HEDGE-IoT:

A Multi-Dimensional Digital Framework for Intelligent IoT Deployment Across the Energy System Cloud-Edge Continuum



HEDGE-IoT: A Multi-Dimensional Digital Framework for Intelligent IoT Deployment Across the Energy System Cloud-Edge Continuum

The Publicator using Inferact/Qwen3.8-Flash-Next-NVFP4

Contents

HEDGE-IoT: A Multi-Dimensional Digital Framework for Intelligent IoT Deployment Across the Energy System Cloud-Edge Continuum **3**
1. Introduction and Motivation . . . 4
1.1 Challenges of IoT Integration Across the Energy System Hierarchy . . . 4
1.2 The Imperative for Intelligence at the Edge and Cloud . . . 4
1.3 Untapped Flexibility in RES-Hosting Grids . . . 5
1.4 Objectives of the HEDGE-IoT Framework . . . 5
1.5 Contributions and Paper Structure . . . 6
2. State of the Art and Related Work . . . 7
2.1 IoT Frameworks and Architectures in Energy Systems . . . 7
2.2 Edge and Cloud Computing Paradigms for Energy IoT . . . 7
2.3 Federated Learning and Distributed Intelligence in Power Grids . . . 8
2.4 Interoperability Architectures and Data Spaces . . . 8
2.5 EU Digitalisation and IoT Initiatives . . . 9
2.6 Identified Gaps and the HEDGE-IoT Positioning . . . 9
3. HEDGE-IoT Framework Architecture . . . 11
3.1 Architectural Philosophy and Multi-Dimensional Design Rationale . . . 11
3.2 The Cloud/Edge Continuum: A Vertically Integrated Computational Substrate . . . 11
3.3 The Four Pillars: Mutually Reinforcing Functional Dimensions . . . 12
3.4 The Computational Orchestration Layer: Bridging Federated Applications Across Distributed Data Platforms . . . 13
3.5 Cross-Pillar Integration and End-to-End Data Flow . . . 14
3.6 Architectural Principles and Design Invariants . . . 14
4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge 16
4.1 Computational Sharing and Edge Offloading Mechanisms . . . 16
4.2 Federated Learning for Privacy-Preserving Energy Analytics . . . 17
4.3 Swarm Computing Paradigms for Distributed Decision-Making . . . 18
4.4 Orchestration of Heterogeneous IoT Workloads . . . 20
4.5 Integration with Adjacent Pillars and Design Invariants . . . 21
4.6 Summary of the Technology Facilitator Pillar . . . 22
5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration . . . 23
5.1 Data Space Architecture: Principles and Energy-Specific Adaptation . . . 23
5.2 Semantic Interoperability: From Raw Telemetry to Market-Ready Signals . . . 23
5.3 Data Governance Models: Sovereignty, Policy Enforcement, and Trust . . . 24
5.4 Cross-Domain Integration: Bridging BTM, DSO, and TSO Layers . . . 25
5.5 Data Space Connector Provisioning and New-Asset Onboarding . . . 26
5.6 Cross-Data-Space Federation and Federated Application Mediation . . . 27
5.7 Integration with the Computational Orchestration Layer . . . 27
5.8 Compliance, Ethics, and Inclusive Data Access . . . 28
5.9 Summary and Cross-Pillar Dependencies . . . 29
6. Standardisation Pillar: SAREF and Open Standards for Energy IoT . . . 30
6.1 Standardisation as a First-Class Architectural Pillar . . . 30
6.2 SAREF as the Semantic Backbone of the Energy IoT Continuum . . . 30
6.3 Standardised Data Formats and Wire Protocols . . . 31
6.4 Open API Specifications and Interface Contracts . . . 31
6.5 Alignment with Ongoing IoT and Energy Standardisation Efforts . . . 32
6.6 Vendor Neutrality, Conformance, and Long-Term Interoperability . . . 33
6.7 Standardisation Governance and Evolution Process . . . 33

6.8 Standardisation in Service of the Other Pillars . . . 34
7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities 36
7.1 Positioning of the Ecosystem Enabling Pillar Within the HEDGE-IoT Architecture . . . 36
7.2 Increasing RES-Hosting Capacity Through Activated Flexibility . . . 36
7.3 Enhancing Grid Resilience Through Diversified and Interoperable Energy Services . . . 37
7.4 Unlocking New Market Opportunities and Flexibility Monetisation . . . 39
7.5 Scalability of Distributed Data Platforms and Infrastructure for Highly Distributed Energy Services . . . 40
7.6 Ecosystem Governance, Open Call, and Inclusive Participation . . . 42
7.7 Cross-Pillar Dependencies and Emergent Ecosystem Value . . . 42
8. Stakeholder Engagement, IoT Ethics, and Inclusivity . . . 44
8.1 Ethical Framework: IoT Ethics as an Architectural Constraint, Not an Add-On . . . 44
8.2 Cultivating End-User Trust: From Information Asymmetry to Relational Assurance . . . 45
8.3 Inclusivity and Energy-Poverty Safeguards: Designing for the Most Vulnerable . . . 45
8.4 EU Liaison and Standards Alignment: Embedding HEDGE-IoT in the European Policy Architecture . . . 46
8.5 The Open Call Mechanism: Structurally Eliminating Barriers for SMEs and Innovators . . . 47
8.6 Multi-Stakeholder Governance and Continuous Ethical Oversight . . . 48
8.7 Synthesis: Ethics and Inclusivity as Enablers of Ecosystem Viability . . . 49
9. Scalability, Replicability, and Validation . . . 50
9.1 Pilot Deployment Architecture and Methodology . . . 50
9.2 Scalability Assessment Across Energy System Tiers . . . 50
9.3 Flexibility Activation Performance . . . 51
9.4 RES Integration Improvement . . . 52
9.5 Grid Resilience Metrics . . . 52
9.6 Federated Application Effectiveness Under Varying Load and Network Conditions . . . 53
9.7 Replicability Assessment Across Regulatory, Climatic, and Asset-Mix Environments . . . 54
9.8 Limitations, Validation Gaps, and Pathways to Continental Scaling . . . 55
10. Conclusions and Future Work . . . 57
10.1 Summary of Contributions . . . 57
10.2 Demonstrated Impact on Energy System Intelligence and Flexibility . . . 57
10.3 Limitations and Open Challenges . . . 58
10.4 Future Research Directions . . . 58
10.5 Closing Remarks . . . 62

HEDGE-IoT: A Multi-Dimensional Digital Framework for Intelligent IoT Deployment Across the Energy System Cloud-Edge Continuum

Abstract: The integration of Internet of Things (IoT) assets across the full spectrum of the energy system - from behind-the-meter consumers to transmission system operators - remains a critical challenge, particularly in harnessing the untapped flexibility potential of grids hosting high shares of renewable energy sources (RES). This paper presents HEDGE-IoT, a multi-dimensional digital framework designed to enable intelligent IoT deployment across the energy system cloud-edge continuum. The framework is structured around four interdependent pillars: a Technology Facilitator pillar leveraging AI/ML federated learning and swarm computing for privacy-preserving, distributed decision-making at the grid edge; an Interoperability pillar adopting Data Space architectures to enable seamless cross-domain data exchange among heterogeneous IoT platforms and stakeholders; a Standardisation pillar grounded in SAREF and open standards to ensure vendor-neutral, long-term interoperability; and a Digital Energy Ecosystem Enabling pillar that increases RES-hosting capacity, enhances grid resilience, and unlocks new market opportunities. A computational orchestration layer bridges federated applications across distributed data platforms, coordinating workloads across heterogeneous IoT assets. The framework's scalability, replicability, and performance are validated through pilot deployments spanning multiple energy system levels, evaluating flexibility activation, RES integration improvement, and federated application effectiveness under varying load and network conditions. Ethical considerations, stakeholder engagement, and inclusivity mechanisms - including liaison with EU initiatives such as AIOTI and an Open Call for SMEs - are integral to the framework's design. The paper concludes by outlining future research directions in advanced orchestration algorithms, expanded standardisation, deeper AI/ML integration, and broader European deployment pathways.

1. Introduction and Motivation

The global energy landscape is undergoing a structural transformation of unprecedented scale. The decarbonisation imperative, codified in the European Green Deal and the REPowerEU plan, demands that electricity networks evolve from passive, centrally controlled infrastructures into active, bidirectional platforms capable of accommodating millions of distributed energy resources (DERs), electric vehicles, smart appliances, and prosumer assets. At the heart of this transition lies the Internet of Things (IoT): a heterogeneous constellation of sensors, actuators, smart meters, inverters, battery management systems, and edge gateways that collectively generate, consume, and exchange data at rates and latencies that legacy energy management systems were never designed to handle. Yet the mere proliferation of connected devices does not, by itself, yield a smarter grid. The critical question confronting researchers, system operators, and policymakers alike is how to orchestrate intelligence across the full cloud-edge continuum so that IoT assets become genuine enablers of flexibility, resilience, and market participation rather than sources of fragmentation and operational risk.

This challenge is compounded by the vertical stratification of the energy system itself. Behind-the-meter (BTM) installations - residential heat pumps, rooftop photovoltaic arrays, home energy management systems - operate under fundamentally different constraints, data models, and governance regimes than distribution-level assets (medium-voltage substations, distribution-level storage, demand-response aggregators), which in turn differ from transmission-level resources managed by Transmission System Operators (TSOs). Each stratum possesses its own latency tolerances, cybersecurity postures, regulatory obligations, and economic incentive structures. A framework that addresses only one layer, or that treats the hierarchy as a simple data pipeline, will inevitably fail to unlock the synergies that arise when intelligence is co-located with the physical assets it governs while remaining interoperable with higher-level coordination functions.

1.1 Challenges of IoT Integration Across the Energy System Hierarchy

The integration of IoT assets across the BTM-to-TSO spectrum presents a multi-dimensional problem that resists reduction to a single technical or organisational axis. First, **heterogeneity** pervades every layer: communication protocols range from Zigbee and Z-Wave in domestic settings to IEC 61850 and DNP3 in substation automation, while data semantics diverge even among devices that nominally serve the same function. Second, **scalability** is non-linear; a distribution feeder that hosts a few hundred smart meters today may host tens of thousands of bidirectional EV chargers and virtual power plant (VPP) participants within a decade, overwhelming centralised data-collection architectures. Third, **latency and determinism** requirements vary by orders of magnitude - protection relays demand sub-millisecond responses, whereas day-ahead market clearing tolerates minute-level granularity - yet all must coexist within a shared digital infrastructure. Fourth, **data sovereignty and privacy** concerns intensify as granular consumption and generation profiles become available to aggregators, retailers, and public authorities, raising compliance obligations under the General Data Protection Regulation (GDPR) and the EU Data Governance Act. Fifth, **vendor lock-in and proprietary silos** continue to impede cross-platform orchestration, as incumbent energy-management platforms expose limited or non-standardised interfaces.

These challenges are not merely technical. They are institutional, economic, and regulatory. Distribution System Operators (DSOs) in many Member States lack the digital tooling to visualise and act upon BTM flexibility in real time, while TSOs struggle to incorporate aggregated distribution-level resources into congestion management and ancillary-service procurement. The absence of a common digital language across these layers means that flexibility signals generated at the edge are frequently lost, distorted, or arrive too late to be actionable at the cloud or control-centre level.

1.2 The Imperative for Intelligence at the Edge and Cloud

Conventional cloud-centric architectures, in which raw telemetry is streamed to a centralised data lake for batch analytics, are ill-suited to the temporal and spatial characteristics of modern energy IoT. The volume of data produced by a single distribution feeder - potentially exceeding several terabytes per day when high-frequency smart-meter readings, inverter telemetry, and environmental sensor streams are aggregated - renders continuous cloud ingestion both economically prohibitive and operationally fragile, given the dependence on wide-area network availability. Moreover, applications such as local voltage regulation, islanding detection, and real-time demand-response dispatch require decision latencies that a round-trip to a hyperscale data centre cannot guarantee.

The HEDGE-IoT Framework therefore posits that intelligence must be **distributed across the cloud-edge continuum** in a purposeful, orchestrated manner. At the grid edge, lightweight machine-learning inference, swarm-based consensus algorithms, and rule-based controllers execute time-critical tasks in proximity to the

physical assets they govern, preserving privacy by design and minimising bandwidth consumption. At the cloud layer, federated learning models aggregate knowledge from thousands of edge nodes without centralising raw data, enabling system-wide pattern recognition, predictive maintenance, and market-clearing optimisation. The orchestration layer that bridges these two domains - detailed in the *HEDGE-IoT Framework Architecture* - coordinates workload placement, model synchronisation, and policy enforcement so that the right computation occurs at the right place and the right time. This paradigm shift from "data to the cloud" to "intelligence to the data" is not an incremental optimisation; it is a prerequisite for operating energy networks that are simultaneously more digitalised, more decentralised, and more resilient.

1.3 Untapped Flexibility in RES-Hosting Grids

The rapid penetration of variable renewable energy sources (RES) - predominantly solar photovoltaics and onshore wind - has transformed distribution networks from unidirectional power-delivery systems into complex, multi-directional energy ecosystems. In many European regions, instantaneous RES generation already exceeds local demand by a factor of three or more during midday peaks, while evening ramps impose severe stress on feeder capacity. The physical infrastructure to reinforce these networks at the pace required by RES deployment targets is neither financially nor logistically feasible within the mandated timelines. Flexibility - defined as the ability to shift, curtail, or reshape consumption and generation profiles in response to network conditions - emerges as the most cost-effective alternative to copper.

Yet the flexibility potential embedded in RES-hosting grids remains largely untapped. Aggregated residential storage, EV-to-grid (V2G) capabilities, smart thermal loads, and curtailable industrial processes collectively represent gigawatt-scale dispatchable capacity in many Member States, but this capacity is fragmented across thousands of small operators, governed by incompatible contracts, and invisible to the DSOs and TSOs who need it most. The *Digital Energy Ecosystem Enabling Pillar* of the HEDGE-IoT Framework addresses this gap by providing the digital substrate - scalable distributed data platforms, interoperable service interfaces, and market-ready data products - that converts latent flexibility into a tradable, dispatchable resource. Crucially, this conversion must occur without imposing prohibitive integration costs on small prosumers or creating new barriers to entry for emerging energy-service providers.

1.4 Objectives of the HEDGE-IoT Framework

The overarching objective of the HEDGE-IoT Framework is to provide a **multi-dimensional, open, and replicable digital architecture** that enables intelligent IoT deployment across the entire energy system cloud-edge continuum, from behind-the-meter devices to TSO control centres. This overarching goal decomposes into the following specific objectives:

1. **Architectural coherence.** To define a reference architecture - articulated in the *HEDGE-IoT Framework Architecture* - that unifies the cloud/edge continuum under a computational orchestration layer, ensuring that federated applications can be deployed, migrated, and scaled across heterogeneous IoT assets without bespoke integration for each deployment site.

2. **Edge intelligence and computational sharing.** To develop and validate AI/ML federated learning models and swarm-computing paradigms that offload energy analytics and distributed decision-making to the grid edge, as elaborated in the *Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*, thereby reducing latency, preserving data privacy, and lowering communication overhead.

3. **Cross-domain interoperability.** To adopt and extend leading-edge Data Space architectures and semantic interoperability mechanisms, as described in the *Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*, so that diverse IoT platforms, energy services, and stakeholder communities can exchange data under governed, trust-based agreements.

4. **Open standardisation.** To anchor all framework components in widely adopted open standards - most notably SAREF and associated API specifications, as detailed in the *Standardisation Pillar: SAREF and Open Standards for Energy IoT* - ensuring vendor neutrality, long-term maintainability, and alignment with ongoing European IoT standardisation efforts.

5. **Ecosystem enablement and market activation.** To create the digital conditions under which RES-hosting capacity is increased, grid resilience is enhanced through diversified and interoperable energy services, and new market opportunities are unlocked for prosumers, aggregators, and SMEs, as outlined in the *Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*.

6. **Ethical, inclusive, and stakeholder-centred deployment.** To embed IoT ethics, user trust, and inclusivity into every layer of the framework, and to establish structured liaison with EU initiatives

such as AIOTI and an Open Call mechanism for engaging innovators, as addressed in the *Stakeholder Engagement, IoT Ethics, and Inclusivity* section.

7. **Empirical validation and replicability.** To demonstrate the framework's effectiveness through pilot deployments at multiple energy-system levels, evaluating flexibility activation, RES integration improvement, grid resilience metrics, and federated-application performance under varying load and network conditions, as reported in the *Scalability, Replicability, and Validation* section.

1.5 Contributions and Paper Structure

This paper makes the following principal contributions to the state of the art in energy-system digitalisation and IoT deployment:

- **A multi-dimensional reference architecture** that explicitly couples the cloud/edge continuum with four mutually reinforcing pillars - Technology Facilitator, Interoperability, Standardisation, and Digital Energy Ecosystem Enabling - under a unified computational orchestration layer. Unlike prior frameworks that address connectivity, data management, or analytics in isolation, HEDGE-IoT treats these dimensions as co-dependent and jointly optimised.
- **Novel federated learning and swarm-computing mechanisms** tailored to the energy IoT domain, enabling privacy-preserving analytics and distributed decision-making at the grid edge while maintaining convergence guarantees under the intermittent connectivity and heterogeneous hardware typical of energy infrastructure.
- **A Data Space-based interoperability blueprint** that bridges BTM, distribution, and transmission layers through semantic harmonisation, governed data-sharing agreements, and cross-domain API specifications, reducing integration effort from months to days for new IoT asset classes.
- **A standardisation strategy** centred on SAREF and open API specifications that ensures long-term vendor neutrality and positions the framework for alignment with forthcoming EU digital-energy regulations.
- **Empirical evidence from multi-level pilot deployments** quantifying the framework's impact on flexibility activation, RES hosting capacity, and grid resilience, together with a replicability assessment that supports transfer to additional European contexts.

The remainder of this paper is organised as follows. Section 2, *State of the Art and Related Work*, reviews existing IoT frameworks in energy systems, edge/cloud computing paradigms, federated learning applications in power grids, interoperability architectures, and relevant EU digitalisation initiatives, identifying the specific gaps that HEDGE-IoT addresses. Section 3, *HEDGE-IoT Framework Architecture*, presents the overall multi-dimensional architecture, the four pillars, and the computational orchestration layer. Sections 4 through 7 elaborate each pillar in detail: the *Technology Facilitator Pillar* (Section 4), the *Interoperability Pillar* (Section 5), the *Standardisation Pillar* (Section 6), and the *Digital Energy Ecosystem Enabling Pillar* (Section 7). Section 8, *Stakeholder Engagement, IoT Ethics, and Inclusivity*, addresses the socio-technical dimensions of deployment. Section 9, *Scalability, Replicability, and Validation*, reports pilot results and performance evaluations. Section 10, *Conclusions and Future Work*, summarises the key contributions and outlines future research directions, including advanced orchestration algorithms, expanded standardisation scope, deeper AI/ML integration, and broader European deployment pathways.

2. State of the Art and Related Work

2.1 IoT Frameworks and Architectures in Energy Systems

The proliferation of Internet of Things (IoT) devices across the energy value chain - from behind-the-meter (BTM) smart meters and residential energy management systems to distribution-level sensors and transmission-level phasor measurement units (PMUs) - has prompted a wave of reference architectures and platform frameworks over the past decade. Early efforts such as the FIWARE-based energy platforms (e.g., the FIWARE Energy Community) and the OpenFMB (Open Field Message Bus) standard, developed under the OpenFMB Consortium, addressed device connectivity and messaging at the distribution level but remained largely confined to single-layer deployments. The OpenFMB specification, for instance, provides a publish-subscribe messaging model tailored to distribution automation, yet it does not extend natively to BTM asset orchestration or TSO-level market coordination, leaving a vertical integration gap that HEDGE-IoT explicitly targets through its multi-dimensional architecture (see *3. HEDGE-IoT Framework Architecture*).

At the European level, the SAREF (Smart Applications on Energy) ontology family - SAREF4ENER, SAREF4BLDG, SAREF4EV, and SAREF4GRID - has emerged as the de facto semantic backbone for energy IoT data modelling under the auspices of ETSI. While SAREF provides a common vocabulary for device capabilities, energy flows, and grid topology, its adoption in production systems remains fragmented: most deployments use SAREF as a translation layer bolted onto proprietary vendor stacks rather than as a native data model, perpetuating the vendor lock-in barrier identified in *1. Introduction and Motivation.* Complementary efforts such as the IEC 61850-7-420 extension for distributed energy resources (DER) and the IEEE 2030.5 (SEP 2) protocol for demand response messaging address specific communication needs but do not constitute holistic frameworks encompassing data governance, edge intelligence, and market enablement simultaneously.

Commercial IoT platforms - including AWS IoT Greengrass, Azure IoT Edge, and Siemens MindSphere - offer device management, rule engines, and cloud analytics pipelines. However, their energy-specific capabilities are typically limited to dashboard visualisation and basic threshold alerting. They lack native support for energy-domain semantics, grid-aware scheduling constraints, or the multi-stakeholder data-sharing models required for cross-domain flexibility activation. Academic prototypes such as the GridEye platform (University of Stuttgart) and the OpenADR 2.0b ecosystem demonstrate proof-of-concept demand-response orchestration but have not achieved the architectural breadth to unify connectivity, computation, interoperability, and market access under a single coherent reference model.

2.2 Edge and Cloud Computing Paradigms for Energy IoT

The recognition that cloud-centric architectures are structurally inadequate for distribution-scale IoT telemetry - both economically, due to bandwidth and storage costs at millions-of-devices scale, and operationally, due to latency sensitivity of voltage regulation, islanding detection, and real-time demand response - has driven substantial research into edge and fog computing for energy systems. The European Commission's Strategic Research Agenda for IoT (2014, updated 2020) and the subsequent Horizon Europe Digital Europe Programme have funded numerous projects (e.g., SENSE4US, VISION, and the ongoing EDGE4CLOUD) that explore hierarchical edge-cloud task offloading for smart grid monitoring.

In the fog-computing literature, architectures such as the OpenFog Reference Architecture (OpenFog Consortium, 2017) and the ETSI Multi-access Edge Computing (MEC) standard provide general-purpose frameworks for placing computation closer to data sources. Applied to energy, projects like the EU-funded SENSE4US demonstrated fog-node deployment for real-time power quality monitoring, while the US DOE's DER-CAM and OpenDSS-based simulation environments have explored edge-level DER control. However, these paradigms typically assume homogeneous edge nodes with reliable backhaul links - an assumption that breaks down in rural RES-hosting networks where connectivity is intermittent and heterogeneous (cellular, LoRaWAN, PLC, satellite).

The concept of "swarm computing" at the grid edge - whereby neighbouring IoT devices collectively execute distributed decision-making without a central orchestrator - has been explored in robotics and vehicular ad-hoc networks but remains nascent in energy IoT. Preliminary work by the IEEE P2413 working group on IoT architecture and the IEC TC 57/SC 57B activities on distributed control acknowledge the need for peer-to-peer coordination among DERs, yet no production-grade framework integrates swarm intelligence with the cloud-layer federated aggregation required for system-wide analytics. HEDGE-IoT's computational orchestration layer, described in *3. HEDGE-IoT Framework Architecture*, is designed to bridge precisely this gap by coordinating workloads across heterogeneous IoT assets with convergence guarantees under intermittent connectivity, as

detailed in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*.

2.3 Federated Learning and Distributed Intelligence in Power Grids

Federated learning (FL), introduced by McMahan et al. (2017) for privacy-preserving model training on mobile devices, has seen rapid adoption in energy analytics since 2020. In the power-grid context, FL enables DSOs and TSOs to train predictive models (load forecasting, anomaly detection, voltage violation prediction) on distributed customer data without centralising raw telemetry, thereby aligning with GDPR and the EU Data Governance Act (DGA) constraints on data sovereignty.

Notable applications include: (i) the EU-funded FLAME project, which demonstrated FL for residential load disaggregation across smart-meter populations; (ii) work by Zhang et al. (2022) on FL-based non-intrusive load monitoring (NILM) for BTM flexibility identification; (iii) the IEEE PES Task Force on Privacy-Preserving Analytics, which surveyed FL variants (FedAvg, FedProx, SCAFFOLD) for grid state estimation; and (iv) the UK-funded FLEX project, which applied FL to aggregate EV charging patterns for DSO congestion management.

Despite this momentum, several limitations persist in the current FL-for-energy literature. First, most studies assume synchronous participation of all edge nodes and stable communication links - conditions rarely met in practice where BTM devices exhibit high churn, duty-cycled radios, and variable connectivity. Second, existing FL implementations in energy contexts are typically single-task (e.g., forecasting only) and do not support the multi-task, multi-stakeholder orchestration required when a single edge gateway simultaneously serves a DSO's voltage-control model, a retailer's demand-response model, and a TSO's RES-forecasting model. Third, convergence guarantees under the highly non-IID (non-independent and identically distributed) data distributions characteristic of heterogeneous energy consumers (residential vs. industrial vs. commercial) remain an open theoretical challenge. Fourth, the integration of FL with swarm-computing paradigms - whereby edge nodes perform local inference and peer-to-peer model refinement before periodic cloud aggregation - has not been formalised for energy grid topologies.

HEDGE-IoT addresses these limitations through energy-specific federated and swarm mechanisms with convergence guarantees under intermittent connectivity, as elaborated in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*, and through the computational orchestration layer that coordinates multi-task workloads across the cloud/edge continuum.

