

ORIGINAL RESEARCH

Open access

A Scalable Grid-Intelligence Architecture for Federated Foundation Models across Distribution Utilities, Microgrids, and EV Aggregators

Elena Petrova^{1*}, Ivan Georgiev¹

Abstract

Grid intelligence increasingly depends on artificial intelligence, yet operational data remains distributed across utilities, microgrids, and electric vehicle aggregators. Federated foundation models offer a pathway for shared learning while preserving privacy, operational autonomy, and data sovereignty. Existing grid AI architectures are not designed to federate foundation models across multiple stakeholder classes with different data assets, control objectives, and regulatory obligations. The resulting fragmentation limits the ability of the power sector to build reusable intelligence across distribution, local energy systems, and mobile flexibility resources. This article proposes a scalable grid-intelligence architecture that federates a foundation model across distribution utilities, microgrid operators, and EV aggregators. The architecture integrates actor-specific layers with a shared federated foundation model layer and a governance/security layer. The article develops a conceptual architecture using layered functions, interface descriptions, model-update flows, and governance responsibilities. Validation pathways are outlined as design checks, staged pilots, and operational readiness assessments without presenting empirical results. The architecture defines how each actor contributes local knowledge, receives model improvements, and preserves control over proprietary data. The tables summarize utility, microgrid, and governance functions that support scalable, privacy-respecting federation. The proposed architecture offers a blueprint for collaborative AI in the power grid. It enables shared intelligence across diverse grid stakeholders while keeping raw data local and governance responsibilities explicit.

Keywords Microgrid, Federated foundation models, Scalability, Grid intelligence, Distribution utility, EV aggregator

*Correspondence:

Elena Petrova
elena.petrova@gmail.com

¹ Department of Energy Informatics and Optimization, Faculty of Engineering, University of Sofia, Sofia, Bulgaria

Introduction

Grid operation is moving from centrally planned control toward coordinated intelligence across distribution utilities, microgrids, and flexible demand resources. Federated learning has emerged as a general paradigm for collaborative model development without centralizing raw data, making it relevant to grid domains where operational records, customer information, and asset telemetry are institutionally separated [1]. In this context, a scalable architecture must support coordination among actors that differ in size, infrastructure, and decision authority. The architecture proposed here treats each actor as a sovereign participant in a shared learning system rather than as a passive data provider.

Foundation models extend the value of federated learning because they can encode reusable representations that are adapted to many downstream tasks. Recent power-sector discussions of foundation models indicate that transformer-based and large-language-model approaches can support forecasting, situational awareness, engineering assistance,

and cross-domain reasoning when aligned with grid data and operational constraints [2]. However, such models are difficult to deploy in critical infrastructure when training data is commercially sensitive, geographically distributed, and governed by different ownership rules. A federated foundation model architecture therefore needs both technical learning mechanisms and institutional safeguards.

The power grid already contains examples of federated learning for load forecasting, cyberattack detection, anomaly detection, and EV-related tasks, showing that the underlying learning concept is compatible with energy applications [3]. Yet these studies are typically task-specific and do not define a common architecture for utilities, microgrids, and EV aggregators to share a reusable model backbone. This article moves from individual applications toward an architecture in which local adaptation, prompt-based use, and federated model updating coexist. The goal is to create a shared intelligence layer that supports many grid-edge functions while respecting stakeholder boundaries.

The proposed article is conceptual and architecture-driven rather than experimental. It builds on federated machine learning principles that distinguish local training, server-side aggregation, and privacy-preserving collaboration [4], while adapting those principles to the operational structure of modern distribution systems. The contribution is a layered architecture with a utility layer, a microgrid layer, an EV aggregator layer, a federated foundation model layer, and a governance/security layer. Each layer defines data responsibilities, model interfaces, scalability constraints, and accountability mechanisms.

