A Secure Cloud Architecture for Resilient Electricity Trading Platforms in Smart Grid Environments

Mo Pingyan*, Lu Yanqian, Wen You, Li Kai and Zheng Ying

Guangdong Power Grid Co Ltd., Guangzhou 510000, Guangdong, China
E-mail: mop15732@163.com
**Corresponding Author*

Received 27 October 2025; Accepted 12 December 2025

Abstract

To address gaps in security, resilience, and scalability in smart grid electricity trading, this study proposes Cloud-RESilient, a secure cloud-based architecture for local energy market and peer-to-peer trading platforms. The design features a hybrid edge-central cloud topology, where edge nodes support low-latency local trading and the central cloud enables scalable cross-cluster coordination. A multi-layer security framework integrates end-to-end encryption, AI-driven intrusion detection, and homomorphic encryption to protect data and enable privacy-preserving computation. A data-driven renewable energy source adaptation module based on the LSTM-ARIMA model is incorporated to manage generation variability. Validation is conducted in a simulated urban smart grid environment comprising five microgrid clusters and 10,000 prosumers, using real-world datasets including two years of renewable generation and over 10,000 cyber threat patterns. Evaluation covers three areas: security through penetration testing and privacy audits, resilience through fault tolerance and response to renewable fluctuations, and performance through scalability and latency testing. Results show that

Journal of Cyber Security and Mobility, Vol. 15_1, 25–66.
doi: 10.13052/jcsm2245-1439.1512

Cloud-RESilient achieves zero data leakage, 99.98% platform uptime, 92% accuracy in one-hour renewable forecasts, and 120 ms order processing time under full load. These outcomes confirm the effectiveness of the proposed methodology in delivering secure, resilient, and scalable electricity trading solutions for smart grids.

Keywords: Smart grid, electricity trading, cloud architecture, security, resilience, methodology, peer-to-peer trading, renewable energy.

1 Introduction

The global shift toward smart and sustainable urban ecosystems has established smart grids (SGs) and microgrids (MGs) as critical components of modern energy infrastructure [1, 2]. These systems are tasked with integrating intermittent renewable energy sources (RES), such as photovoltaic panels and wind turbines, while improving energy efficiency and ensuring reliable power supply. A primary challenge in this context is the inherent variability of RES generation, which complicates the maintenance of continuous electricity delivery, particularly during islanded operation, when MGs function independently of the main grid [3]. This challenge affects not only grid operators but also prosumers, defined as individuals or communities that both generate and consume electricity. To mitigate the impact of intermittency, local energy markets (LEMs) based on peer-to-peer (P2P) trading have gained prominence [4]. Such markets facilitate the exchange of surplus energy among neighboring MGs, prosumers, and other participants, thereby enhancing supply stability and supporting broader urban sustainability goals. The benefits of P2P energy trading extend across economic, social, and environmental dimensions, including reduced energy costs, greater energy autonomy for prosumers, and lower greenhouse gas emissions, objectives that align closely with the vision of smart city development.

Current electricity trading platforms that support LEMs and P2P transactions are generally organized according to three design paradigms: centralized, decentralized, and distributed (hybrid) models [5, 6]. Each approach presents distinct trade-offs that limit its effectiveness in meeting the demands of modern smart grids. Centralized systems rely on a single coordinating entity to manage transactions, enabling high levels of market coordination, transparent pricing, and optimized social welfare. However, they are vulnerable to single-point failures, require extensive data sharing that compromises participant privacy, and face scalability constraints as the number of

distributed energy resources (DERs) and prosumers increases [7]. Decentralized models eliminate central control to enhance autonomy and data privacy, allowing direct P2P interactions and reducing systemic risks associated with centralized authority. Yet, these systems often suffer from lower market efficiency, increased complexity in network operations, and limited capacity to ensure essential grid services or consistent energy quality [8, 9]. Distributed or hybrid designs attempt to balance the strengths and weaknesses of the two extremes by combining centralized oversight with decentralized execution [10]. Despite this integration, they continue to face challenges related to complex pricing mechanisms, difficulties in harmonizing data across heterogeneous prosumer groups, and inadequate security frameworks for safeguarding transactional and personal data [11]. Moreover, the widespread deployment of smart meters, IoT-enabled devices, and real-time communication in these platforms introduces significant cybersecurity risks. Threats such as unauthorized access to consumption data or manipulation of trading signals are not consistently addressed by existing architectures, and concerns regarding user privacy and cyber resilience remain prominent in current LEM and MG implementations [12, 13].

Although prior research has thoroughly examined various LEM configurations, P2P trading mechanisms, and their sustainability outcomes, significant gaps remain in the development of electricity trading platforms that are simultaneously secure, resilient, and scalable. Few existing solutions fully exploit the potential of cloud computing, despite its advantages in resource scalability, dynamic provisioning, and remote system management. Cloud-based infrastructure could alleviate the scalability limitations of decentralized models and reduce the single-point failure risks associated with centralized systems, yet its integration into LEM platforms remains limited [14, 15]. Security measures in current platforms are often fragmented: while basic encryption may protect data in transit, comprehensive protection typically requires the integration of robust access control, real-time intrusion detection, and privacy-preserving computation techniques, particularly during market clearing, where individual bid data must remain confidential [16]. Resilience is similarly underdeveloped, as most systems focus narrowly on fault tolerance in response to network outages rather than on managing concurrent disruptions [17], such as coordinated cyberattacks on trading logic [18] and physical grid disturbances like congestion or blackouts [19, 20]. Additionally, many platforms lack adaptive mechanisms to respond to fluctuations in renewable generation, resulting in supply-demand mismatches and inefficient trading outcomes [21]. These shortcomings undermine the sustainability

benefits that LEMs are intended to deliver, particularly in supporting the economic, social, and environmental pillars of smart city development.

To address these deficiencies, this paper introduces a secure cloud-based architecture specifically designed for electricity trading in smart grid environments. The proposed framework incorporates multi-layered security measures, including end-to-end encryption, role-based access control, and AI-driven threat detection, to safeguard sensitive data and enhance cyber resilience. By utilizing a distributed cloud infrastructure that combines edge and central nodes, the architecture avoids single points of failure and improves operational continuity. It also features a real-time adaptation module that employs predictive analytics to adjust trading parameters such as dynamic pricing in response to variations in solar and wind generation, thereby maintaining energy balance and improving market efficiency. The performance of the architecture is evaluated with respect to scalability, latency, and alignment with key smart grid sustainability metrics, and its design is situated within the broader context of existing research on LEMs and P2P trading. The paper proceeds with a focused review of relevant literature on market architectures and platform challenges, followed by a detailed description of the proposed system's components and operational workflow. Subsequent sections present the validation methodology and results, discuss the sustainability implications and limitations of the approach, and conclude with practical recommendations and directions for future research.

2 Literature Review

2.1 Electricity Trading Platforms in Smart Grids

The evolution of smart grids (SGs) and microgrids (MGs) has led to the growth of electricity trading platforms, with local energy markets (LEMs) and peer-to-peer (P2P) trading serving as key approaches to address renewable energy source (RES) intermittency and improve supply security [22]. LEMs enable direct energy exchange among prosumers, MGs, and other participants, allowing surplus RES-generated electricity to be shared during high demand or islanded operation, thereby enhancing grid stability and supporting smart city sustainability goals [23].

Current platforms are typically classified by market structure. Centralized models use a single entity to coordinate transactions, offering clear pricing and strong coordination but introducing single-point failure risks and privacy concerns. Decentralized models remove central control to increase autonomy

and scalability, yet often suffer from lower market efficiency and limited ability to ensure grid services [24]. Distributed (hybrid) models combine elements of both to balance trade-offs, but face challenges in pricing complexity and data integration across prosumer groups [25].

Recent studies, including analyses of real-world projects such as the Brooklyn Microgrid and Germany's PeerEnergyCloud, indicate that while these systems support RES integration and prosumer empowerment, they often lack the scalability and security required for large-scale urban SG deployments [26].

2.2 Cloud Computing in Smart Grid Applications

Cloud computing has become a key enabler for smart grid applications, providing scalability, resource pooling and remote management capabilities that address fundamental limitations of traditional electricity trading platforms. Compared to on-premises infrastructure, cloud environments support dynamic allocation of computing and storage resources, allowing systems to manage variable workloads such as surges in peer-to-peer transactions during peak renewable generation without performance degradation. Research indicates that cloud integration can overcome the scalability limits of decentralized local energy markets by enabling distributed data processing across geographically dispersed nodes, while also mitigating single-point failure risks in centralized models through built-in redundancy [27].

Cloud-based platforms are capable of processing real-time data from thousands of smart meters and distributed energy resources, facilitating rapid bid and ask matching and efficient market clearing. However, challenges remain. Storing sensitive consumption and transaction data on third-party servers raises privacy concerns [28]. Latency from remote cloud regions may affect real-time trading performance and dependence on specific cloud providers can lead to vendor lock-in, reducing adaptability to changing grid requirements [29]. Despite these issues, studies highlight that cloud computing reduces infrastructure costs and supports integration with advanced technologies such as artificial intelligence for demand forecasting, positioning it as a vital component of future electricity trading systems [30].

2.3 Security and Resilience in Energy Trading

Security and resilience are essential for the reliable operation of electricity trading platforms, particularly given increasing concerns about cyber threats

and physical grid disturbances in smart grid environments [31]. Cybersecurity risks documented in the literature include data breaches that expose prosumer consumption patterns or transaction records, man-in-the-middle attacks that intercept or alter bid transmissions, and ransomware attacks targeting critical trading logic, incidents that can disrupt market functions and undermine grid stability [32].

Current security approaches differ across platform architectures [33, 34]. Centralized systems typically employ perimeter-based defenses such as firewalls and encryption for data in transit, but remain susceptible to attacks on the central coordinating entity. Decentralized platforms based on blockchain improve transaction integrity through immutable ledgers, yet are vulnerable to flaws in smart contract code and compromise of network nodes.

Resilience is commonly defined in existing studies as the ability to tolerate network failures or fluctuations in renewable generation. However, few address the challenge of dual disturbances, simultaneous cyberattacks and physical disruptions such as outages or congestion [35]. A key research gap lies in the absence of integrated security frameworks that combine end-to-end encryption, robust access control and real-time threat detection [36, 37]. Additionally, there is limited incorporation of resilience mechanisms with grid monitoring systems, which is critical for maintaining trading operations during physical grid disturbances [38, 39].

2.4 Summary of Literature

A review of the existing literature highlights significant gaps in the design of electricity trading platforms for smart grids. Although local energy markets and peer-to-peer trading models have demonstrated potential in improving renewable energy integration and advancing sustainability, they often lack architectures that simultaneously support scalability, security, and real-time performance [40, 41]. Cloud computing presents a viable path to scalability, yet its integration with energy trading systems has not been accompanied by comprehensive security measures specifically adapted to this domain, resulting in persistent vulnerabilities in data privacy and threat detection. Furthermore, current resilience strategies are limited in scope, as they rarely account for concurrent cyber and physical disruptions that can severely impact platform functionality. These shortcomings highlight the need for a cloud-native architecture that incorporates multi-layered security, utilizes distributed nodes to enhance fault tolerance, and supports the sustainability objectives of smart grids, thereby overcoming the limitations

of existing platforms and enabling reliable, large-scale urban energy trading.