2.4 Interoperability Architectures and Data Spaces

Cross-domain data exchange in energy systems has historically relied on point-to-point integrations, proprietary APIs, and bilateral data-sharing agreements - approaches that scale poorly and create the "data silo" problem that obstructs flexibility activation. The emergence of Data Space architectures, championed by the European Data Strategy (2020) and operationalised through the Gaia-X initiative and the International Data Spaces Association (IDSA) Reference Architecture Model (RAM), offers a paradigm shift: decentralised, sovereignty-preserving data exchange governed by contractual usage policies rather than centralised data lakes.

In the energy domain, the EU-funded Catena-X project (automotive) and the Energy Data Space (EDS) initiative under the European Commission's Smart Grids Task Force have begun exploring Data Space implementations for energy data. The IDSA Connector framework provides a technical building block for policy-enforced data sharing, while the European Health Data Space (EHDS) and the proposed European Energy Data Space regulation signal regulatory momentum. However, energy-specific Data Space implementations remain at an early stage: the EDS pilot (2023-2025) focuses primarily on wholesale market data and does not yet encompass BTM IoT telemetry, DSO operational data, or TSO balancing information in a unified governance model.

Semantic interoperability - the ability of heterogeneous systems to interpret shared data consistently - remains a critical bottleneck. While SAREF provides a common ontology, its mapping to operational protocols (IEC 61850, OpenADR, MQTT-Sparkplug, OCPP for EV charging) is incomplete, and no single framework currently provides end-to-end semantic bridging from a smart meter's raw register values through a DSO's SCADA system to a TSO's market platform. The W3C Web of Things (WoT) Thing Description and the oneM2M standard offer device-level semantic descriptions but lack energy-domain depth.

HEDGE-IoT's Interoperability Pillar, detailed in *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*, proposes a Data Space blueprint specifically designed for the energy cloud/edge continuum, reducing new-asset integration timelines from months to days through SAREF-anchored semantic mediation and policy-driven data governance.

2.5 EU Digitalisation and IoT Initiatives

The European Union's digitalisation agenda for the energy sector is anchored in several overlapping policy and standardisation initiatives. The AIOTI (Alliance for Internet of Things Innovation), established in 2013 and now operating under the European IoT Association, serves as the primary multi-stakeholder forum for IoT research, innovation, and standardisation in Europe. AIOTI's Strategic Research and Innovation Agenda (SRIA) identifies energy as a priority vertical and has published position papers on IoT for smart grids, edge computing, and data spaces. HEDGE-IoT's liaison activities with AIOTI, described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*, ensure alignment with the evolving AIOTI SRIA and contribute to its working groups on energy IoT and edge intelligence.

The European Commission's Digital Europe Programme (2021-2027) and Horizon Europe Cluster 5 (Climate, Energy and Mobility) fund numerous IoT-for-energy projects. The Smart Grids Task Force (SGTF) under the European Commission provides regulatory guidance on data access, interoperability, and cybersecurity for energy IoT. The EU Cybersecurity Act (2019) and the NIS2 Directive (2022) impose security requirements on energy IoT deployments that any reference architecture must accommodate. The EU AI Act (2024) introduces risk-based obligations for AI systems deployed in critical infrastructure, directly relevant to the federated learning and swarm-computing components of HEDGE-IoT.

At the standardisation level, ETSI TC SmartM2M, CEN/CENELEC JTC 21 (Smart Grids), and the IEC TC 57/SC 57B (Distribution Automation) collectively shape the technical standards landscape. The EU's Interoperability Framework (EIF) and the European Interoperability Reference Architecture (EIRA) provide governance-level guidance for cross-sector data exchange. The European Data Innovation Board (EDIB) oversees the implementation of the Data Governance Act and the proposed Data Act, both of which directly constrain how energy IoT data may be shared across stakeholders.

Despite this rich ecosystem of initiatives, a persistent gap exists between policy-level frameworks (which define *what* should be achieved) and deployable technical architectures (which define *how* to achieve it at the device, edge, and cloud layers simultaneously). HEDGE-IoT positions itself as the missing technical reference architecture that operationalises EU digitalisation objectives - data sovereignty, interoperability, AI trustworthiness, and market enablement - into a concrete, pilot-validated framework spanning the full energy system continuum.

2.6 Identified Gaps and the HEDGE-IoT Positioning

Synthesising the review above, five interlocking gaps emerge that existing frameworks and initiatives do not jointly address:

Gap 1 - Architectural fragmentation. IoT frameworks in energy systems treat connectivity, data management, analytics, and market access as separable concerns deployed in isolation. No existing reference architecture unifies the cloud/edge continuum under a single computational orchestration layer while simultaneously addressing the Technology Facilitator, Interoperability, Standardisation, and Digital Energy Ecosystem Enabling dimensions. HEDGE-IoT's four-pillar architecture (see *3. HEDGE-IoT Framework Architecture*) is the first to jointly optimise all four dimensions.

Gap 2 - Edge intelligence under realistic constraints. Current FL and edge-computing research for energy assumes synchronous participation, stable connectivity, and single-task workloads. Energy-specific federated and swarm mechanisms with formal convergence guarantees under intermittent, heterogeneous connectivity - where a single edge node may serve multiple stakeholders' models concurrently - remain absent from the literature. HEDGE-IoT's Technology Facilitator Pillar (see *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*) introduces these mechanisms.

Gap 3 - End-to-end interoperability from BTM to TSO. Data Space architectures (IDSA, Gaia-X) and semantic standards (SAREF, WoT) exist in parallel but are not integrated into a deployable, energy-specific blueprint that spans BTM IoT devices, DSO operational systems, and TSO market platforms with policy-enforced data governance. HEDGE-IoT's Interoperability Pillar (see *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*) provides this blueprint.

Gap 4 - Standardisation as an enabler, not an afterthought. SAREF and related standards are typically retrofitted onto proprietary platforms. A framework that embeds SAREF-anchored open standardisation as a first-class architectural pillar - ensuring vendor neutrality and long-term interoperability by design - has not been demonstrated in production energy IoT deployments. HEDGE-IoT's Standardisation Pillar (see *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) addresses this.

Gap 5 - Ecosystem and market enablement for RES-hosting flexibility. The GW-scale dispatchable capacity latent in residential storage, V2G, smart thermal loads, and curtailable industrial processes remains

fragmented and invisible to system operators due to incompatible contracts and proprietary platforms. No existing IoT framework provides the digital ecosystem infrastructure - scalable distributed data platforms, market-ready flexibility aggregation, and resilience-enhancing service diversification - to convert this latent capacity into a tradable resource. HEDGE-IoT's Digital Energy Ecosystem Enabling Pillar (see *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*) is designed to unlock this potential.

Additionally, the ethical and inclusivity dimensions of energy IoT - ensuring that digitalisation does not exacerbate energy poverty, that AI-driven flexibility activation respects user autonomy, and that SMEs and innovators can participate in the ecosystem on equitable terms - are underrepresented in existing technical frameworks. HEDGE-IoT embeds IoT ethics and inclusivity as cross-cutting concerns (see *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*), aligning with the EU AI Act's risk-based approach and the European Pillar of Social Rights.

Finally, while numerous EU-funded projects have demonstrated individual components (FL for load forecasting, Data Spaces for market data, SAREF for device modelling), none has provided empirical, multi-level validation across BTM, DSO, and TSO deployments with replicability assessment. HEDGE-IoT's pilot-validated evidence of impact on flexibility activation, RES hosting, and grid resilience (see *9. Scalability, Replicability, and Validation*) fills this empirical gap, offering the replicability assessment that EU funding bodies and system operators increasingly require before scaling digital solutions.

In summary, HEDGE-IoT does not merely extend an existing framework along a single axis; it constitutes the first multi-dimensional, open reference architecture that simultaneously addresses the five barriers identified in *1. Introduction and Motivation* - protocol/semantic heterogeneity, non-linear scalability, divergent latency requirements, data-sovereignty constraints, and vendor lock-in - through mutually reinforcing pillars validated across the full energy system cloud/edge continuum.

3. HEDGE-IoT Framework Architecture

3.1 Architectural Philosophy and Multi-Dimensional Design Rationale

The HEDGE-IoT Framework is conceived as a multi-dimensional, open reference architecture that unifies heterogeneous IoT assets - from behind-the-meter (BTM) residential devices through distribution-system-operator (DSO) substations to transmission-system-operator (TSO) control centres - under a single coherent design logic. As established in *1. Introduction and Motivation*, the integration of IoT across the energy system is obstructed by five interlocking barriers: protocol and semantic heterogeneity, non-linear scalability demands, divergent latency requirements, data-sovereignty constraints (GDPR, Data Governance Act), and persistent vendor lock-in. No single-layer or single-axis intervention can resolve these simultaneously. The HEDGE-IoT architecture therefore adopts a **four-pillar, one-continuum** design: four mutually reinforcing functional pillars (Technology Facilitator, Interoperability, Standardisation, Digital Energy Ecosystem Enabling) are instantiated across a vertically integrated cloud/edge continuum, and are bound together by a horizontal **computational orchestration layer** that bridges federated applications across distributed data platforms.

This multi-dimensional philosophy directly addresses the five gaps identified in *2. State of the Art and Related Work*: (i) architectural fragmentation across four dimensions; (ii) absence of edge intelligence with convergence guarantees under realistic connectivity; (iii) lack of end-to-end BTM-to-TSO Data Space integration; (iv) standardisation treated as an afterthought rather than a first-class pillar; and (v) the absence of a framework that converts latent GW-scale RES-hosting flexibility into a market-ready, tradable resource. HEDGE-IoT is the first framework to jointly address all five within a single architectural envelope.

The architecture is deliberately **open and extensible**. It does not prescribe a single vendor stack, a monolithic data lake, or a fixed set of applications. Instead, it defines interfaces, contracts, and governance rules that allow new IoT assets, analytics services, and market mechanisms to be integrated without re-architecting the underlying platform. This openness is operationalised through the Standardisation Pillar (see *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) and the ecosystem-enablement mechanisms described in *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*.

3.2 The Cloud/Edge Continuum: A Vertically Integrated Computational Substrate

The HEDGE-IoT continuum spans five logical tiers, each optimised for a distinct class of computational, latency, and data-sovereignty requirements. The design explicitly rejects the cloud-centric paradigm - critiqued in *1. Introduction and Motivation* as economically prohibitive at distribution-scale telemetry volumes and operationally fragile under network variability - in favour of **co-locating intelligence with physical assets at the edge**, reserving the cloud for federated model aggregation, system-wide analytics, and market coordination.

Tier 0 - Device / Sensor Layer. At the lowest tier, heterogeneous IoT assets (smart meters, PMUs, EVSEs, heat pumps, battery storage units, industrial PLCs, solar inverters) generate raw telemetry at sampling rates ranging from sub-second (voltage, frequency) to minutes (energy consumption, state-of-charge). Connectivity is inherently heterogeneous: cellular (4G/5G), LoRaWAN, power-line communication (PLC), Wi-Fi, and satellite links coexist within a single distribution feeder. The architecture treats this heterogeneity as a first-class design constraint rather than an exception, in contrast to edge/fog paradigms such as OpenFog or ETSI MEC that presuppose homogeneous edge nodes with reliable backhaul (see *2. State of the Art and Related Work*).

Tier 1 - Edge / Gateway Layer. Edge gateways - deployed at BTM aggregation points, DSO secondary substations, and TSO regional nodes - host lightweight inference engines, local rule engines, and swarm-computing agents. Time-critical applications (voltage regulation, islanding detection, real-time demand response) execute entirely within this tier, tolerating no cloud round-trip latency. The gateway layer is the primary locus for the **Technology Facilitator Pillar** (see *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*), where federated learning (FL) participants perform local model updates and swarm agents execute peer-to-peer coordination without centralised coordination overhead.

Tier 2 - Fog / Regional Aggregation Layer. Regional fog nodes aggregate model gradients, anonymised feature vectors, and operational summaries from multiple edge gateways within a DSO or TSO control area. This tier performs cross-feeder coordination, congestion-aware load balancing of computational workloads, and preliminary anomaly detection. It also enforces **data-sovereignty policies** mandated by the GDPR and the Data Governance Act: raw telemetry never leaves the jurisdictional boundary of the DSO; only differentially private or federated artefacts ascend to the cloud.

Tier 3 - Cloud / Platform Layer. The cloud tier hosts the **Interoperability Pillar** (see *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*) and the **Digital Energy Ecosystem**

Enabling Pillar (see *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*). Functions include federated model aggregation with convergence guarantees under intermittent connectivity, system-wide analytics, market-clearing engines, Data Space connectors, and the semantic translation layer that converts raw register values into market-ready flexibility signals. The cloud is explicitly *not* a centralised data lake; it is a federated, policy-enforced federation of distributed data platforms.

Tier 4 - Ecosystem / Market Layer. The outermost tier encompasses external stakeholders: energy exchanges, aggregators, regulatory authorities, SMEs, and end-users. Interaction with this tier is mediated exclusively through the Data Space governance model and the open API contracts defined by the Standardisation Pillar, ensuring that no single platform operator retains unilateral control over data flows or market access.

The five tiers are not rigidly separated; the computational orchestration layer (Section 3.4) dynamically allocates workloads across tiers based on latency budgets, data-sensitivity classifications, and real-time network conditions. A voltage-regulation inference task, for example, may execute entirely at Tier 1 under normal conditions but be escalated to Tier 2 during a regional contingency requiring cross-feeder coordination.

3.3 The Four Pillars: Mutually Reinforcing Functional Dimensions

The HEDGE-IoT architecture is organised around four pillars that are **mutually reinforcing rather than sequentially dependent**. Each pillar addresses a distinct dimension of the integration challenge, yet the value of any single pillar is contingent on the presence of the other three. The following subsections summarise the architectural role of each pillar; detailed technical treatments are provided in Sections 4 through 7.

3.3.1 Technology Facilitator Pillar

The Technology Facilitator Pillar provides the **computational substrate** for intelligence at the grid edge. It encompasses:

- **Federated Learning (FL) engines** that enable privacy-preserving model training across BTM devices without raw-data centralisation. The HEDGE-IoT FL design extends beyond the synchronous, single-task assumptions of prior energy-grid FL work (see *2. State of the Art and Related Work*) by incorporating asynchronous participation protocols, non-IID-aware aggregation (e.g., FedProx, SCAFFOLD variants), and convergence guarantees under the intermittent, heterogeneous connectivity characteristic of RES-hosting distribution networks.
- **Swarm-computing paradigms** for distributed, peer-to-peer decision-making among edge agents. Swarm agents coordinate local flexibility activation, islanding detection, and microgrid formation without a central controller, exploiting the grid's physical topology as a communication and coordination substrate.
- **Workload orchestration** that maps AI/ML tasks to the optimal tier (Tier 0-3) based on latency, energy, and privacy constraints.

This pillar is elaborated in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge.*

3.3.2 Interoperability Pillar

The Interoperability Pillar provides the **data-exchange and semantic-bridging infrastructure** that allows heterogeneous IoT platforms, energy services, and stakeholders to communicate. Its architectural components include:

- **Data Space connectors** aligned with IDSA/Gaia-X reference architectures, instantiated as energy-specific blueprints that span BTM IoT telemetry, DSO operational data, and TSO market platforms under a unified, policy-enforced governance model.
- **Semantic translation layers** that map raw device register values (e.g., Modbus, IEC 61850, OpenADR) to SAREF-anchored ontologies and, ultimately, to market-ready flexibility signals.
- **Cross-domain integration contracts** that define data-sharing terms, consent management, and purpose limitation across stakeholder boundaries.

This pillar is elaborated in *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration.*

3.3.3 Standardisation Pillar

The Standardisation Pillar ensures **long-term interoperability and vendor neutrality** by anchoring all platform interfaces, data formats, and API specifications to widely adopted open standards. Its architectural role is threefold:

- **SAREF-anchored semantic standardisation**: All data exchanged across the continuum is expressed in SAREF (Smart Applications REFerence) ontologies, eliminating the retrofitted, proprietary semantic layers that perpetuate vendor lock-in in existing platforms (see *2. State of the Art and Related Work*).
- **Open API and protocol specifications**: RESTful and event-driven APIs, message-broker contracts (e.g., MQTT, AMQP), and device-onboarding procedures are standardised to reduce new-asset integration from months to days.
- **Alignment with ongoing IoT standardisation efforts**: The pillar maintains active liaison with AIOTI, the Smart Grids Task Force, and CEN/CENELEC/ETSI working groups to ensure forward compatibility with evolving EU digitalisation mandates (AI Act, NIS2, Data Governance Act).

This pillar is elaborated in *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*.

3.3.4 Digital Energy Ecosystem Enabling Pillar

The Digital Energy Ecosystem Enabling Pillar translates the technical capabilities of the other three pillars into **tangible energy-system outcomes**: increased RES-hosting capacity, enhanced grid resilience, and new market opportunities. Its architectural components include:

- **Flexibility activation pipelines** that convert latent, fragmented capacity (residential storage, V2G, smart thermal loads, curtailable industrial processes) into aggregated, dispatchable, and tradable resources visible to DSOs and TSOs.
- **Resilience orchestration** that leverages diversified, interoperable energy services to maintain critical loads during contingencies, islanding events, and extreme weather.
- **Scalable distributed data platforms** that support highly distributed energy services without centralised bottlenecks, enabling the framework to scale from single-feeder pilots to national deployments (validated in *9. Scalability, Replicability, and Validation*).
- **Market-access mechanisms** (Open Call, SME/innovator onboarding) that lower barriers to entry for new ecosystem participants, as further described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*.

This pillar is elaborated in *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*.

3.4 The Computational Orchestration Layer: Bridging Federated Applications Across Distributed Data Platforms

The **computational orchestration layer** is the horizontal binding element of the HEDGE-IoT architecture. It sits logically between the four pillars and the cloud/edge continuum, providing the runtime infrastructure that coordinates federated applications, manages workload placement, and enforces cross-pillar consistency. Without this layer, the four pillars would remain parallel but unintegrated - precisely the fragmentation criticised in *2. State of the Art and Related Work*.

The orchestration layer comprises four functional sub-components:

3.4.1 Workload Scheduler and Placement Engine. The scheduler receives application descriptors (latency budget, data-sensitivity class, computational footprint, connectivity assumptions) from federated applications and maps them to the optimal tier and node within the continuum. Placement decisions are multi-objective: minimising end-to-end latency for time-critical tasks, minimising energy consumption of the computational substrate itself, and respecting data-sovereignty constraints that prohibit cross-jurisdictional data movement. The scheduler is aware of the heterogeneous connectivity landscape (cellular, LoRaWAN, PLC, satellite) and can gracefully degrade or migrate workloads when a link degrades, addressing the unrealistic network assumptions of prior edge/fog architectures.

3.4.2 Federated Application Broker. The broker manages the lifecycle of federated applications - FL training jobs, swarm-coordination tasks, analytics pipelines - across distributed data platforms. It handles participant registration, model-version negotiation, gradient/parameter exchange, and convergence monitoring. Critically, the broker operates **across Data Space boundaries**: an FL job may involve BTM devices governed

by one Data Space connector, DSO operational data governed by another, and TSO market data governed by a third, with the broker enforcing each connector's usage policies at every data-access point. This cross-Data-Space federation is a distinguishing architectural feature of HEDGE-IoT relative to existing FL frameworks that assume a single, homogeneous data domain.

3.4.3 Semantic and Policy Mediator. The mediator ensures that data flowing between pillars and between tiers is semantically consistent and policy-compliant. It invokes the SAREF-anchored translation services defined by the Standardisation Pillar to normalise heterogeneous data representations, and it enforces the governance rules (consent, purpose limitation, retention, access control) defined by the Interoperability Pillar's Data Space model. The mediator operates as a lightweight, distributed service mesh rather than a centralised gateway, preserving the edge-autonomy principle.

3.4.4 Orchestration State and Observability Plane. A distributed state store tracks the current placement of all federated applications, the health and connectivity status of edge/fog/cloud nodes, the version lineage of shared models, and the provenance of data artefacts traversing Data Space connectors. This observability plane feeds the scalability and replicability assessment described in *9. Scalability, Replicability, and Validation* and provides the empirical evidence base for the framework's multi-level pilot validation.

3.5 Cross-Pillar Integration and End-to-End Data Flow

The four pillars and the orchestration layer are not independent modules; they interact through well-defined interfaces that create **emergent architectural properties** greater than the sum of individual pillar capabilities. Three representative cross-pillar flows illustrate this integration:

Flow A - Privacy-Preserving Flexibility Discovery (Pillars 1 → 2 → 3 → 4). A swarm-computing agent at Tier 1 (Technology Facilitator) detects a cluster of EVSEs with available V2G capacity. The agent publishes a SAREF-encoded flexibility descriptor (Standardisation Pillar) through a Data Space connector (Interoperability Pillar), which applies the relevant usage policy and consent check. The descriptor is consumed by a market-clearing engine at Tier 3 (Digital Energy Ecosystem Enabling), which aggregates it with descriptors from other feeders and submits a bid to the TSO balancing market. The computational orchestration layer ensures that the swarm agent's local computation completes within the market's pre-qualification latency window and that no raw charging-session data leaves the BTM boundary.

Flow B - Federated Congestion Forecasting (Pillars 1 → 2 → 3 → 4). An FL training job (Technology Facilitator) is orchestrated across 200 edge gateways in a DSO control area. The broker (orchestration layer) manages asynchronous participation under intermittent LoRaWAN and cellular links, applying non-IID-aware aggregation. Model gradients are exchanged through Data Space connectors (Interoperability Pillar) with differential-privacy noise injection. The resulting congestion-forecast model is published as a SAREF-annotated service (Standardisation Pillar) and consumed by a DSO resilience dashboard (Digital Energy Ecosystem Enabling) that pre-positions reactive power resources ahead of predicted congestion events.

Flow C - New-Asset Onboarding (Pillars 3 → 2 → 1 → 4). A new smart-heat-pump model is introduced by an SME via the Open Call mechanism (Digital Energy Ecosystem Enabling). The Standardisation Pillar provides the SAREF device profile and API contract; the Interoperability Pillar provisions a Data Space connector with the appropriate governance policy; the orchestration layer registers the device as a swarm-computing participant and an FL data contributor (Technology Facilitator). The end-to-end onboarding is completed in days rather than months, directly addressing the integration-latency gap identified in *2. State of the Art and Related Work*.

3.6 Architectural Principles and Design Invariants

The HEDGE-IoT architecture is governed by a set of design invariants that constrain all implementation choices and ensure coherence across the four pillars and the cloud/edge continuum:

1. **Edge-autonomy by default.** Intelligence and decision-making reside at the lowest tier capable of satisfying the application's latency and accuracy requirements. Cloud resources are invoked only for federated aggregation, system-wide analytics, and market coordination.
2. **Data never leaves its sovereign boundary unprocessed.** Raw telemetry is transformed, aggregated, or differentially private-protected before crossing a Data Space connector or a jurisdictional boundary, in compliance with the GDPR and the Data Governance Act.
3. **Open standards as the default interface.** All inter-component and inter-stakeholder interfaces are expressed in SAREF ontologies and open API specifications. Proprietary extensions are permitted only as opt-in, versioned add-ons that do not fracture the common semantic layer.

4. **Graceful degradation under connectivity loss.** Every federated application must continue to operate (at reduced fidelity) when one or more participants become unreachable, reflecting the intermittent connectivity of RES-hosting distribution networks.
5. **Ethical and inclusive design as architectural constraints, not add-ons.** IoT ethics, energy-poverty safeguards, and SME/innovator access mechanisms are embedded in the orchestration layer's policy engine and the Data Space governance model, as elaborated in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*.
6. **Empirical validation at every tier.** Architectural claims are substantiated through multi-level pilot deployments and replicability assessments (see *9. Scalability, Replicability, and Validation*), ensuring that the framework's theoretical properties hold under real-world load, network, and stakeholder conditions.

These invariants collectively ensure that the HEDGE-IoT Framework is not merely a conceptual reference model but a **deployable, verifiable, and extensible** architecture capable of transforming heterogeneous IoT assets into coordinated, privacy-preserving, market-ready flexibility resources across the full energy-system continuum.

4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge

4.1 Computational Sharing and Edge Offloading Mechanisms

The Technology Facilitator Pillar operationalises the architectural principle established in *3. HEDGE-IoT Framework Architecture* - that intelligence must be co-located with physical assets at the edge (Tier 1) while the cloud (Tier 3) is reserved for federated model aggregation, system-wide analytics, and market coordination. This section details the computational sharing mechanisms that make such co-location feasible across the heterogeneous, intermittently connected IoT estate characteristic of RES-hosting distribution networks.

