3,L_4\}$, where: L_1 is Edge IDS, L_2 is Gateway IDS, L_3 is Network IDS/IPS and L_4 is AI/Data Integrity IDS. One can assume that each layer generates a normalized alert score $A_i(t)$ that represents the detection confidence level associated with a potential security event at time t. The global alert score follows a weighted aggregation model:

$$A_{global}(t) = \sum_{i=1}^{4} w_i \cdot A_i(t) \tag{11}$$

where: $A_{global}(t)$ is the aggregated alert score, $A_i(t)$ is the alert score produced by layer L_i with the corresponding weight wi. These weights should be normalized according to $\Sigma^4_{(i=1)} w_i=1$. This weighting scheme is quite similar to many data fusion schemes for multimodal biometric systems (Soviany et al., 2026). It reveals the reliability and detection capability of each layer. AI-level integrity monitoring may get a higher weight when protecting the prediction pipeline, while edge-level alerts may have lower confidence due to limited computational power and other resource-related constraints. An intrusive case is confirmed when the aggregated alert score exceeds a predefined correlation threshold: $A_{global}(t)>\delta$, where δ is the correlation threshold defined by the security policy.

This aggregation mechanism allows to detect multi-stage attacks, where individual detection signals may remain below their respective thresholds but all together indicate intrusive activity. A distributed attack scenario may involve the following detection cases:

- *at the edge layer* (L_1): unusual sensor behavior;
- *at the gateway layer* (L_2): increased communication frequency;
- *at the network IDS layer* (L_3): suspicious traffic patterns;
- *at the AI layer* (L_4): abnormal prediction error.

Individually, these anomalies may seem to be non-malicious; however, their correlation significantly increases the probability of an ongoing attack.

The detection reliability could be further improved by using a temporal correlation window. In this approach (to be explored in a future work), the aggregation of the alerts should be done within a time slot τ such as: $\tau=[t-\Delta t,t]$ in which Δt represents the correlation time window. The correlated alert score can therefore be defined as:

$$A_{correl}(t) = \sum_{i=1}^{4} w_i \cdot \left(\frac{1}{\Delta t} \int_{t-\Delta t}^{t} A_i(\tau)\, d\tau \right) \tag{12}$$

This temporal aggregation should be able to detect slow multi-stage attacks evolving over extended periods. The cross-layer correlation model could ensure: reduction of false positive alerts, improved detection of distributed or stealthy attacks, unified security monitoring across the entire IoT–AI infrastructure.

The design options for scalability are: the gateway-level models should be locally trained, only the parameters should be shared, raw

data remains local, as in Federated Learning (Olanrewaju-George& Pranggono, 2025). This approach supports GDPR protection and ensures a reduced centralized attack surface.

The integration of alert signals from all layers allows MLC-IDS architecture to ensure a comprehensive defense-in-depth strategy aligned with the PRECONERG security-by-design requirement.

Evaluation Metrics: KPI for IDS Evaluation

The following metrics will be used for evaluation during the development process, with respect to the fixed targets-KPI (Key Performance Indicators): detection performance metrics- TPR (True Positive Rate), FPR, AUC (Area-Under-the Receiver Operating Characteristic Curve), Precision, Accuracy, F_1-score; *operational metrics* -MTTD (Mean Time To Detect the incident), MTTR (Mean Time To Response), Latency overhead (ms), CPU overhead (%), Energy overhead (mWh/day per device); *AI Integrity metrics* - Data poisoning detection rate, False rejection of legitimate data, Model drift alert accuracy. The evaluations will be presented in future case studies during the PRECONERG platform development.

TECHNICAL FRAMEWORK AND EXPERIMENTAL VALIDATION OF THE HYBRID IDS ARCHITECTURE

The initial experimentation is done for the ML component of the hybrid IDS architecture.

Dataset Description

The experimental validation of the proposed hybrid IDS was conducted using the UNSW-NB15 benchmark dataset (Moustafa and Slay, 2015). The dataset was generated in the Cyber Range Lab at UNSW Canberra using the IXIA PerfectStorm traffic generator, combining legitimate modern network traffic with synthetically generated contemporary cyber-attacks.