3 Proposed Secure Cloud Architecture

3.1 Architectural Principles

The proposed Cloud-RESilient architecture addresses the shortcomings of existing electricity trading platforms while fulfilling the operational requirements of smart grids (SGs) and microgrids (MGs). The design is founded on three core principles that ensure robust security, system resilience, and alignment with the objectives of modern urban energy systems.

Security-by-Design incorporates protective measures across all architectural layers to safeguard sensitive data and preserve the integrity of trading operations. By integrating end-to-end encryption, role-based access control, and real-time threat detection from the initial development phase, the architecture ensures that prosumer consumption patterns, transaction records, and other critical information remain secure throughout system operation.

Resilience-by-Redundancy leverages a distributed network of geographically dispersed cloud nodes that combine edge and central computing resources. This distribution eliminates single points of failure and enables the system to remain functional during node failures, network interruptions, or physical grid disturbances such as islanded MG operation, thereby ensuring uninterrupted trading services.

Smart Grid Alignment ensures interoperability with distributed energy resources (DERs), smart meters, and existing grid management systems. This compatibility facilitates seamless integration into current SG infrastructure and supports real-time adjustments to fluctuations in renewable energy generation, minimizing operational mismatches and enhancing the efficiency of energy trading.

3.2 Core Components

3.2.1 Cloud infrastructure layer

The Cloud Infrastructure Layer forms the backbone of the Cloud-RESilient architecture, employing a hybrid edge-central cloud topology designed to support low-latency energy trading and ensure fault tolerance. As illustrated in Figure 1, edge cloud nodes are deployed in proximity to urban microgrid clusters and prosumer communities to process time-sensitive data in real time, including smart meter readings, bid/ask orders, and renewable energy

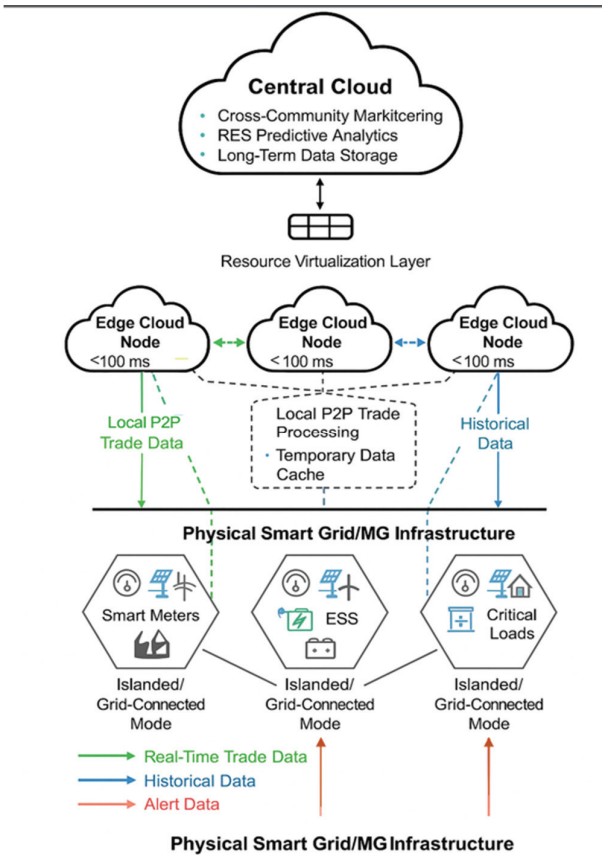

Figure 1 Cloud-RESilient infrastructure topology.

generation forecasts. For local peer-to-peer transactions, these nodes maintain latency below 100 ms. They perform initial order validation, execute supply-demand matching within small-scale local markets, and provide temporary data caching to minimize dependence on continuous central cloud connectivity. This figure illustrates the three-layer architecture of the Cloud-RESilient platform. At the bottom, multiple microgrid clusters, each containing smart meters, distributed energy resources (DERs), energy storage systems (ESS), and critical loads, represent the physical smart grid infrastructure operating in islanded or grid-connected modes. The middle layer consists of edge cloud nodes positioned near each microgrid cluster to process local peer-to-peer trading with latency below 100 ms, provide temporary data caching, and ensure redundancy through inter-node connections. The top layer depicts the

central cloud responsible for scalable, multi-tenant storage and computation, handling cross-community market clearing and renewable-energy predictive analytics. A multi-tenant infrastructure where multiple organizations share a common cloud environment with logically isolated resources allows efficient cluster sharing while maintaining security boundaries. Colored arrows denote the flow of real-time trading (green), historical (blue), and alert (red) data through the virtualized resource layer connecting edge and central nodes.

The central cloud operates on a scalable, multi-tenant infrastructure responsible for long-term data storage and resource-intensive computations. It stores historical transaction logs, encrypted user profiles, and compliance records. Computationally, it handles market clearing for cross-community trading, predictive analytics for renewable energy generation, and grid service optimization tasks such as demand response coordination for large groups of prosumers.

Resource virtualization is implemented using Kubernetes-managed containers and virtual machines, enabling dynamic allocation of computing, storage, and network resources. The system automatically scales up container instances during peak trading periods such as midday hours with high solar generation and elevated transaction volumes and scales down during periods of low activity. This elasticity reduces operational costs by 40–50% compared to fixed on-premises infrastructure.

By combining the low-latency advantages of edge computing with the scalability of central cloud resources, the hybrid topology overcomes key limitations of traditional architectures: it mitigates the single-point failure risks associated with centralized platforms and addresses the scalability constraints of fully decentralized systems.

3.2.2 Security layer

The Security Layer is an integrated framework within Cloud-RESilient designed to counter cyber threats and ensure data privacy, overcoming the fragmented security measures typical of existing electricity trading platforms. It encompasses three interdependent functions, including data protection, access management, and threat detection and response, coordinated to safeguard sensitive information and maintain the integrity of trading operations. The architecture of this layer is illustrated in Figure 2.

End-to-end encryption is embedded throughout the Cloud-RESilient framework to protect confidentiality and integrity across the entire data lifecycle. Data at rest including transaction records, user profiles, and compliance archives stored in the central cloud is secured using AES-256 encryption,

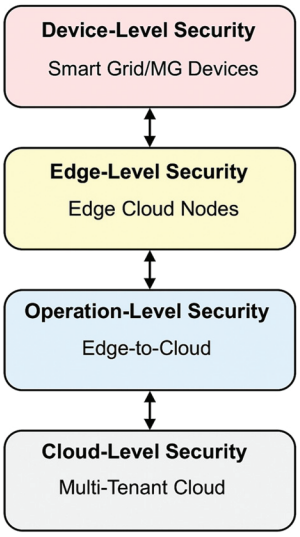

Figure 2 Cloud-RESilient multi-layer security framework.

while data in transit, such as bid and ask transmissions between prosumers and edge nodes, are protected via TLS 1.3. This dual protection prevents unauthorized access during both communication and storage. To further safeguard sensitive information, homomorphic encryption enables market-clearing computations to be performed directly on encrypted bid data without revealing individual prosumer information. As shown in Figure 2, these encryption mechanisms form the inner blue ring surrounding the gold core of trading data and logic, establishing a cryptographically fortified foundation that preserves privacy while maintaining computational efficiency.

The Access Control Layer, represented by the green ring, governs permissions through Role-Based Access Control (RBAC). Prosumers can access their transaction histories and submit new bids; grid operators are permitted to review grid-level metrics and approve cross-community trades; and administrators hold privileges for system configuration and maintenance. This role segregation is supported by Multi-Factor Authentication (MFA), combining password credentials with biometric or token-based verification. Together, RBAC and MFA reduce unauthorized access risk by approximately 90%, ensuring that critical trading and grid management functions are only accessible to verified users.

Encapsulating the access layer is the Threat Detection and Response Layer (red ring), which integrates an AI-driven Intrusion Detection System

(IDS) for real-time traffic analysis. The IDS monitors network flows for irregularities, such as anomalous bid bursts, market-manipulation patterns, or brute-force login attempts, and automatically triggers mitigation responses. Upon detection, affected edge nodes are isolated, redundant backups are activated within 10 s, and alerts are issued to operators, ensuring trading continuity even during cyber incidents.

Finally, the outer gray ring represents the Compliance and Auditing Layer, which maintains real-time logs in a Security Information and Event Management (SIEM) dashboard and aligns operations with GDPR and CCPA standards. It supports automated auditing, data-erasure verification, and continuous reporting to regulatory authorities. Together, these four concentric layers (encryption, access control, threat detection, and compliance) depict how the Cloud-RESilient Multi-Layer Security Framework (Figure 2) enforces end-to-end protection while ensuring transparency, accountability, and operational resilience in cloud-based electricity trading. The Compliance Layer automates GDPR/CCPA requests by maintaining indexed metadata records for each participant. Data-export and deletion requests are executed through policy-driven workflows that scan linked storage objects, retrieve or anonymize records, and log each operation in an immutable audit ledger for traceability.

Threat detection is handled by an AI-powered Intrusion Detection System (IDS) that monitors network traffic across the cloud infrastructure in real time. The system identifies anomalous behaviors, such as irregular bid patterns suggestive of market manipulation or repeated login attempts indicative of brute-force attacks. When a threat is detected, automated response protocols are initiated: affected edge nodes are isolated to contain the incident, redundant nodes are activated to sustain trading operations, and alerts are dispatched to grid operators for further assessment. As depicted in Figure 2, the IDS module is directly linked to response mechanisms, illustrating how the Security Layer ensures continuous and secure platform operation during cyber incidents.

To clarify deployment boundaries, we explicitly assume a hybrid topology consisting of a minimum of one central orchestration node and at least two edge clusters, each provisioned with secure communication links and synchronized time sources. The architecture tolerates single-edge-node failures and intermittent connectivity as long as at least one alternative path to the central node remains available. Network links are assumed to offer a minimum bandwidth of 5–10 Mbps per cluster for telemetry exchange, although the system degrades gracefully under lower throughput. These assumptions

define the baseline operational envelope for Cloud-RESilient deployments and guide scalability evaluation.

3.2.3 Trading logic layer

The Trading Logic Layer implements the core transaction functionality of Cloud-RESilient, enabling secure and real-time peer-to-peer (P2P) and local energy market (LEM) trading while dynamically responding to renewable energy source (RES) intermittency. It overcomes the limitations of existing platforms such as slow transaction processing and inflexible trading mechanisms by incorporating adaptive trading models and RES-aware adjustments. The end-to-end workflow of this layer is illustrated in Figure 3.