Scalability Problem

Data volume and velocity across grid actors

Distribution utilities operate with high-volume streams from smart meters, supervisory control systems, outage management platforms, distributed energy resource management systems, and field-device telemetry. Federated learning surveys emphasize that scale becomes difficult when many participants generate updates at different speeds and under unequal communication conditions [5]. In a grid setting, this problem is intensified by the difference between slow planning data, near-real-time operational telemetry, and event-driven alerts. A scalable architecture must therefore separate raw data handling from model-update exchange and avoid assuming that all actors can participate synchronously.

Heterogeneity of data types and tasks

Grid actors produce heterogeneous data because utilities observe feeders and assets, microgrids observe local generation and storage, and EV aggregators observe mobile charging behavior and flexibility commitments. Personalized federated learning for residential load forecasting shows that non-identical data distributions are a central challenge when energy users and sites follow different operating patterns [6]. The architecture must support heterogeneous tasks such as forecasting, anomaly detection, congestion awareness, flexibility estimation, and operator assistance without forcing every actor into one uniform data schema. This requirement motivates a foundation model layer that can host shared representations while allowing local task adapters.

Privacy and regulatory constraints

Power-system data can reveal customer behavior, asset vulnerabilities, commercial strategies, and operational weaknesses, so privacy and sovereignty are not optional design features. Security and privacy reviews of federated learning identify differential privacy, secure aggregation, and attack resilience as core requirements rather than add-on protections [7]. In multi-stakeholder energy systems, privacy also includes institutional control over who may access models, updates, logs, and derived insights. The architecture must therefore keep raw data inside each actor layer while using governed model updates as the primary collaboration channel.

The need for a federated approach

A centralized model-development process would require utilities, microgrid operators, and aggregators to transfer sensitive operational data into a common repository, which is difficult to justify technically or institutionally. Federated learning provides an alternative by allowing participants to train locally and share only selected model information, a design already explored for smart-grid security and forecasting applications [8]. For foundation models, the federated approach is especially valuable because a common model backbone can be adapted repeatedly across different grid tasks. The scalability problem is therefore not simply data size; it is the need to coordinate many sovereign learning environments.

Architecture Requirements

Functional requirements

The architecture must allow each actor to ingest local data, preprocess it, adapt the shared model, evaluate local behavior, and exchange approved updates with the federation layer. Federated learning applications in smart grids show that practical deployments must connect learning functions to operational tasks such as forecasting, detection, and decision support [9]. Functionally, the system must support local fine-tuning, prompt-based inference, model adapter management, and controlled update submission. These capabilities should be available without requiring each actor to expose raw measurements or internal operating rules.

Scalability and communication requirements

Scalability requires the architecture to operate across thousands of grid-edge nodes, multiple organizational domains, and uneven communication links. Communication-efficient federated systems are needed because repeated transfer of large model updates can create operational burden, especially when edge devices and local controllers have limited bandwidth [1]. The architecture should use hierarchical coordination, local caching, update compression, and scheduled participation windows to reduce unnecessary communication. These mechanisms allow DSOs, microgrids, and EV aggregators to participate at different cadences while remaining part of one shared intelligence process.

Privacy and security requirements

Privacy requirements must be built into the learning protocol through local data retention, secure aggregation, and controlled disclosure of model updates. Smart-grid studies on privacy-preserving distributed learning demonstrate that federated approaches can be aligned with attack detection and load forecasting without exposing participant datasets [10]. Security requirements must also address malicious updates, model poisoning, inference leakage, and unauthorized model access. The architecture therefore treats privacy and security as cross-layer functions that shape data interfaces, model interfaces, and governance interfaces.

Governance and compliance requirements

Governance is required because the architecture spans organizations with different mandates, contractual obligations, and risk tolerances. Blockchain-empowered federated learning literature highlights the need for accountable coordination, verifiable update histories, and mechanisms for trusted collaboration among distributed parties [11]. For grid intelligence, governance must define who may join the federation, what model versions may be deployed, how updates are audited, and how incidents are escalated. Compliance must be operationalized through access policies, model documentation, audit trails, and review procedures rather than treated as a separate administrative layer.

Figure 1 illustrates the proposed layered architecture in which distribution utilities, microgrid operators, and EV aggregators retain local data sovereignty while contributing governed model updates to a shared federated foundation model supported by security, privacy, auditability, and version-control mechanisms.