The dataset contains 49 extracted flow-based features derived from approximately 100 GB of raw packet capture data processed using Argus and Bro-IDS tools. It includes 9 attack families - Fuzzers, Analysis, Backdoors, DoS, Exploits, Generic, Reconnaissance, Shellcode, and Worms - alongside normal traffic. For the present study, predefined training (175,341 records) and testing (82,332 records) partitions were employed, ensuring strict separation between learning and evaluation stages. The class distribution reflects realistic network imbalance conditions, with attack traffic representing approximately 2/3 of the dataset.

The UNSW-NB15 dataset was selected due to its realistic traffic composition and diversity of modern attack scenarios, which make it more representative of contemporary network environments compared to legacy intrusion detection datasets. Recent studies confirm that UNSW-NB15 remains one of the most widely used benchmarks to evaluate ML-based intrusion detection models in modern networks and IoT environments (Ring et al., 2019).

Proposed Multi-Layer IDS Architecture (basic features)

The proposed IDS will be embedded within an IoT-based intelligent energy prediction infrastructure and follows a multi-layer defensive model designed to provide resilience across system boundaries. Such layered cybersecurity architectures are widely adopted in IoT environments in order to distribute detection capabilities across devices, gateways, and cloud platforms, improving system resilience against distributed attacks and against a large variety of intrusive attacks in IoT networks (Ashraf et al., 2022).

At the device layer, lightweight behavioral validation mechanisms monitor telemetry frequency and payload consistency. The gateway layer performs protocol inspection, rate-limiting enforcement, and flow aggregation analysis. The perimeter layer hosts the hybrid detection engine combining supervised and anomaly-based models. At the application and AI level, the integrity validation monitors deviations from predictions and potential attempts to contaminate data that affect the

energy forecasting engine. This approach ensures that the detection is not isolated but integrated directly into the AI data flow, protecting both cybersecurity and forecast reliability.

Machine Learning Models

To validate the hybrid detection strategy, 2 complementary models were considered.

Supervised Detection – Random Forest

The supervised component was implemented using a balanced RF classifier. Class weighting was applied to address dataset imbalance and reduce bias toward the dominant attack class. RF models continue to be widely adopted for intrusion detection due to their ability to handle high-dimensional traffic features and maintain strong detection accuracy in complex network environments (Sahani et al., 2023).

The model achieved: Accuracy: 0.68; Precision (attack): 0.73; Recall (attack): 0.67; F1-score (attack): 0.70; ROC-AUC: 0.81. These results suggest that there is enough space for further performance improvement through various optimization techniques in order to achieve stronger discriminative capability between malicious and non-malicious traffic flows.

Anomaly-Based Detection – Isolation Forest

To enhance robustness against zero-day threats, an anomaly-based detection layer is considered using Isolation Forest. This model isolates statistical outliers without relying on labelled attack data. The anomaly-based model achieved a ROC-AUC of 0.79, demonstrating meaningful separation between normal and anomalous traffic patterns.

The similar performance of both models supports the architectural choice of a hybrid detection mechanism.

Experimental Evaluation

ROC Analysis

The ROC (Receiver Operating Characteristic) comparison between the supervised RF and the anomaly-based Isolation Forest is presented in Figure 3. This figure illustrates ROC curves for both detection models. RF classifier achieves a ROC-AUC of 0.81, outperforming the Isolation Forest model (ROC-AUC = 0.79) across most false positive rate thresholds. The supervised model demonstrates higher discriminative power, particularly in low false-positive regions critical for operational environments. However, the anomaly-based model maintains competitive performance, supporting the architectural choice of a hybrid detection framework capable of addressing both known and previously unseen attack patterns.

The RF model generally performed better than the anomaly-based model across most false positive rate thresholds. Its advantage was especially noticeable in areas with a low false positive rate, which are particularly important for practical, real-world use. The ROC-AUC score of 0.81 indicates that the model can effectively distinguish between different traffic classes.

The Isolation Forest model achieved a ROC-AUC score of 0.79, suggesting that it is capable of identifying unusual patterns in the data, including behaviors that do not necessarily match previously known attack signatures.