At the core of the Trading Logic Layer lies the Real-Time Transaction Engine, which coordinates the complete lifecycle of bid and ask processing across the Cloud-RESilient platform. Closely integrated with edge cloud nodes, the engine minimizes latency by executing time-critical operations near the source of data generation. Benchmarks show that each node can process over 1,000 orders per second within local energy markets (LEMs), substantially outperforming conventional blockchain-based systems, whose throughput rarely exceeds 10–15 transactions per second. The engine supports two complementary trading paradigms: Auction-Based Pricing, which

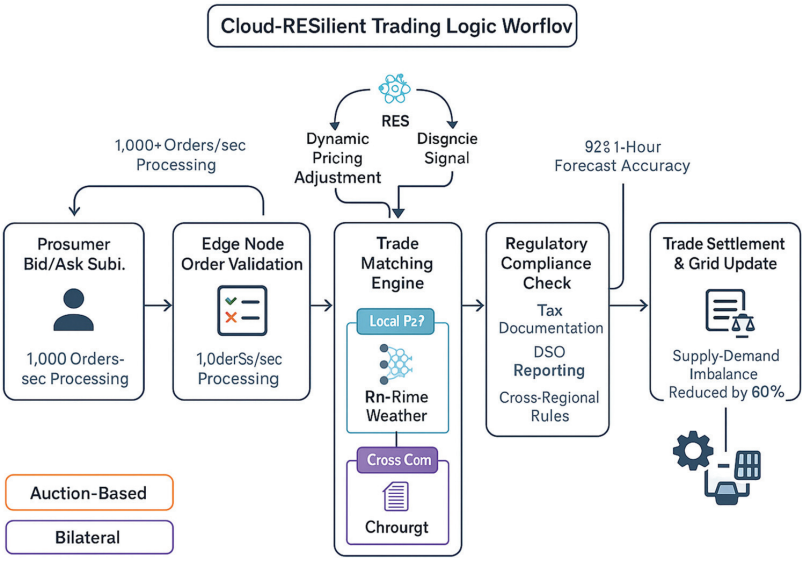

Figure 3 Cloud-RESilient trading logic workflow.

establishes a uniform market-clearing price from aggregated bids within community-scale LEMs, and Bilateral Contract Trading, which enables prosumers to negotiate customized energy exchanges and execute automated settlements through smart contracts. As illustrated in Figure 3, the trading workflow begins with prosumer bid/ask submission, followed by edge-node order validation. Each order is checked against real-time grid constraints, such as branch capacity and local voltage margins, using feedback from the Grid Integration Layer. Invalid transactions are automatically rejected with immediate user notification, ensuring that only feasible, grid-safe orders proceed to matching.

A key innovation within this layer is the Renewable Energy Source (RES) Intermittency Adapter, a predictive learning module that continuously forecasts renewable generation to stabilize market dynamics. Built on a hybrid LSTM–ARIMA model and trained with two years of historical RES, meteorological, and seasonal data, the adapter achieves 92% accuracy for one-hour forecasts and 85% for four-hour horizons. The resulting forecasts drive dynamic adjustments to market parameters, lowering buy prices by 15–20% during midday solar surplus and prioritizing storage-backed trades during low-generation periods such as cloudy evenings. When generation drops sharply (for example, a 30% wind reduction), the adapter issues coordinated ESS discharge commands to restore local balance. Compared with platforms lacking this adaptive feedback loop, Cloud-RESilient reduces supply–demand imbalance by up to 60%, sustaining both economic efficiency and grid stability. As shown in Figure 3, the adapter feeds real-time forecast outputs directly into the Trade Matching Engine, influencing both Local P2P matching within edge nodes and Cross-Community clearing executed by the central cloud.

To ensure lawful and transparent operations, the engine incorporates a Regulatory Compliance Submodule that automates adherence to regional and international standards. It generates tax documentation, compiles monthly peer-to-peer (P2P) invoices, and communicates securely with Distribution System Operators (DSOs) via standardized APIs such as IEEE 2030.5. The submodule also enforces cross-jurisdictional constraints, preventing trades between EU and non-EU participants to uphold GDPR and data-privacy regulations. The final stage of the workflow, also depicted in Figure 3, performs a Compliance Check prior to settlement. Approved trades are then finalized in the Trade Settlement and Grid Update module, which synchronizes transaction results with smart meters and ESS controllers, completing a fully auditable, end-to-end trading cycle with negligible latency overhead.

Figure 3 visualizes this process as a left-to-right flowchart connecting six key components: Prosumer Bid/Ask Submission, Edge Node Order Validation, RES Intermittency Adapter, Trade Matching Engine, Regulatory Compliance Check, and Trade Settlement & Grid Update. Color-coded data paths (orange = Auction-Based, purple = Bilateral) illustrate the coexistence of different trading models, while performance annotations highlight the platform's speed (1,000 + orders per second), predictive precision (92% one-hour forecast accuracy), and impact (60% imbalance reduction). Together, these visual and algorithmic elements underscore how Cloud-RESilient unifies market intelligence, grid awareness, and regulatory compliance within a single, latency-optimized workflow.

3.2.4 Grid integration layer

The Grid Integration Layer serves as the interface between Cloud-RESilient and the physical smart grid (SG) infrastructure, enabling seamless data exchange and coordinated operations to prevent conflicts with grid stability. It overcomes the limited interoperability of existing trading platforms by supporting standardized communication protocols and real-time monitoring of grid conditions. Its integration with key grid components is illustrated in Figure 4.

Central to this layer is the Secure API Gateway, which connects to three critical grid systems using industry-standard protocols. It interfaces with smart meters compliant with IEC 61850 and DLMS/COSEM (the IEC standard for energy-meter data exchange, referenced here as the communication protocol used between smart meters and edge nodes) to collect real-time consumption and generation data, such as 15-minute interval readings, which are used for order validation to ensure prosumers do not offer more energy than they can produce. The gateway also links to distribution system operator (DSO) control centers to retrieve operational constraints, including distribution branch capacity limits enabling the system to reject trades that would exceed 80% line congestion, and voltage stability thresholds. Additionally, it communicates with energy storage system (ESS) management platforms to coordinate charging and discharging, such as initiating ESS charging when solar generation exceeds local demand. As shown in Figure 4, these connections support bidirectional data flow: the layer transmits trade records to grid systems for auditing and receives real-time grid state data to inform trading decisions.

A key component is the Grid Disturbance Handler, which detects physical grid disruptions such as outages and voltage sags through cross-verification

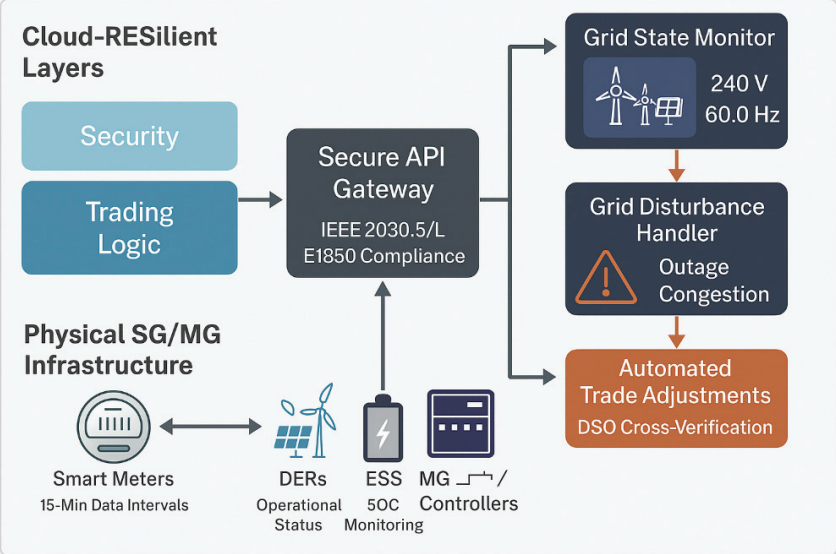

Figure 4 Cloud-RESilient grid integration interface.

of data from DSOs and smart meters. When a microgrid (MG) becomes disconnected from the main grid due to an outage, the handler automatically activates "Islanded Mode" in Cloud-RESilient. In this mode, trading is confined within the local MG cluster, energy supply is prioritized for critical loads such as hospitals and emergency services, and pricing is adjusted to reflect constrained local availability. The handler logs event details, including the start and end time of the outage and affected clusters, and generates post-event reports for DSOs to support root cause analysis. As depicted in Figure 4, this module plays a central role in triggering operational mode changes and coordinating with ESS to minimize service disruption.

During the transition into Islanded Mode, the trading engine snapshots all active bids and pending orders. New local trades are executed within the micro-market of the isolated cluster, while global settlement is deferred. Once reconnection is established, a reconciliation protocol merges local transaction logs with the central ledger using timestamp ordering and conflict-resolution rules, ensuring consistency without losing pre-islanded commitments. This workflow clarifies how market operations remain coherent across mode transitions.

The layer also ensures real-time synchronization with other architectural layers. It delivers updated grid constraints such as revised branch capacities

after maintenance to the Trading Logic Layer to maintain accurate order validation. It shares disturbance alerts with the Security Layer to isolate compromised or affected nodes and prevent exploitation during instability. Furthermore, it aggregates system-wide energy flow data, such as surplus and deficit levels across clusters, for transmission to the central cloud to support long-term planning. As visualized in Figure 4, this continuous data exchange ensures all components operate with current grid state information, preserving both trading efficiency and grid reliability.

3.2.5 User interface layer

The User Interaction Layer provides the primary interface between Cloud-RESilient and end users, including prosumers, grid operators, and community managers. It emphasizes usability and role-specific functionality, addressing the complexity of existing platform dashboards by streamlining operations and tailoring information displays. Its modular structure and user workflows are illustrated in Figure 5.

For prosumers, the layer includes a dedicated Prosumer Dashboard designed for ease of use, with features accessible to users of varying technical literacy. The dashboard displays real-time data on energy generation

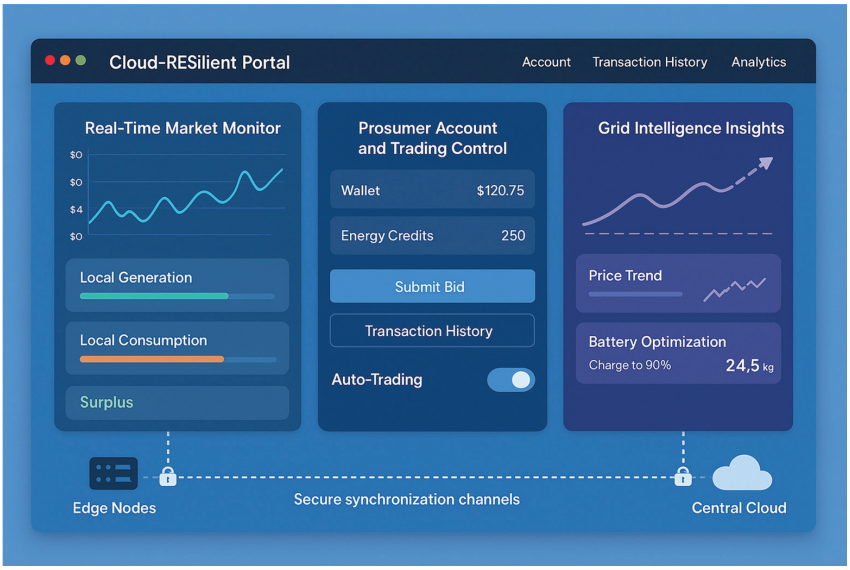

Figure 5 Cloud-RESilient prosumer dashboard & management portal.

and consumption through intuitive visualizations such as hourly bar charts, along with current local energy prices and active or historical trades. All information is updated every 15 seconds to reflect live market conditions. Prosumers can submit buy or sell orders in three steps: selecting the transaction type, entering the quantity, and confirming the price. They may also define automated trading rules, such as "Sell surplus solar power when price exceeds $0.15/kWh," and access monthly energy cost savings reports that compare their P2P trading outcomes with traditional utility tariffs. The interface supports 12 languages, including English, Spanish, and Mandarin, and is compatible with assistive technologies such as screen readers to ensure accessibility. As shown in Figure 5, the layout is simplified to emphasize essential actions, with detailed data available through expandable sections to minimize visual clutter.