4.1.1 Tiered Computational Offloading Model

HEDGE-IoT defines a three-level offloading taxonomy aligned with the five-tier continuum of *3. HEDGE-IoT Framework Architecture*:

- **Intra-device offloading (Tier 0 → Tier 1):** Resource-constrained sensors and actuators (smart meters, thermal controllers, EVSE units) delegate inference tasks to co-located edge gateways when local compute (typically <1 TOPS NPU) is insufficient for the required model complexity. The offloading decision is governed by a lightweight cost model that weighs inference latency, energy budget, and network availability against the marginal accuracy gain of executing the task locally.
- **Inter-edge offloading (Tier 1 → Tier 1):** Edge gateways within the same feeder or substation cluster share computational capacity through a peer-to-peer resource pool. A gateway with surplus GPU/NPU cycles (e.g., a DSO substation controller during off-peak hours) can host inference workloads for neighbouring residential aggregators operating under constrained hardware. This mechanism directly addresses the non-linear scalability demands identified in *1. Introduction and Motivation*, preventing the need for uniform hardware upgrades across thousands of BTM nodes.
- **Edge-to-cloud offloading (Tier 1 → Tier 3):** Only aggregated model parameters, gradient updates, and anonymised feature statistics traverse the edge-to-cloud boundary. Raw telemetry never leaves its sovereign boundary unprocessed, in strict adherence to design invariant (ii) of *3. HEDGE-IoT Framework Architecture*. Cloud offloading is reserved for global model aggregation, cross-feeder pattern mining, and market-signal generation - tasks whose computational and data-volume requirements exceed the aggregate capacity of any single edge cluster.

4.1.2 Heterogeneous Connectivity as a First-Class Constraint

Unlike the edge/fog paradigms catalogued in *2. State of the Art and Related Work* (OpenFog, ETSI MEC, SENSE4US), which presuppose homogeneous edge nodes with reliable backhaul, HEDGE-IoT treats the coexistence of cellular (4G/5G), LoRaWAN, power-line communication (PLC), and satellite links as a design constraint rather than an exception. The computational sharing layer implements:

- **Connectivity-aware task partitioning:** Workload schedulers decompose AI/ML pipelines into sub-tasks annotated with connectivity requirements (bandwidth, latency, jitter tolerance). A voltage-regulation inference requiring <10 ms round-trip is pinned to the nearest Tier 1 gateway reachable via PLC, while a 15-minute-ahead congestion forecast tolerating 200 ms latency may traverse a 5G link to a Tier 2 fog node.
- **Store-and-forward gradient buffering:** Under intermittent LoRaWAN or satellite links (duty-cycle constrained to <1% airtime), FL gradient updates are buffered locally and transmitted opportunistically during connectivity windows. The buffer implements a priority queue that ensures convergence-critical updates (e.g., from under-represented non-IID clients) are transmitted before redundant ones.
- **Graceful degradation under connectivity loss:** Consistent with design invariant (iv), edge nodes continue executing locally trained models and swarm-based local decisions during extended outages. Upon reconnection, a delta-synchronisation protocol reconciles model drift without requiring full retraining, bounding the accuracy loss to a pre-defined tolerance (typically <2% MAPE for load forecasting).

4.1.3 Resource-Aware Model Adaptation

To accommodate the extreme hardware heterogeneity across Tier 0 and Tier 1 devices - from 32-bit MCU-based smart plugs to 8-core ARM edge servers - HEDGE-IoT employs a family of model-adaptation strategies:

- **Progressive model distillation:** A teacher model trained at Tier 3 is progressively distilled into student variants of decreasing complexity (full-precision → INT8 quantised → binary-weight) tailored to the compute envelope of each target device class. The distillation pipeline is orchestrated by the federated application broker (see *3. HEDGE-IoT Framework Architecture*) and triggered automatically upon new-asset registration.
- **Dynamic model switching:** Edge gateways maintain a model zoo indexed by accuracy-latency-energy trade-offs. A runtime controller selects the appropriate model variant based on real-time resource telemetry (CPU/GPU utilisation, thermal state, battery level for battery-backed nodes) and the criticality of the inference task.
- **Hardware-aware neural architecture search (NAS):** For new deployment scenarios lacking a pre-trained model, a lightweight NAS loop executes at Tier 2 fog nodes, generating architectures constrained to the target device's memory, compute, and power budgets. The search space is bounded by SAREF-annotated task ontologies (see *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) to ensure architectural validity for energy-domain workloads.

4.2 Federated Learning for Privacy-Preserving Energy Analytics

Federated learning (FL) constitutes the primary mechanism through which HEDGE-IoT satisfies the data-sovereignty constraints (GDPR, Data Governance Act) identified in *1. Introduction and Motivation* while enabling cross-stakeholder analytics. Unlike the FL applications reviewed in *2. State of the Art and Related Work* - which assume synchronous participation, stable links, and single-task models - HEDGE-IoT's FL architecture is engineered for the production-grade robustness gap identified therein.

4.2.1 Energy-Specific FL Topology

The FL topology mirrors the physical grid hierarchy:

- **Local aggregation (Tier 1):** Within a single BTM household or a DSO substation cluster, a local FL server (the edge gateway) aggregates gradients from Tier 0 devices. This intra-site aggregation provides a first layer of differential privacy: individual device contributions are never visible beyond the gateway boundary.
- **Regional aggregation (Tier 2):** Multiple Tier 1 gateways within a feeder or distribution substation contribute to a Tier 2 fog-level aggregator. This tier handles cross-site pattern extraction (e.g., feeder-level congestion signatures) without exposing per-household telemetry.
- **Global aggregation (Tier 3):** The cloud-level FL orchestrator aggregates model updates from Tier 2 nodes across DSO and TSO governance boundaries. This tier produces system-wide models (national load profiles, cross-border congestion predictors) while preserving the sovereignty of each governance domain.

This hierarchical topology ensures that no single entity - neither a DSO nor a TSO nor a cloud provider - ever possesses the raw data of all participants, satisfying the design invariant that data never leaves its sovereign boundary unprocessed.

4.2.2 Convergence Guarantees Under Non-IID and Intermittent Connectivity

The convergence gap identified in *2. State of the Art and Related Work* is addressed through three complementary mechanisms:

- **FedProx-E (Energy-adapted FedProx):** The proximal term in FedProx is augmented with a grid-topology-aware regularisation that penalises model divergence proportional to the electrical distance between contributing nodes. Nodes on the same feeder (electrically proximate) are expected to exhibit similar load patterns; the regularisation term encodes this prior, accelerating convergence under non-IID distributions by 30-45% relative to vanilla FedAvg in pilot evaluations (see *9. Scalability, Replicability, and Validation*).
- **Asynchronous partial participation with staleness compensation:** Given that 15-40% of BTM devices may be offline at any given aggregation round (due to LoRaWAN duty cycles, PLC noise, or scheduled maintenance), the global aggregator accepts partial updates and applies a staleness-aware weighting scheme. Gradients from nodes that have been offline for k rounds are down-weighted by a factor of $\widehat{\ }k$ ($0 < < 1$), preventing stale updates from destabilising the global model while preserving the information content of delayed contributions.

- **Topology-aware client selection:** The federated application broker selects participating clients per round based on (i) connectivity reliability (historical uptime per link type), (ii) data diversity (ensuring representation across consumer archetypes: residential, commercial, industrial, EVSE), and (iii) electrical topology (ensuring spatial coverage across feeders to avoid geographic bias in the learned model).

4.2.3 Privacy Mechanisms Beyond Gradient Exchange

FL alone does not guarantee privacy; gradient inversion attacks can reconstruct training data from shared updates. HEDGE-IoT layers additional protections:

- **Secure aggregation with homomorphic encryption:** Tier 1 and Tier 2 aggregators compute weighted averages of encrypted gradients using Paillier or BFV schemes, ensuring that no aggregator observes individual contributions. The encryption overhead is bounded to <15% of communication cost for typical model sizes (<50 MB) encountered in energy analytics.
- **Differential privacy with adaptive noise scheduling:** Gaussian noise is injected into aggregated updates with a privacy budget (,) calibrated per data sensitivity class. Residential consumption data (high sensitivity) receives 1.0; aggregated feeder-level statistics (lower sensitivity) permit 5.0. The noise schedule adapts dynamically: early training rounds (where model parameters are far from convergence) tolerate higher noise, while late rounds reduce noise to preserve final accuracy.
- **Federated feature selection:** Rather than sharing full gradient vectors, clients transmit only the indices and magnitudes of the top-k most informative features per round. This reduces the attack surface for gradient inversion while simultaneously compressing communication by 60-80% - critical for LoRaWAN-constrained links.

4.2.4 Representative FL Applications in the Energy Domain

Application	Tier	Model Class	Privacy Requirement	Latency Budget
Residential load forecasting	$1 \rightarrow 2$	Temporal Fusion Transformer (TFT)	1.0, secure agg.	15 min (batch)
Non-intrusive load monitoring (NILM)	$0 \rightarrow 1$	1D-CNN + attention	Local-only (no cross-site)	<100 ms
Feeder-level congestion prediction	$2 \rightarrow 3$	Graph Neural Network (GNN)	3.0, homomorphic agg.	5 min
Cross-DSO flexibility discovery	3 (global)	Variational autoencoder	5.0, DP-SGD	1 h
EV charging pattern clustering	$1 \rightarrow 2$	Federated k-means + GMM	2.0, secure agg.	30 min

These applications are orchestrated through the federated application broker described in *3. HEDGE-IoT Framework Architecture*, which manages model lifecycle, client registration, and cross-Data-Space federation (see *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*).

4.3 Swarm Computing Paradigms for Distributed Decision-Making

While FL addresses the *learning* dimension of edge intelligence, swarm computing addresses the *decision-making* dimension - enabling real-time, distributed coordination among heterogeneous IoT assets without a central controller. This paradigm fills the gap identified in *2. State of the Art and Related Work*, where the integration of FL with swarm-computing paradigms for peer-to-peer edge coordination has not been formalised for energy grid topologies.

4.3.1 Swarm Architecture for Grid-Edge Coordination

HEDGE-IoT's swarm computing layer implements a bio-inspired coordination protocol adapted to the physical constraints of distribution networks:

- **Stigmergic communication via shared state registers:** Rather than direct peer-to-peer messaging (which scales poorly beyond ~100 nodes), swarm agents (smart inverters, battery storage units, thermal loads, EVSEs) coordinate through shared state registers hosted on Tier 1 edge gateways. Each agent reads the current "pheromone" state (e.g., feeder voltage profile, aggregate flexibility available, congestion risk index) and adjusts its local control policy accordingly. This indirect communication pattern is robust to node churn and intermittent connectivity.
- **Gradient-based swarm consensus:** For tasks requiring agreement among agents (e.g., distributed voltage regulation across a 500-node feeder), HEDGE-IoT employs a gossip-based consensus protocol with convergence guarantees under bounded message delays. Each agent exchanges its local control vector with

a randomly selected subset of neighbours (fan-out $f = 3$-5) per iteration. The protocol converges to a Nash equilibrium of the underlying distributed optimisation problem within *O(log N / log f)* rounds, where N is the number of participating agents.

- **Hierarchical swarm nesting:** Swarms are nested to mirror grid topology: a *micro-swarm* comprises devices within a single household or building; a *meso-swarm* aggregates micro-swarms within a feeder; a *macro-swarm* spans multiple feeders under a substation. Coordination between swarm levels follows a leader-election protocol where the most computationally capable node (typically a DSO substation controller) assumes the meso-swarm coordinator role, but this role is revocable and migratable to preserve resilience.

4.3.2 Swarm-FL Synergy: Learning-Decision Coupling

The critical innovation of HEDGE-IoT's Technology Facilitator Pillar is the formal coupling of FL and swarm computing into a unified edge-intelligence loop:

1. **FL informs swarm policy:** The global FL model (aggregated at Tier 3) provides the *prior* for swarm decision-making. For example, a congestion-prediction GNN trained via FL outputs a per-feeder congestion probability that seeds the initial control vectors of the voltage-regulation swarm.
2. **Swarm execution generates training data:** The realised control actions and their outcomes (measured voltage deviations, actual power flows) constitute the *ground-truth labels* for the next FL training round. This closes the learning loop without requiring manual annotation.
3. **Swarm feedback modulates FL participation:** Agents whose local model predictions consistently diverge from swarm-observed outcomes are flagged by the federated application broker for targeted retraining or model replacement, preventing the accumulation of stale or biased local models.

This coupling ensures that the system's predictive accuracy (FL) and its operational responsiveness (swarm) co-evolve, addressing the production-grade robustness gap identified in *2. State of the Art and Related Work*.

4.3.3 Distributed Decision-Making Use Cases

- **Real-time voltage regulation (Tier 1 swarm):** Smart inverters and battery storage units within a feeder execute a distributed model-predictive control (MPC) loop via the gossip consensus protocol. Each agent solves a local optimisation (minimise voltage deviation subject to inverter rating and battery SOC constraints) and exchanges boundary conditions with neighbours. Convergence to a feasible operating point is achieved within 2-5 consensus rounds (<500 ms), satisfying the sub-second latency requirement for islanding detection and voltage regulation identified in *1. Introduction and Motivation*.
- **Peer-to-peer energy trading (Tier 1-2 swarm):** Prosumers with surplus PV generation and consumers with flexible demand negotiate bilateral trades through a swarm-based double-auction mechanism. The auction clearing is distributed: each agent submits a bid vector to its meso-swarm coordinator, which computes a market-clearing price via a distributed ADMM (Alternating Direction Method of Multipliers) solver. No central market operator is required at the BTM/DSO level, reducing single points of failure and enabling micro-market operation even during backhaul outages.
- **Distributed demand response (Tier 1-2 swarm):** A DSO issues a flexibility request (e.g., "shed 2 MW within 15 minutes on Feeder 7"). The request propagates to the meso-swarm, which decomposes it into per-agent curtailment targets via a distributed knapsack solver. Each agent (smart thermal load, EVSE, industrial process) independently determines its feasible curtailment profile and commits via a two-phase protocol. The swarm guarantees aggregate compliance within $\pm 5\%$ of the requested capacity, with graceful degradation (proportional curtailment) if individual agents become unreachable.

4.3.4 Resilience and Byzantine Fault Tolerance

Energy grid operations cannot tolerate single-node failures or malicious behaviour. HEDGE-IoT's swarm layer incorporates:

- **Byzantine fault tolerance (BFT) in consensus:** The gossip protocol is augmented with a BFT variant (PBFT-inspired) that tolerates up to $f = (N\text{-}1)/3$ faulty or adversarial agents per swarm. Compromised agents (e.g., a spoofed smart meter injecting false voltage readings) are detected via cross-validation against neighbouring agents' observations and excluded from subsequent consensus rounds.

- **Swarm reconfiguration under node loss:** If a meso-swarm coordinator fails, a new leader is elected within one consensus round via a Raft-inspired protocol. Micro-swarm members temporarily operate in autonomous mode (executing their last-known control policy) until coordination is restored, ensuring no interruption to time-critical grid functions.
- **Adversarial robustness of FL models:** The FL pipeline incorporates adversarial training (FGSM, PGD attacks) on the global model to ensure that a compromised client cannot poison the global model beyond a bounded accuracy degradation. The federated application broker monitors per-client gradient norms and flags outliers for investigation.

4.4 Orchestration of Heterogeneous IoT Workloads

The computational orchestration layer described in *3. HEDGE-IoT Framework Architecture* - comprising the workload scheduler/placement engine, the federated application broker, the semantic and policy mediator, and the distributed observability plane - serves as the binding element that prevents the Technology Facilitator Pillar from remaining an isolated analytics stack. This section details how orchestration coordinates AI/ML and swarm workloads across heterogeneous IoT assets.

4.4.1 Workload Scheduler and Placement Engine

The placement engine solves a multi-objective optimisation problem for each incoming AI/ML or swarm task:

- **Objective function:** Minimise a weighted combination of (i) end-to-end inference latency, (ii) energy consumption of the execution platform, (iii) communication cost (bytes × link cost), and (iv) a fairness term ensuring no single edge node is overloaded.
- **Constraints:** Hardware capability (compute, memory, accelerator type), connectivity class (PLC vs. 5G vs. LoRaWAN), data-sovereignty policy (task must execute within a specified governance boundary), and real-time deadline.
- **Solution approach:** A two-phase heuristic: (1) a constraint-filtering phase eliminates infeasible placements; (2) a genetic-algorithm-based search over the remaining feasible set converges to a near-optimal placement within <100 ms for workloads up to 10,000 tasks. For larger fleets, a distributed auction mechanism (each edge node bids its available capacity) achieves comparable quality in <500 ms.
- **Multi-task, multi-stakeholder co-scheduling:** A single edge gateway may simultaneously host a DSO congestion-prediction model, a TSO flexibility-discovery FL job, and a BTM NILM inference. The scheduler enforces resource isolation (cgroups, GPU MIG slicing) and priority preemption: safety-critical tasks (voltage regulation, protection) preempt analytics tasks, which in turn preempt batch training jobs.

4.4.2 Federated Application Broker

The federated application broker manages the full lifecycle of FL jobs across governance boundaries:

- **Job registration and discovery:** Analytics providers (DSOs, TSOs, aggregators, research institutions) register FL jobs via a declarative specification (model architecture, data requirements, privacy parameters, convergence criteria). The broker matches jobs to eligible client pools based on data availability, connectivity class, and policy compliance.
- **Cross-Data-Space federation:** A single FL job may span BTM, DSO, and TSO Data Spaces (see *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*). The broker mediates the federation by translating between Data Space connector protocols, enforcing per-domain policy constraints (e.g., "residential data may contribute gradients but not raw features"), and managing the secure aggregation pipeline across governance boundaries - without a centralised gateway, preserving edge autonomy.
- **Model versioning and A/B deployment:** The broker maintains a model registry with version lineage. New model candidates (produced by FL training or NAS) are deployed to a canary subset of edge nodes; the distributed observability plane monitors performance metrics (accuracy, latency, energy) before promoting the candidate to full fleet deployment. Rollback is automatic upon metric regression.

4.4.3 Semantic and Policy Mediation

Heterogeneous IoT assets expose data in incompatible formats and semantic vocabularies. The orchestration layer's semantic mediator:

- **Maps raw register values to SAREF-annotated ontologies** (see *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*), enabling FL clients with different sensor suites (e.g., a smart meter exposing 15-min interval active power vs. a smart inverter exposing 1-s instantaneous power) to contribute to the same FL model through a common feature space.
- **Enforces policy-as-code constraints** derived from the Data Space governance model (see *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*). Policies govern which data attributes may participate in which FL jobs, under which privacy parameters, and for which purpose. The mediator intercepts all gradient exchanges and blocks non-compliant contributions.
- **Resolves semantic conflicts in swarm coordination:** When swarm agents from different vendors report "available flexibility" using different units, time horizons, or confidence intervals, the mediator normalises these into a canonical SAREF FlexibilityProfile representation before the swarm consensus protocol operates on them.

4.4.4 Distributed Observability Plane

The orchestration layer's observability plane provides the feedback loop necessary for adaptive orchestration:

- **Per-task telemetry:** Each AI/ML inference, FL round, and swarm consensus iteration emits structured telemetry (latency, accuracy, energy, communication volume) to a distributed metrics store. No raw data is centralised; only operational metadata traverses the edge-to-cloud boundary.
- **Anomaly detection on orchestration metrics:** A lightweight FL-trained anomaly detector (running at Tier 2) identifies degraded workloads (e.g., a model whose accuracy has drifted >5% over the last 10 rounds, or a swarm whose consensus latency has doubled). Alerts trigger automatic remediation: model retraining, client re-selection, or swarm reconfiguration.
- **SLA monitoring and enforcement:** The observability plane tracks compliance with service-level agreements defined per application class (e.g., "voltage regulation: 99.9% of decisions within 500 ms"). Persistent SLA violations trigger escalation to the workload scheduler for re-placement or to the federated application broker for model replacement.

4.5 Integration with Adjacent Pillars and Design Invariants

The Technology Facilitator Pillar does not operate in isolation; its efficacy depends on tight coupling with the other three pillars of *3. HEDGE-IoT Framework Architecture*.

4.5.1 Coupling with the Interoperability Pillar

FL and swarm workloads are instantiated within Data Space governance boundaries. The federated application broker (4.4.2) is the technical interface through which cross-Data-Space FL federation is realised. Data Space connectors (see *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*) provide the policy-enforced data access that FL clients require, while the swarm's shared state registers are exposed as Data Space assets with access-control policies. This coupling ensures that the computational sharing mechanisms of 4.1 operate within the legal and governance envelope defined by the Interoperability Pillar.

4.5.2 Coupling with the Standardisation Pillar

SAREF-anchored ontologies (see *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) provide the semantic substrate for: (i) the feature-space normalisation performed by the semantic mediator (4.4.3); (ii) the declarative FL job specifications registered with the federated application broker; (iii) the canonical flexibility profiles exchanged in swarm coordination (4.3.3); and (iv) the model metadata (architecture, training data provenance, privacy parameters) stored in the model registry. Without SAREF anchoring, the orchestration layer would require bespoke adapters for every new IoT asset class, reintroducing the vendor lock-in barrier identified in *1. Introduction and Motivation*.

4.5.3 Coupling with the Digital Energy Ecosystem Enabling Pillar

The FL and swarm mechanisms described herein are the technical enablers for the market-ready flexibility activation targeted by *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*. Specifically:

- FL-derived congestion forecasts and flexibility-discovery models provide the *informational foundation* for DSO/TSO flexibility markets.

- Swarm-based distributed demand response and P2P trading provide the *execution mechanism* through which aggregated BTM flexibility is delivered to the market.
- The orchestration layer's SLA monitoring provides the *performance evidence* required for market settlement and regulatory compliance.

4.5.4 Adherence to Design Invariants

All mechanisms in this section are constrained by the six design invariants of *3. HEDGE-IoT Framework Architecture*:

Invariant	Implementation in Technology Facilitator Pillar
(i) Edge-autonomy by default	Swarm decisions execute locally; cloud involvement is optional and asynchronous
(ii) Data never leaves sovereign boundary unprocessed	FL exchanges only encrypted gradients; raw telemetry stays at Tier 0/1
(iii) Open standards as default interface	SAREF ontologies, open FL protocols (FedML, Flower), open swarm APIs
(iv) Graceful degradation under connectivity loss	Store-and-forward buffering, autonomous swarm operation, model zoo fallback
(v) Ethical and inclusive design	DP noise calibrated per sensitivity; fairness term in scheduler; SME-accessible model zoo
(vi) Empirical validation at every tier	Distributed observability plane; pilot metrics reported in *9. Scalability, Replicability, and Validation*

4.6 Summary of the Technology Facilitator Pillar

The Technology Facilitator Pillar provides the computational substrate upon which the HEDGE-IoT Framework's intelligence operates. By combining tiered computational offloading (4.1), production-grade federated learning with convergence guarantees under non-IID and intermittent conditions (4.2), bio-inspired swarm computing for real-time distributed decision-making (4.3), and a multi-objective orchestration layer that coordinates heterogeneous workloads across governance boundaries (4.4), the pillar transforms the edge from a passive data-collection point into an active intelligence node. The formal coupling of FL and swarm paradigms (4.3.2) and the tight integration with the Interoperability, Standardisation, and Ecosystem Enabling pillars (4.5) ensure that this edge intelligence is privacy-preserving, semantically coherent, market-ready, and resilient - addressing the five interlocking barriers of *1. Introduction and Motivation* and the production-grade robustness gap of *2. State of the Art and Related Work* in a unified, deployable architecture.

5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration

5.1 Data Space Architecture: Principles and Energy-Specific Adaptation

The Interoperability Pillar operationalises the HEDGE-IoT Framework's commitment to open, sovereign, and policy-enforced data exchange by adopting and extending the International Data Spaces Association (IDSA) Reference Architecture Model (RAM) and the Gaia-X ecosystem principles into a deployable, energy-specific blueprint. Unlike generic Data Space implementations that treat data as an undifferentiated commodity, the HEDGE-IoT Data Space architecture recognises that energy-system data carries unique operational, regulatory, and physical constraints: telemetry from a residential smart meter is subject to GDPR and the Data Governance Act (DGA); DSO operational data is governed by NIS2 security directives and national grid codes; TSO market data is regulated under REMIT and national energy-market legislation. A single, flat governance model cannot accommodate these divergent regimes without either over-constraining innovation or under-protecting critical infrastructure.

The HEDGE-IoT Data Space architecture therefore introduces a **layered sovereignty model** in which each tier of the five-tier cloud/edge continuum defined in *3. HEDGE-IoT Framework Architecture* constitutes a distinct Data Space domain with its own governance authority, policy engine, and connector infrastructure. Tier 0 (Device/Sensor) and Tier 1 (Edge) nodes collectively form the **Behind-the-Meter (BTM) Data Space**, governed by the data subject (prosumer, household, or industrial asset owner) through consent-based usage policies. Tier 2 (Local Aggregation) constitutes the **DSO Operational Data Space**, governed by the distribution system operator under NIS2 and national regulatory frameworks. Tier 3 (Cloud/Regional) forms the **TSO and Market Data Space**, governed by transmission operators and market participants under REMIT and national energy-market rules. Tier 4 (Ecosystem/Market) constitutes the **Cross-Domain Federation Space**, a meta-governance layer that mediates policy conflicts and enables cross-Data-Space transactions without collapsing sovereign boundaries.

This layered design directly addresses the gap identified in *2. State of the Art and Related Work*, where IDSA/Gaia-X architectures and SAREF/WoT semantic standards exist as parallel but unintegrated efforts. The HEDGE-IoT blueprint is the first to span BTM IoT telemetry, DSO operational data, and TSO market platforms under a unified yet federated governance model, eliminating the need for a centralised data lake while preserving the edge-autonomy-by-default invariant established in *3. HEDGE-IoT Framework Architecture*.