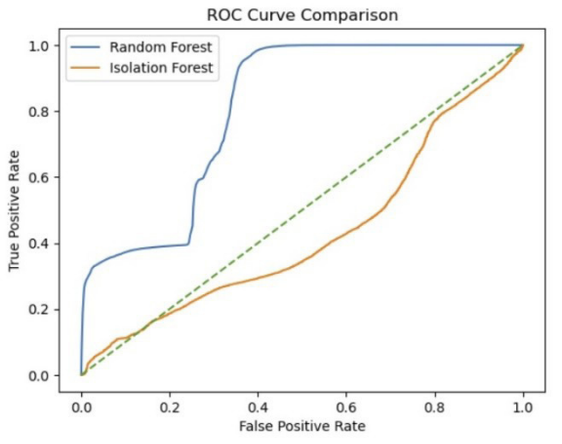

Figure 3. *ROC curve comparison between the supervised RF classifier and the anomaly-based Isolation Forest model using the UNSW-NB15 dataset*

Confusion Matrix Analysis

The detailed classification breakdown of the supervised model is illustrated in Figure 4. This shows the classification results of the supervised detection layer. The matrix highlights how predictions are distributed between true positives (attacks correctly detected), true negatives (normal traffic correctly identified), false positives, and false negatives.

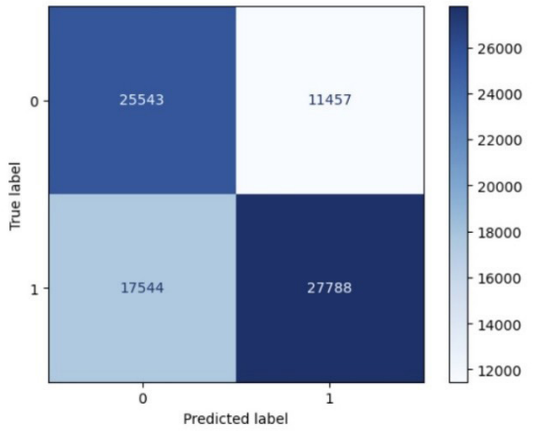

Figure 4. *Confusion Matrix of the Balanced RF classifier evaluated on the UNSW-NB15 testing dataset*

The results suggest a fairly balanced detection performance, which matches the reported Precision (0.73) and Recall (0.67) values. The moderate number of false negatives indicates that the model is able to detect most attacks effectively, while the acceptable level of false positives shows a reasonable trade-off. This balance is typical in security-focused IoT environments, where it is usually more important to detect potential threats even if it means generating a slightly higher number of alerts.

The confusion matrix highlights the distribution of: True Positives (correctly detected attacks), True Negatives (correctly identified normal traffic), False Positives (non-malicious traffic flagged as an attack), and False Negatives (missed attacks).

The results show that the balanced training approach helped to reduce bias in the classification process. The moderate FPR is still considered acceptable in critical infrastructure environments, where the detection of potential attacks is more important than avoiding occasional alerts for normal activity. From an IoT security standpoint, reducing false negatives is especially important. If malicious traffic is not detected, it could spread across connected energy devices and interfere with predictive optimization processes.

Feature Importance and Interpretability

To enhance transparency and architectural alignment, feature importance analysis was conducted using the RF model, an approach commonly employed to interpret ML decisions in cybersecurity analytics. The ranking of features is presented in Figure 5. This ensures a certain amount of explainability, as a step towards using XAI for the intrusion detection in IoT environments. The figure presents the most important traffic attributes with the strongest influence on performance, contributing to supervised attack detection. The feature importance analysis identified several network traffic attributes that significantly contribute to the detection performance of the supervised model. The most important features include *sttl* (Source Time-To-Live), *ct_state_ttl*, *dload* (destination load), *rate* (flow rate), and *sload* (source load). These variables capture both

packet propagation characteristics and traffic intensity patterns, which represent critical indicators of abnormal network behavior.