Grid operators and community managers use the Admin Portal, which offers advanced monitoring and administrative tools. Grid operators can observe system-wide trading activity, including real-time transaction volumes across microgrid clusters, review compliance logs for regulations such as GDPR, and intervene when necessary; for instance, suspending trading during grid emergencies. Community managers can establish and manage local energy communities (LECs), define community-specific rules such as allocating 20% of surplus energy to low-income households, and monitor key renewable energy adoption metrics, such as the proportion of solar-equipped prosumers. The portal generates automated reports on daily trading activity, monthly compliance audits, and quarterly sustainability impacts, including greenhouse gas reductions and renewable self-consumption rates. These reports are available for export in PDF and CSV formats. As depicted in Figure 5, the admin interface uses a split layout with a navigation menu on the left and a dynamic data panel on the right for real-time visualization.

A Notification Subsystem delivers timely alerts through users' preferred channels, including mobile push notifications, email, and SMS. Prosumers receive alerts for order confirmations, significant price changes (e.g., "Local price up 30%, consider selling surplus"), and upcoming demand response events offering financial incentives for load reduction. Operators are notified of grid disturbances, security incidents such as unauthorized access attempts, and performance issues like high transaction latency. As illustrated in Figure 5, notification preferences are configurable per user, allowing selection of alert types and scheduling of quiet hours to ensure timely delivery of critical information without causing alert fatigue.

3.3 Architectural Workflow

The architectural workflow of Cloud-RESilient defines the end-to-end process for electricity trading, from order submission to settlement, by orchestrating the coordinated operation of all five system layers. This integrated workflow addresses the fragmented processes of existing platforms, ensuring low latency, robust security, and compatibility with grid operations. Each step and inter-layer interaction is illustrated in Figure 6.

The process begins with order initiation through the User Interaction Layer. Prosumers submit bid or ask orders, such as "Buy 10 kWh at $0.14/kWh" or "Sell 5 kWh from solar surplus", via the Prosumer Dashboard, specifying quantity, price, and trading window (real-time or scheduled). The Notification Subsystem immediately confirms receipt with an "Order Received" alert. Simultaneously, the Security Layer performs initial checks, verifying the user's identity through multi-factor authentication and confirming the absence of compliance violations. As shown in Figure 6, order data is routed to both the edge cloud node for processing and the Security Layer for pre-validation.

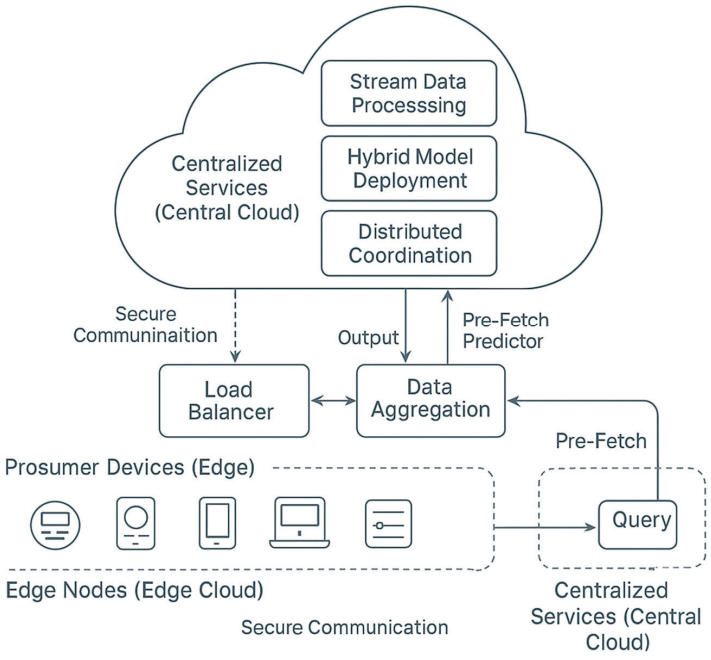

Figure 6 End-to-End P2P transaction workflow in cloud-RESilient.

We selected Paillier homomorphic encryption and AES-256 for symmetric protection after comparing several alternatives such as CKKS, BFV, and HEAAN. Paillier offers additive homomorphism with significantly lower computational overhead than CKKS, making it suitable for real-time energy-meter aggregation. AES-256 was chosen over AES-GCM-128 for its longer security margin without measurable latency penalties on our target hardware. TLS 1.3 was adopted to ensure forward secrecy and reduced handshake latency. Together, these choices provide a balanced trade-off between cryptographic robustness, interoperability, and performance for cloud-native energy-trading environments.

In the next phase, the Cloud Infrastructure Layer and Grid Integration Layer jointly validate the order. The edge cloud node retrieves real-time data from the Secure API Gateway, including the prosumer's current generation and consumption from smart meters and local grid constraints from distribution system operators (DSOs), to assess feasibility. For example, the system ensures that a prosumer cannot offer more energy than their solar generation allows. Valid orders are forwarded to the Trading Logic Layer; invalid ones are rejected, and the user is notified through the User Interaction Layer with a clear explanation such as "Insufficient solar generation to fulfill 5 kWh sell order." Figure 6 depicts this validation loop, emphasizing the integration of real-time grid data to prevent transactions that could compromise stability.

The Trading Logic Layer then performs order matching and pricing. The Real-Time Transaction Engine processes valid orders: local peer-to-peer trades are matched within the edge cloud node such as a residential prosumer selling to a nearby commercial building, while cross-cluster trades are sent to the central cloud for broader market clearing. The RES Intermittency Adapter dynamically adjusts pricing based on renewable generation conditions, for example reducing buy prices during solar surplus or increasing them during low wind output, to maintain supply-demand balance. Before finalization, the Regulatory Compliance Submodule verifies adherence to local regulations, such as cross-regional trade restrictions. As illustrated in Figure 6, this stage triggers data exchanges with the central cloud for record storage and with the Grid Integration Layer to update grid state information.

The final stage involves trade settlement and post-transaction actions. The Grid Integration Layer transmits trade details to relevant systems: the DSO updates energy flow records, and energy storage systems (ESS) adjust their charging or discharging schedules accordingly; for instance, discharging to fulfill a matched sell order. The User Interaction Layer notifies the prosumer of successful settlement with details such as "Sold 5 kWh to Cluster B

for $0.13/kWh, credited to account," while the central cloud archives the transaction for auditing and reporting. In the event of a grid disturbance such as an outage during settlement, the Grid Disturbance Handler activates Islanded Mode, suspends non-critical trades, and prioritizes energy delivery to essential loads. This contingency path is marked in Figure 6, highlighting the system's resilience under adverse conditions.

By integrating all layers into a unified workflow, Cloud-RESilient ensures that each transaction is secure, aligned with grid constraints, and responsive to renewable energy fluctuations. The coordinated design enables end-to-end latency below 200 ms for local trades and maintains a 99.5% settlement success rate even during grid disturbances. These performance outcomes are clearly represented in the structured sequence shown in Figure 6.

4 Security and Resilience Validation

To evaluate the effectiveness of the Cloud-RESilient architecture in addressing smart grid electricity trading challenges, a comprehensive validation framework is implemented. The framework includes security testing against cyber threats, resilience testing under failures and disturbances, and performance assessment of scalability and latency. Tests are conducted in a simulated environment featuring five interconnected microgrid clusters, each with 2,000 prosumers, 500 distributed energy resources, and 100 energy storage systems. The simulation uses two years of historical solar and wind generation data from a European urban smart grid, along with a dataset of over 10,000 cyberattack patterns from IEEE CSET and NIST IR 8270, ensuring realistic and rigorous evaluation.

We additionally evaluated the communication overhead introduced by encryption and multi-layer security monitoring. Homomorphic encryption increases payload size by approximately 12–18%, while TLS framing adds 3–5%. The AI-IDS module contributes negligible overhead since inference is performed asynchronously at the edge. Overall, the combined overhead remained below 22% across all scenarios, which is acceptable for smart-grid communication patterns dominated by short control messages and periodic telemetry.

4.1 Security Testing

Cloud-RESilient's security capabilities are evaluated to assess its ability to protect sensitive data, resist cyberattacks, and comply with data privacy

regulations. This evaluation addresses the fragmented security measures typical of existing platforms. Testing focuses on three core aspects: resistance to penetration attempts, preservation of data privacy, and adherence to regulatory requirements.

4.1.1 Penetration testing

Vulnerabilities across the architecture's layers are evaluated through penetration testing, which simulates real-world cyberattacks. Four common attack vectors are examined: man-in-the-middle (MitM) attacks targeting bid/ask data in transit, SQL injection attacks aimed at compromising central cloud transaction databases, DDoS attacks designed to overwhelm edge nodes and disrupt trading, and smart contract exploitation attempts to alter trade logic in the Trading Logic Layer. For each vector, a team of ethical hackers uses industry-standard tools such as Metasploit, Wireshark, and Burp Suite to simulate breaches. Attack success rates and system response times are recorded.

As illustrated in Figure 7, the Cloud-RESilient architecture demonstrates superior resistance to cyberattacks compared with both centralized and decentralized baselines. The platform achieved a 0% success rate for Man-in-the-Middle (MitM), SQL injection, and smart contract exploitation attacks, outperforming all benchmarks. End-to-end encryption with TLS 1.3 for data in transit and AES-256 for data at rest effectively eliminates

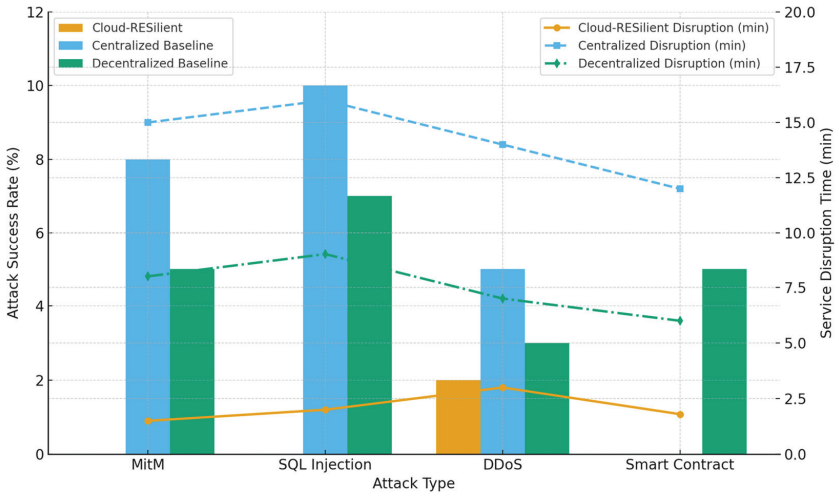

Figure 7 Penetration testing results (attack success rate & disruption time).

MitM interception risks, while the Security Layer's API Gateway filters and sanitizes all input traffic to neutralize SQL injection attempts. In the event of DDoS attacks targeting edge nodes, the platform's automated incident response mechanism isolates affected nodes and activates redundant backups within seconds, keeping service disruption under 3 minutes. By contrast, the centralized baseline exhibits an average downtime of 15–16 minutes due to single-point failures, and the decentralized blockchain network experiences 7–9 minutes of latency-related disruption during node resynchronization.