Figure 1. Layered Federated Foundation-Model Architecture for Multi-Actor Grid Intelligence

Utility Layer

Data sources and preprocessing

The utility layer represents the distribution system operator or distribution utility that manages feeders, substations, grid-edge devices, outages, and regulated customer-facing operations. Local preprocessing converts smart-meter data, feeder measurements, topology records, outage logs, distributed energy resource telemetry, and cybersecurity events into formats suitable for model adaptation. False-data-injection detection research using secure federated deep learning illustrates why utilities need local preprocessing that preserves operational sensitivity while enabling collaborative learning [12]. In this architecture, preprocessing also tags data by operational context so that the foundation model can distinguish planning, monitoring, protection, and customer-service tasks.

Local AI tasks and model usage

The utility layer uses the federated foundation model for feeder-level forecasting, anomaly detection, outage-context interpretation, voltage-control support, and operator decision assistance. Federated reinforcement learning for decentralized voltage control indicates that learning-based grid support can be distributed across control areas without requiring one centralized training environment [13]. The utility does not surrender operational authority to the model; instead, it uses local adapters and prompts to translate model outputs into decision-support artifacts. Human operators and existing control systems remain responsible for final actions in safety-critical workflows.

Interaction with federated model

The utility layer interacts with the federated foundation model by downloading approved model versions, adapting them locally, evaluating them against internal validation checks, and submitting governed updates. Applied Energy work on incentive edge-based federated learning for grid attack detection shows that utilities can participate in collaborative learning while maintaining local operational boundaries [14]. In the proposed architecture, utility updates may include adapter weights, task-specific refinements, metadata summaries, and security attestations rather than raw telemetry. These interactions are routed through the governance/security layer before contributing to the shared model state.

Expected benefits

The utility layer benefits from a shared foundation model because knowledge learned from diverse feeders, operating conditions, and participant types can improve local adaptability without direct data exchange. Privacy-preserving forecasting work shows the value of learning across residential and grid-edge environments while preserving participant

confidentiality [15]. At the architecture level, the utility gains a reusable model interface for many tasks instead of maintaining isolated AI pipelines for each function. **Table 1** summarizes the utility layer’s data sources, AI tasks, compute resources, and interaction points with the federated foundation model.

Table 1. Utility Layer Functions and Federated Model Interaction

Utility-layer element	Architectural role	Representative local data	Local AI task	Compute resource	Federated model interaction
Feeder operations interface	Connects real-time distribution monitoring to local model use	Feeder telemetry, switch states, voltage measurements, outage signals	Situational awareness and anomaly triage	Utility control-center servers and secured edge gateways	Receives approved model versions and submits validated operational adapters
Customer and metering interface	Converts customer-linked measurements into privacy-governed model inputs	Smart-meter intervals, service events, demand patterns	Load forecasting and customer-impact interpretation	Meter data management environment with privacy controls	Applies differential privacy before update release where required
DER coordination interface	Aligns distributed generation and storage observations with grid constraints	Solar output, storage status, inverter events, connection requests	DER hosting support and congestion awareness	DER management platform and local analytics nodes	Shares task-specific refinements without exposing asset-level commercial data
Cyber-operational interface	Links grid cybersecurity signals with operational state	Alarms, authentication logs, false-data indicators, device events	Intrusion and false-data-injection detection	Security operations environment and isolated test sandboxes	Sends security-screened updates with audit metadata
Model governance interface	Ensures utility participation follows policy and compliance obligations	Validation logs, approval records, version histories	Model review and deployment authorization	Governed model registry and access-control service	Accepts only approved federation rounds and records model provenance

Microgrid Layer

Microgrid-specific data and constraints

The microgrid layer represents campus systems, community energy systems, industrial microgrids, and remote grids that combine local generation, storage, controllable loads, and islanding capability. These actors generate data that is more localized than utility data but often richer in controllable flexibility and operational context. Federated renewable-energy research emphasizes that distributed energy applications face site-specific data conditions and privacy needs, making microgrids natural participants in federated architectures [16]. The architecture must therefore allow each microgrid to retain local control policies while contributing learning signals to a broader grid-intelligence model.