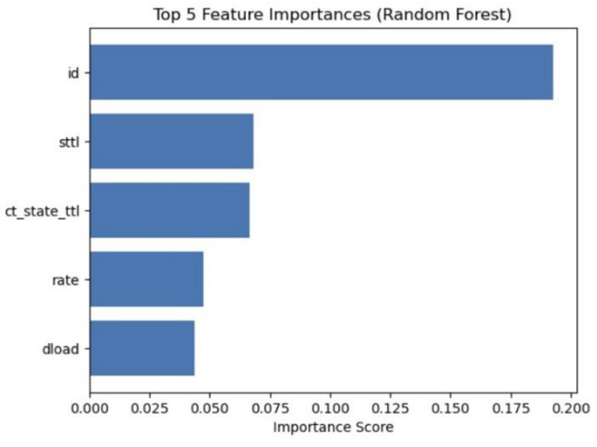

Figure 5. *Top-ranked feature importance scores derived from RF classifier*

The dominance of TTL (Time-To-Live)-related metrics, such as *sttl* and *ct_state_ttl,* together with traffic-load indicators including *dload, sload* and *rate*, confirms that both temporal and volumetric characteristics play a crucial role in distinguishing benign from malicious network behavior. TTL-based attributes provide information about packet propagation and routing paths within the network, making them particularly useful for identifying irregular traffic patterns. Abnormal TTL values may indicate spoofing attempts, packet manipulation, or routing anomalies, which are commonly associated with various intrusion scenarios. Similarly, elevated traffic load and irregular throughput patterns are often linked to scanning activities, volumetric attacks (e.g., DoS), and exploit-driven bursts of network traffic (Paxson, 1999).

These findings are consistent with observations reported in network traffic analysis and intrusion detection research, where TTL-related indicators contribute significantly to identifying suspicious communication patterns in modern network environments (Paxson, 1999). Furthermore, the results validate the behavioral monitoring emphasis at the gateway and perimeter layers of the proposed IDS architecture. The interpretability of feature rankings strengthens the transparency of the model and supports its integration into critical IoT energy infrastructures.

These findings empirically validate the behavioral emphasis of the proposed IDS architecture. The combined interpretation of Figures 3, 4, and 5 demonstrates strong alignment between model-level performance, feature-level interpretability, and architectural design objectives.

Security Performance Composite Index (SPCI)

To assess the overall effectiveness of the proposed security framework beyond isolated ML evaluation metrics, the obtained results should be integrated into a composite indicator named SPCI (Security Performance Composite Index). This index will provide a system-level perspective by aggregating detection accuracy with operational and architectural considerations. It will be fully defined and explained in a future work.

DP (Detection Performance) component is calculated as the arithmetic mean of the ROC-AUC and F1-score metrics, capturing both discrimination capability and classification balance:

DP = (ROC-AUC + F1-score)/2 resulting in DP = (0.81 + 0.70)/2 = 0.755. When combined with indicators reflecting operational efficiency and AI integrity stability, the target for the resulting composite score is to exceed 0.80, suggesting a high degree of architectural robustness and a well-balanced trade-off between cybersecurity effectiveness and system performance. SPCI will be properly defined and evaluated using extended datasets from the real application/ platform process development.

Consequently, SPCI is expected to serve as a holistic evaluation metric that should align detection accuracy, operational feasibility, and AI reliability, offering a comprehensive perspective on security performance within IoT-enabled energy infrastructure environments.

Architectural Implications for IoT Energy Systems

The experimental findings have direct implications for IoT-based intelligent energy infrastructures.

Firstly, the importance of TTL and flow-rate features confirms that gateway-level traffic aggregation and behavioral monitoring are essential for detection. Secondly, the validation of the hybrid model shows that combining supervised classification with anomaly detection improves the system's ability to handle both known attacks and new, previously unseen threats. Finally, the integration of intrusion detection directly into the AI data stream helps to protect predictive models against risks such as data poisoning, traffic manipulation, or adversarial interference, which could otherwise affect energy demand forecasts or optimization strategies. Taken together, the alignment between dataset analysis, model performance results, interpretability assessment, and system architecture suggests that the proposed IDS will work as an integral layer of cybersecurity within the smart energy IoT ecosystem, rather than simply acting as an external monitoring tool.