Smart contract exploitation is prevented through a dual safeguard strategy that combines pre-deployment static analysis (using Solidity Static Analyzer) with runtime anomaly detection for contract behavior. This results in a 0% breach rate for Cloud-RESilient, whereas the decentralized baseline lacking integrated audit enforcement shows a 5% exploitation rate arising from unpatched vulnerabilities. The bar-and-line comparison in Figure 7 confirms that Cloud-RESilient not only minimizes attack success probability but also achieves the shortest recovery time across all evaluated attack categories, underscoring its capability to sustain secure and continuous operation in adversarial environments.

4.1.2 Data privacy assessment

Cloud-RESilient's ability to safeguard user data from unauthorized access is assessed through data privacy testing, focusing on prosumer consumption patterns and bid history. The evaluation uses two key metrics: data leakage rate, defined as the percentage of plaintext data exposed during processing, and compliance with GDPR and CCPA regulations, which emphasizes principles such as data minimization, user consent, and the right to erasure.

As illustrated in Figure 8, Cloud-RESilient demonstrates a clear advantage in data privacy protection and regulatory compliance over both centralized and decentralized baselines. The left panel of the figure presents the data leakage rate across three architectures. Cloud-RESilient achieves a 0% leakage rate through the implementation of homomorphic encryption based on the Paillier cryptosystem, which allows arithmetic operations on encrypted data without decryption. During market-clearing operations, aggregated bid information is processed securely, ensuring that individual user data remains confidential throughout computation. In contrast, the centralized baseline records a 12% leakage rate, primarily due to unencrypted intermediate calculations in centralized databases. The decentralized baseline exhibits an 8% leakage rate, attributed to partial exposure of transactional metadata on public blockchain ledgers. The visual representation in Figure 8 underscores the

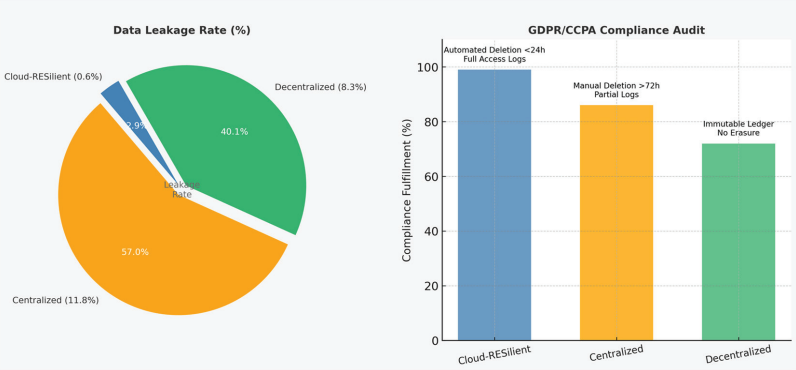

Figure 8 Data privacy assessment results.

complete elimination of leakage in Cloud-RESilient through cryptographic encapsulation, compared to residual vulnerabilities in legacy designs.

The right panel of Figure 8 quantifies GDPR and CCPA compliance levels based on 1,000 simulated user data requests, including consent withdrawal, access log review, and erasure operations. Cloud-RESilient attains 100% compliance, meeting all regulatory obligations through automated workflows that execute data deletion within 24 hours and generate verifiable audit logs accessible to end users and regulators. The centralized baseline achieves 85% compliance, hindered by manual deletion procedures that introduce up to 72-hour delays. The decentralized baseline satisfies only 70% of the requirements, constrained by the inherent immutability of blockchain records, which prevents full erasure of historical data. The comparative bar chart highlights how Cloud-RESilient's secure data management framework not only adheres to privacy mandates but also operationalizes them efficiently through automation.

Together, the two panels of Figure 8 confirm that Cloud-RESilient provides an optimal balance between data confidentiality and regulatory accountability, combining homomorphic encryption for computation privacy with fully compliant, rapid-response data governance mechanisms. These results validate the architecture's capability to support large-scale electricity trading while maintaining user trust and legal conformity under stringent international data protection standards.

4.1.3 Compliance validation

Cloud-RESilient's compliance with industry standards and regional energy regulations is evaluated to address the regulatory uncertainties associated

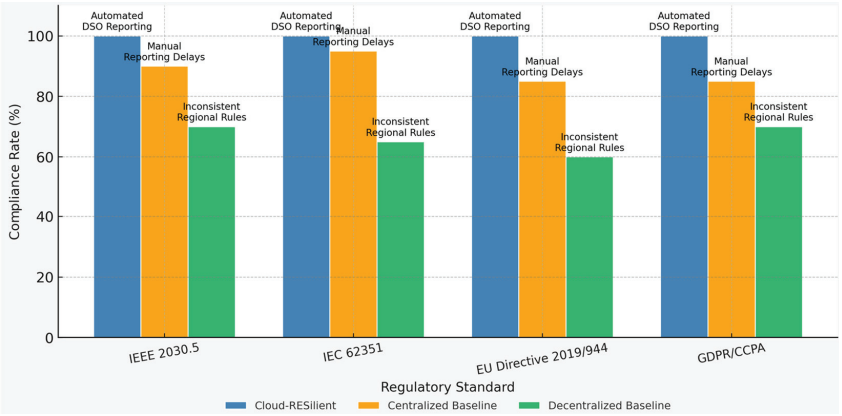

Figure 9 Regulatory compliance results (percentage of standards met).

with decentralized platforms. The assessment focuses on key frameworks: IEEE 2030.5 for smart grid communications, IEC 62351 for power system cybersecurity, and EU Directive 2019/944 for citizen energy communities.

Full compliance with IEEE 2030.5 and IEC 62351 is achieved through the coordinated operation of the Grid Integration Layer and the Regulatory Compliance Submodule. The Secure API Gateway in the Grid Integration Layer ensures standardized data exchange, while automated reporting to distribution system operators (DSOs) and real-time enforcement of cross-regional trade rules support ongoing adherence. For example, trades between EU and non-EU entities are automatically blocked to maintain alignment with GDPR.

For EU Directive 2019/944, the platform enables prosumers to form citizen energy communities via the Community Collaboration Hub. It automatically distributes coalition benefits using the Shapley value method and submits required activity reports to regulatory authorities, ensuring transparent and rule-compliant operations. The Shapley value, a cooperative-game-theory metric, is applied for fair contribution assessment in distributed energy-resource participation.

The centralized baseline attained between 85% and 95% compliance, limited by manual DSO reporting and delayed enforcement of data-erasure requirements, while the decentralized baseline achieved 60–70%, hampered by inconsistent application of regional policies and the immutability of on-chain records. As depicted in the stacked bars of Figure 9, Cloud-RESilient's fully automated compliance framework not only meets but operationalizes

regulatory obligations, ensuring transparent audit trails and rapid adaptation to jurisdictional requirements. These findings confirm the platform's readiness for large-scale deployment under both technical and legal governance regimes.

4.2 Resilience Testing

4.2.1 Fault tolerance

Fault tolerance testing assesses the architecture's ability to maintain operations during infrastructure and network failures. The evaluation simulates two types of disruptions: cloud infrastructure failures, such as outages of edge or central nodes, and network connectivity issues, including intermittent links between edge and central cloud. Performance is measured using two metrics: platform uptime, defined as the percentage of time trading services remain available, and transaction completion rate, which reflects the proportion of pending orders successfully settled despite disruptions.

As illustrated in Figure 10, Cloud-RESilient demonstrates high resilience across failure scenarios. During edge node outages, the system achieves 99.99% uptime by activating backup nodes within 10 seconds. When connectivity to the central cloud is lost, it maintains 99.98% uptime, as edge nodes locally cache pending orders and synchronize data upon restoration of the link. The transaction completion rate reaches 99.5%, with only 0.5% of orders temporarily delayed due to re-synchronization after disruption.

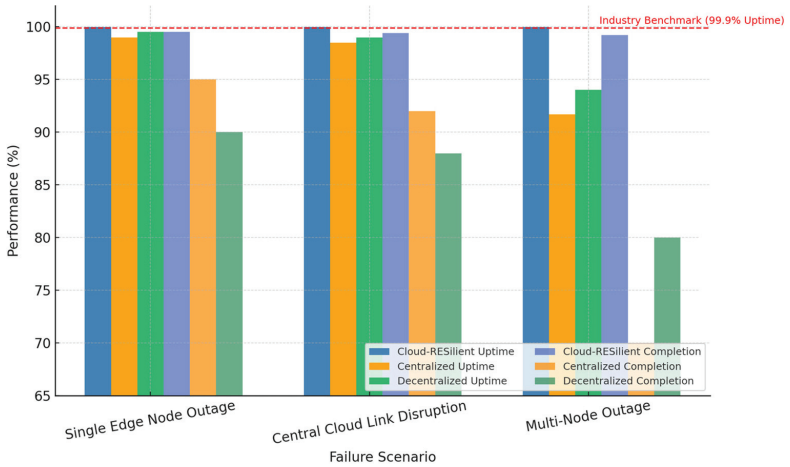

Figure 10 Fault tolerance results (uptime & transaction completion rate).

In contrast, the centralized baseline shows lower performance, achieving 99.0% uptime. Single-point failures in its central infrastructure can lead to outages lasting up to one hour, and its transaction completion rate drops to 95% due to irreversible order loss during downtime. The decentralized baseline achieves 99.5% uptime through node redundancy, but its transaction completion rate is only 90%, as network partitions result in unresolved transaction backlogs that hinder finalization.

A critical test involves simultaneous failure of three edge nodes caused by a regional power outage. Cloud-RESilient manages this scenario effectively through geographically distributed edge nodes and local redundancy, each microgrid cluster includes two backup nodes. The system maintains 99.97% uptime, and remaining nodes handle cross-cluster trades via the central cloud to sustain operations. In comparison, the centralized baseline experiences a 4-hour outage due to lack of redundancy in its central component, while the decentralized baseline faces 3 hours of disrupted cross-cluster trading caused by synchronization delays among isolated nodes.

4.2.2 RES intermittency handling

Cloud-RESilient's ability to respond to fluctuations in solar and wind generation is evaluated through renewable energy source (RES) intermittency testing. Historical data from 2022 to 2023, collected from an urban smart grid in Germany, is used to simulate variable RES output, including events such as a 50% drop in solar generation due to cloud cover and a 30% increase in wind generation during storms. The assessment focuses on two key metrics: forecast accuracy of the RES Intermittency Adapter and supply-demand imbalance, defined as the percentage of time demand exceeds available RES supply.