AI for local control and optimization

Microgrids use the shared foundation model to support forecasting, local dispatch planning, resilience assessment, islanding preparation, and operator explanation. Autoencoder-based federated anomaly detection in distributed power systems shows how local energy systems can benefit from shared learning while detecting abnormal behavior in their own operating environments [17]. In the proposed architecture, the foundation model is not a replacement for certified protection or control equipment; it is a local intelligence layer that supports planning and supervisory decisions. The microgrid operator can fine-tune local adapters for site-specific objectives such as resilience, cost management, or renewable utilization.

Federated model contribution and usage

The microgrid contributes to the federation through locally trained updates that encode patterns from distributed generation, storage operation, demand flexibility, and islanded operation. Federated anomaly detection studies using recurrent and autoencoder models demonstrate that distributed sites can learn from one another without sharing raw operational sequences [18]. The microgrid receives shared model improvements that may improve local recognition of rare events, renewable variability, and abnormal operating states. Its updates are filtered through privacy, security, and governance checks before being accepted by the federated foundation model layer.

Scalability to many microgrids

Scalability across many microgrids requires the architecture to support lightweight participation, intermittent connectivity, and heterogeneous local objectives. Interpretable federated load forecasting at the edge shows that local energy actors may need models that are both resource-aware and understandable to operators [19]. The proposed architecture therefore allows microgrids to participate through edge gateways, local model caches, adapter-only updates, and flexible synchronization schedules. **Table 2** summarizes the microgrid layer's data, AI tasks, and integration points within the federated grid-intelligence architecture.

Table 2. Microgrid Layer Functions and Integration

Microgrid-layer element	Architectural role	Representative local data	Local AI task	Integration point	Scalability consideration
Local energy resource interface	Represents generation, storage, and controllable load behavior	Solar generation, battery state, controllable load schedules, inverter events	Renewable forecasting and flexibility estimation	Microgrid energy management system	Supports adapter-only participation when compute is limited
Resilience and islanding interface	Captures operating modes under grid-connected and islanded conditions	Islanding events, reserve status, critical-load priorities, restoration logs	Resilience planning and abnormal-state recognition	Local supervisory controller	Allows asynchronous updates after disrupted communication periods
Site-operations interface	Connects model use to operator workflows	Maintenance records, operator notes, alarms, local constraints	Decision support and operator explanation	Local control-room dashboard	Keeps human approval before operational deployment
Privacy-preserving	Converts local learning into	Local gradients or adapters, validation	Federated contribution	Secure federation	Prevents raw operational histories

learning interface	governed federation updates	summaries, privacy parameters	and local adaptation	gateway	from leaving the microgrid
Multi-microgrid coordination interface	Enables learning across diverse sites without direct data exchange	Aggregated model metadata, site-type descriptors, participation status	Cross-site pattern transfer	Federated foundation model layer	Supports many microgrids through hierarchical or scheduled synchronization

EV Aggregator Layer

EV fleet data and flexibility

The EV aggregator layer represents operators that coordinate charging infrastructure, vehicle fleets, customer charging preferences, mobility schedules, and market-facing flexibility offers. EV charging-load prediction research using federated learning shows that charging data is highly distributed across stations, users, and service territories, making centralized data pooling difficult [20]. In the proposed architecture, the aggregator keeps vehicle-level and customer-level data within its own systems while exposing only governed model updates to the federation. This allows the shared foundation model to learn from EV flexibility patterns without directly accessing private mobility or charging histories.

AI for charging and grid services

The EV aggregator uses the foundation model to support charging-demand forecasting, congestion-aware scheduling, flexibility estimation, station-occupancy prediction, and grid-service coordination. Adaptive federated learning for charging-station occupancy prediction demonstrates that EV infrastructure can benefit from collaborative learning while adapting to local station behavior [21]. The architecture treats the aggregator as both a flexibility provider and a learning participant, so model outputs must remain compatible with utility constraints and customer commitments. AI assistance is therefore embedded into scheduling and coordination workflows rather than treated as an autonomous dispatch authority.