CONCLUSIONS AND FUTURE DIRECTIONS

This paper introduces a Multi-Layer Correlated IDS tailored for intelligent energy prediction systems in non-residential buildings. It is specifically designed for the PRECONERG multi-layer architecture. The proposed model integrates 4 functional layers: (1) lightweight statistical anomaly detection at IoT edge devices, (2) protocol-aware gateway detection using Isolation Forest and One-Class SVM models for MQTT/CoAP traffic profiling, (3) hybrid signature- and ML-based network IDS/IPS with SIEM-driven cross-layer alert correlation, and (4) AI-level data integrity monitoring incorporating trust scoring and concept drift detection to mitigate data poisoning attacks. A cross-layer alert fusion function is defined to reduce false positives while preserving detection sensitivity under IoT resource constraints. By integrating lightweight edge detection, protocol-aware gateway monitoring, hybrid network IDS, and AI integrity protection, the model ensures end-to-end resilience. The proposed framework enhances both cybersecurity level and ML prediction reliability for use-cases belonging to energy consumption in smart buildings. This supports the advance beyond traditional building automation security models. The future research directions include large-scale pilot validation, AI-driven adaptive thresholding, federated anomaly detection, Zero Trust dynamic identity integration. The more specific works will include several case studies with synthetic and operational data that will be collected during the platform development process.

ACKNOWLEDGEMENTS

This work was supported by 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

Al Maawi, K.N. & Abdullah, M. (2025). A review on intrusion detection systems for MQTT in IoT environments. *International Journal of Safety and Security Engineering.* 15(8), 1733-1744. https://doi.org/10.18280/ijsse.150818 [Accessed 16th February 2026].

Albulayhi, K., Smadi, A.A., Sheldon, F.T. & Abercrombie, R.K. (2021). IoT Intrusion Detection Taxonomy, Reference Architecture, and Analyses. *Sensors.* 21, 6432. https://doi.org/10.3390/s21196432

Ashraf, E., Areed, N., Salem, H., Abdelhady, E., & Farouk, A. (2022). IoT based intrusion detection systems from the perspective of machine and deep learning: a survey and comparative study. *Delta University Scientific Journal.* 5(2), 367-386, https://www.researchgate.net/publication/366757863_IoT_Based_Intrusion_Detection_Systems_from_The_Perspective_of_Machine_and_Deep_Learning_A_Survey_and_Comparative_Study/references [Accessed 6th February 2026].

Awad, O., Hazim, L., Jasim, A. & Ata, O. (2024). Enhancing IIoT security with Machine Learning and Deep Learning for intrusion detection, *Malaysian Journal of Computer Science.* 37(2), 139-153, https://doi:10.22452/mjcs.vol37no2.3. [Accessed 5th February 2026].

Baker, K. (2025) *Lateral Movement Explained, Cybersecurity 101: The Fundamentals of Cybersecurity.* https://www.crowdstrike.com/en-us/cybersecurity-101/cyberattacks/lateral-movement/ [Accessed 17th February 2026].

Baz, M. (2022) SEHIDS: Self Evolving Host-Based Intrusion Detection System for IoT Networks. *Sensors.* 22, 6505. https://doi.org/10.3390/s22176505. [Accessed 5th February 2026].

Dikii, D., Arustamov, S. & Grishentsev, A. (2021) DoS attacks detection in MQTT networks. *Indonesian Journal of Electrical Engineering and Computer Science.* 21(1), 601-608. http://doi.org/10.11591/ijeecs.v21.i1.pp601-608. [Accessed 16th February 2026].

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. https://doi.org/10.1155/2020/8860598. [Accessed 16th February 2026].

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.* https://doi.org/10.1145/3626232.3653271. [Accessed 16th February 2026].

Ghadami, R. (2025) An intrusion detection system in the Internet of Things with deep learning and an improved arithmetic optimization algorithm (AOA) and sine cosine algorithm (SCA). *Scientific Reports.* 15, 38156, https://doi.org/10.1038/s41598-025-22074-3. [Accessed 4th February 2026].