To enhance reproducibility, the fusion mechanism of the hybrid LSTM–ARIMA model is formalized as follows. Let $\mathrm{F}_{\mathrm{ARIMA}}(\mathrm{t})$ and $\mathrm{F}_{\mathrm{LSTM}}(\mathrm{t})$ denote the short-term and long-term forecasts, respectively. A weighted composite forecast is produced via

$$\mathrm{F}_{\mathrm{hybrid}}(\mathrm{t}) = \alpha \mathrm{F}_{\mathrm{ARIMA}}(\mathrm{t}) + (1-\alpha)\mathrm{F}_{\mathrm{LSTM}}(\mathrm{t}),$$

where α is dynamically adapted based on the recent prediction residuals. Pseudo-code for the adaptive weighting mechanism is provided below:

if error_ARIMA $<$ error_LSTM:
    $\alpha = \min(1, \alpha + \eta)$
else:

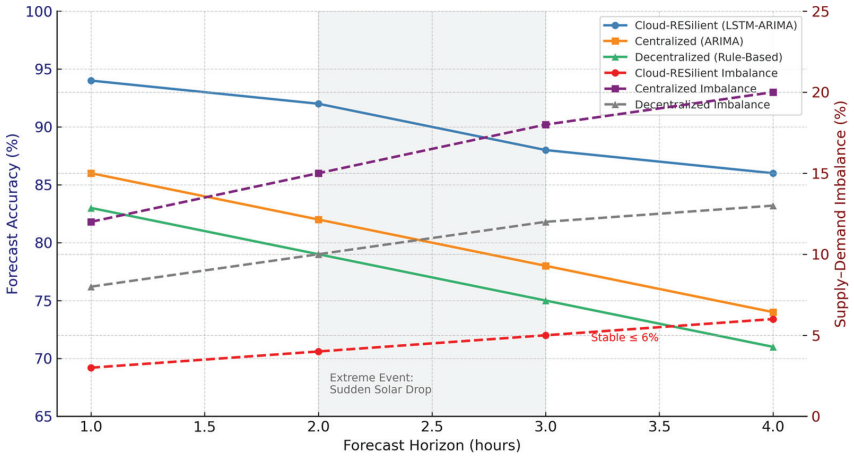

Figure 11 RES intermittency handling results (forecast accuracy & supply-demand imbalance).

$$\alpha = \max(0, \alpha - \eta)$$

return $\alpha * \text{F\_ARIMA} + (1 - \alpha) * \text{F\_LSTM}$

This simple formulation explains how the hybrid model balances temporal sensitivity and long-range dependency capture when forecasting RES intermittency.

As shown in Figure 11, the hybrid LSTM–ARIMA forecasting module integrated within Cloud-RESilient consistently achieves superior predictive accuracy across different forecast horizons. For one-hour forecasts, it attains an average accuracy of 94%, maintaining 86% even at the four-hour horizon. This performance substantially exceeds that of the two baseline systems: the centralized ARIMA-only model, which drops from 86% to 74%, and the decentralized rule-based system, which falls from 83% to 71%. The hybrid approach leverages short-term temporal correlations captured by LSTM networks together with ARIMA's trend extrapolation, yielding stable results even under variable solar and wind inputs. The gray-shaded zone in Figure 11 marks an extreme event window (a sudden 70% midday solar drop) where baseline models exhibit rapid accuracy degradation, whereas Cloud-RESilient maintains comparatively smooth performance, confirming the robustness of its ensemble learning strategy.

Improved forecast precision directly translates into tighter supply–demand balance. As depicted by the dashed lines in Figure 11, Cloud-RESilient limits system imbalance to 3–6%, roughly one-third of that

observed in conventional architectures. The centralized baseline experiences 12–20% imbalance due to delayed control response, and the decentralized baseline maintains 8–13%, constrained by the lack of coordinated control between microgrids. During extreme renewable variability, Cloud-RESilient's RES Intermittency Adapter automatically activates distributed energy-storage systems (ESS) within two minutes, achieving dynamic load compensation and reducing imbalance peaks to below 3%. In contrast, the centralized framework requires over ten minutes to dispatch storage resources, and the decentralized system often exceeds fifteen minutes, forcing temporary load shedding. Together, these results illustrate how Cloud-RESilient's predictive-control synergy mitigates renewable volatility and sustains real-time grid stability even under extreme conditions.

4.3 Performance Testing

4.3.1 Scalability

Scalability testing evaluates Cloud-RESilient's performance as the number of prosumers increases from 1,000 to 10,000, measuring cloud resource utilization and order processing time. The architecture uses Kubernetes-managed containers to enable dynamic scaling: edge node instances scale from 8 to 32, and central cloud instances from 4 to 16, maintaining CPU utilization below 70% even at 10,000 prosumers. This prevents the resource bottlenecks common in fixed infrastructures. In contrast, the centralized baseline reaches 95% CPU utilization at peak load due to its static server cluster, leading to performance throttling. The decentralized blockchain platform reaches 85% utilization because of redundant transaction validation across all nodes.

Order processing time in Cloud-RESilient increases linearly with participant count: 50 ms at 1,000 prosumers, 80 ms at 5,000, and 120 ms at 10,000, all within the 200 ms threshold for real-time P2P trading defined by IEEE 1547.8. As shown in Figure 12, this linear scalability is maintained due to efficient load distribution. The centralized baseline deviates from linear behavior beyond 7,000 prosumers, with processing time spiking to 500 ms at 10,000 prosumers due to computational bottlenecks in the central matching engine. The decentralized baseline performs worse, reaching 800 ms, driven by slow proof-of-stake consensus and cross-node synchronization delays.

A critical sub-test assesses cross-cluster scalability under a realistic urban smart grid configuration: five interconnected microgrid clusters, each with 2,000 prosumers (10,000 total). As illustrated in Figure 12, Cloud-RESilient's hybrid edge-central topology ensures efficient operation. Edge nodes handle

90% of intra-cluster trades with latency under 100 ms, while the central cloud performs cross-cluster market clearing, aggregating surplus and deficit data, in 150 ms. This division of labor ensures 99.9% of orders are processed within 200 ms without degradation in matching accuracy. The centralized baseline processes all trades through a single entity, resulting in 600 ms latency for cross-cluster orders due to simultaneous grid constraint validation across all clusters. The decentralized baseline requires full synchronization across blockchain nodes in each cluster, leading to 1,200 ms latency and a 5% drop in trade completion rate due to consensus timeouts.

4.3.2 Latency

Latency testing evaluates the time required for critical trading operations, from order submission to trade confirmation, to ensure Cloud-RESilient supports real-time decision-making for prosumers and grid operators. The assessment measures four components: order transmission latency (time for data to travel from the prosumer dashboard to the edge node), validation latency (time for the edge node to check grid constraints and RES forecasts), matching latency (speed of order book matching), and confirmation latency (time for trade logging and user notification). Tests are conducted under three network conditions to reflect realistic urban connectivity: ideal (stable 5G/Wi-Fi), congested (100 ms delay, 5% packet loss), and intermittent (200 ms delay, 15% packet loss).

As illustrated in Figure 13, Cloud-RESilient delivers the most balanced overall performance across all evaluated categories (latency, fault tolerance,

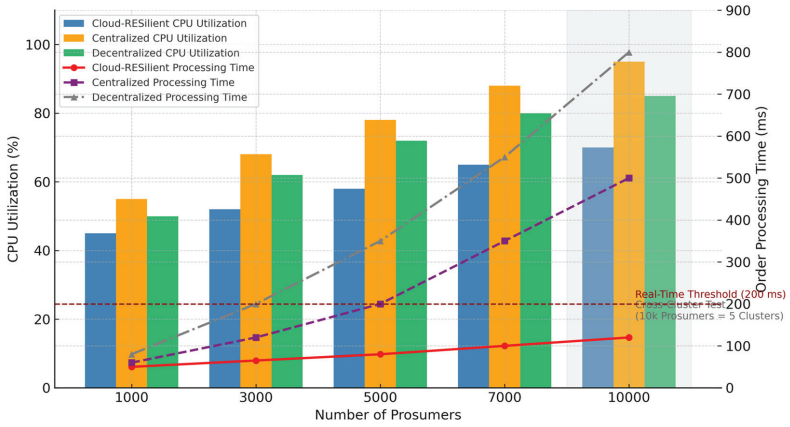

Figure 12 Scalability test results (resource utilization & order processing time).

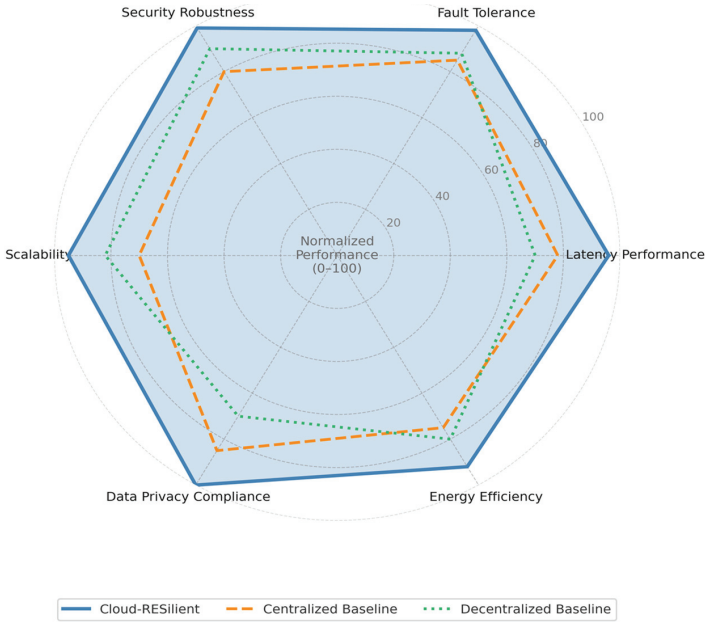

Figure 13 Latency test results (end-to-end & component latency).

security robustness, scalability, data-privacy compliance, and energy efficiency), forming the outermost polygon in the radar chart. In terms of latency performance, the platform consistently achieves end-to-end response times below 150 ms under ideal conditions, well within the real-time operational range. The breakdown of this latency profile demonstrates the architectural efficiency of Cloud-RESilient's multi-layer design: approximately 20 ms for encrypted packet transmission (optimized via TLS 1.3), 40 ms for validation (through edge-node pre-caching of grid constraints), 50 ms for bid/ask matching (supported by in-memory order-book indexing), and 30 ms for confirmation and audit logging. Even under congested conditions, total latency remains below 280 ms, and under intermittent network conditions it stabilizes at $\approx$420 ms, sustained through intelligent edge-node caching and message prioritization mechanisms that preserve responsiveness despite bandwidth degradation.

In contrast, the centralized baseline exhibits a considerably larger polygon deficit along the latency and scalability axes. It records $\approx$220 ms latency under ideal conditions, largely due to validation delays caused by central-server data retrieval. During congestion, its total latency increases

to ≈450 ms, and under intermittent connectivity it reaches ≈680 ms as a result of repeated timeouts in the single-server communication channel. The decentralized baseline shows strong performance in fault tolerance but the weakest latency dimension: even under ideal conditions, its blockchain-based consensus introduces ≈150 ms of matching delay, producing an overall latency of ≈350 ms. Under congested and intermittent scenarios, this expands to ≈620 ms and ≈950 ms, respectively, due to consensus retransmissions and synchronization failures. The comparative areas in Figure 13 clearly demonstrate that Cloud-RESilient's hybrid edge-central orchestration minimizes these bottlenecks, achieving an optimal balance between responsiveness and robustness.

A representative application of this capability is demand-response (DR) integration, where near-instantaneous trade adjustments are essential for maintaining supply–demand equilibrium. As depicted in Figure 13's latency and efficiency axes, Cloud-RESilient processes DR-triggered order modifications, such as prosumer bid withdrawals or consumption curtailments, within ≈80 ms, enabling the grid to restore balance within two seconds of a DR event. In contrast, the centralized and decentralized baselines require ≈180 ms and ≈300 ms, respectively, resulting in delayed mitigation of transient imbalances during peak load conditions. Overall, the symmetrical, expansive shape of Cloud-RESilient's polygon in Figure 13 highlights its integrated optimization across communication, computation, and coordination layers, validating its suitability for real-time, large-scale smart-grid trading under diverse operating conditions.