Each Data Space domain is instantiated through a **Data Space Connector (DSC)** - a software component that mediates all inbound and outbound data flows, enforces usage policies, maintains audit trails, and exposes a standardised interface to the computational orchestration layer. The DSC implements the IDSA Connector specification (version 4.0+) with energy-specific extensions: a grid-topology-aware policy engine that evaluates data-sharing requests against physical network constraints (e.g., a DSO may permit aggregated voltage data sharing but reject per-feeder granularity if it reveals critical-infrastructure topology), a temporal validity window tied to market settlement periods, and a data-quality attestation mechanism that certifies telemetry provenance and calibration status before exchange.

5.2 Semantic Interoperability: From Raw Telemetry to Market-Ready Signals

Protocol-level interoperability - ensuring that a Modbus TCP register value and an MQTT JSON payload can be parsed by the same application - is necessary but insufficient for the cross-domain integration demanded by the HEDGE-IoT Framework. The deeper challenge, identified as the first of five interlocking barriers in *1. Introduction and Motivation*, is **semantic heterogeneity**: the same physical quantity (e.g., active power at a household connection point) is represented differently across vendor platforms, national metering standards, and market-clearing systems. A smart meter from Vendor A may report instantaneous active power in watts with a 1-second timestamp; a DSO SCADA system may report 15-minute average active power in kilowatts with a settlement-period-aligned timestamp; a TSO balancing-market platform may require a 5-minute forward flexibility signal in megawatts with a bidirectional sign convention. Without semantic bridging, these representations are mutually unintelligible, and the flexibility they encode remains invisible to downstream consumers.

The HEDGE-IoT semantic interoperability layer is anchored in the **Smart Applications REFERENCE (SAREF) ontology family**, as mandated by the Standardisation Pillar (*6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*). SAREF provides the canonical vocabulary for energy-system entities (Device, Meter, Battery, ElectricVehicle, HeatPump, GridConnectionPoint), their properties (activePower, reactivePower, stateOfCharge, tariff), and their relationships (isMeasuredBy, isConnectedTo, isPartOf). The HEDGE-IoT extension introduces three additional ontological layers:

1. **The Telemetry-to-Signal Transformation Ontology (TSTO).** This layer defines deterministic and probabilistic transformation rules that map raw sensor register values to semantically enriched, market-ready signals. For example, a rule might specify that a sequence of 1-second instantaneous active-power readings from a Tier 0 smart meter is aggregated into a 15-minute average, converted to the DSO's reporting convention, annotated with a confidence interval derived from the meter's accuracy class, and tagged with a SAREF `saref:Measurement` instance linked to the physical `saref:GridConnectionPoint`. The TSTO is expressed as a set of SPARQL-constructible transformation pipelines, executable by the semantic mediator within the computational orchestration layer.
2. **The Flexibility Taxonomy Ontology (FTO).** This layer classifies latent flexibility resources (residential storage, V2G, smart thermal loads, curtailable industrial processes) into a common taxonomy that spans BTM, DSO, and TSO representations. The FTO defines flexibility attributes - duration, ramp rate, direction (up/down), activation latency, minimum bid size, and reliability class - that are vendor-neutral and market-agnostic. A V2G-capable electric vehicle's flexibility is described identically whether it is discovered by a BTM aggregator, reported to a DSO congestion-management system, or bid into a TSO balancing market.
3. **The Cross-Domain Context Ontology (CDCO).** This layer captures the operational context in which data is exchanged: the physical grid topology (feeder, substation, zone), the temporal context (settlement period, day-ahead market window, real-time balancing window), the regulatory context (applicable national grid code, DGA consent scope), and the trust context (data-sharing agreement identifier, purpose limitation). The CDCO ensures that a data-sharing decision made in one Data Space domain is interpretable and enforceable in another.

The semantic mediator component of the computational orchestration layer (described in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*) executes these ontological mappings at runtime. When a federated learning job spans BTM, DSO, and TSO governance boundaries, the semantic mediator normalises feature vectors to SAREF-anchored representations before gradient exchange, ensuring that a model trained on Vendor A's BTM telemetry can be evaluated against Vendor B's DSO operational data without manual re-engineering. This mechanism is the technical foundation for the new-asset onboarding acceleration (from months to days) reported in *3. HEDGE-IoT Framework Architecture.*

5.3 Data Governance Models: Sovereignty, Policy Enforcement, and Trust

The Data Governance Act (DGA), the General Data Protection Regulation (GDPR), the NIS2 Directive, and the AI Act collectively impose a complex, multi-layered compliance landscape on energy-system data exchange. The HEDGE-IoT Data Space architecture translates these regulatory requirements into machine-executable **usage policies** that are evaluated at the point of data exchange by each Data Space Connector, rather than relying on post-hoc legal enforcement or trust-based agreements.

The governance model is structured around three interlocking mechanisms:

Policy-as-Code Enforcement. Each data asset published to a Data Space Connector is accompanied by a machine-readable usage policy expressed in the Open Digital Contract (ODC) format, extended with energy-specific policy predicates. A residential smart-meter telemetry stream, for instance, may carry the following policy constraints: (i) purpose limitation - data may be used exclusively for demand-response optimisation and voltage regulation, not for commercial profiling; (ii) temporal restriction - raw telemetry is available for 72 hours post-collection, after which only aggregated statistics (hourly averages, peak values) remain accessible; (iii) geographic restriction - data may not be processed outside the EU/EEA; (iv) consent revocation - upon withdrawal of the data subject's consent, all downstream copies must be purged within 24 hours, with cryptographic attestation of deletion; (v) derivative-data ownership - model parameters derived from federated training on this data remain the property of the data subject unless explicitly licensed. The DSC's policy engine evaluates every access request against these constraints in <5 ms, denying non-compliant requests and logging the decision for audit.

Data-Sovereignty Boundary Enforcement. The design invariant that "data never leaves its sovereign boundary unprocessed" (established in *3. HEDGE-IoT Framework Architecture*) is operationalised through a **processing-at-source** requirement. Raw BTM telemetry is never transmitted to a DSO or TSO Data Space in its native form. Instead, the Tier 1 edge gateway executes the TSTO transformation pipeline locally, producing semantically enriched, aggregated, and (where required) differentially private outputs before the DSC releases them to the next Data Space domain. For federated learning workloads, only model gradients or aggregated parameters traverse the BTM-to-DSO boundary, and only encrypted model updates traverse the DSO-to-TSO boundary, as detailed in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing*

at the Grid Edge. The DSC enforces this boundary by inspecting outbound payloads against a schema that rejects raw-register-level data exceeding the permitted granularity.

Trust and Provenance Framework. Cross-domain data exchange in energy systems requires not only policy compliance but also verifiable trust in data quality, provenance, and the identity of the data provider. The HEDGE-IoT trust framework implements a **decentralised identity and attestation layer** based on W3C Verifiable Credentials (VCs) and Decentralised Identifiers (DIDs). Each Data Space Connector, each IoT asset, and each data-processing service is issued a DID by a national or sectoral trust anchor (e.g., a DSO's PKI, a national energy regulator's certification authority). Data assets are accompanied by a Verifiable Credential attesting to their provenance (which physical asset generated the data, under which calibration certificate, at which timestamp), their quality classification (accuracy class, completeness ratio, anomaly score), and their governance classification (personal data, operational data, market data). The receiving DSC validates the VC against the issuer's DID document before accepting the data, preventing spoofed or degraded telemetry from entering downstream analytics or market-clearing processes.

This governance model directly addresses the data-sovereignty constraints identified in *1. Introduction and Motivation* and the absence of a unified policy-enforced governance model noted in *2. State of the Art and Related Work*. It also provides the technical substrate for the IoT ethics and inclusivity requirements discussed in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*, ensuring that consent, purpose limitation, and data-subject rights are enforced architecturally rather than procedurally.

5.4 Cross-Domain Integration: Bridging BTM, DSO, and TSO Layers

The energy system's physical hierarchy - behind-the-meter assets, distribution networks, transmission networks, and wholesale markets - maps onto a governance hierarchy in which no single entity possesses complete visibility or authority. A residential battery's state of charge is known to the prosumer and, if shared, to the BTM aggregator; the DSO knows the aggregate load on the feeder to which the battery is connected but not the individual battery's SOC; the TSO knows the net injection at the transmission-distribution interface but not the composition of that injection. The HEDGE-IoT Interoperability Pillar bridges these layers through a **progressive abstraction and aggregation protocol** that preserves information fidelity at each level while respecting governance boundaries.

The cross-domain integration protocol operates as follows:

BTM-to-DSO Bridge. The BTM Data Space Connector publishes aggregated, semantically enriched flexibility signals to the DSO Data Space Connector at a granularity determined by the DSO's operational need and the data subject's consent scope. For voltage-regulation applications, the DSO receives per-feeder aggregate flexibility (total available MW up/down, aggregate ramp rate, aggregate reliability class) rather than per-household data. For congestion-management applications, the DSO may receive per-substation-zone flexibility with a 15-minute resolution. The TSTO transformation pipeline at the Tier 1 gateway performs the aggregation, and the DSC enforces the granularity constraint. The DSO Data Space Connector, in turn, publishes network-state information (feeder loading, voltage profiles, congestion forecasts) back to the BTM Data Space, enabling edge-level swarm-computing decisions (as described in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*) to be informed by network context without exposing individual consumer data.

DSO-to-TSO Bridge. The DSO Data Space Connector publishes aggregated flexibility capacity and network-state summaries to the TSO Data Space Connector at the transmission-distribution interface. The granularity is determined by the TSO's market-clearing requirements: for day-ahead markets, 15-minute resolution per bidding zone; for real-time balancing, 5-minute resolution per control area. The DSO retains authority over the aggregation methodology and the selection of which BTM flexibility pools contribute to the published signal, ensuring that the TSO receives a market-ready flexibility product without visibility into individual consumer assets. The TSO Data Space Connector publishes system-wide balancing signals, RES forecast updates, and congestion alerts back to the DSO Data Space, enabling DSO-level orchestration of distributed flexibility resources.

Cross-Domain Federation for Federated Applications. The most architecturally demanding integration scenario is a federated learning or swarm-computing job that spans all three governance domains simultaneously - for example, a congestion-forecasting model trained on BTM load patterns, DSO feeder-state histories, and TSO RES generation forecasts. The computational orchestration layer's federated application broker (described in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*) mediates this cross-Data-Space federation by: (i) decomposing the global model into domain-specific sub-models, each trained within its own Data Space boundary; (ii) orchestrating the exchange of encrypted model updates through the respective DSCs, with each DSC enforcing its domain's usage policy on the outbound update; (iii)

performing secure aggregation at the boundary between domains, ensuring that no single entity reconstructs another domain's raw training data; and (iv) maintaining a distributed audit trail that records every cross-domain data movement for regulatory compliance.

This progressive abstraction protocol ensures that the GW-scale latent flexibility identified in *1. Introduction and Motivation* - aggregated residential storage, V2G, smart thermal loads, curtailable industrial processes - becomes visible and actionable at the DSO and TSO levels without requiring any single entity to possess a complete, centralised view of all distributed assets. The result is a **federated visibility model** that converts fragmented, invisible flexibility into a coordinated, market-ready resource, directly enabling the ecosystem-enablement objectives of *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*.

5.5 Data Space Connector Provisioning and New-Asset Onboarding

A critical operational requirement for the Interoperability Pillar is the ability to onboard new IoT assets - whether a new generation of smart meters, a novel EV charging platform, or an industrial heat-pump controller - into the Data Space ecosystem without re-architecting the underlying platform. The HEDGE-IoT Framework reduces new-asset integration from the months-long cycles typical of proprietary platform onboarding to a **days-long, self-service provisioning workflow**, as reported in *3. HEDGE-IoT Framework Architecture*.

The onboarding workflow comprises five automated stages:

1. **Asset Registration and Identity Issuance.** The new asset (or its manufacturer/aggregator) registers with the relevant Data Space Connector, presenting a manufacturer's attestation (a Verifiable Credential issued by a recognised certification body) and a SAREF-compliant device description. The DSC issues a Decentralised Identifier (DID) to the asset and provisions a lightweight connector agent on the asset's edge gateway (Tier 1) or, for constrained Tier 0 devices, on the nearest Tier 1 gateway.
2. **Semantic Mapping and Validation.** The asset's native data model (vendor-specific register maps, proprietary JSON schemas, or legacy protocol payloads) is automatically mapped to SAREF ontologies using the semantic mediator's mapping library. The mapping is validated against the TSTO transformation rules to ensure that the asset's telemetry can be converted to market-ready signals. If the asset's data model cannot be fully mapped (e.g., a novel sensor type not yet covered by SAREF), the system flags the gap for the Standardisation Pillar's ontology-extension process (*6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) and provisions a provisional mapping with a confidence score.
3. **Policy Configuration and Consent Binding.** The asset's data-sharing policies are configured through a policy-as-code interface, binding the asset's telemetry to the applicable governance constraints (GDPR consent scope, DGA data-altruism provisions, NIS2 security requirements, national grid-code reporting obligations). For BTM assets, the data subject (prosumer) is presented with a human-readable consent interface that maps to the machine-readable policy, ensuring that the ethical and inclusivity requirements of *8. Stakeholder Engagement, IoT Ethics, and Inclusivity* are satisfied at the point of onboarding.
4. **Connector Provisioning and Interoperability Testing.** The DSC provisions the asset's connector agent with the appropriate policy engine, semantic mapping, and trust credentials. An automated interoperability test suite validates that the asset's telemetry can be ingested by the BTM Data Space, transformed through the TSTO pipeline, and consumed by at least one downstream application (e.g., a demand-response aggregator or a voltage-regulation swarm controller) without manual intervention.
5. **Orchestration-Layer Registration and SLA Binding.** The newly onboarded asset is registered with the computational orchestration layer's workload scheduler, making it available for federated learning participation, swarm-computing coordination, and flexibility-market bidding. A Service Level Agreement (SLA) is established, specifying the asset's availability, data-quality guarantees, and response-time commitments. The distributed observability plane begins monitoring the asset's telemetry quality and connector health.

This workflow is designed to be **vendor-neutral and extensible**: a new asset type requires only a SAREF-compliant description and a policy configuration, not a custom integration project. The Open Call mechanism described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity* provides SMEs and innovators with sandbox access to the Data Space infrastructure, enabling them to prototype and validate new IoT services against the HEDGE-IoT interoperability stack before committing to full production integration.

5.6 Cross-Data-Space Federation and Federated Application Mediation

The HEDGE-IoT Framework's most architecturally distinctive interoperability capability is **cross-Data-Space federation**: the ability to execute a single computational task - whether a federated learning job, a swarm-computing coordination cycle, or a flexibility-market clearing process - across multiple Data Space domains without collapsing their governance boundaries or introducing a centralised gateway. This capability is essential for the energy-specific federated and swarm mechanisms described in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*, where a single FL job may span BTM residential nodes, DSO regional aggregators, and TSO cloud analytics platforms.

Cross-Data-Space federation is mediated by the **Federated Application Broker (FAB)**, a component of the computational orchestration layer that operates as a policy-aware, stateless coordinator. The FAB does not store, process, or transit raw data; it orchestrates the exchange of encrypted model updates, swarm state registers, and market signals between Data Space Connectors, enforcing each domain's usage policy at the boundary. The federation protocol proceeds as follows:

- **Job Specification and Policy Negotiation.** A federated application (e.g., a congestion-forecasting model) is specified as a SAREF-anchored job descriptor that declares the required input data categories (expressed in SAREF terms), the desired output (a model, a forecast, a control signal), and the governance constraints (maximum data granularity, permitted processing locations, retention periods). The FAB broadcasts the job specification to the relevant Data Space Connectors, which evaluate it against their local policies and return a participation offer (or a policy-conflict notification).
- **Secure Channel Establishment.** Upon successful policy negotiation, the FAB establishes pairwise secure channels (mutually authenticated TLS 1.3 with DID-based identity verification) between the participating DSCs. For FL workloads, the channels carry homomorphically encrypted model updates; for swarm workloads, they carry signed state-register deltas; for market workloads, they carry digitally signed bid/offer messages.
- **Federated Execution with Boundary Enforcement.** The FL or swarm execution proceeds as described in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*, with the critical addition that every cross-domain data movement is intercepted by the sending DSC, evaluated against the sending domain's usage policy, and logged. If a policy violation is detected (e.g., a model update that, through gradient inversion, could reveal individual consumer data beyond the permitted differential-privacy budget), the DSC blocks the transmission and alerts the FAB.
- **Result Aggregation and Policy-Compliant Release.** The final model, forecast, or control signal is aggregated at the boundary of the requesting domain's Data Space. The requesting DSC evaluates the result against its own usage policy before releasing it to the consuming application. For example, a TSO congestion-forecasting model trained on BTM and DSO data is released to the TSO's market-clearing system only after the TSO DSC verifies that the model's output does not constitute personal data under GDPR and does not reveal critical-infrastructure topology under NIS2.

This federation protocol ensures that the computational orchestration layer's cross-Data-Space capability - identified as a critical binding element in *3. HEDGE-IoT Framework Architecture* - operates without a centralised gateway, preserving the edge-autonomy-by-default invariant and the data-sovereignty constraints mandated by EU regulation.

5.7 Integration with the Computational Orchestration Layer

The Interoperability Pillar does not operate as a standalone data-management subsystem; it is architecturally inseparable from the computational orchestration layer and the other three pillars of the HEDGE-IoT Framework. The Data Space Connectors, semantic mediator, and policy engines are the **data-plane substrate** upon which the orchestration layer's control-plane functions - workload scheduling, federated application brokering, and distributed observability - operate.

The integration is structured around three interfaces:

Orchestration-to-Data-Space Interface. The workload scheduler/placement engine queries the Data Space Connectors for available data assets, their governance constraints, and their quality classifications before placing a federated learning or swarm-computing task. The placement decision is constrained not only by latency, energy, and communication costs (as described in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*) but also by data-sovereignty constraints: a task requiring BTM residential data cannot be placed on a cloud node outside the EU/EEA, and a task requiring DSO operational

data cannot be placed on a node lacking NIS2 security certification. The scheduler treats governance constraints as hard constraints in its multi-objective optimisation, not as soft preferences.

Data-Space-to-Orchestration Interface. The Data Space Connectors publish real-time telemetry about data availability, quality, and policy status to the distributed observability plane. If a BTM Data Space Connector detects that a data subject has revoked consent, it notifies the orchestration layer, which triggers the removal of that subject's data from all active FL training sets and the retraining of affected models. If a DSO Data Space Connector detects a data-quality anomaly (e.g., a metering device drifting beyond its accuracy class), it flags the affected data streams, and the orchestration layer adjusts the FL client-selection algorithm to exclude degraded data sources.

Cross-Pillar Semantic Coupling. The semantic mediator within the orchestration layer is the shared component that serves both the Interoperability Pillar (mapping heterogeneous sensor vocabularies to SAREF ontologies for cross-vendor data exchange) and the Standardisation Pillar (*6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*, providing the semantic substrate for feature normalisation, job specification, and model metadata). The same SAREF-anchored ontology graph is used by the Ecosystem Enabling Pillar (*7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*) to describe flexibility resources in market-ready terms, ensuring that a flexibility signal generated by a BTM swarm controller is directly interpretable by a TSO balancing-market platform without intermediate translation.

This tight coupling ensures that the four pillars remain mutually reinforcing rather than parallel but unintegrated, fulfilling the architectural coherence objective stated in *1. Introduction and Motivation*.

5.8 Compliance, Ethics, and Inclusive Data Access

The Interoperability Pillar's data-governance architecture is designed to satisfy not only the technical and regulatory requirements of the DGA, GDPR, NIS2, and AI Act but also the ethical and inclusivity principles that the HEDGE-IoT Framework embeds as architectural constraints rather than add-ons (as established in *3. HEDGE-IoT Framework Architecture* and elaborated in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*).

Three design principles govern the ethical dimension of the Interoperability Pillar:

Purpose Limitation by Architecture. The policy-as-code enforcement mechanism ensures that data shared for one purpose (e.g., voltage regulation) cannot be repurposed for another (e.g., commercial profiling or insurance risk assessment) without a new, explicit consent transaction. The Data Space Connector's policy engine evaluates the purpose field of every access request against the asset's registered purpose limitations, and the distributed audit trail provides a tamper-evident record of all data uses. This architectural enforcement is stronger than procedural compliance, as it operates at the data-plane level and cannot be bypassed by application-layer workarounds.

Energy-Poverty Safeguards in Data Exchange. The HEDGE-IoT Data Space architecture includes a **vulnerability-classification layer** that identifies data subjects in energy-poverty situations (based on self-declared status, DSO-identified arrears patterns, or social-service referrals, all subject to strict purpose limitation). For these subjects, the policy engine automatically applies enhanced protections: raw telemetry is never shared beyond the BTM Data Space; only aggregated, anonymised flexibility signals are published to DSO and TSO domains; and the data subject retains the right to opt out of all data-sharing arrangements without loss of access to essential energy services. These safeguards are implemented as default policy templates within the DSC, not as optional configurations, ensuring that vulnerable consumers are protected by architecture rather than by individual choice.

Inclusive Access for SMEs and Innovators. The Data Space architecture is designed to be accessible to small and medium-sized enterprises (SMEs) and independent innovators, not only to incumbent utilities and large technology vendors. The Open Digital Contract format is open-source and royalty-free; the SAREF ontologies are publicly available; the Data Space Connector reference implementation is provided under an Apache 2.0 licence. The Open Call mechanism (described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*) provides SMEs with sandbox access to a production-equivalent Data Space environment, enabling them to develop, test, and validate new IoT services, analytics applications, and market mechanisms against the HEDGE-IoT interoperability stack before committing to full integration. The new-asset onboarding workflow (Section 5.5) is designed for self-service operation, eliminating the need for costly, vendor-mediated integration projects that historically excluded smaller participants.

Algorithmic Transparency and Explainability. Where Data Space-mediated data exchange feeds into AI/ML applications (federated learning models, swarm-computing control policies), the governance framework requires that the data provenance, feature-selection rationale, and model-metadata be recorded as Verifiable

Credentials alongside the model artefacts. This ensures that a DSO or TSO consuming a federated congestion forecast can trace the model's training data back to its source Data Space domains, verify the differential-privacy budget applied, and assess the model's suitability for its operational context. This transparency requirement aligns with the AI Act's high-risk-system provisions and supports the trust-building objectives of *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*.

5.9 Summary and Cross-Pillar Dependencies

The Interoperability Pillar provides the data-plane foundation upon which the HEDGE-IoT Framework's computational, standardisation, and ecosystem-enablement capabilities are built. By adopting and extending IDSA/Gaia-X Data Space architectures into an energy-specific, layered-sovereignty blueprint, the Pillar resolves the protocol/semantic heterogeneity and data-sovereignty barriers identified in *1. Introduction and Motivation* and the absence of an end-to-end BTM-to-TSO Data Space integration noted in *2. State of the Art and Related Work*. The semantic interoperability layer, anchored in SAREF and extended with the TSTO, FTO, and CDCO ontologies, bridges the gap between raw IoT telemetry and market-ready flexibility signals. The policy-as-code governance model enforces GDPR, DGA, NIS2, and AI Act compliance at the data-plane level, while the trust and provenance framework ensures verifiable data quality across domain boundaries.

The Pillar's cross-domain integration protocol and cross-Data-Space federation mechanism are architecturally inseparable from the computational orchestration layer and the Technology Facilitator Pillar (*4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*): FL and swarm workloads are instantiated within Data Space governance boundaries, and the semantic mediator provides the shared ontological substrate for feature normalisation, job specification, and model metadata. The Standardisation Pillar (*6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) supplies the canonical vocabularies and API specifications that make the Data Space's semantic interoperability possible, while the Ecosystem Enabling Pillar (*7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*) consumes the market-ready flexibility signals produced by the Data Space's progressive abstraction protocol. The ethical and inclusivity safeguards embedded in the governance model operationalise the commitments of *8. Stakeholder Engagement, IoT Ethics, and Inclusivity* at the architectural level.

Together, these mechanisms transform the HEDGE-IoT Data Space from a theoretical interoperability reference into a deployable, policy-enforced, vendor-neutral data-exchange infrastructure that reduces new-asset integration from months to days, enables cross-domain federated computation without centralised gateways, and converts the GW-scale latent flexibility of RES-hosting grids into a coordinated, privacy-preserving, market-ready resource.

6. Standardisation Pillar: SAREF and Open Standards for Energy IoT

6.1 Standardisation as a First-Class Architectural Pillar

The HEDGE-IoT Framework elevates standardisation from a retrofitted compliance exercise - characteristic of existing energy IoT platforms where SAREF adoption is grafted onto proprietary stacks, perpetuating vendor lock-in (as catalogued in *2. State of the Art and Related Work*) - to a first-class architectural pillar with binding influence on every interface, data artefact, and governance contract within the system. This positioning is a direct consequence of design invariant (iii) established in *3. HEDGE-IoT Framework Architecture*: open standards as the default interface. The Standardisation Pillar does not merely publish a conformance checklist; it prescribes the ontological substrate, wire-format contracts, API surface, and lifecycle governance rules through which all four pillars, all five tiers of the cloud/edge continuum, and all external stakeholders exchange information.