Hejazi, M., & Singh, Y. P. (2013) One-Class Support Vector Machines Approach to Anomaly Detection. *Applied Artificial Intelligence*. 27(5), 351–366. https://doi.org/10.1080/08839514.2013.785791.

HiveMQ Team (2026) *MQTT Topics, Wildcards, & Best Practices – MQTT Essentials: Part 5*, https://www.hivemq.com/blog/mqtt-essentials-part-5-mqtt-topics-best-practices/#:~:text=Here%20are%2 0some%20examples%20of,located%20on%20the%20ground%20floor. [Accessed 18th February 2026].

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

Katole, H. & Pattanshetti, T.R. (2025) Study on Supervised Anomaly Detection Model for MQTT-Based IoT data for DoS attacks. *International Journal for Multidisciplinary Research.* 7(2). https://doi.org/10.36948/ijfmr.2025.v07i02.39144. [Accessed 4th February 2026].

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. https://doi.org/10.3390/electronics13163229. [Accessed 4th February 2026].

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. https://doi.org/10.1109/MSP.2020.2975749.

Matlab Help Center (n.d.) Anomaly Detection with Isolation Forest https://www.mathworks.com/help/stats/anomaly-detection-with-isolation-forest.html [Accessed 30th January 2026].

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. https://doi.org/10.1016/j.iot.2020.100227. [Accessed 16th February 2026].

Moustafa, N. & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). *Military Communications and Information Systems Conference (MilCIS).* 1–6. https://doi.org/10.1109/MilCIS.2015.7348942 [Accessed 16th February 2026].

Olanrewaju-George, B. & Pranggono, B. (2025) Federated learning-based intrusion detection system for the internet of things using unsupervised and supervised deep learning models, *Cyber Security and Applications.* 3, 100068. https://doi.org/10.1016/j.csa.2024.100068. [Accessed 4th February 2026].

Paxson, V. (1999) Bro: a system for detecting network intruders in real-time. *Computer Networks.* 31(23–24), 2435-2463. https://doi.org/10.1016/S1389-1286(99)00112-7. [Accessed 4 February 2026].

Ring, M., Wunderlich, S., Scheuring, D., Landes, D. & Hotho, A. (2019) A survey of network-based intrusion detection datasets. *Computers & Security.* 86, 147–167. https://doi.org/10.1016/j.cose.2019.06.005. [Accessed 18 February 2026].

Rose, S., Borchert, O., Mitchell, S. & Connelly, S. (2020) Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207. [Accessed 18th February 2026].

Sadhwani, S., Modi, U.K., Muthalagu, R., & Pawar, P.M. (2024) SmartSentry: Cyber Threat Intelligence in Industrial IoT. *IEEE Access.* 12, 34720-34740. https://doi.org/10.1109/ACCESS. 2024.3371996. [Accessed 18th February 2026].

Sahani, N., Zhu, R., Cho, J.-H. & Liu, C.-C. (2023) Machine Learning-based Intrusion Detection for Smart Grid Computing: A Survey. ACM Transactions on Cyber-Physical Systems, 7(2), 1-31. https://doi.org/10.1145/3578366. [Accessed 16th February 2026].

Sanjalawe, Y., Fraihat, S., Al-E'mari, S. & Makhadmed, S. N., (2026) A review of artificial intelligence-based intrusion detection in industrial internet of things. *Discover Internet of Things.* https://doi.org/10.1007/s43926-026-00285-y [Accessed 3th March 2026].

Soviany, S., Gheorghe, C.-G., Gheorghe-Moisii, M. (2026) Multimodal biometric authentication of identical twin users in e-learning platforms: Security architecture. *In: International Conference on Virtual Learning*, ISSN 2971-9291, ISSN-L 1844-8933, vol. 21, pp. 445-460, 2026. https://doi.org/10.58503/icvl-v21y202638

Tudora, E., Soviany, S., Bica, O. (2025) Security Architecture for Data Protection in Intelligent Energy Prediction Systems for Non- Residential Buildings. *Romanian Cyber Security Journal.* 7(2), 13-27. https://doi.org/10.54851/v7i2y202502.