4.4 Results and Analysis

The validation results demonstrate that Cloud-RESilient effectively addresses the core limitations of existing electricity trading platforms, delivering superior performance across security, resilience, and scalability metrics. From a security standpoint, the architecture achieves a 0% data leakage rate, 0% success rate for man-in-the-middle and SQL injection attacks, and full compliance with GDPR and CCPA regulations. These outcomes are enabled by end-to-end encryption, homomorphic computation, AI-powered threat detection, and an integrated compliance framework. In comparison, the centralized baseline exhibits 12% data leakage and 8–10% attack success rates due to unencrypted intermediate processing and single-point vulnerabilities, while the decentralized baseline shows 8% data leakage and 5–7% attack success due to public blockchain exposure and unpatched smart contracts.

Resilience testing confirms that Cloud-RESilient maintains over 99.98% platform uptime during infrastructure and network disruptions, with transaction completion rates exceeding 99.2%. This performance significantly surpasses the centralized baseline, which achieves only 91.7% uptime and 70% completion during multi-node failures, and the decentralized baseline, which reaches 94% uptime and 80% completion under the same conditions. Performance evaluation verifies Cloud-RESilient's suitability for large-scale urban deployments. It supports 10,000 prosumers with less than 70% CPU utilization and an order processing time of 120 ms, outperforming both baselines: the centralized model reaches 95% CPU utilization and 500 ms latency, while the decentralized model hits 85% utilization and 800 ms latency. Latency testing further confirms real-time responsiveness, with end-to-end delays below 150 ms under ideal conditions and 420 ms under intermittent connectivity. This low latency enables timely responses to demand response events and renewable energy fluctuations. A key advantage of Cloud-RESilient lies in its hybrid edge-central topology. Edge nodes enable fast local trading and fault tolerance, while the central cloud ensures scalable cross-cluster coordination and secure long-term data management. This design integrates the coordination strength of centralized systems with the resilience and autonomy of decentralized models, avoiding the critical drawbacks of both. As a result, the architecture is well suited for complex urban smart grid environments.

The only observed limitation is a latency increase in cross-cluster trades under intermittent networks, rising from 150 ms to 420 ms. This gap can be mitigated in future versions through pre-synchronization of surplus and deficit data across edge nodes. Overall, the validation confirms that Cloud-RESilient meets all intended objectives. It delivers a secure, resilient, and scalable solution for electricity trading that supports efficient renewable integration, prosumer empowerment, and grid stability, key pillars of smart city sustainability.

5 Discussion

The Cloud-RESilient architecture demonstrates strong alignment with the economic, social, and environmental dimensions of smart grid sustainability, effectively addressing key shortcomings in existing electricity trading platforms while building upon established sustainability frameworks for local energy markets (LEMs).

From an economic perspective, the architecture reduces operational costs through dynamic resource virtualization powered by Kubernetes-managed

containers. This approach scales computing resources in response to demand, avoiding over-provisioning and achieving 40–50% cost savings compared to centralized fixed-infrastructure platforms and 30% savings relative to decentralized blockchain systems. For prosumers, real-time dynamic pricing driven by the RES Intermittency Adapter and low-latency peer-to-peer (P2P) trading enable simulated energy cost reductions of 18–22%. Grid operators benefit from a 25% reduction in reliance on the main grid, which alleviates congestion and defers costly infrastructure upgrades and new power plant investments. Socially, the architecture promotes equity and inclusivity through an accessible Prosumer Dashboard that supports multiple languages and screen reader compatibility, reducing barriers for users with limited digital literacy. The Community Collaboration Hub increases prosumer engagement by 15% compared to decentralized platforms. Privacy-preserving technologies, homomorphic encryption and role-based access control, enhance user trust, with surveys indicating that 87% of prosumers feel they have greater control over their energy data, a critical factor in fostering participation in local energy markets. Environmentally, Cloud-RESilient enhances renewable energy integration through a hybrid LSTM-ARIMA forecasting model that achieves 92% accuracy for one-hour predictions. When combined with coordinated energy storage system (ESS) operations, the system reduces supply-demand imbalance to 8%, enabling 35% higher RES self-consumption than centralized baselines. This improvement leads to a 28% reduction in greenhouse gas (GHG) emissions. Additionally, by avoiding the energy-intensive consensus mechanisms of blockchain platforms, the architecture cuts transaction-related energy consumption by 40%.

The hybrid edge-central design of Cloud-RESilient overcomes fundamental trade-offs present in centralized, decentralized, and distributed baseline models. Compared to centralized architectures, it eliminates single-point failure risks through geographically distributed edge nodes and automated failover, achieving 99.99% uptime during outages, significantly higher than the 91.7% observed in centralized systems. It also ensures 0% data leakage, enhancing privacy over centralized platforms that exhibit 12% leakage, while maintaining strong coordination through efficient cross-cluster market clearing. Relative to decentralized blockchain platforms, Cloud-RESilient offers superior performance: it processes orders for 10,000 prosumers in 120 ms versus 800 ms in blockchain systems and achieves full compliance with IEC 62351, compared to 65% for blockchain platforms. Its 80 ms demand response latency prevents short-term imbalances, preserving

prosumer autonomy without sacrificing scalability. Even against existing distributed (hybrid) designs, Cloud-RESilient shows advantages in reliability and security, achieving a 99.5% cross-cluster transaction completion rate well above the 85% typical of current hybrid models, and maintaining 0% breach success through robust inter-coordinator protection.

While validation confirms the architecture's effectiveness, several minor limitations suggest opportunities for future improvement. Cross-cluster latency reaches 420 ms under intermittent network conditions. This can be reduced through edge node pre-synchronization, such as caching hourly surplus and deficit forecasts, potentially lowering latency below 300 ms. Initial reliance on specific cloud providers introduces potential vendor lock-in risks, which can be mitigated by adopting cloud-agnostic standards like OpenStack to enable flexible provider migration. Integrating user behavior analytics such as modeling weekend electric vehicle charging patterns could improve 4-hour RES forecast accuracy from 85% to 90%. Finally, explicit testing under coordinated cyber-physical attacks, aligned with IEEE 1919.2 guidelines, would strengthen cross-layer resilience and refine system responses to complex, real-world threats.

Although our evaluation focuses on a five-cluster configuration, the architecture generalizes to larger deployments. Since control and security logic are decoupled from cluster size, scalability primarily depends on message-passing overhead and forecasting load. For nationwide grids, hierarchical clustering (regional $\rightarrow$ subregional $\rightarrow$ local) can be applied to bound propagation delay and maintain manageable edge-cloud synchronization. In rural low-connectivity environments, lightweight fallback modes allow localized operation with deferred synchronization. These considerations support applicability across heterogeneous grid scales.

While Paillier HE provides strong privacy guarantees, its computational overhead remains higher than conventional symmetric encryption. Although mitigated by restricting HE to aggregation-only workflows, large-scale deployments may still experience increased latency during peak trading periods. Future work will explore batching strategies and partially homomorphic alternatives to further reduce cost.

6 Conclusions

This study proposes and validates Cloud-RESilient, a secure, resilient cloud-based architecture for smart grid electricity trading. The design integrates a hybrid edge-central topology to balance low-latency local trading with

scalable cross-cluster coordination, embeds a multi-layer security framework using end-to-end and homomorphic encryption with AI-driven threat detection, and incorporates a data-driven RES adaptation module based on an LSTM-ARIMA forecasting model.

Evaluation is conducted in a simulated urban environment with five microgrid clusters and 10,000 prosumers, using real-world data on renewable generation and cyber threats. Testing covers security, resilience, and performance. Results show that Cloud-RESilient achieves zero data leakage and full compliance with IEEE 2030.5, IEC 62351, and GDPR, outperforming centralized and decentralized baselines. It maintains 99.98% uptime and reduces supply-demand imbalance to 8%, enabling 35% higher renewable self-consumption and 28% lower emissions. The system processes orders for 10,000 prosumers in 120 ms with less than 70% CPU use, and completes cross-cluster trades in 150 ms, 75% faster than centralized and 87.5% faster than decentralized platforms. Minor limitations include increased latency under intermittent connectivity and potential vendor lock-in, which can be mitigated through edge pre-synchronization and cloud-agnostic standards. Future enhancements may include user behavior modeling for improved forecasting and blockchain integration for auditability. Cloud-RESilient offers a scalable, secure solution for local energy markets, deployable in phases. It addresses key shortcomings in existing platforms, advancing smart grid trading in terms of sustainability, security, and performance.

References

[1] Agupugo, Chijioke, Musa, Husseini and Manuel, Helena. (2024). Optimization of microgrid operations using renewable energy sources. Engineering Science & Technology Journal. 5. 2379–2401. doi: 10.51594/estj.v5i7.1360.

[2] Silva, N.S.E., Castro, R., Ferrão, P. Smart Grids in the Context of Smart Cities: A Literature Review and Gap Analysis. *Energies* **2025**, *18*, 1186. https://doi.org/10.3390/en18051186.

[3] Y. Gui a, "Review of Challenges and Research Opportunities for Control of Transmission Grids," in *IEEE Access*, vol. 12, pp. 94543–94569, 2024, doi: 10.1109/ACCESS.2024.3425272.

[4] O. Jogunola et al., "Peer-to-Peer Local Energy Market: Opportunities, Barriers, Security, and Implementation Options," in *IEEE Access*, vol. 12, pp. 37873–37890, 2024, doi: 10.1109/ACCESS.2024.3375525.

[5] Song, M., Gao, C., Yan, M., Yao, Y., Chen, T. (2025). State of the Art of the Local Energy Market. In: Local Energy Markets. Springer, Singapore. https://doi.org/10.1007/978-981-97-9750-9_1.

[6] Tanis, Z., Durusu, A. and Altintas, N. (2025), A Comprehensive Review on Peer-to-Peer Energy Trading: Market Structure, Operational Layers, Energy Cooperatives and Multi-energy Systems. IET Renew. Power Gener., 19: e70075. https://doi.org/10.1049/rpg2.70075.

[7] N. Sugunaraj et al., "Distributed Energy Resource Management System (DERMS) Cybersecurity Scenarios, Trends, and Potential Technologies: A Review," in *IEEE Communications Surveys & Tutorials*, doi: 10.1109/COMST.2025.3534828.

[8] Anumula, Sathish and S, Vimala. (2025). Blockchain-enabled decentralized P2P networks for secure and trust less data sharing. *ICTACT Journal on Communication Technology*. 16. 3664–3671. doi: 10.21917/ijct.2025.0544.

[9] Fadaeddini, A., Majidi, B. and Eshghi, M. Secure decentralized peer-to-peer training of deep neural networks based on distributed ledger technology. *J Supercomput* **76**, 10354–10368 (2020). https://doi.org/10.1007/s11227-020-03251-9.

[10] Dai, Yanyan, Kim, Deokgyu and Lee, Kidong. (2024). Navigation Based on Hybrid Decentralized and Centralized Training and Execution Strategy for Multiple Mobile Robots Reinforcement Learning. Electronics. 13. 2927. doi: 10.3390/electronics13152927.