The strategic rationale is threefold. First, the five interlocking integration barriers identified in *1. Introduction and Motivation* - protocol/semantic heterogeneity, non-linear scalability demands, divergent latency requirements, data-sovereignty constraints, and persistent vendor lock-in - cannot be resolved by any single protocol or proprietary middleware. Only a layered, standards-anchored approach simultaneously addresses semantic, syntactic, and governance heterogeneity. Second, the GW-scale latent flexibility in RES-hosting grids remains invisible to DSOs and TSOs precisely because heterogeneous IoT assets speak incompatible vocabularies and expose non-interoperable interfaces; standardisation is the precondition for converting fragmented capacity into a coordinated, tradable resource. Third, the framework's commitment to vendor neutrality and long-term deployability (design invariant vi: empirical validation at every tier) demands that no implementation choice create a dependency on a single vendor's roadmap, pricing model, or end-of-life schedule.

The Standardisation Pillar operationalises these imperatives through four interlocking mechanisms: (i) a SAREF-anchored semantic ontology stack that provides the shared vocabulary for all data exchange; (ii) standardised data formats and wire protocols that ensure syntactic interoperability across heterogeneous connectivity technologies; (iii) open API specifications that define the contract surface between the computational orchestration layer, Data Space connectors, and external platforms; and (iv) a governance and evolution process that keeps the standards stack aligned with the pace of IoT, energy, and AI standardisation at European and international levels.

6.2 SAREF as the Semantic Backbone of the Energy IoT Continuum

The Smart Applications REFERENCE (SAREF) ontology, developed under ETSI TC SmartM2M and maintained through the ETSI TS 103 264 series, constitutes the semantic backbone of the HEDGE-IoT Framework. SAREF provides a vendor-neutral, domain-agnostic upper ontology for smart-device data, with energy-specific extensions (SAREF4ENERGY, SAREF4BLDG, SAREF4EV, SAREF4CITY) covering the full spectrum of assets from behind-the-meter smart meters and heat pumps to electric-vehicle chargers and distributed generation inverters. Within HEDGE-IoT, SAREF is not an optional annotation layer but the mandatory semantic substrate through which every data artefact, model specification, and market signal is expressed.

The framework extends the base SAREF ontology through three energy-specific ontological layers, introduced in *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration* and elaborated here in their standardisation context:

Telemetry-to-Signal Transformation Ontology (TSTO). The TSTO defines deterministic and probabilistic transformation pipelines that convert raw register-level telemetry (e.g., 1-second instantaneous power readings from a smart meter) into aggregated, confidence-annotated, SAREF-tagged signals suitable for DSO operational analytics or TSO market participation. Each transformation step is expressed as a SAREF-compatible class hierarchy with explicit input/output type constraints, temporal aggregation rules, and quality-attestation predicates. The TSTO ensures that a flexibility signal produced by a Tier 1 edge gateway in a BTM Data Space is semantically identical - regardless of the underlying sensor vendor or communication protocol - to the signal consumed by a Tier 3 TSO market-clearing engine.

Flexibility Taxonomy Ontology (FTO). The FTO classifies heterogeneous flexibility resources - residential battery storage, vehicle-to-grid (V2G) units, smart thermal loads, curtailable industrial processes, and distributed generation - into a vendor-neutral, market-agnostic taxonomy. Each flexibility class inherits from SAREF's `saref:Device` and `saref:Function` hierarchies and adds energy-specific attributes: dispatchability envelope (minimum/maximum power, ramp rate, duration), response-time class, location granularity, and market-participation eligibility. The FTO is the semantic contract that enables the Ecosystem Enabling Pillar

(*7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*) to aggregate heterogeneous resources into homogeneous, bid-ready flexibility portfolios without bespoke per-vendor mapping.

Cross-Domain Context Ontology (CDCO). The CDCO captures the physical, temporal, regulatory, and trust context in which a data asset or computational task operates. It extends SAREF's `saref:Location` and `saref:Time` classes with grid-topology references (feeder, substation, bidding zone), regulatory jurisdiction markers (national grid code, NIS2 classification, REMIT applicability), and trust-level annotations (W3C Verifiable Credential references, differential-privacy budget status). The CDCO is the semantic mechanism through which the policy-as-code enforcement described in *5. Interoperability Pillar* evaluates access requests in <5 ms: predicates are expressed as CDCO class restrictions, and the Data Space Connector's policy engine resolves them against the requesting entity's Verifiable Credentials.

All three ontological layers are published as W3C-compliant OWL 2 DL ontologies with SHACL shapes for instance-level validation. The semantic mediator within the computational orchestration layer (described in *4. Technology Facilitator Pillar*) executes runtime mappings from heterogeneous sensor vocabularies to these SAREF-anchored ontologies, enabling cross-vendor federated learning participation and swarm coordination without manual re-engineering. Critically, the ontologies are versioned under a semantic-versioning scheme (major.minor.patch) aligned with ETSI's SAREF release cadence, ensuring backward compatibility for deployed assets while permitting incremental extension.

6.3 Standardised Data Formats and Wire Protocols

Semantic interoperability is necessary but insufficient; the Standardisation Pillar also prescribes syntactic and transport-level standards that ensure data artefacts are parseable, serialisable, and transmissible across the heterogeneous connectivity landscape of RES-hosting distribution networks (cellular, LoRaWAN, PLC, satellite, Ethernet).

Structured data formats. All data artefacts exchanged within and across Data Space domains conform to JSON-LD (W3C Recommendation) with embedded SAREF context URIs, enabling both human-readable inspection and machine-parseable semantic resolution. For high-frequency telemetry streams (Tier 0 → Tier 1), the framework adopts a compact binary encoding - Protocol Buffers (proto3) with a SAREF-derived schema registry - reducing payload size by 60-80% relative to raw JSON while preserving full semantic fidelity through schema-ID references. For model artefacts (federated learning parameters, swarm policy configurations), the framework standardises on the ONNX (Open Neural Network Exchange) format for model serialisation and the MLflow model registry schema for metadata, ensuring that a model trained on a Tier 1 edge node is loadable by any Tier 3 cloud inference engine regardless of the underlying ML framework (PyTorch, TensorFlow, JAX).

Wire protocols and transport standards. The framework does not mandate a single transport protocol but defines a conformance profile per connectivity class: MQTT 5.0 with SAREF-tagged topic hierarchies for constrained IoT devices (LoRaWAN gateways, PLC modems); HTTPS/REST with OpenAPI 3.1 specifications for Tier 1 → Tier 2 and Tier 2 → Tier 3 exchanges; and gRPC with Protocol Buffers for latency-critical intra-edge swarm coordination (Tier 1 Tier 1 peer-to-peer). All transports are secured via mutually authenticated TLS 1.3 with DID-based identity binding, as specified in the cross-Data-Space federation protocol of *5. Interoperability Pillar*. The choice of transport is governed by the workload scheduler's connectivity-aware task partitioning (described in *4. Technology Facilitator Pillar*), which selects the protocol profile matching the available link characteristics without application-level reconfiguration.

Time synchronisation and ordering. Energy-system applications impose strict temporal semantics: voltage-regulation decisions must be timestamped to sub-millisecond precision, market-clearing signals must carry settlement-period validity windows, and federated learning gradients must be ordered within a bounded staleness window. The Standardisation Pillar mandates IEEE 1588 Precision Time Protocol (PTP) at Tier 0-1 for sub-microsecond synchronisation, NTPv4 with PTP fallback at Tier 2, and logical-clock ordering (Lamport timestamps augmented with hybrid logical clocks) for cross-Data-Space message sequencing where physical clock synchronisation is infeasible. All timestamps are expressed in ISO 8601 with explicit timezone and leap-second handling, and are embedded as SAREF `saref:hasTimestamp` properties on every data artefact.

6.4 Open API Specifications and Interface Contracts

The Standardisation Pillar defines the complete API surface through which the computational orchestration layer, Data Space connectors, edge gateways, and external platforms interact. These specifications are published as OpenAPI 3.1 documents and AsyncAPI 2.6 specifications, versioned under the same semantic-versioning scheme as the ontologies, and made available through a public, royalty-free specification repository.

Orchestration-layer APIs. The workload scheduler exposes a RESTful interface for task submission, placement queries, and status polling, with a gRPC streaming endpoint for real-time telemetry ingestion from the distributed observability plane. The federated application broker exposes a job-specification API (accepting SAREF-annotated FL job descriptors with per-participant policy constraints), a participation-registration API (DID-authenticated), and a result-release API (policy-gated). The semantic mediator exposes a mapping-resolution API that accepts heterogeneous sensor vocabularies and returns SAREF-anchored class instances, and a model-metadata API that registers and queries ONNX model artefacts with their training-data provenance (Verifiable Credential references).

Data Space Connector APIs. Each DSC exposes a standardised interface for: asset registration and DID issuance; policy configuration and consent binding (accepting Open Digital Contract documents); data publication and subscription (with SAREF-tagged payload validation); cross-domain federation negotiation (mutually authenticated TLS 1.3 handshake with DID-based identity); and audit-log retrieval. The five-stage self-service onboarding workflow described in *5. Interoperability Pillar* is fully automatable through these APIs, requiring only a SAREF-compliant device description (expressed as a JSON-LD document conforming to the FTO class hierarchy) and a policy configuration (expressed as an ODC with CDCO predicates).

Edge-gateway and device APIs. Tier 1 edge gateways expose a local REST/MQTT interface for Tier 0 device registration, local model inference, swarm-state register access, and connectivity-status reporting. The stigmergic coordination protocol (shared state registers) and gossip-based consensus mechanism (described in *4. Technology Facilitator Pillar*) are specified as a set of SAREF-annotated key-value operations with defined consistency semantics (eventual consistency for swarm state, strong consistency for BFT consensus rounds). Device manufacturers implement a minimal conformance profile - SAREF device description, MQTT 5.0 publish/subscribe with SAREF topic hierarchy, and TLS 1.3 mutual authentication - to achieve plug-and-play integration without vendor-specific SDKs.

External-platform integration APIs. The framework exposes standardised interfaces for integration with external energy-market platforms (e.g., EPEX SPOT, Nord Pool, national balancing-market systems), DSO/TSO SCADA and ADMS systems (IEC 61850-7-420 distributed energy resources model), and regulatory reporting systems (REMIT transaction reporting, GDPR data-subject access requests). These interfaces are expressed as SAREF-to-domain-specific mapping adapters, ensuring that the framework's internal SAREF-anchored representation is the single source of truth while external systems consume their native formats.

6.5 Alignment with Ongoing IoT and Energy Standardisation Efforts

The Standardisation Pillar is not a static artefact but a living alignment mechanism that tracks, adopts, and contributes to the evolving landscape of European and international IoT, energy, and AI standardisation. This alignment is essential to ensure that HEDGE-IoT deployments remain interoperable with the broader ecosystem over a 10-15-year asset lifecycle, during which individual standards will undergo multiple revisions.

ETSI and SAREF evolution. The framework tracks ETSI TC SmartM2M's SAREF release cadence (currently TS 103 264-1 through -12, with SAREF4ENERGY extensions under active development) and participates in the ETSI SAREF Task Force. HEDGE-IoT's TSTO, FTO, and CDCO ontologies are submitted as candidate extensions to the SAREF ecosystem, ensuring that energy-specific semantic constructs benefit from community review and are available to other SAREF adopters. The framework's ontology versioning scheme is synchronised with ETSI's release numbering to prevent semantic drift.

CEN/CENELEC and IEC energy standards. The FTO taxonomy is aligned with the CEN/CENELEC CLC/TC 57X (Smart Grids) work on distributed energy resource models and the IEC 61850-7-420 (Distributed Energy Resources) information model. The framework's DSO Operational Data Space interfaces conform to IEC 61968 (Common Information Model) for utility data exchange, and the TSO Market Data Space interfaces align with IEC 62357 (Common Information Model for power systems) and the ENTSO-E transparency platform data specifications. Where SAREF and IEC CIM overlap (e.g., device identification, measurement semantics), the framework defines explicit crosswalk mappings maintained as versioned SHACL shapes.

IDSA, Gaia-X, and Data Space standardisation. The Data Space architecture described in *5. Interoperability Pillar* is aligned with the IDSA Reference Architecture Model (RAM) v3.0 and Gaia-X Trust Framework, and tracks the ongoing work of the IDSA Technical Committee on connector specifications, policy language (ODRL-based), and federation protocols. The framework's Open Digital Contract format is a superset of the IDSA Contract Template, extended with energy-specific predicates (settlement-period validity, grid-code compliance, NIS2 classification). The framework contributes its energy-specific policy predicates and grid-topology-aware policy engine design to the IDSA working groups, ensuring that future IDSA releases incorporate energy-domain requirements.

AIOTI and EU digitalisation policy alignment. The Standardisation Pillar maintains active liaison with the Alliance for IoT Innovation (AIOTI) and its Energy IoT Cluster, as referenced in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*. The framework's standards stack is mapped against the requirements of the Data Governance Act, the AI Act (particularly the transparency and record-keeping obligations for high-risk AI systems in critical infrastructure), the NIS2 Directive (cybersecurity and incident-reporting interfaces), and the Digital Europe Programme's strategic objectives for sovereign data infrastructure. This mapping is maintained as a living traceability matrix, updated at each regulatory milestone, ensuring that standardisation decisions anticipate rather than react to regulatory changes.

WoT and device-level interoperability. At the device and edge-gateway level, the framework aligns with the W3C Web of Things (WoT) Thing Description (TD) specification for device discovery and metadata exchange. SAREF device descriptions are serialised as WoT TD documents with SAREF context URIs embedded in the `@context` field, enabling WoT-compliant discovery agents to locate and interpret HEDGE-IoT assets without framework-specific tooling. The WoT Binding Templates for MQTT, CoAP, and HTTP are adopted as the transport-level conformance profiles, ensuring that a WoT-compliant edge gateway can participate in HEDGE-IoT swarm coordination and FL participation without custom integration.

6.6 Vendor Neutrality, Conformance, and Long-Term Interoperability

The Standardisation Pillar operationalises vendor neutrality through a multi-layered conformance and certification regime that prevents any single implementation from becoming a de facto standard while ensuring that all conformant implementations achieve genuine interoperability.

Conformance profiles and certification tiers. The framework defines three conformance tiers: (i) *Semantic Conformance* - the implementation correctly serialises, deserialises, and validates data artefacts against the SAREF-anchored ontology stack (TSTO, FTO, CDCO) using published SHACL shapes; (ii) *Interface Conformance* - the implementation exposes and consumes the OpenAPI/AsyncAPI specifications for its role (device, edge gateway, DSC, orchestration component, external platform) with full request/response schema compliance; and (iii) *Behavioural Conformance* - the implementation exhibits the specified operational semantics (e.g., <5 ms policy evaluation, <500 ms swarm convergence, $<2\%$ MAPE delta-synchronisation accuracy loss, 24-hour consent-revocation purge attestation). Conformance is verified through a publicly available test suite (conformance test vectors, negative test cases, and performance benchmarks) and, for production deployments, through third-party certification by an accredited body.

Vendor-neutral reference implementations. The framework publishes open-source reference implementations of each component (edge gateway firmware, DSC software, orchestration-layer services, semantic mediator) under Apache 2.0 licensing. These reference implementations are not the only conformant path; they serve as the normative behavioural specification against which commercial implementations are validated. The reference implementations are maintained in a public repository with continuous integration against the conformance test suite, ensuring that the specification and the reference code remain synchronised.

Anti-lock-in architectural guarantees. The framework's architecture enforces vendor neutrality structurally: (i) no component holds a monopoly on semantic interpretation - the semantic mediator is a stateless, horizontally scalable service that any conformant implementation can instantiate; (ii) data artefacts are self-describing (JSON-LD with embedded SAREF context) and do not require vendor-specific parsers; (iii) model artefacts are in ONNX format, loadable by any conformant inference engine; (iv) policy documents are in ODRL/ODC format, evaluable by any conformant policy engine; and (v) the orchestration layer's workload scheduler treats all compute resources as fungible, preventing any single edge-cloud provider from becoming a single point of dependency. Migration between conformant implementations requires only re-registration of DIDs and re-binding of policy documents - no data reformatting, no model retraining, no application re-engineering.

Long-term interoperability and standards evolution. The framework's semantic-versioning scheme (major.minor.patch) and the ETSI SAREF release cadence provide a structured path for standards evolution without breaking deployed assets. Major version increments (e.g., SAREF 3.0 $\rightarrow$ 4.0) trigger a migration-assessment workflow: the semantic mediator maintains backward-compatible mapping adapters for at least two major versions, and the conformance test suite includes cross-version interoperability tests. The framework's governance process (described in Section 6.7) includes a standards-impact assessment for every proposed change, evaluating the effect on deployed assets, in-flight FL jobs, and active market contracts. This ensures that the 10-15-year asset lifecycle of energy IoT devices is not disrupted by standards evolution.

6.7 Standardisation Governance and Evolution Process

The Standardisation Pillar is governed by a transparent, multi-stakeholder process that balances the need for stability (deployed assets must not be disrupted) with the need for evolution (new asset classes, new market

mechanisms, new regulatory requirements must be accommodated).

Governance structure. The standards stack is maintained by a Standards Working Group (SWG) comprising representatives from DSOs, TSOs, technology vendors, academic institutions, consumer organisations, and EU standardisation bodies (ETSI, CEN/CENELEC, IDSA). The SWG operates under a public charter with defined decision-making rules (consensus-seeking with documented dissent), a quarterly meeting cadence, and a public issue tracker. Changes to the ontology stack, API specifications, or conformance profiles follow a Request-for-Comment (RFC) process with a minimum 60-day public review period, extended to 120 days for changes affecting deployed production systems.

Change management and impact assessment. Every proposed change undergoes a structured impact assessment covering: (i) semantic backward compatibility (can existing data artefacts be reinterpreted under the new specification?); (ii) interface compatibility (do existing API consumers require modification?); (iii) performance impact (does the change affect the <5 ms policy-evaluation budget, the <500 ms swarm-convergence target, or the $<2\%$ MAPE accuracy-loss bound?); (iv) regulatory alignment (does the change maintain compliance with GDPR, DGA, NIS2, AI Act, and applicable grid codes?); and (v) ecosystem impact (how many deployed assets, active FL jobs, and market contracts are affected?). The assessment is published alongside the RFC, and the SWG's decision (adopt, defer, reject) is recorded with rationale.

Contribution and extension mechanism. The framework's openness - prescribing interfaces, contracts, and governance rules rather than a monolithic stack, as established in *3. HEDGE-IoT Framework Architecture* - extends to the standards themselves. Any stakeholder may propose a new ontology extension (e.g., a SAREF4HEATPUMP extension for a novel thermal-storage technology), a new API endpoint (e.g., a peer-to-peer energy-trading interface for a new market design), or a new conformance profile (e.g., a satellite-connectivity profile for remote rural deployments). Proposals are evaluated by the SWG against the framework's design invariants and the existing ontology architecture. Accepted extensions are incorporated into the next minor or major release, with full conformance-test coverage and migration guidance.

Liaison with external standardisation bodies. The SWG maintains formal liaison relationships with ETSI TC SmartM2M (SAREF), CEN/CENELEC CLC/TC 57X (Smart Grids), IEC TC 57 (Power Systems Management), the IDSA Technical Committee (Data Spaces), the W3C WoT Interest Group, and the AIOTI Energy IoT Cluster. These relationships ensure bidirectional information flow: HEDGE-IoT's energy-specific requirements inform upstream standards development, and upstream standards releases are assessed for their impact on the HEDGE-IoT stack within one release cycle. The framework's contributions to ETSI SAREF (TSTO, FTO, CDCO as candidate extensions) and to IDSA (energy-specific policy predicates, grid-topology-aware policy engine design) exemplify this bidirectional engagement.

Open Call and SME integration pathway. The standardisation governance process is explicitly designed to be accessible to SMEs and innovators, consistent with the Open Call mechanism described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*. The conformance test suite, reference implementations, and API specifications are available without registration or licensing fees. The SWG reserves two seats for SME/innovator representatives, and the RFC process includes a dedicated "SME impact" assessment ensuring that proposed changes do not impose disproportionate compliance costs on smaller participants. The sandbox environment (production-equivalent Data Space) provides SMEs with a risk-free pathway to validate their SAREF-compliant device descriptions and API implementations before production deployment, reducing the integration barrier that historically excluded smaller participants from energy-IoT ecosystems.

6.8 Standardisation in Service of the Other Pillars

The Standardisation Pillar is architecturally inseparable from the other three pillars, providing the semantic and syntactic substrate upon which they operate. This coupling is not incidental but a deliberate design choice ensuring that standardisation is not a parallel, afterthought activity but the connective tissue that prevents the four pillars from remaining "parallel but unintegrated" (as characterised in *2. State of the Art and Related Work*).

The Technology Facilitator Pillar (*4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*) depends on the Standardisation Pillar for: SAREF-anchored feature normalisation (ensuring that a feature vector produced by a Tier 1 gateway in one BTM Data Space is semantically identical to one produced in another, enabling cross-vendor FL participation); SAREF-tagged job specifications (the federated application broker accepts FL job descriptors expressed as SAREF class instances with CDCO policy constraints); and ONNX/MLflow model metadata (ensuring that a model trained under one FL round is interpretable and loadable by any conformant inference engine in subsequent rounds).

The Interoperability Pillar (*5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*)

depends on the Standardisation Pillar for: the TSTO, FTO, and CDCO ontologies that define the semantic content of every data asset flowing through Data Space connectors; the Open Digital Contract format (ODRL superset) that encodes policy-as-code predicates; the WoT Thing Description format for device discovery; and the conformance profiles that ensure all DSC implementations exhibit the specified <5 ms policy-evaluation and <24-hour consent-revocation behaviours.

The Digital Energy Ecosystem Enabling Pillar (*7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*) depends on the Standardisation Pillar for: the FTO taxonomy that classifies heterogeneous flexibility resources into market-agnostic, bid-ready categories; the SAREF-to-IEC 61850-7-420 and SAREF-to-IEC 62357 crosswalk mappings that translate internal flexibility signals into DSO/TSO-native formats; and the REMIT-compliant transaction-reporting interfaces that ensure market participation meets regulatory obligations.

This tight coupling ensures that the Standardisation Pillar is not a documentation exercise but an operational dependency: a change to the ontology stack propagates through the semantic mediator to all FL job specifications, all DSC policy evaluations, and all market-signal transformations within a single release cycle. The framework's design invariants - particularly "open standards as the default interface" and "data never leaves its sovereign boundary unprocessed" - are enforced at the standards level, making compliance a structural property of conformant implementations rather than a procedural obligation subject to audit.

7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities

7.1 Positioning of the Ecosystem Enabling Pillar Within the HEDGE-IoT Architecture

The Digital Energy Ecosystem Enabling Pillar constitutes the fourth and outermost pillar of the HEDGE-IoT Framework, translating the technical capabilities established by the Technology Facilitator, Interoperability, and Standardisation Pillars into tangible economic, operational, and societal value across the energy system. Where the preceding three pillars construct the computational, semantic, and governance substrate - federated learning and swarm computing at the grid edge (*4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*), layered-sovereignty Data Spaces with progressive abstraction protocols (*5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*), and SAREF-anchored ontological and interface standards (*6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) - the Ecosystem Enabling Pillar defines the market mechanisms, service portfolios, resilience architectures, and scalability pathways through which heterogeneous IoT assets become coordinated, privacy-preserving, market-ready flexibility resources.

This pillar operationalises the framework's central economic proposition articulated in *1. Introduction and Motivation*: that the GW-scale dispatchable capacity embedded in aggregated residential storage, vehicle-to-grid (V2G) systems, smart thermal loads, and curtailable industrial processes - collectively representing the most cost-effective alternative to physical network reinforcement - can be converted from a latent, fragmented, and contractually locked resource into a tradable, coordinated, and continuously optimised asset class. The conversion is not achieved through a single intervention but through the emergent coupling of all four pillars: FL-derived forecasts and swarm-based execution provide the informational and operational foundations; Data Space governance and progressive abstraction ensure that market signals respect data sovereignty and consent boundaries; SAREF-to-IEC/REMIT crosswalks render flexibility signals legible to existing market-clearing platforms; and the computational orchestration layer coordinates cross-Data-Space market processes without a centralised gateway.

The pillar addresses three interdependent ecosystem outcomes: (i) measurable increases in RES-hosting capacity through activated flexibility; (ii) enhanced grid resilience through diversified, interoperable, and geographically distributed energy services; and (iii) the unlocking of new market opportunities for prosumers, aggregators, SMEs, and system operators. Each outcome is underpinned by the scalability characteristics of the distributed data platforms and infrastructure that constitute the framework's operational backbone.

7.2 Increasing RES-Hosting Capacity Through Activated Flexibility

7.2.1 From Latent to Activated Flexibility: The Conversion Pipeline

The fundamental constraint on RES-hosting capacity in distribution networks is not the absence of flexibility but its invisibility, fragmentation, and contractual inaccessibility. As established in *1. Introduction and Motivation*, aggregated residential storage, V2G, smart thermal loads, and curtailable industrial processes collectively represent GW-scale dispatchable capacity in many Member States, yet this capacity remains locked behind incompatible contracts, proprietary platforms, and protocol/semantic heterogeneity. The Ecosystem Enabling Pillar defines a four-stage conversion pipeline that transforms latent flexibility into a continuously available, market-ready resource:

Stage 1 - Discovery and Classification. IoT assets across Tier 0 (Device/Sensor) and Tier 1 (Edge Gateway) register their technical capabilities - power rating, ramp rate, duration, availability windows, storage state-of-charge, thermal inertia, and V2G bidirectional capability - through the five-stage self-service onboarding workflow described in *5. Interoperability Pillar*. The Flexibility Taxonomy Ontology (FTO), anchored in SAREF and maintained under the governance of the Standards Working Group (*6. Standardisation Pillar*), classifies each resource into a vendor-neutral, market-agnostic taxonomy. This classification is not a one-time event: the Telemetry-to-Signal Transformation Ontology (TSTO) continuously updates availability and capability annotations as operational conditions evolve (e.g., a V2G vehicle's state-of-charge trajectory, a thermal load's pre-heating window).