Federated model participation

The EV aggregator participates in the federation by locally adapting the shared model to charging sessions, fleet behavior, station utilization, and grid-service response records. Security-enhanced federated transformer learning for vehicle-to-grid dispatch indicates that transformer-based models can support EV-grid coordination under non-identical participant data conditions [22]. In this architecture, the aggregator may submit adapter updates, charging-context embeddings, validation summaries, and security attestations after internal review. The federated model returns improved representations that help the aggregator align charging flexibility with distribution-grid needs.

Dynamic and mobile nature of EVs

EV participation differs from utility and microgrid participation because vehicles are mobile, charging behavior is user-driven, and flexibility availability changes rapidly across time and location. Vertical federated learning for charging-station load prediction in coupled transportation and distribution systems illustrates the need to coordinate information across mobility and power-network domains without collapsing them into one centralized dataset [23]. The architecture therefore allows EV aggregators to maintain mobility-sensitive data boundaries while sharing model improvements through governed federation interfaces. **Table 3** summarizes the EV aggregator layer's data, flexibility services, and interaction with the federated foundation model.

Table 3. EV Aggregator Layer Functions and Federated Model Interaction

EV aggregator-layer element	Architectural role	Representative local data	Flexibility service	Model interaction	Scalability consideration
Charging-session interface	Converts station and vehicle charging records into local model inputs	Plug-in times, charging duration, energy requests, station utilization	Charging-demand forecasting	Fine-tunes local adapters for station and fleet behavior	Supports large numbers of charging points through local aggregation
Mobility-context interface	Represents movement-linked constraints without exposing personal travel histories	Route patterns, depot schedules, availability windows, anonymized location classes	Flexibility availability estimation	Shares privacy-filtered update signals rather than raw mobility traces	Handles changing vehicle membership and location-dependent availability
Grid-service interface	Connects aggregator flexibility to distribution-grid needs	Curtailment requests, vehicle-to-grid events, response confirmations, feeder constraints	Congestion support and demand response	Aligns local scheduling models with utility-facing operational constraints	Enables coordination across many feeders and market zones
Customer-preference interface	Preserves service commitments and user constraints	Departure targets, charging preferences, opt-in status, service agreements	Customer-aware charging control	Uses local prompting and policy adapters to respect user commitments	Keeps customer-sensitive data inside the aggregator environment
Federation gateway	Manages update release and model-version intake	Adapter updates, validation summaries, participation logs, security checks	Governed learning participation	Sends approved updates and receives certified model versions	Allows asynchronous participation during high-mobility operating periods

Federated Foundation Model Layer

Choice of foundation model

The federated foundation model layer hosts a reusable grid-intelligence model that can support forecasting, anomaly interpretation, operational summarization, and actor-specific adaptation. Reviews of large language models for energy systems show that foundation-model capabilities must be aligned with energy-domain constraints, terminology, and safety expectations before being embedded into operational workflows [24]. The architecture therefore favors a transformer-based foundation model pre-trained on broad grid, energy, and operational text-and-time-series representations, then adapted locally by each participant. Its purpose is not to replace existing grid applications but to provide a common representation layer for many downstream tasks.

Federated learning protocol

The federated learning protocol defines how actor layers receive a model version, adapt it locally, evaluate it under internal checks, and submit approved updates. General federated machine learning principles distinguish the

collaboration mechanism from the ownership of raw data, which is essential for this architecture's multi-actor setting [4]. The protocol supports local fine-tuning, prompt-template refinement, adapter updates, and privacy-filtered model contributions. It also allows participation to vary by actor type, since utilities, microgrids, and EV aggregators have different compute environments and operational schedules.