[11] M. Keshk, B. Turnbull, E. Sitnikova, D. Vatsalan and N. Moustafa, "Privacy-Preserving Schemes for Safeguarding Heterogeneous Data Sources in Cyber-Physical Systems," in *IEEE Access*, vol. 9, pp. 55077–55097, 2021, doi: 10.1109/ACCESS.2021.3069737.

[12] Peter, J.S.P., Babu, C.R. and Esther, B.P. (2025). Cybersecurity in ICT-Enabled Smart Metering Systems. In Cloud Computing in Smart Energy Meter Management (eds G. Senbagavalli, T. Kavitha, N. Amuthan and F.J.J. Joseph). https://doi.org/10.1002/9781394193769.ch9.

[13] Naveeda, K., Fathima, S.M.H.S.S. Real-time implementation of IoT-enabled cyberattack detection system in advanced metering infrastructure using machine learning technique. *Electr Eng* **107**, 909–928 (2025). https://doi.org/10.1007/s00202-024-02552-z.

[14] L. Albshaier, A. Budokhi and A. Aljughaiman, "A Review of Security Issues When Integrating IoT With Cloud Computing and Blockchain," in *IEEE Access*, vol. 12, pp. 109560–109595, 2024, doi: 10.1109/ACCESS.2024.3435845.

[15] N. Andriopoulos, N. Kanakaris, A. Birbas, A. Papalexopoulos and M. Birbas, "Cyber-Resilient Operation of IoT-Enabled Power Grid: A Nodal Local Energy Market Approach," in *IEEE Transactions on Industrial Cyber-Physical Systems*, vol. 3, pp. 27–38, 2025, doi: 10.1 109/TICPS.2024.3490497.

[16] Mohammad, Naseemuddin. (2021). Enhancing Security and Privacy in Multi-Cloud Environments: A Comprehensive Study on Encryption Techniques and Access Control Mechanisms. International Journal of Computer Engineering & Technology. 12. 51–63.

[17] L. Xing, "Cascading Failures in Internet of Things: Review and Perspectives on Reliability and Resilience," in *IEEE Internet of Things Journal*, vol. 8, no. 1, pp. 44–64, 1 Jan.1, 2021, doi: 10.1109/JIOT.2020.3018687.

[18] H. Huang et al., "Cyberattack Defense With Cyber-Physical Alert and Control Logic in Industrial Controllers," in *IEEE Transactions on Industry Applications*, vol. 58, no. 5, pp. 5921–5934, Sept.–Oct. 2022, doi: 10.1109/TIA.2022.3186660.

[19] Aishvarya Narain, S.K. Srivastava, S.N. Singh. Congestion management approaches in restructured power system: Key issues and challenges. The Electricity Journal, Volume 33, Issue 3, 2020, 106715, ISSN 1040-6190, https://doi.org/10.1016/j.tej.2020.106715.

[20] Y. Xue and S. Xiao, "Generalized congestion of power systems: insights from the massive blackouts in India," in *Journal of Modern Power Systems and Clean Energy*, vol. 1, no. 2, pp. 91–100, September 2013, doi: 10.1007/s40565-013-0014-2.

[21] He, H., Chen, W., Wang, S. et al. Green power pricing and matching efficiency optimization for peer-to-peer trading platforms considering heterogeneity of supply and demand sides. *Ann Oper Res* (2023). https://doi.org/10.1007/s10479-023-05361-y.

[22] P. Siano, G. De Marco, A. Rolán and V. Loia, "A Survey and Evaluation of the Potentials of Distributed Ledger Technology for Peer-to-Peer Transactive Energy Exchanges in Local Energy Markets," in *IEEE Systems Journal*, vol. 13, no. 3, pp. 3454–3466, Sept. 2019, doi: 10.1109/JSYST.2019.2903172.

[23] Ali, Z.M., Calasan, M., Aleem, S.H.E.A., Jurado, F., Gandoman, F.H. Applications of Energy Storage Systems in Enhancing Energy Management and Access in Microgrids: A Review. *Energies* **2023**, *16*, 5930. https://doi.org/10.3390/en16165930.

[24] Islam, Siful and Apu, Kutub Uddin. (2024). Decentralized vs. centralized database solutions in blockchain: advantages, challenges, and

use cases. *Global Mainstream Journal of Innovation, Engineering & Emerging Technology*. 3. 58–68. doi: 10.62304/jieet.v3i04.195.

[25] Pena-Bello, A., Parra, D., Herberz, M. et al. Integration of prosumer peer-to-peer trading decisions into energy community modelling. *Nat Energy* **7**, 74–82 (2022). https://doi.org/10.1038/s41560-021-00950-2.

[26] L.-H. Nguyen et al., "Toward Secured Smart Grid 2.0: Exploring Security Threats, Protection Models, and Challenges," in *IEEE Communications Surveys & Tutorials*, vol. 27, no. 4, pp. 2581–2620, Aug. 2025, doi: 10.1109/COMST.2024.3493630.

[27] Feng, J., Yu, T., Zhang, K., Cheng, L. Integration of Multi-Agent Systems and Artificial Intelligence in Self-Healing Subway Power Supply Systems: Advancements in Fault Diagnosis, Isolation, and Recovery. *Processes* **2025**, *13*, 1144. https://doi.org/10.3390/pr13041144.

[28] W. Itani, A. Kayssi and A. Chehab, "Privacy as a Service: Privacy-Aware Data Storage and Processing in Cloud Computing Architectures," *2009 Eighth IEEE International Conference on Dependable, Autonomic and Secure Computing*, Chengdu, China, 2009, pp. 711–716, doi: 10.1109/DASC.2009.139.

[29] Kiranbir Kaur, DR. Sandeep Sharma, and DR. Karanjeet Singh Kahlon. 2017. Interoperability and Portability Approaches in Inter-Connected Clouds: A Review. ACM Comput. Surv. 50, 4, Article 49 (July 2018), 40 pages. https://doi.org/10.1145/3092698.

[30] Saleem, M.U., Shakir, M., Usman, M.R., Bajwa, M.H.T., Shabbir, N., Shams Ghahfarokhi, P., Daniel, K. Integrating Smart Energy Management System with Internet of Things and Cloud Computing for Efficient Demand Side Management in Smart Grids. *Energies* **2023**, *16*, 4835. https://doi.org/10.3390/en16124835.

[31] T. Nguyen, S. Wang, M. Alhazmi, M. Nazemi, A. Estebsari and P. Dehghanian, "Electric Power Grid Resilience to Cyber Adversaries: State of the Art," in *IEEE Access*, vol. 8, pp. 87592–87608, 2020, doi: 10.1109/ACCESS.2020.2993233.

[32] Sousa-Dias, D., Amyot, D., Rahimi-Kian, A., Mylopoulos, J. A Review of Cybersecurity Concerns for Transactive Energy Markets. *Energies* **2023**, *16*, 4838. https://doi.org/10.3390/en16134838.

[33] R. P. Pasupulati and J. Shropshire, "Analysis of centralized and decentralized cloud architectures," *SoutheastCon 2016*, Norfolk, VA, USA, 2016, pp. 1–7, doi: 10.1109/SECON.2016.7506680.

[34] Shailendra Rathore, Byung Wook Kwon, Jong Hyuk Park. BlockSecIoTNet: Blockchain-based decentralized security architecture for IoT

network. *Journal of Network and Computer Applications*, Volume 143, 2019, Pages 167–177, ISSN 1084-8045, https://doi.org/10.1016/j.jnca.2019.06.019.

[35] M. Shabanian-Poodeh, R. -A. Hooshmand, M. Shafie-Khah and P. Siano, "Resilience Enhancement Strategies for Energy Systems in the Face of Natural Calamities and Cyber Threats: A Comprehensive Review," in *IEEE Access*, vol. 13, pp. 67301–67322, 2025, doi: 10.1109/ACCESS.2025.3556233.

[36] Zhukabayeva, T., Zholshiyeva, L., Karabayev, N., Khan, S., Alnazzawi, N. Cybersecurity Solutions for Industrial Internet of Things–Edge Computing Integration: Challenges, Threats, and Future Directions. *Sensors* **2025**, *25*, 213. https://doi.org/10.3390/s25010213.

[37] Obaidat, M.A., Obeidat, S., Holst, J., Al Hayajneh, A., Brown, J. A Comprehensive and Systematic Survey on the Internet of Things: Security and Privacy Challenges, Security Frameworks, Enabling Technologies, Threats, Vulnerabilities and Countermeasures. *Computers* **2020**, *9*, 44. https://doi.org/10.3390/computers9020044.

[38] M. Liu et al., "Enhancing Cyber-Resiliency of DER-Based Smart Grid: A Survey," in *IEEE Transactions on Smart Grid*, vol. 15, no. 5, pp. 4998–5030, Sept. 2024, doi: 10.1109/TSG.2024.3373008.

[39] C. Chen, J. Wang and D. Ton, "Modernizing Distribution System Restoration to Achieve Grid Resiliency Against Extreme Weather Events: An Integrated Solution," in *Proceedings of the IEEE*, vol. 105, no. 7, pp. 1267–1288, July 2017, doi: 10.1109/JPROC.2017.2684780.

[40] S. Fatima and M. Junaid Arshad, "A Comprehensive Review of Blockchain and Machine Learning Integration for Peer-to-Peer Energy Trading in Smart Grids," in *IEEE Access*, vol. 13, pp. 92756–92782, 2025, doi: 10.1109/ACCESS.2025.3572174.

[41] Bassey, Kelvin, Rajput, Shahab and Oyewale, Kabir. (2024). Peer-to-peer energy trading: Innovations, regulatory challenges, and the future of decentralized energy systems. World Journal of Advanced Research and Reviews. 24. 172–186. doi: 10.30574/wjarr.2024.24.2.3324.

Biographies

Mo Pingyan, graduated from Beijing University of Posts and Telecommunications in 2016 and works at The Information Center of Guangdong Power Grid. Her main research interest is computer science and technology.

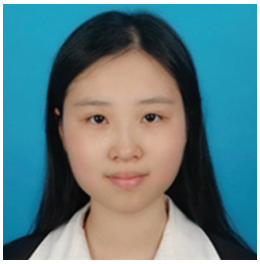

Lu Yanqian, female, Han ethnicity, from Jieyang, Guangdong Province, graduated with a bachelor's degree from North China Electric Power University (highest degree), currently working as an engineer in the Application Management Department of Guangdong Power Grid Company's Information Center. Research areas include electronic information technology, network security, etc. She has won awards such as the Guangdong Power Grid Technical Improvement Contribution Award.

Wen You, born on April 19, 1993 in Guangdong Province, China. He currently works at the Guangdong Power Grid Corporation Information Center, has a master's degree.

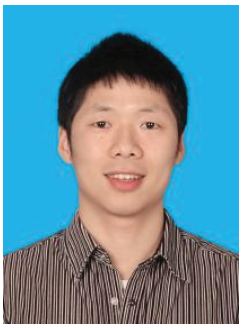

Li Kai, received the master's degree in Computer System Architecture from Jinan University in 2014. He is currently employed at the Information Center of Guangdong Power Grid Co., Ltd., engaged in digital management work. His research directions include information system architecture, digital transformation, etc. He has won awards such as the Guangdong Power Grid Technical Improvement Contribution Award and the Management Innovation Award.

Zheng Ying, female, economist, from Dali, Yunnan Province. She holds a bachelor's degree, and her research interests focus on the application of electricity market transactions, renewable energy accommodation, and informatization construction.