Stage 2 - Aggregation and Progressive Abstraction. The progressive abstraction and aggregation protocol, enforced by Data Space Connectors (DSCs) at each governance boundary, publishes flexibility signals at granularity determined by operational need and consent scope. Per-feeder aggregates flow to DSO Operational Data Spaces; per-bidding-zone summaries flow to TSO and Market Data Spaces; per-household data never leaves the BTM Data Space. This federated visibility model, a direct implementation of design invariant

(ii) from *3. HEDGE-IoT Framework Architecture* - "data never leaves its sovereign boundary unprocessed" - ensures that system operators gain the aggregate flexibility visibility required for RES-hosting decisions without any single entity possessing a complete centralised view of individual consumer behaviour.

Stage 3 - Predictive Forecasting and Swarm-Based Execution. FL models trained through the hierarchical topology (local → regional → global aggregation) described in *4. Technology Facilitator Pillar* generate probabilistic flexibility forecasts that account for non-IID data distributions, intermittent connectivity, and consumer-archetype diversity. These forecasts feed the swarm-computing coordination layer, where stigmergic protocols and gossip-based consensus enable 500+ node feeders to reach voltage-regulation and congestion-management equilibria within <500 ms. The co-evolutionary learning-decision loop - where FL models provide the predictive prior for swarm control policies and swarm execution outcomes generate ground-truth labels for the next FL training round - ensures that flexibility activation accuracy improves continuously without centralised data collection.

Stage 4 - Market Integration and Settlement. SAREF-to-IEC 62357 and SAREF-to-REMIT crosswalk mappings, maintained as part of the standards stack (*6. Standardisation Pillar*), translate the framework's internal SAREF-anchored flexibility representation into the native formats consumed by existing balancing markets, capacity markets, and local flexibility markets. The Open Digital Contract (ODC) attached to each flexibility signal encodes settlement terms, temporal validity windows tied to settlement periods, and derivative-data ownership, enabling automated market participation without bespoke integration projects.

7.2.2 Quantified Impact on RES-Hosting Capacity

The conversion pipeline yields measurable increases in RES-hosting capacity through three mechanisms. First, **congestion relief through distributed flexibility activation**: swarm-based voltage regulation and congestion management at Tier 1 edge gateways defer or eliminate the need for physical network reinforcement (conductor upgrades, transformer replacements, STATCOM installations) that would otherwise be triggered by RES export peaks. Pilot deployments (detailed in *9. Scalability, Replicability, and Validation*) demonstrate that activated flexibility at the feeder level can absorb 15-30% of peak RES export without network reinforcement, directly translating into deferred capital expenditure for DSOs.

Second, **curtailment reduction through predictive flexibility scheduling**: FL-derived forecasts of RES generation and demand, combined with swarm-based dispatch of aggregated storage and thermal loads, reduce RES curtailment by shifting consumption into surplus-generation windows. The hierarchical FL topology mirrors the physical grid hierarchy, ensuring that forecasts are spatially and temporally relevant to the specific feeder or bidding zone in which curtailment risk is highest.

Third, **inertia and frequency-response provision through distributed resources**: the FTO taxonomy explicitly classifies resources capable of providing synthetic inertia, primary frequency response, and secondary reserve. Swarm-computing execution, with BFT consensus tolerating up to $(N-1)/3$ faulty or adversarial agents, enables sub-second response times that meet TSO frequency-control requirements without reliance on synchronous generation. This capability is particularly critical in RES-dominated systems where the declining synchronous generation fleet erodes natural system inertia.

7.2.3 RES-Hosting Capacity as a Continuous, Self-Optimising Resource

Unlike static network reinforcement, activated flexibility is a continuous, self-optimising resource whose capacity evolves with the IoT asset base, consumer behaviour, and RES penetration trajectory. The computational orchestration layer's workload scheduler treats flexibility activation as a multi-objective optimisation problem (latency, energy, communication, fairness) with constraint filtering, re-solving placement and dispatch decisions as new assets onboard, connectivity conditions change, or market signals shift. The distributed observability plane provides per-task telemetry and SLA enforcement, enabling DSOs and TSOs to monitor realised flexibility delivery against contracted commitments and to trigger targeted retraining of FL models when forecast accuracy degrades. This continuous optimisation loop ensures that RES-hosting capacity grows in lockstep with the expanding IoT asset base rather than requiring periodic re-planning cycles.

7.3 Enhancing Grid Resilience Through Diversified and Interoperable Energy Services

7.3.1 Resilience as a Multi-Dimensional, Distributed Property

Grid resilience in the context of HEDGE-IoT is not a single metric but a multi-dimensional property encompassing: (i) **robustness** - the ability to maintain service continuity under component failures, cyber-attacks, and extreme weather events; (ii) **redundancy** - the availability of alternative resources and communication paths

when primary assets or links are compromised; (iii) **recovery speed** - the time required to restore normal operation following a disruption; and (iv) **adaptive capacity** - the ability to reconfigure resource allocation and service delivery in response to evolving threat landscapes. The Ecosystem Enabling Pillar enhances all four dimensions through the diversification and interoperability of energy services enabled by the framework's architecture.

7.3.2 Service Diversification Through Heterogeneous IoT Assets

The framework's explicit treatment of heterogeneous connectivity (cellular, LoRaWAN, PLC, satellite) as a first-class design constraint, rather than an exception, is the foundational resilience mechanism. In a RES-hosting distribution network, no single communication technology is universally available or universally reliable: LoRaWAN provides low-bandwidth coverage in rural areas but is vulnerable to interference; cellular networks offer higher bandwidth but are susceptible to congestion during extreme events; PLC operates over the power line itself but is disrupted by line faults; satellite links provide last-resort connectivity but with elevated latency. The computational sharing layer's connectivity-aware task partitioning and store-and-forward gradient buffering ensure that edge nodes continue executing locally trained models and swarm decisions during connectivity loss, with delta-synchronisation bounding accuracy loss to <2% MAPE upon reconnection.

Service diversification extends beyond communication to the energy resources themselves. The FTO taxonomy's vendor-neutral classification of residential storage, V2G, smart thermal loads, curtailable industrial processes, and distributed generation ensures that resilience services are not concentrated in a single asset class or vendor ecosystem. A cyber-attack targeting a specific V2G platform, for example, degrades only the V2G subset of flexibility; residential storage and thermal loads continue to provide congestion relief and voltage support through their independent Data Space governance boundaries. The graceful degradation invariant (design invariant iv) ensures that the loss of any subset of assets or services reduces system performance proportionally rather than catastrophically.

7.3.3 Interoperability as a Resilience Multiplier

The layered-sovereignty Data Space architecture (*5. Interoperability Pillar*) provides a structural resilience advantage over monolithic or vendor-locked platforms. Because each governance domain (BTM, DSO Operational, TSO/Market, Cross-Domain Federation) is instantiated through an independent DSC with a grid-topology-aware policy engine, the compromise or failure of one domain does not cascade to others. A DSO Operational Data Space experiencing a NIS2-reportable incident can be isolated through policy-as-code enforcement at the DSC boundary, preventing lateral movement into BTM Data Spaces or TSO market platforms while preserving the operational autonomy of edge nodes.

Cross-Data-Space federation, mediated by the Federated Application Broker (FAB), provides an additional resilience layer: because FL jobs, swarm-computing cycles, and market-clearing processes span governance boundaries through policy-aware, stateless coordination rather than a centralised gateway, there is no single point of failure whose compromise would halt all cross-domain operations. If a DSC in one domain becomes unavailable, the FAB's stateless protocol allows in-flight tasks to be re-routed through alternative federation paths, and edge nodes continue executing locally trained models and swarm decisions independently.

7.3.4 Resilience Services Portfolio

The Ecosystem Enabling Pillar defines a portfolio of resilience services that can be procured, contracted, and executed through the framework's market mechanisms:

- **Islanding detection and microgrid formation**: Swarm-computing execution at Tier 1 gateways, informed by FL-derived network-state models, enables sub-second detection of islanding conditions and autonomous formation of microgrids from available distributed generation, storage, and controllable loads. The BFT consensus protocol ensures that islanding decisions are made by a quorum of edge nodes, preventing single-node failures or adversarial behaviour from triggering spurious islanding.
- **Black-start and restoration sequencing**: FL models trained on historical restoration data generate optimised black-start sequences that prioritise critical loads, coordinate distributed generation ramp-up, and manage inrush currents from aggregated storage. The progressive abstraction protocol ensures that restoration commands are issued at the appropriate governance level (feeder-level to DSO, zone-level to TSO) without exposing individual consumer data.
- **Cyber-resilience through decentralised trust**: The decentralised trust framework based on W3C Verifiable Credentials and Decentralised Identifiers provides cryptographic provenance, quality classification, and governance classification for every data asset. Spoofed or degraded telemetry is detected and blocked

at the DSC policy engine before entering downstream analytics or market-clearing processes, preventing adversarial data injection from propagating through the ecosystem.

- **Physical-resilience augmentation through predictive maintenance**: FL models trained on IoT sensor telemetry (vibration, temperature, partial discharge, insulation resistance) predict component degradation and failure risk at the asset level. Predictive maintenance scheduling, coordinated through the workload scheduler's multi-objective optimisation, reduces unplanned outages and extends asset lifecycles, directly enhancing the physical robustness dimension of resilience.

7.3.5 Resilience Metrics and Continuous Assessment

The distributed observability plane provides real-time telemetry on resilience-relevant indicators: edge-node availability, swarm-consensus convergence times, FL model accuracy under degraded connectivity, DSC policy-enforcement latency, and cross-Data-Space federation task completion rates. These indicators are aggregated (at appropriate abstraction levels) into resilience dashboards accessible to DSOs and TSOs through their respective Data Spaces, enabling continuous assessment of system resilience posture and identification of emerging vulnerabilities before they manifest as service disruptions.

7.4 Unlocking New Market Opportunities and Flexibility Monetisation

7.4.1 Market Architecture: From Centralised Clearing to Federated, Multi-Level Participation

The HEDGE-IoT Framework does not propose a replacement for existing energy market structures but rather a digital layer that enables new participants, new products, and new clearing mechanisms to operate within and alongside existing markets. The market architecture enabled by the Ecosystem Enabling Pillar operates at three levels:

Local flexibility markets (Tier 1-2): At the distribution level, the framework enables local flexibility markets in which DSOs procure congestion-relief, voltage-support, and RES-absorption services from aggregated BTM resources. The FTO taxonomy provides the standardised product definitions; the progressive abstraction protocol ensures that DSOs see per-feeder aggregates without per-household data; and the ODC-based settlement mechanism automates payment and compliance verification. Swarm-computing execution enables real-time dispatch of procured flexibility, with BFT consensus ensuring that dispatch commands are validated by a quorum of edge nodes before actuation.

Regional and national balancing markets (Tier 2-3): At the transmission level, aggregated flexibility from multiple DSO domains is presented to TSOs and balancing-market operators through the TSO and Market Data Space. SAREF-to-IEC 62357 crosswalk mappings translate flexibility signals into the formats consumed by existing balancing-market platforms (e.g., XBID, PICASSO, MARI). The Cross-Domain Context Ontology (CDCO) captures the physical, temporal, regulatory, and trust context required for TSOs to assess the deliverability and quality of flexibility bids. FL-derived forecasts provide the probabilistic confidence intervals that market operators require for risk-adjusted clearing.

Peer-to-peer and community energy markets (Tier 0-1): At the prosumer level, the framework enables P2P energy trading and community energy schemes in which households, EV owners, and small commercial consumers trade surplus generation, storage capacity, and demand flexibility directly. Swarm-computing execution operates even during backhaul outages, ensuring that local trading continues during network disruptions. The BTM Data Space governance model, with its consent-based data sharing and vulnerability-classification safeguards, ensures that P2P trading does not compromise data-subject rights or expose energy-poverty-affected consumers to adverse market dynamics.

7.4.2 New Market Products and Revenue Streams

The framework's semantic and computational infrastructure enables the creation of market products that are infeasible under existing platform architectures:

- **Granularity-tiered flexibility products**: The progressive abstraction protocol enables the same underlying flexibility resource to be marketed at different granularities to different buyers - a residential storage unit's aggregate availability is sold as a feeder-level congestion-relief product to the DSO, while its contribution to a bidding-zone frequency-response pool is sold to the TSO, without either buyer seeing the other's data or the consumer's individual profile.
- **Resilience-as-a-service contracts**: DSOs and TSOs can procure resilience services (islanding capability, black-start support, cyber-resilience assurance) as contracted products from distributed IoT asset

pools, with SLA enforcement through the distributed observability plane and penalty/reward mechanisms encoded in ODCs.

- **Predictive analytics-as-a-service**: FL models trained across multiple DSO domains (without raw data leaving any domain) generate congestion forecasts, RES-generation predictions, and demand-response potential estimates that can be licensed to market participants, aggregators, and network planners as subscription services.
- **Data-space participation fees and connector-as-a-service**: The DSC infrastructure, with its policy-as-code enforcement and progressive abstraction capabilities, can be offered as a managed service to smaller DSOs, aggregators, and market participants that lack the resources to deploy and operate their own Data Space infrastructure, creating a recurring revenue stream for infrastructure providers while lowering barriers to ecosystem participation.

7.4.3 Lowering Barriers to Market Entry for SMEs and Innovators

A critical ecosystem-enabling function is the reduction of market-entry barriers for SMEs, innovators, and new market participants. The five-stage self-service onboarding workflow (*5. Interoperability Pillar*) reduces new-asset integration from months to days, requiring only a SAREF-compliant device description and a policy configuration rather than a custom integration project. The Open Call mechanism provides SMEs and innovators with sandbox access to a production-equivalent Data Space environment, eliminating the vendor-mediated integration costs that historically excluded smaller participants. The conformance test suite, reference implementations, and API specifications are available without registration or licensing fees, with two Standards Working Group seats reserved for SME/innovator representatives (*6. Standardisation Pillar*).

The framework's vendor-neutral architecture - enforced through the three-tier conformance regime (Semantic, Interface, Behavioural) and the fungibility of compute resources in the workload scheduler - ensures that no single vendor can capture the value chain. An SME developing a novel V2G optimisation algorithm can deploy it as an FL job or swarm policy within the framework without negotiating integration contracts with a platform vendor, a DSO, or a TSO. The Federated Application Broker mediates the job across governance boundaries; the semantic mediator maps the algorithm's inputs and outputs to SAREF ontologies; and the ODC-based settlement mechanism handles revenue distribution. This structural openness is the precondition for a competitive, innovative ecosystem rather than a consolidated oligopoly.

7.4.4 Regulatory Alignment and Market Compliance

All market mechanisms enabled by the Ecosystem Enabling Pillar are designed for compliance with the EU regulatory landscape: REMIT for market transparency and reporting; the Electricity Market Design Regulation for consumer protection and active participation rights; the Data Governance Act and GDPR for data sharing and consent; the AI Act for algorithmic transparency and risk classification; and NIS2 for operational security. The living traceability matrix maintained by the Standards Working Group maps the framework's market mechanisms against these regulatory requirements, updated at each regulatory milestone. The ODC-based settlement mechanism encodes regulatory compliance predicates (purpose limitation, temporal restriction, geographic restriction, consent revocation with 24-hour purge attestation) as machine-readable constraints evaluated at the DSC policy engine in <5 ms, ensuring that market participation is compliant by construction rather than by procedural audit.

7.5 Scalability of Distributed Data Platforms and Infrastructure for Highly Distributed Energy Services

7.5.1 The Scalability Imperative: From Thousands to Millions of Assets

The energy system's IoT asset base is growing at a rate that renders centralised architectures economically and operationally untenable. A single DSO domain may encompass 10 -10 BTM IoT devices; a national TSO domain may aggregate data from 10 -10 distributed resources. Streaming raw telemetry from this asset base to centralised data lakes is economically prohibitive at distribution-scale volumes and operationally fragile under network variability, as established in *1. Introduction and Motivation*. The Ecosystem Enabling Pillar's scalability strategy is therefore predicated on the distributed data platforms and infrastructure defined by the framework's architecture, not on the assumption of ever-increasing centralised compute and storage capacity.

7.5.2 Horizontal Scalability of Data Space Infrastructure

The layered-sovereignty Data Space architecture scales horizontally through three mechanisms:

Domain-level horizontal scaling. Each governance domain (BTM, DSO Operational, TSO/Market, Cross-Domain Federation) can be instantiated as multiple independent DSC instances, each serving a subset of the domain's data assets. A DSO Operational Data Space serving 500,000 BTM devices can be partitioned into 50 DSC instances, each serving 10,000 devices, with the grid-topology-aware policy engine ensuring that cross-partition queries are resolved through the progressive abstraction protocol rather than raw-data joins. Adding capacity requires provisioning additional DSC instances and registering them with the FAB - no re-architecture, no data migration, no application re-engineering.

Cross-domain federation scaling. The FAB's stateless coordination protocol scales to arbitrary numbers of participating Data Spaces because each federation task is self-contained: job specification, policy negotiation, secure channel establishment, federated execution, and result release are executed per-task without persistent cross-domain state. The O(log N / log f) convergence of gossip-based consensus in swarm-computing cycles ensures that coordination overhead grows logarithmically with the number of participating edge nodes, not linearly.

Edge-tier scaling through inter-edge resource pooling. The intra-device and inter-edge offloading mechanisms (*4. Technology Facilitator Pillar*) enable Tier 1 edge gateways to pool compute resources dynamically, preventing the need for uniform hardware upgrades across thousands of BTM nodes. A gateway experiencing a transient compute load (e.g., during a swarm-consensus cycle for a 500-node feeder) can offload sub-tasks to neighbouring gateways with available capacity, governed by the workload scheduler's four-objective optimisation. This peer-to-peer resource pooling scales the effective compute capacity of the edge tier without centralised orchestration.

7.5.3 Vertical Scalability: From Pilot to Continental Deployment

The framework's architecture supports vertical scaling from single-feeder pilots to continental deployments through the hierarchical structure of the cloud/edge continuum. Tier 0 (Device/Sensor) and Tier 1 (Edge Gateway) scale with the physical IoT asset base; Tier 2 (DSO Operational) scales with the number of DSO domains; Tier 3 (Cloud/TSO) scales with the number of TSO and market domains; Tier 4 (Ecosystem/Market) scales with the number of cross-domain federation participants. Each tier scales independently, governed by its own Data Space and policy engine, with the computational orchestration layer providing the binding coordination.

The FL hierarchical topology (local → regional → global aggregation) mirrors this physical hierarchy, ensuring that model training scales with the asset base without requiring any single entity to possess all raw data. The asynchronous partial participation mechanism, tolerating 15-40% offline clients, ensures that FL convergence is robust to the heterogeneous and intermittent connectivity characteristic of large-scale RES-hosting networks. The topology-aware client selection mechanism ensures spatial and consumer-archetype diversity in each FL round, preventing model bias toward well-connected or homogeneous asset clusters.

7.5.4 Infrastructure Scalability: Compute, Storage, and Communication

Compute scalability. The workload scheduler's multi-objective placement optimisation (latency, energy, communication, fairness) with constraint filtering and genetic-algorithm search (<100 ms for 10,000 tasks) ensures that compute workloads are distributed across the edge-cloud continuum in proportion to available capacity, not concentrated at the cloud tier. The fungibility of compute resources - enforced by the vendor-neutral conformance regime - means that adding compute capacity at any tier (a new edge gateway, a new cloud instance, a new DSC) immediately increases the system's effective capacity without re-architecting the orchestration layer.

Storage scalability. The "data never leaves its sovereign boundary unprocessed" invariant ensures that storage requirements scale with the number of governance domains, not with the total volume of raw telemetry. Raw IoT telemetry is stored at Tier 0-1 (device and edge), aggregated signals at Tier 2 (DSO), and market-ready flexibility signals at Tier 3 (TSO/Market). The progressive abstraction protocol ensures that each tier stores only the granularity required for its operational and market functions, preventing the storage explosion that would result from centralising all raw telemetry.

Communication scalability. The Protocol Buffers (proto3) transport conformance profile with a SAREF-derived schema registry provides 60-80% payload reduction for high-frequency telemetry, directly reducing the bandwidth requirements of edge-to-cloud and inter-domain exchanges. The federated feature selection mechanism in FL reduces communication overhead by 60-80% by transmitting only the most informative feature gradients. Store-and-forward gradient buffering for LoRaWAN and satellite links decouples communication scheduling from computation scheduling, preventing communication bottlenecks from stalling FL training or swarm execution.

7.5.5 Scalability Under Adverse Conditions

The framework's scalability is not contingent on ideal network conditions. The graceful degradation invariant (design invariant iv) ensures that the system continues to operate - albeit at reduced performance - under connectivity loss, node failure, and policy-engine overload. Edge nodes continue executing locally trained models and swarm decisions during extended outages; DSCs queue and batch cross-domain data movements during communication congestion; the workload scheduler re-optimises placement when compute capacity is lost. The distributed observability plane provides real-time telemetry on degradation status, enabling operators to assess the system's resilience posture and trigger contingency procedures before service-level thresholds are breached.

7.6 Ecosystem Governance, Open Call, and Inclusive Participation

7.6.1 Multi-Stakeholder Ecosystem Governance

The Ecosystem Enabling Pillar establishes a governance model for the digital energy ecosystem that is transparent, multi-stakeholder, and resistant to capture by any single participant class. The Standards Working Group (DSOs, TSOs, vendors, academia, consumer organisations, EU bodies) operates under a public charter with RFC processes (60-120-day review) and structured impact assessments, ensuring that ecosystem rules - market product definitions, flexibility taxonomy extensions, policy predicates, settlement mechanisms - evolve through a transparent, consensus-based process rather than through unilateral vendor or operator decisions.

The living traceability matrix, mapping the framework's ecosystem mechanisms against GDPR, DGA, NIS2, AI Act, and applicable grid-code requirements, is publicly accessible and updated at each regulatory milestone. This transparency ensures that ecosystem participants - particularly SMEs, consumer organisations, and energy-poverty advocacy groups - can verify that market mechanisms and data-sharing arrangements remain compliant with evolving regulatory requirements without relying on vendor self-certification.

7.6.2 The Open Call Mechanism as an Ecosystem Entry Point

The Open Call mechanism, described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*, provides SMEs, innovators, and research institutions with sandbox access to a production-equivalent Data Space environment. This mechanism serves a dual ecosystem-enabling function: it lowers the technical and financial barriers to ecosystem participation (eliminating the vendor-mediated integration costs that historically excluded smaller participants), and it provides a controlled environment for testing novel energy services, FL models, swarm policies, and market mechanisms before production deployment. The Open Call's integration with the five-stage self-service onboarding workflow means that a successful sandbox deployment can be promoted to production with minimal additional effort, accelerating the innovation-to-deployment cycle.

7.6.3 Inclusive Participation and Energy-Poverty Safeguards

The Ecosystem Enabling Pillar embeds inclusive participation as an architectural constraint, not an optional configuration. The vulnerability-classification layer automatically applies enhanced protections for energy-poverty-affected data subjects: raw telemetry never shared beyond BTM; aggregated anonymised signals only; opt-out without loss of essential services. These safeguards, operationalised at the data-plane level through the DSC policy engine (*5. Interoperability Pillar*), ensure that market participation and flexibility activation do not impose disproportionate burdens on vulnerable consumers.

The framework's support for community energy schemes and P2P trading at Tier 0-1 enables collective participation models in which energy-poverty-affected households can pool their flexibility resources and share market revenues, reducing individual exposure to market volatility while maintaining data sovereignty within the BTM Data Space. Algorithmic transparency requirements - mandating that data provenance, feature-selection rationale, and model metadata be recorded as Verifiable Credentials alongside model artefacts - enable consumer organisations and regulatory authorities to audit whether FL models and swarm policies systematically disadvantage specific consumer archetypes.

7.7 Cross-Pillar Dependencies and Emergent Ecosystem Value

7.7.1 Architectural Inseparability

The Ecosystem Enabling Pillar is architecturally inseparable from the other three pillars, consuming their outputs and providing the economic and operational context that justifies their investment. The Technology Facilitator Pillar's FL-derived forecasts and swarm-based execution constitute the informational and operational foundations for market-ready flexibility activation; the Interoperability Pillar's progressive abstraction protocol

and Data Space governance produce the market-ready flexibility signals that the Ecosystem Enabling Pillar consumes; and the Standardisation Pillar's FTO taxonomy and SAREF-to-IEC/REMIT crosswalks render those signals legible to existing market-clearing platforms. This tight coupling ensures that the four pillars remain mutually reinforcing rather than parallel but unintegrated, fulfilling the framework's core proposition stated in *1. Introduction and Motivation*.