Aggregation and synchronization

Aggregation and synchronization determine how local updates become a shared model improvement without allowing one actor to dominate the federation. Secure and efficient edge-cloud federated learning for smart grids demonstrates the relevance of edge-cloud collaboration when learning must scale across distributed power-system environments [8]. In the proposed architecture, aggregation may occur through hierarchical coordinators, regional federation nodes, or a governed central model service, depending on institutional arrangements. Synchronization is designed to be tolerant of intermittent participation, delayed updates, and actor-specific validation timelines.

Handling heterogeneous data distributions

The federated foundation model layer must handle differences among feeder telemetry, microgrid operations, EV charging behavior, and cybersecurity signals. Personalized federated methods for building-level electricity forecasting show why non-identical data distributions require actor-specific adaptation rather than one uniform model behavior [25]. The architecture addresses heterogeneity through local adapters, task-specific prompts, metadata-aware update review, and optional personalization layers. This design allows the shared backbone to capture transferable grid knowledge while leaving local specialization under participant control.

Model update and versioning

Model update and versioning functions ensure that every deployed model state is identifiable, reviewable, and reversible. Research on privacy-preserving federated learning against poisoning attacks in smart grids highlights the need to inspect updates before they influence shared model behavior [26]. The architecture therefore maintains version histories, participant update records, rollback pathways, and compatibility checks across actor layers. Versioning also supports staged deployment, where a model can be tested locally before being promoted into operational workflows.

Governance/Security Layer

Access control and authorization

The governance/security layer defines who may join the federation, which model functions they may access, and what kinds of updates they may submit. Federated learning security surveys emphasize that access control is central to reducing privacy leakage, unauthorized participation, and manipulation risks [7]. In this architecture, utilities, microgrids, and EV aggregators receive role-specific permissions that reflect their operational responsibilities and contractual rights. Authorization applies not only to data and models but also to prompts, adapters, audit records, and deployment approvals.

Privacy guarantees and data sovereignty

Privacy guarantees are implemented through local data retention, secure aggregation, differential privacy where appropriate, and participant control over update release. Privacy-preserving heterogeneous federated learning for electricity theft detection demonstrates how energy applications can combine collaborative learning with protection for sensitive participant data [27]. Data sovereignty is treated as an architectural principle: each actor decides which local assets are eligible for learning, which updates are releasable, and which validation evidence must accompany participation. The governance layer documents these decisions so privacy is enforceable rather than merely assumed.

Auditability and model traceability

Auditability requires each model update, model version, deployment approval, and incident response action to be traceable. Blockchain-supported federated learning research shows how distributed trust mechanisms can support verifiable collaboration when participants do not share a single administrative domain [11]. In this architecture, audit trails record update provenance, validation outcomes, access events, and version transitions across the utility, microgrid, EV aggregator, and foundation-model layers. Traceability enables regulators, operators, and participating organizations to reconstruct how a model entered service and which controls governed its deployment.

Figure 2 maps the governed lifecycle through which local model adaptations from utilities, microgrids, and EV aggregators are validated, screened, aggregated, versioned, deployed, monitored, and auditable throughout the federated grid-intelligence system.




Figure 2. Governed Model-Update, Versioning, and Auditability Workflow in Federated Grid Intelligence

Adversarial robustness and incident response

Adversarial robustness addresses poisoning, false data injection, model inversion, unauthorized prompting, and attempts to manipulate grid-facing decisions. Federated false-data-injection detection and distributed anomaly-detection studies show that smart-grid learning systems must be designed with cyber-operational threats in mind [28]. The architecture therefore includes update screening, anomaly checks, adversarial testing, rollback procedures, and incident escalation paths before model changes affect operational use. Table 4 summarizes the governance mechanisms, security protocols, and auditability features that protect the federated grid-intelligence architecture.

Table 4. Governance and Security Mechanisms

Governance/security mechanism	Architectural purpose	Protected asset	Operational implementation	Auditability feature
Role-based access control	Restricts model, prompt, adapter, and registry access by participant role	Model interfaces and governance records	Utility, microgrid, aggregator, auditor, and administrator permission profiles	Access events are logged with actor identity and approved role
Secure aggregation	Prevents the federation coordinator from inspecting	Local model updates and participant learning signals	Cryptographic aggregation service between actor gateways and federation layer	Aggregation round records show participants and acceptance status

	individual participant updates			
Differential privacy policy	Limits the disclosure risk of sensitive operational or customer-linked data	Customer behavior, asset state, and local operating patterns	Privacy parameters selected by actor policy and task sensitivity	Privacy settings are attached to update metadata
Model-update screening	Detects malicious, unstable, or policy-violating contributions	Shared foundation model and downstream adapters	Automated anomaly checks, validation gates, and manual review for high-risk updates	Rejected and accepted updates retain review explanations
Version registry and rollback	Maintains approved model states and recovery paths	Deployed model versions and operational continuity	Certified model registry with staged promotion and rollback authority	Version lineage records link updates to deployed releases
Incident-response workflow	Coordinates response to suspected model or security compromise	Federation trust, operational safety, and participant confidence	Escalation rules, temporary suspension, forensic review, and recovery approval	Incident timelines preserve actions, decisions, and responsible parties

Conclusion

This article proposed a scalable grid-intelligence architecture for federating foundation models across distribution utilities, microgrid operators, and EV aggregators. The architecture separates participant responsibilities into actor-specific layers while connecting them through a shared federated foundation model layer. It is designed to preserve data sovereignty while still enabling collaborative learning across heterogeneous grid environments. The result is a conceptual blueprint for shared intelligence without raw data centralization.

The utility layer provides feeder-scale operational knowledge, the microgrid layer contributes localized energy-system intelligence, and the EV aggregator layer contributes mobile flexibility knowledge. Each actor retains control over its data, local model adaptation, validation procedures, and deployment decisions. The federated foundation model layer converts these separate learning environments into a coordinated system for reusable grid intelligence. This layered design addresses the practical reality that modern power systems are distributed institutionally as well as physically.

The governance/security layer is essential because the architecture operates in a safety-critical and multi-organizational environment. Access control, privacy guarantees, audit trails, model versioning, and incident response are not external compliance activities but core architectural functions. These mechanisms ensure that collaboration remains accountable, reviewable, and resilient to manipulation. The architecture therefore combines technical federation with institutional trust management.

Future work should move from conceptual specification to reference implementations, controlled pilots, and operational readiness assessments. Pilot deployments should test the interfaces among utilities, microgrids, aggregators, model registries, and governance services without compromising participant data. Additional work should refine validation procedures for foundation-model behavior in critical grid contexts. The long-term objective is a trustworthy collaborative AI ecosystem that improves grid operation while respecting the autonomy and obligations of every participant.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 29 Sep 2025

Revised: 05 Nov 2025

Accepted: 14 Dec 2025

Published online: 20 March 2026

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Process Mag*. 2020;37(3):50-60.
- Hamann HF, Gjorgiev B, Brunswiler T, Martins LS, Puech A, Varbella A, et al. Foundation models for the electric power grid. *Joule*. 2024;8(12):3245-58.
- Husnoo MA, Anwar A, Hosseinzadeh N, Islam SN, Mahmood AN, Doss R. A secure federated learning framework for residential short-term load forecasting. *IEEE Trans Smart Grid*. 2023;15(2):2044-55.
- Yang Q, Liu Y, Chen T, Tong Y. Federated machine learning: Concept and applications. *ACM Trans Intell Syst Technol (TIST)*. 2019;10(2):1-9.
- Kairouz P, McMahan HB. Advances and open problems in federated learning. *Found Trends Machine Learn*. 2021;14(1-2):1-210.
- Qu X, Guan C, Xie G, Tian Z, Sood K, Sun C, et al. Personalized federated learning for heterogeneous residential load forecasting. *Big Data Min Anal*. 2023;6(4):421-32.
- Mothukuri V, Parizi RM, Pouriyeh S, Huang Y, Dehghantanha A, Srivastava G. A survey on security and privacy of federated learning. *Future Gener Comput Syst*. 2021;115:619-40.