7.7.2 Emergent Ecosystem Value

The value of the Ecosystem Enabling Pillar exceeds the sum of its individual market mechanisms, resilience services, and scalability pathways. Three emergent properties arise from the cross-pillar integration:

First, the self-reinforcing flexibility flywheel. As more IoT assets onboard through the self-service workflow, the aggregate flexibility pool grows, improving FL model accuracy (more training data, greater spatial diversity) and swarm-computing effectiveness (more nodes for consensus, more resources for dispatch). Improved accuracy and effectiveness increase the value of flexibility in local and balancing markets, attracting further asset onboarding. The progressive abstraction protocol ensures that this flywheel operates without any single entity accumulating a centralised data monopoly, preserving the ecosystem's openness and competitiveness.

Second, the resilience-diversification feedback loop. As the ecosystem's service portfolio diversifies across asset classes, communication technologies, and governance domains, the system's resilience improves, reducing the risk premium that DSOs and TSOs must price into flexibility procurement. Lower risk premiums increase the economic attractiveness of flexibility activation relative to physical reinforcement, accelerating the conversion of latent flexibility into activated capacity and further diversifying the resource base.

Third, the standardisation-network-effect virtuous cycle. As more participants adopt the SAREF-anchored ontologies, conformance profiles, and API specifications, the network effect increases the value of standardisation for all participants, attracting further adoption and reducing the marginal cost of integration for new entrants. The Standards Working Group's transparent governance process and the Open Call mechanism ensure that this virtuous cycle is not captured by incumbent vendors but remains open to SMEs, innovators, and new market participants.

7.7.3 Validation and Replicability of Ecosystem Outcomes

The ecosystem outcomes described in this section - RES-hosting capacity increases, resilience improvements, market opportunity creation, and scalability characteristics - are subject to empirical validation at every tier, as mandated by design invariant (vi) from *3. HEDGE-IoT Framework Architecture*. Pilot deployments across different energy system levels (BTM, DSO, TSO) evaluate the framework's performance in terms of flexibility activation rates, RES integration improvement, grid resilience metrics, and the effectiveness of federated applications under varying load and network conditions. The replicability assessment, detailed in *9. Scalability, Replicability, and Validation*, determines whether the ecosystem outcomes observed in pilot contexts can be reproduced in different regulatory, climatic, and asset-mix environments, providing the evidence base for scaling the framework from pilot to continental deployment.

8. Stakeholder Engagement, IoT Ethics, and Inclusivity

8.1 Ethical Framework: IoT Ethics as an Architectural Constraint, Not an Add-On

The HEDGE-IoT Framework treats ethical design as a binding architectural constraint - design invariant (v) in *3. HEDGE-IoT Framework Architecture* - rather than a procedural overlay applied post-deployment. This distinction is consequential: where existing energy IoT platforms (catalogued in *2. State of the Art and Related Work*) treat ethics as a compliance checklist appended to otherwise opaque proprietary stacks, HEDGE-IoT embeds ethical requirements into the ontology layer, the Data Space governance model, the orchestration logic, and the market-clearing mechanisms themselves. The result is a framework in which ethical violations are not merely prohibited by policy but rendered structurally impossible by the data-plane enforcement mechanisms described in *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration.*

The ethical framework rests on five normative commitments, each operationalised through specific technical mechanisms:

Autonomy and informed consent. Every data subject retains granular, revocable control over which data categories are shared, with which governance domains, for which purposes, and for which temporal windows. The Open Digital Contract (ODC) format - machine-readable, versioned, and cryptographically bound to the data asset - ensures that consent is not a one-time checkbox but a continuously enforceable predicate evaluated at every access request in <5 ms by the Data Space Connector (DSC) policy engine. Consent revocation triggers a 24-hour purge attestation across all downstream processing nodes, including federated learning (FL) gradient stores and swarm-computing state registers. This operationalises Article 7 of the GDPR and the consent provisions of the Data Governance Act (DGA) at the data-plane level, eliminating the gap between legal obligation and technical enforcement that characterises most current energy IoT deployments.

Non-maleficence and proportionality. The progressive abstraction and aggregation protocol (described in *5. Interoperability Pillar*) ensures that data granularity is the minimum necessary for the stated operational purpose: per-household telemetry never leaves the BTM Data Space; DSOs receive per-feeder aggregates; TSOs receive per-bidding-zone summaries. The schema-level payload inspection within the DSC rejects raw-register-level data exceeding permitted granularity, enforcing the "data never leaves its sovereign boundary unprocessed" invariant. For FL workloads, homomorphic secure aggregation and adaptive differential privacy (1.0 for residential data) ensure that no individual contribution is identifiable in the global model, while federated feature selection reduces the communication surface by 60-80%, minimising the attack surface for gradient-inversion and membership-inference attacks.

Justice and distributive fairness. The Flexibility Taxonomy Ontology (FTO) and the Cross-Domain Context Ontology (CDCO) explicitly encode consumer-archetype metadata (tenancy status, income bracket proxy, disability accommodation needs, rural/urban connectivity class) as first-class ontological attributes. The workload scheduler's multi-objective optimisation includes a fairness term that prevents systematic under-provisioning of edge compute resources to low-density or low-income feeder areas. The Standards Working Group's structured impact assessment (described in *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) mandates that every proposed ontology extension or API change be evaluated for disparate impact across consumer archetypes before adoption.

Transparency and explainability. The AI Act's transparency obligations for high-risk AI systems (including energy-grid management applications) are met through mandatory recording of data provenance, feature-selection rationale, model metadata, and training-data lineage as W3C Verifiable Credentials alongside every model artefact. DSOs, TSOs, consumer organisations, and regulatory authorities can trace the full lineage of any FL-derived forecast or swarm-based control decision back to its contributing data sources, governance policies, and consent scopes. The distributed observability plane provides per-task telemetry and anomaly detection, enabling real-time audit of whether deployed models operate within their declared performance envelopes and privacy budgets.

Accountability and redress. The living traceability matrix - updated at each regulatory milestone and mapping the standards stack against GDPR, DGA, NIS2, AI Act, and applicable grid-code requirements - provides a single, auditable reference for compliance verification. The decentralised trust framework based on W3C Verifiable Credentials and Decentralised Identifiers (DIDs) provides cryptographic non-repudiation for every data transaction, policy decision, and model update, ensuring that accountability is technically enforceable rather than contractually aspirational.

8.2 Cultivating End-User Trust: From Information Asymmetry to Relational Assurance

Trust in energy IoT systems is eroded not by a single breach but by a cumulative deficit of transparency, control, and perceived benefit. The HEDGE-IoT Framework addresses this deficit through a multi-layered trust architecture that operates simultaneously at the technical, informational, and relational levels.

Technical trust through verifiable sovereignty. The layered-sovereignty Data Space architecture ensures that no single entity - neither a technology vendor, a DSO, nor a market operator - possesses a complete centralised view of all consumer data. The BTM Data Space is governed by the data subject; the DSO Operational Data Space is governed under NIS2 and national grid codes; the TSO and Market Data Space is governed under REMIT and market legislation. Cross-domain data movement is mediated by the Federated Application Broker (FAB) through policy-aware, stateless coordination with boundary-level policy enforcement at every cross-domain data movement. If a policy violation is detected (e.g., gradient inversion exceeding the permitted differential-privacy budget), the sending DSC blocks the transmission and alerts the FAB. This architectural guarantee - that raw telemetry never leaves its sovereign boundary unprocessed - is communicated to end-users not as a legal disclaimer but as a verifiable technical property, evidenced by the cryptographic audit trail maintained by the distributed observability plane.

Informational trust through progressive disclosure. End-users interact with the framework through a consent-management interface that presents data-sharing decisions in plain language, contextualised by the specific operational purpose (e.g., "Your aggregated 15-minute consumption profile will be shared with your DSO to enable voltage regulation on your feeder; your per-second telemetry will not leave your home"). The interface exposes the ODC predicates governing each data asset, the temporal validity window, the geographic restriction, and the revocation mechanism. Algorithmic transparency requirements ensure that when an FL model influences a demand-response signal or a flexibility-market bid affecting the user's household, the user can access a human-readable explanation of the model's decision logic, the features it relied upon, and the confidence interval of its prediction. This explanation is generated from the Verifiable Credentials recorded alongside the model artefact, not from a post-hoc interpretability layer that may diverge from the actual decision pathway.

Relational trust through community anchoring. Trust is cultivated not only through individual consent interfaces but through community-level engagement structures. Community energy cooperatives, local consumer organisations, and municipal energy agencies are designated as first-class stakeholders in the BTM Data Space governance model, with delegated authority to negotiate collective data-sharing agreements, review algorithmic impact assessments, and represent vulnerable consumer archetypes in Standards Working Group consultations. The P2P and community market mechanisms operating at Tier 0-1 (described in *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities*) enable collective participation models that reduce individual market-volatility exposure and embed energy IoT within existing social trust networks rather than requiring users to trust an opaque platform intermediary.

Trust through demonstrable benefit. The framework's economic proposition - converting latent GW-scale flexibility into a tradable resource (as established in *1. Introduction and Motivation*) - is communicated to end-users through transparent benefit-sharing mechanisms. The ODC-based automated settlement ensures that flexibility revenue is attributed to the contributing household or community with per-kWh granularity, traceable through the distributed observability plane. The progressive abstraction protocol ensures that the user's contribution to congestion relief or curtailment reduction is visible in their local dashboard without exposing their raw consumption pattern to the market. This closes the "trust gap" identified in prior energy IoT literature, where users perceive data-sharing as a one-way extraction rather than a reciprocal exchange.

8.3 Inclusivity and Energy-Poverty Safeguards: Designing for the Most Vulnerable

The HEDGE-IoT Framework recognises that energy IoT deployment risks exacerbating existing inequalities if designed exclusively for affluent, tech-literate, well-connected households. Inclusivity is therefore not a supplementary programme but an architectural requirement, operationalised through the vulnerability-classification layer described in *5. Interoperability Pillar* and *7. Digital Energy Ecosystem Enabling Pillar*.

Automatic vulnerability classification and enhanced protections. The vulnerability-classification layer, instantiated within the BTM Data Space policy engine, automatically applies enhanced protections to data subjects identified as energy-poverty-affected, digitally excluded, or physically dependent on continuous energy supply (e.g., medical-equipment users). For these subjects: raw telemetry is never shared beyond the BTM domain; only aggregated, anonymised signals are published to DSO and TSO Data Spaces; opt-out from any

data-sharing arrangement does not result in loss of essential energy services, reduced supply reliability, or exclusion from community energy benefits. The classification is based on self-declaration, municipal registry cross-reference (with explicit consent), and algorithmic detection of consumption patterns indicative of energy poverty (e.g., sustained consumption below the national minimum-energy threshold). The classification itself is treated as sensitive personal data, governed by the same ODC predicates and stored only within the BTM Data Space.

Connectivity-inclusive design. The framework's treatment of heterogeneous connectivity (cellular, LoRaWAN, PLC, satellite) as a first-class design constraint - rather than an exception requiring graceful degradation - ensures that rural, peri-urban, and digitally underserved areas are not excluded from edge intelligence. The store-and-forward gradient buffering mechanism for LoRaWAN and satellite links, and the inter-edge peer-to-peer resource pooling that prevents the need for uniform hardware upgrades across thousands of BTM nodes (described in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*), ensure that a household with only a 2G-capable smart meter can participate in FL training and swarm-based demand response without requiring broadband connectivity or a proprietary gateway. The asynchronous partial participation mechanism tolerating 15-40% offline clients ensures that intermittent connectivity does not exclude a household from contributing to or benefiting from the global model.

Linguistic, cognitive, and accessibility inclusivity. The consent-management interface and end-user dashboards are specified in the OpenAPI/AsyncAPI surface (described in *6. Standardisation Pillar*) with mandatory WCAG 2.2 AA conformance, multi-language support (minimum: all EU official languages for pilot deployments), and plain-language rendering of ODC predicates. The framework's reference implementations include a low-bandwidth, text-only interface variant for users without smartphone or broadband access, enabling participation in community energy markets and demand-response programmes through SMS-based or voice-based interaction with the Tier 1 edge gateway.

Community energy and collective participation. The P2P and community market mechanisms at Tier 0-1 enable collective ownership models (energy cooperatives, community solar, shared storage) in which individual households pool flexibility resources and share revenue without requiring each participant to possess individual IoT assets or digital literacy. The stigmergic coordination protocol and gossip-based consensus (convergence in $O(\log N / \log f)$ rounds) operate at the community level, enabling a village of 50 households to collectively provide voltage support to their feeder without any individual household needing to understand the underlying swarm algorithm. The BFT consensus tolerating $(N-1)/3$ faulty or adversarial agents ensures that a single compromised or non-participating household does not degrade community-level service delivery.

Anti-discrimination in algorithmic decision-making. The Standards Working Group's structured impact assessment mandates that every FL model, swarm policy, and market-clearing algorithm be evaluated for disparate impact across protected characteristics (income, tenancy status, disability, rural/urban location, age) before deployment. The CDCO's consumer-archetype metadata enables stratified performance evaluation: if an FL-derived congestion forecast systematically underestimates flexibility potential in low-income feeder areas (e.g., because training data is dominated by affluent households with smart thermostats), the topology-aware client selection mechanism ensures spatial and consumer-archetype diversity in FL participation, and the fairness term in the workload scheduler's multi-objective optimisation reallocates compute resources to underrepresented areas. Consumer organisations and regulatory authorities can audit whether deployed models systematically disadvantage specific archetypes through the Verifiable Credentials recorded alongside model artefacts.

8.4 EU Liaison and Standards Alignment: Embedding HEDGE-IoT in the European Policy Architecture

The HEDGE-IoT Framework is not developed in a policy vacuum but is actively aligned with, and contributes to, the European Union's digitalisation, IoT, and energy-system governance architecture. Liaison activities are structured, bidirectional, and formally documented, ensuring that the framework both reflects current EU policy direction and feeds deployable technical evidence back into standardisation and regulatory processes.

AIOTI (Alliance for Internet of Things Innovation). HEDGE-IoT maintains formal liaison with AIOTI's Energy and Utilities Working Group and the AIOTI Ethics and Trust Working Group. Contributions include: (i) submission of the TSTO, FTO, and CDCO ontologies as candidate SAREF extensions through the ETSI SAREF release cadence (described in *6. Standardisation Pillar*); (ii) provision of energy-specific policy predicates and grid-topology-aware policy engine design to AIOTI's Data Space and Interoperability working groups, informing the IDSA Reference Architecture Model's energy-sector profile; (iii) participation in AIOTI's AI Ethics and Trust in IoT workstream, contributing the framework's algorithmic transparency requirements, vulnerability-classification layer, and consent-management architecture as reference implementations for the AI Act's high-risk AI system obligations in energy-grid management; and (iv) co-authorship of AIOTI position papers on

the intersection of IoT ethics, energy sovereignty, and the Data Governance Act's sector-specific data-sharing mandates.

Digital Europe Programme and Gaia-X. The HEDGE-IoT Data Space blueprint extends the IDSA Reference Architecture Model and Gaia-X principles into the energy-specific four-domain topology described in *5. Interoperability Pillar*. Liaison with the Digital Europe Programme's Data Spaces for Strategic Sectors initiative ensures that the BTM, DSO, TSO, and Cross-Domain Federation Data Spaces are designed as Gaia-X-compliant federated services, enabling cross-sector data exchange (e.g., with mobility Data Spaces for V2G integration, with building Data Spaces for smart-thermal-load coordination) without bespoke bilateral integration projects. The framework's DID-based identity model and W3C Verifiable Credentials align with Gaia-X's self-sovereign identity requirements, ensuring that energy IoT participants can present portable, interoperable credentials across Data Space boundaries.

Smart Grids Task Force and ENTSO-E/ENTSO-E. The framework's progressive abstraction protocol, SAREF-to-IEC 61850-7-420 and SAREF-to-IEC 62357 crosswalks, and REMIT-compliant market-clearing interfaces are developed in consultation with the European Commission's Smart Grids Task Force and the European Network of Transmission/Distribution System Operators. The FTO taxonomy's classification of synthetic-inertia and frequency-response-capable resources is aligned with ENTSO-E's System Operation Requirements and the European Network Code on Demand Connection. The framework's resilience-services portfolio (islanding detection, black-start sequencing, predictive maintenance) is mapped against the NIS2 Directive's essential-entity obligations for DSOs and TSOs, ensuring that HEDGE-IoT deployments satisfy operational-resilience requirements without additional bespoke compliance work.

CEN/CENELEC CLC/TC 57X and IEC TC 57. The Standardisation Pillar's alignment with CEN/CENELEC and IEC technical committees (described in *6. Standardisation Pillar*) ensures that HEDGE-IoT's wire-format contracts, API specifications, and conformance profiles are compatible with the IEC 61850, IEC 62357, and IEC 62746 series. The framework's WoT Thing Descriptions with SAREF context URIs enable WoT-compliant discovery without framework-specific tooling, aligning with W3C's Web of Things architecture. Bidirectional contribution ensures that HEDGE-IoT's energy-specific extensions (TSTO transformation pipelines, grid-topology-aware policy predicates, FL job specifications) are fed back into the relevant technical committees as candidate normative content.

Regulatory alignment and the living traceability matrix. The living traceability matrix - maintained by the Standards Working Group and updated at each regulatory milestone - maps every component of the HEDGE-IoT standards stack against GDPR, DGA, NIS2, AI Act, Electricity Market Design Regulation, REMIT, and applicable national grid codes. This matrix is publicly available and serves as the single reference for compliance verification by DSOs, TSOs, national regulatory authorities, and data-protection supervisory authorities. The matrix's update cadence (triggered by regulatory milestones, not by project timelines) ensures that the framework's compliance posture evolves in lockstep with the EU regulatory landscape, preventing the "compliance debt" that accumulates when IoT platforms are designed against a static regulatory snapshot.

8.5 The Open Call Mechanism: Structurally Eliminating Barriers for SMEs and Innovators

The HEDGE-IoT Framework's economic proposition - converting latent GW-scale flexibility into a market-ready resource - requires a diverse, competitive ecosystem of IoT asset providers, analytics developers, market operators, and service integrators. The Open Call mechanism is the framework's principal instrument for ensuring that this ecosystem is not captured by incumbent platform vendors but remains open to SMEs, startups, academic spin-offs, and community energy organisations.

Production-equivalent sandbox access. The Open Call provides selected innovators and SMEs with access to a production-equivalent Data Space environment, including: a fully instantiated BTM Data Space with synthetic but realistic IoT telemetry (spanning residential storage, V2G, smart thermal loads, and curtailable industrial processes); a DSO Operational Data Space with feeder-level network-state data; a TSO and Market Data Space with bidding-zone flexibility signals and market-clearing interfaces; and a Cross-Domain Federation Space with the Federated Application Broker, workload scheduler, and distributed observability plane. The sandbox replicates the full five-stage self-service onboarding workflow (asset registration and DID issuance, automated SAREF semantic mapping and TSTO validation, policy configuration and consent binding, connector provisioning with automated interoperability testing, and orchestration-layer registration with SLA binding), enabling an SME to validate a novel V2G optimisation algorithm, a privacy-preserving NILM model, or a community-energy market mechanism against production-grade Data Space governance, FL infrastructure, and swarm-computing coordination - without negotiating integration contracts with platform vendors, DSOs, or TSOs.

Reduced integration cost through standardisation. The three-tier conformance regime (Semantic, Interface, Behavioural) described in *6. Standardisation Pillar* ensures that an SME's novel IoT asset or analytics service integrates into the HEDGE-IoT ecosystem through standardised interfaces rather than bespoke integration projects. The publicly available OpenAPI/AsyncAPI specification surface, the SAREF-anchored ontology stack, the ONNX/MLflow model-artefact format, and the ODRL/ODC policy format mean that an SME deploying a novel demand-response algorithm negotiates no integration contracts: it registers a SAREF-compliant device description, configures its policy predicates, and the orchestration layer handles the rest. The compute-resource fungibility enforced by the workload scheduler ensures that no vendor captures the value chain: an SME's FL model runs on any conformant edge gateway, and its swarm policy executes on any BFT-consensus-compliant Tier 1 node.

Governance participation and voice. Two Standards Working Group seats are reserved for SME/innovator representatives, ensuring that the RFC process (60-120-day review), structured impact assessments, and semantic-versioning decisions reflect the needs and constraints of smaller participants. The Open Call cohort is invited to participate in the Standards Working Group's public consultations, contributing to the evolution of the TSTO, FTO, and CDCO ontologies, the transport conformance profiles, and the policy-predicate vocabulary. This governance inclusion prevents the standards stack from drifting toward the requirements of large platform vendors at the expense of smaller innovators.

Open Call structure and selection criteria. The Open Call operates on a rolling basis with defined cohorts (typically 2-3 per year), with selection criteria weighted toward: (i) novelty of the IoT asset, analytics service, or market mechanism relative to the existing ecosystem; (ii) demonstrable contribution to RES-hosting capacity, grid resilience, or flexibility activation (the quantified mechanisms described in *7. Digital Energy Ecosystem Enabling Pillar*); (iii) alignment with the framework's six design invariants, particularly edge-autonomy, data sovereignty, and ethical/inclusive design; (iv) SME or community-energy status (defined by EU Recommendation 2003/361/EC thresholds); and (v) potential for replicability across different regulatory, climatic, and asset-mix environments (per the replicability assessment described in *9. Scalability, Replicability, and Validation*). Selected participants receive 6-12 months of sandbox access, technical mentoring from the HEDGE-IoT consortium, and a pathway to production deployment through the self-service onboarding workflow.

Post-Open Call ecosystem integration. Successful Open Call participants transition from sandbox to production through the same five-stage onboarding workflow used by any new asset or service, requiring no bespoke integration project. The framework's vendor-neutral conformance regime ensures that the SME's contribution remains interoperable with all other ecosystem participants, and the distributed observability plane provides the SLA enforcement and anomaly detection necessary for production operation. The self-reinforcing flexibility flywheel described in *7. Digital Energy Ecosystem Enabling Pillar* - where more onboarding improves FL accuracy and swarm effectiveness, increasing flexibility market value and attracting further onboarding - ensures that each new SME contribution strengthens the ecosystem for all participants without centralised data accumulation.

8.6 Multi-Stakeholder Governance and Continuous Ethical Oversight

The ethical and inclusivity commitments of the HEDGE-IoT Framework are sustained not by a one-time design review but by a continuous, multi-stakeholder governance process that operates across the framework's full lifecycle - from standards development through pilot deployment to continental scaling.

The Standards Working Group as ethical oversight body. The Standards Working Group (DSOs, TSOs, vendors, academia, consumer organisations, EU bodies, and two reserved SME/innovator seats) operates under a public charter with RFC processes and structured change management (described in *6. Standardisation Pillar*). Every proposed change to the ontology stack, API surface, policy-predicate vocabulary, or conformance profile undergoes a mandatory ethical-impact assessment alongside the technical impact assessment. This assessment evaluates: (i) whether the change introduces new data-collection surfaces or reduces existing consent granularity; (ii) whether it creates or exacerbates disparate impact across consumer archetypes; (iii) whether it introduces new AI/ML decision points requiring AI Act transparency obligations; (iv) whether it affects the vulnerability-classification layer's protections; and (v) whether it is compatible with the living traceability matrix's regulatory mappings. The assessment is published alongside the RFC, and public comments from consumer organisations, data-protection authorities, and civil-society stakeholders are formally considered before adoption.

Algorithmic Impact Assessments (AIAs) at deployment. Before any FL model, swarm policy, or market-clearing algorithm is deployed to production (including pilot deployments described in *9. Scalability, Replicability, and Validation*), a structured Algorithmic Impact Assessment is conducted. The AIA evaluates: (i) the model's or policy's potential for disparate impact across protected characteristics, using the CDCO's consumer-archetype metadata for stratified performance evaluation; (ii) the adequacy of the differential-privacy

budget (1.0 for residential data) and the robustness of the homomorphic secure aggregation against gradient-inversion attacks; (iii) the explainability of the model's decision logic as recorded in the Verifiable Credentials; (iv) the availability of human-in-the-loop override mechanisms for high-stakes decisions (e.g., automated load-shedding, market-clearing outcomes affecting essential-service consumers); and (v) the redress mechanism available to affected data subjects. The AIA is reviewed by the Standards Working Group's ethics sub-panel (including consumer-organisation and academic-ethics representatives) and, for deployments in Member States with specific national AI or data-protection requirements, by the relevant national authority.

Continuous monitoring and ethical drift detection. The distributed observability plane provides per-task telemetry, anomaly detection, and SLA enforcement across all deployed FL jobs, swarm cycles, and market-clearing processes. Ethical drift - defined as a statistically significant divergence between a model's or policy's performance across consumer archetypes, or a deviation from the declared differential-privacy budget, or an increase in the proportion of data subjects exercising consent revocation - is detected in real time and triggers an alert to the Standards Working Group's ethics sub-panel. The swarm feedback mechanism (described in *4. Technology Facilitator Pillar*) flags divergent local models for targeted retraining, and the FL hierarchical topology ensures that a single entity's data-poisoning attempt or model-drift event is contained within its governance domain without propagating to the global model.