- Su Z, Wang Y, Luan TH, Zhang N, Li F, Chen T, et al. Secure and efficient federated learning for smart grid with edge-cloud collaboration. *IEEE Trans Ind Inform*. 2021;18(2):1333-44.
- Zhang Z, Rath S, Xu J, Xiao T. Federated learning for smart grid: A survey on applications and potential vulnerabilities. *ACM Trans Cyber Phys Syst*. 2026;10(1):1-26.
- Tran HY, Hu J, Yin X, Pota HR. An efficient privacy-enhancing cross-silo federated learning and applications for false data injection attack detection in smart grids. *IEEE Trans Inf Forensics Secur*. 2023;18:2538-52.
- Zhu J, Cao J, Saxena D, Jiang S, Ferradi H. Blockchain-empowered federated learning: Challenges, solutions, and future directions. *ACM Comput Surv*. 2023;55(11):1-31.
- Li Y, Wei X, Li Y, Dong Z, Shahidehpour M. Detection of false data injection attacks in smart grid: A secure federated deep learning approach. *IEEE Trans Smart Grid*. 2022;13(6):4862-72.
- Liu H, Wu W. Federated reinforcement learning for decentralized voltage control in distribution networks. *IEEE Trans Smart Grid*. 2022;13(5):3840-3.
- Lin WT, Chen G, Huang Y. Incentive edge-based federated learning for false data injection attack detection on power grid state estimation: A novel mechanism design approach. *Appl Energy*. 2022;314:118828.
- Dong Y, Wang Y, Gama M, Mustafa MA, Deconinck G, Huang X. Privacy-preserving distributed learning for residential short-term load forecasting. *IEEE Internet Things J*. 2024;11(9):16817-28.
- Grataloup A, Jonas S, Meyer A. A review of federated learning in renewable energy applications: Potential, challenges, and future directions. *Energy AI*. 2024;17:100375.
- Kea K, Han Y, Kim TK. Enhancing anomaly detection in distributed power systems using autoencoder-based federated learning. *PLoS One*. 2023;18(8):e0290337.
- Shrestha R, Mohammadi M, Sinaei S, Salcines A, Pampliega D, Clemente R, et al. Anomaly detection based on LSTM and autoencoders using federated learning in smart electric grid. *J Parallel Distrib Comput*. 2024;193:104951.
- Xu K, Guo J. A multi-source feature extraction network for salient object detection. *Neural Comput Appl*. 2023;35(35):24727-42.
- Yin W, Ji J. Prediction of EV charging load based on federated learning. *Energy*. 2025;316:134559.
- Hallak K, Kem O. Adaptive federated learning framework for predicting EV charging stations occupancy. *Int J Transp Sci Technol*. 2025.
<https://doi.org/10.1016/j.ijtst.2025.04.007>.
- Shang Y, Li S. FedPT-V2G: Security enhanced federated transformer learning for real-time V2G dispatch with non-IID data. *Appl Energy*. 2024 ;358:122626.
- Han Q, Li X. A Vertical federated learning method for electric vehicle charging station load prediction in coupled transportation and power distribution systems. *Processes*. 2025;13(2):468.
- Mirshekali H, Shadi MR, Ladan i FG, Shaker HR. A Review of Large Language Models for Energy Systems: Applications, challenges, and future prospects. *IEEE Access*. 2025;13:163162-88.
- Qin D, Wang C, Wen Q, Chen W, Sun L, Wang Y. Personalized federated DARTS for electricity load forecasting of individual buildings. *IEEE Trans Smart Grid*. 2023;14(6):4888-901.

Li X, Wen M, He S, Lu R, Wang L. A privacy-preserving federated learning scheme against poisoning attacks in smart grid. *IEEE Internet Things J*. 2024;11(9):16805-16.

Wen H, Liu X, Lei B, Yang M, Cheng X, Chen Z. A privacy-preserving heterogeneous federated learning framework with class imbalance learning for electricity theft detection. *Appl Energy*. 2025;378:124789.

Jithish J, Alangot B, Mahalingam N, Yeo KS. Distributed anomaly detection in smart grids: a federated learning-based approach. *IEEE Access*. 2023;11:7157-79.