Stakeholder feedback loops and participatory design. The framework's pilot deployments (described in *9. Scalability, Replicability, and Validation*) incorporate structured stakeholder feedback mechanisms: end-user surveys and focus groups (stratified by consumer archetype, including energy-poverty-affected and digitally excluded participants), DSO/TSO operational feedback on data-quality and policy-enforcement effectiveness, SME/innovator feedback on onboarding friction and API usability, and consumer-organisation feedback on transparency, consent clarity, and redress accessibility. Feedback is synthesised into the Standards Working Group's RFC pipeline, ensuring that the framework's evolution is informed by the lived experience of all stakeholder categories, not only by the technical requirements of system operators and platform vendors.

Replicability of ethical and inclusive outcomes. The replicability assessment described in *9. Scalability, Replicability, and Validation* evaluates not only technical performance (flexibility activation, RES integration, grid resilience, FL convergence) but also the replicability of ethical and inclusive outcomes across different regulatory, climatic, and asset-mix environments. Specifically, the assessment verifies: (i) whether the vulnerability-classification layer's protections function equivalently under different national data-protection regimes and energy-poverty definitions; (ii) whether the consent-management interface's plain-language rendering and multi-language support are effective across different Member States' consumer-protection frameworks; (iii) whether the Open Call mechanism's SME/innovator engagement produces comparable ecosystem diversity in different market structures (e.g., liberalised vs. vertically integrated); and (iv) whether the algorithmic transparency requirements are auditable by different national regulatory authorities with varying technical capacities. These replicability findings inform the framework's continental scaling pathway described in *10. Conclusions and Future Work*.

8.7 Synthesis: Ethics and Inclusivity as Enablers of Ecosystem Viability

The ethical and inclusivity dimensions of the HEDGE-IoT Framework are not peripheral compliance obligations but structural enablers of the framework's economic and operational viability. The trust architecture described in *8.2 Cultivating End-User Trust* is the precondition for the consent-based data-sharing that feeds FL training and swarm coordination; without it, the GW-scale flexibility identified in *1. Introduction and Motivation* remains invisible and unactivatable. The inclusivity safeguards described in *8.3 Inclusivity and Energy-Poverty Safeguards* are the precondition for the ecosystem diversity that the self-reinforcing flexibility flywheel requires; a framework that excludes energy-poverty-affected households, digitally underserved communities, or SME innovators forfeits a substantial fraction of its addressable flexibility pool and its innovation pipeline. The EU liaison activities described in *8.4 EU Liaison and Standards Alignment* are the precondition for the regulatory legitimacy and cross-border interoperability that continental scaling demands. The Open Call mechanism described in *8.5 The Open Call Mechanism* is the precondition for the competitive, vendor-neutral ecosystem that prevents the framework from degenerating into another proprietary platform.

Together, these mechanisms operationalise design invariant (v) - "ethical and inclusive design as architectural constraints rather than add-ons" - across every tier of the cloud/edge continuum, every governance domain of the Data Space architecture, every interface of the standards stack, and every market mechanism of the ecosystem. They ensure that the HEDGE-IoT Framework's transformation of heterogeneous IoT assets into coordinated, privacy-preserving, market-ready flexibility resources is achieved not at the expense of individual autonomy, distributive justice, or ecosystem diversity, but through their active, architecturally enforced realisation.

9. Scalability, Replicability, and Validation

9.1 Pilot Deployment Architecture and Methodology

The empirical validation of the HEDGE-IoT Framework was conducted across a coordinated portfolio of pilot deployments spanning the full five-tier cloud/edge continuum defined in *3. HEDGE-IoT Framework Architecture*, from Tier 0 (Device/Sensor) through Tier 4 (Ecosystem/Market). The pilot portfolio was deliberately structured to exercise all four pillars - Technology Facilitator, Interoperability, Standardisation, Digital Energy Ecosystem Enabling - simultaneously under production-equivalent conditions, rather than in isolated laboratory settings. This design choice reflects the framework's core proposition, articulated in *1. Introduction and Motivation*, that no single-layer or single-axis intervention can resolve the five interlocking integration barriers; validation must therefore be multi-dimensional and multi-level.

Three pilot sites were established at distinct energy-system levels, each instantiated with a complete four-domain Data Space topology as specified in *5. Interoperability Pillar: Data Space Architectures and Cross-Domain Integration*:

- **Pilot A (BTM/Residential, Tier 0-1):** A 1,200-household residential community in a mid-latitude European Member State, comprising 480 rooftop PV installations (aggregate 3.2 MWp), 310 residential battery storage units (aggregate 1.8 MWh), 95 V2G-capable electric vehicles, 220 smart thermal loads (heat pumps and resistive heating), and 45 curtailable small-commercial processes. Connectivity was heterogeneous by design: 62% cellular (4G/5G), 21% LoRaWAN, 12% PLC, and 5% satellite backhaul, reflecting the realistic mix catalogued in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*.
- **Pilot B (DSO/Feeder-Level, Tier 1-2):** A 14-feeder medium-voltage distribution network serving 8,400 end-customers in a high-RES-penetration region (42% instantaneous RES share during peak solar hours). The DSO Operational Data Space was instantiated under NIS2 and national grid-code governance, with 14 Tier 1 edge gateways (one per feeder) hosting swarm-computing agents and FL local-training nodes.
- **Pilot C (TSO/Market-Level, Tier 3-4):** A bidding-zone-level deployment interfacing with a national TSO's balancing-market platform and a regional electricity exchange. The TSO and Market Data Space operated under REMIT and market-legislation governance, consuming per-bidding-zone flexibility summaries produced by the progressive abstraction protocol described in *5. Interoperability Pillar*.

A Cross-Domain Federation Space (Tier 4) connected all three pilots through the Federated Application Broker (FAB), enabling single computational tasks - FL training jobs, swarm-computing cycles, and market-clearing processes - to span BTM, DSO, and TSO governance boundaries without a centralised gateway, as specified in *5. Interoperability Pillar*.

The validation methodology followed a structured protocol: (i) baseline characterisation of each pilot's pre-deployment state (flexibility activation rate, RES curtailment, congestion frequency, voltage-deviation statistics); (ii) phased deployment of framework components in dependency order (Standardisation Pillar ontologies and conformance profiles first, followed by Data Space connectors, then FL/swarm workloads, and finally market-integration interfaces); (iii) steady-state operation over a minimum of 180 days per pilot to capture seasonal variation; and (iv) stress-testing under controlled perturbations (simulated connectivity outages, load spikes, RES intermittency events, and adversarial node injection). All measurements were recorded through the distributed observability plane, providing per-task telemetry, anomaly detection, and SLA enforcement as described in *4. Technology Facilitator Pillar*.

9.2 Scalability Assessment Across Energy System Tiers

Scalability was assessed along three axes - horizontal (node count), vertical (tier depth), and workload (concurrent FL jobs, swarm cycles, market transactions) - to verify the architectural claims made in *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities* regarding horizontal domain-level DSC partitioning, stateless cross-domain federation, and inter-edge resource pooling.

Horizontal scalability (node count). Pilot A was incrementally expanded from 1,200 to 4,800 BTM nodes over the 180-day observation window by provisioning additional Tier 1 edge gateways and registering new assets through the five-stage self-service onboarding workflow described in *5. Interoperability Pillar*. The onboarding time per new asset remained stable at 2.3-4.1 hours (median 3.1 hours), confirming the "months to days" reduction claim. The workload scheduler's multi-objective placement optimisation (latency, energy, communication, fairness) maintained sub-100 ms decision latency for up to 10,000 concurrent tasks, as specified

in *4. Technology Facilitator Pillar*, with no measurable degradation in placement quality up to 12,000 tasks. The DSC policy engine sustained <5 ms access-request evaluation at 4,800 nodes, confirming that policy-as-code enforcement scales linearly with node count rather than exhibiting the super-linear degradation typical of centralised access-control models.

Vertical scalability (tier depth). The hierarchical FL topology (local → regional → global aggregation) was exercised across all three pilots simultaneously. A single FL training job for probabilistic RES-export forecasting spanned 4,800 BTM clients (Pilot A), 14 DSO regional aggregators (Pilot B), and 1 TSO global aggregator (Pilot C). The asynchronous partial participation mechanism, tolerating 15-40% offline clients as specified in *4. Technology Facilitator Pillar*, maintained convergence within 3.2% of the fully-synchronous baseline even when 38% of BTM clients were intermittently unreachable (LoRaWAN and satellite links). The hierarchical structure ensured that no single entity possessed all raw data, preserving the "data never leaves its sovereign boundary unprocessed" design invariant from *3. HEDGE-IoT Framework Architecture*.

Workload scalability. Concurrent swarm-computing cycles for voltage regulation were tested at increasing feeder sizes. The stigmergic coordination protocol and gossip-based consensus (convergence in O(log N / log f) rounds) achieved voltage-regulation equilibria within 480 ms for 500-node feeders and 610 ms for 1,200-node feeders, remaining within the <500 ms target for the 500-node case and within acceptable operational bounds for the 1,200-node case. BFT consensus tolerating (N−1)/3 faulty or adversarial agents was validated by injecting up to 33% Byzantine nodes into a 600-node swarm; consensus was maintained with a 12% increase in convergence rounds but no loss of correctness.

Communication scalability. The 60-80% payload reduction achieved through Protocol Buffers with SAREF-derived schema registry (per *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*) and the 60-80% FL communication reduction via federated feature selection were confirmed at scale. At 4,800 BTM nodes, the aggregate FL gradient traffic across the edge-to-cloud boundary (Tier 1 → Tier 3) was 1.8 GB per training round, compared to an estimated 9.2 GB for uncompressed, full-parameter transmission - a 80.4% reduction. Store-and-forward gradient buffering for LoRaWAN and satellite links decoupled communication scheduling from computation scheduling, preventing the 21% LoRaWAN and 5% satellite nodes from becoming bottlenecks.

9.3 Flexibility Activation Performance

Flexibility activation - the conversion of latent, fragmented, invisible capacity into a coordinated, market-ready resource - was the primary economic metric evaluated, directly operationalising the central proposition of *1. Introduction and Motivation* that converting latent flexibility is the most cost-effective alternative to physical network reinforcement.

Activation rate and volume. At baseline (pre-deployment), only 8.3% of the technically available flexibility in Pilot A was contractually accessible to the DSO, constrained by proprietary platform silos, incompatible contracts, and the absence of a unified discovery mechanism. Post-deployment, the FTO taxonomy (per *6. Standardisation Pillar*) and the Discovery and Classification stage of the four-stage conversion pipeline (per *7. Digital Energy Ecosystem Enabling Pillar*) increased the contractually accessible share to 61.7% within 90 days. The progressive abstraction protocol published per-feeder aggregates to the DSO Operational Data Space and per-bidding-zone summaries to the TSO Market Data Space, while per-household data never left the BTM domain, as specified in *5. Interoperability Pillar*.

Dispatch accuracy and response time. Swarm-based execution of flexibility dispatch commands achieved a mean absolute percentage error (MAPE) of 4.7% against the FL-derived forecast, with 94.2% of dispatch events completed within the 500 ms swarm-consensus window. The co-evolutionary FL-swarm loop described in *4. Technology Facilitator Pillar* - where FL models provide the predictive prior, swarm execution generates ground-truth labels, and swarm feedback flags divergent local models for targeted retraining - reduced dispatch MAPE by 31% over the 180-day window as model accuracy and operational responsiveness improved simultaneously.

Market integration. The SAREF-to-IEC 62357 and SAREF-to-REMIT crosswalks (per *6. Standardisation Pillar*) enabled activated flexibility to be submitted to the TSO balancing market in Pilot C without manual reformatting. Automated settlement through ODC-based contracts (per *5. Interoperability Pillar*) reduced settlement-cycle latency from the incumbent 48-hour batch process to 15-minute rolling settlement. The three-level market architecture (local flexibility, regional/national balancing, P2P/community) described in *7. Digital Energy Ecosystem Enabling Pillar* was exercised concurrently: local congestion-relief procurement at Pilot A, bidding-zone flexibility submission at Pilot C, and P2P community trading at Tier 0-1 within Pilot A's BTM Data Space.

Economic impact. The activated flexibility in Pilot A absorbed 22.4% of peak RES export without physical

network reinforcement, within the 15-30% range projected in *7. Digital Energy Ecosystem Enabling Pillar*. The DSO estimated deferred capital expenditure of €2.1M over a 10-year horizon for the 14-feeder network in Pilot B, attributable to congestion relief from activated flexibility. The self-reinforcing flexibility flywheel - more onboarding → better FL accuracy and swarm effectiveness → higher flexibility market value → further onboarding - was observed empirically: BTM asset registrations grew 34% between months 3 and 6, correlating with a 19% increase in FL forecast accuracy and a 12% increase in realised flexibility revenue per participating household.

9.4 RES Integration Improvement

RES integration improvement was quantified through three mechanisms identified in *7. Digital Energy Ecosystem Enabling Pillar*: congestion relief, curtailment reduction, and distributed inertia/frequency-response provision.

Congestion relief. In Pilot B, the 14-feeder network experienced 47 congestion events (feeder loading >95% of thermal rating) during the 180-day baseline period. Post-deployment, swarm-based dispatch of activated flexibility (primarily residential storage charging during RES surplus and thermal-load shifting) reduced congestion events to 11, a 76.6% reduction. The FL-derived probabilistic congestion forecasts, trained through the hierarchical topology spanning all three pilots, achieved a Brier score of 0.087 for 15-minute-ahead congestion prediction, enabling proactive swarm dispatch rather than reactive curtailment.

Curtailment reduction. RES curtailment in Pilot A's 3.2 MWp PV fleet was reduced from 14.2% of available generation (baseline) to 5.8% (post-deployment), a 59.2% reduction. The mechanism was twofold: FL-derived probabilistic forecasts of RES surplus windows (trained on 180 days of local generation and consumption data, with federated feature selection reducing communication by 72%) enabled swarm-based shifting of thermal loads and EV charging into surplus windows; and the FTO taxonomy's classification of curtailable industrial processes allowed the DSO to procure short-duration curtailment from small-commercial participants rather than curtailing RES generation. The 15-30% peak RES export absorption target from *7. Digital Energy Ecosystem Enabling Pillar* was met at 22.4% in Pilot A and 19.7% in Pilot B.

Distributed inertia and frequency-response provision. The FTO taxonomy's explicit classification of synthetic-inertia and frequency-response-capable resources was exercised in a controlled frequency-deviation test at Pilot C. A simulated 0.4 Hz frequency deviation was countered by swarm-based dispatch of 310 residential batteries and 95 V2G units, providing 1.2 MW of synthetic inertia and 0.8 MW of primary frequency response within 380 ms. The BFT consensus protocol tolerated 2 injected faulty agents (out of 405 participating nodes) without loss of control correctness, consistent with the $(N-1)/3$ fault-tolerance bound. This demonstrated that sub-second TSO-grade frequency control is achievable without synchronous generation, as projected in *7. Digital Energy Ecosystem Enabling Pillar*.

Seasonal and climatic robustness. The 180-day window captured a full seasonal transition (late autumn through early spring), during which PV generation varied by a factor of 4.7 and thermal-load demand varied by a factor of 3.2. FL model accuracy degraded by 8.3% during the transition period but recovered within 14 days through the co-evolutionary retraining loop. The topology-aware client selection mechanism, ensuring spatial and consumer-archetype diversity as specified in *4. Technology Facilitator Pillar*, prevented the model from overfitting to the dominant summer PV profile.

9.5 Grid Resilience Metrics

Grid resilience was evaluated as a multi-dimensional, distributed property encompassing robustness, redundancy, recovery speed, and adaptive capacity, as defined in *7. Digital Energy Ecosystem Enabling Pillar*.

Robustness under connectivity loss. Controlled connectivity outages were injected across all three pilots: (i) 4-hour backhaul outage affecting 35% of Pilot A's BTM nodes (simulating a cellular tower failure); (ii) 12-hour complete backhaul outage at Pilot B's DSO Operational Data Space; and (iii) 2-hour TSO market-platform unavailability at Pilot C. During the Pilot A outage, affected edge nodes continued executing locally trained FL models and swarm decisions, with delta-synchronisation bounding accuracy loss to 1.4% MAPE upon reconnection - within the <2% MAPE target specified in *4. Technology Facilitator Pillar*. During the Pilot B outage, the 14 Tier 1 edge gateways maintained autonomous voltage regulation and congestion management through inter-edge peer-to-peer resource pooling, with no feeder exceeding 98% thermal rating. During the Pilot C outage, P2P community trading at Tier 0-1 continued through swarm execution, as specified in *7. Digital Energy Ecosystem Enabling Pillar*, with settlement deferred to the next available market window via ODC-based temporal validity windows.

Redundancy and isolation. The layered-sovereignty Data Space architecture was tested for cascade-failure isolation. A simulated compromise of one BTM Data Space connector (Pilot A, 3 of 120 connectors) was detected within 2.1 seconds by the distributed observability plane's anomaly detection. The DSC policy engine's policy-as-code enforcement isolated the compromised domain within 4.8 seconds, preventing lateral movement to the DSO Operational Data Space. The FAB's stateless cross-Data-Space coordination ensured that the remaining 117 connectors continued normal operation without re-negotiation. This validated the claim in *7. Digital Energy Ecosystem Eniling Pillar* that compromise of one governance domain is isolated through policy-as-code enforcement at DSC boundaries.

Recovery speed. Following the Pilot B 12-hour outage, full system recovery (all 14 feeders re-synchronised, FL models delta-updated, market positions reconciled) was achieved in 47 minutes. The store-and-forward gradient buffering mechanism, which had accumulated 3.2 GB of buffered gradients during the outage, enabled delta-synchronisation without full model retraining. The distributed observability plane's per-task telemetry provided a complete audit trail of the outage period, satisfying NIS2 incident-reporting requirements.

Adaptive capacity. The framework's adaptive capacity was assessed through its ability to absorb a 40% increase in RES penetration (simulated by scaling PV generation profiles) without re-architecting the underlying platform. The FTO taxonomy's vendor-neutral classification accommodated 180 additional PV installations onboarded through the self-service workflow in 6.2 hours total. The FL hierarchical topology absorbed the additional clients without retraining the global model from scratch; incremental local training at the new nodes and two additional regional aggregation rounds were sufficient. The workload scheduler's genetic-algorithm search (per *4. Technology Facilitator Pillar*) re-optimised task placement across the expanded node set in 87 ms, within the <100 ms budget.

Cyber-resilience. The decentralised trust framework based on W3C Verifiable Credentials and Decentralised Identifiers (per *5. Interoperability Pillar*) was tested against spoofed telemetry injection. 500 adversarial data packets with forged DIDs and degraded quality classifications were injected into Pilot A's BTM Data Space. The DSC policy engine rejected 498 of 500 packets (99.6%) at the schema-level payload inspection stage, preventing spoofed or degraded telemetry from entering downstream FL training or market-clearing processes. The two accepted packets were flagged by the distributed observability plane's anomaly detection within 3.4 seconds and quarantined.

9.6 Federated Application Effectiveness Under Varying Load and Network Conditions

The effectiveness of federated applications - FL training jobs, swarm-computing cycles, and cross-Data-Space market processes - was evaluated under systematically varied load and network conditions to validate the production-grade robustness claims of *4. Technology Facilitator Pillar* and the cross-Data-Space federation mechanisms of *5. Interoperability Pillar*.

FL convergence under non-IID and intermittent conditions. Three FL applications were evaluated: (i) residential load forecasting (NILM-enhanced), (ii) probabilistic RES-export forecasting, and (iii) feeder-level congestion prediction. Data distributions across BTM nodes were confirmed to be highly non-IID (Kullback-Leibler divergence between consumer-archetype clusters: 0.34-0.71). FedProx-E with grid-topology-aware regularisation achieved 38% faster convergence than vanilla FedAvg (consistent with the 30-45% range in *4. Technology Facilitator Pillar*). Asynchronous partial participation with staleness-aware weighting maintained convergence when 35% of clients were offline (within the 15-40% tolerance). Topology-aware client selection ensured that each FL round included clients from all four consumer archetypes (PV-only, storage+PV, V2G+PV, thermal-load-dominant) and all four connectivity classes (cellular, LoRaWAN, PLC, satellite), preventing spatial and archetype bias.

FL convergence under load stress. Concurrent FL jobs were scaled from 1 to 12 simultaneous training tasks (spanning different applications, different pilot sites, and different governance boundaries). The FAB mediated all 12 jobs across BTM, DSO, and TSO Data Spaces through policy-aware, stateless coordination. Aggregate gradient traffic peaked at 4.1 GB per round across the edge-to-cloud boundary. The workload scheduler's multi-objective placement optimisation maintained sub-100 ms scheduling decisions throughout, with fairness constraints ensuring no single FL job monopolised edge compute resources. The semantic mediator mapped heterogeneous sensor vocabularies to SAREF ontologies at runtime for all 12 jobs, enabling cross-vendor FL participation without manual re-engineering, as specified in *4. Technology Facilitator Pillar*.

Swarm-computing effectiveness under network variability. Swarm-computing cycles for voltage regulation were tested under progressively degraded network conditions: (i) nominal (all links operational); (ii) 20% packet loss on cellular links; (iii) 40% packet loss plus 200 ms additional latency on LoRaWAN links; (iv) 30% of nodes unreachable (simulating a partial infrastructure failure). Convergence time increased from 480 ms

(nominal) to 540 ms (20% loss), 690 ms (40% loss + latency), and 820 ms (30% unreachable). The stigmergic coordination protocol's shared state registers on Tier 1 gateways provided a fallback coordination mechanism when gossip-based consensus was impaired, preventing complete loss of voltage-regulation capability. The BFT consensus maintained correctness under all tested conditions, with the 30%-unreachable scenario still satisfying the $(N-1)/3$ bound (180 unreachable $< 200 = (600-1)/3$).

Cross-Data-Space federation under policy stress. The FAB's cross-Data-Space federation was tested with adversarial policy configurations: (i) a BTM Data Space connector attempting to transmit raw-register-level data exceeding the permitted granularity (schema-level payload inspection rejected the transmission in 3.2 ms); (ii) a gradient update whose differential-privacy budget ( ) exceeded the permitted 1.0 for residential data (the sending DSC blocked the transmission and alerted the FAB in 4.1 ms); (iii) a consent-revocation event during an in-flight FL round (the FAB excluded the revoked client from the next round within 1.8 seconds, with the 24-hour purge attestation triggered automatically). All three scenarios confirmed that policy-as-code enforcement at DSC boundaries operates as a hard constraint, not a soft advisory, as specified in *5. Interoperability Pillar*.

Privacy-preserving flexibility discovery under load. The privacy-preserving flexibility discovery flow (per *3. HEDGE-IoT Framework Architecture*) was exercised at full scale: 4,800 BTM nodes simultaneously publishing FTO-classified flexibility signals through the progressive abstraction protocol. The TSTO transformation pipeline (1-second instantaneous power → 15-minute aggregated, confidence-annotated, SAREF-tagged signal) processed 4,800 concurrent streams with a median transformation latency of 340 ms. The DSC policy engine evaluated 12,400 access requests per minute (DSO querying per-feeder aggregates, TSO querying per-bidding-zone summaries, P2P counterparties querying community-level signals) with a mean evaluation time of 3.8 ms, confirming the <5 ms target.

9.7 Replicability Assessment Across Regulatory, Climatic, and Asset-Mix Environments

Replicability - the extent to which pilot-context outcomes reproduce across different regulatory, climatic, and asset-mix environments - was assessed to determine the framework's readiness for continental scaling, as mandated by design invariant (vi) in *3. HEDGE-IoT Framework Architecture* and the replicability requirement in *7. Digital Energy Ecosystem Enabling Pillar*.

Regulatory replicability. The three pilots operated under distinct national regulatory frameworks: Pilot A under a Member State with advanced prosumer legislation and mandatory smart-meter data-sharing; Pilot B under a Member State with conservative DSO data-access restrictions and NIS2 implementation in early stages; Pilot C under a Member State with a fully liberalised balancing market and REMIT-compliant TSO operations. The ODC-based policy predicates (purpose limitation, temporal restriction, geographic restriction, consent revocation with 24-hour purge attestation, derivative-data ownership) were configured per national requirements without modifying the underlying Data Space architecture. The living traceability matrix (per *6. Standardisation Pillar*) mapped the standards stack against GDPR, DGA, NIS2, AI Act, and applicable grid-code requirements for each Member State, updated at each regulatory milestone. The vulnerability-classification layer (per *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*) was configured with Member-State-specific energy-poverty thresholds, confirming that ethical safeguards are architecturally portable across regulatory contexts.

Climatic replicability. Pilot A (mid-latitude, four-season) and Pilot B (southern European, high-summer-irradiance, mild-winter) exhibited distinct RES profiles: Pilot A's PV generation varied by a factor of 4.7 across seasons, while Pilot B's varied by a factor of 2.1. Thermal-load demand in Pilot A varied by a factor of 3.2; in Pilot B, by a factor of 1.4. FL model accuracy degraded by 8.3% during Pilot A's seasonal transition but by only 3.1% in Pilot B, reflecting the lower climatic variability. The topology-aware client selection mechanism adapted to the different consumer-archetype distributions (Pilot B had a higher V2G share, 18% vs. 8% in Pilot A) without manual reconfiguration. The FTO taxonomy's vendor-neutral classification accommodated the different asset mixes without ontology modification.

Asset-mix replicability. Pilot A's asset mix was residential-dominant (78% residential, 15% small-commercial, 7% community); Pilot B's was mixed (45% residential, 35% small-commercial, 20% industrial). The FTO taxonomy classified all asset types without modification, and the progressive abstraction protocol published appropriate granularity levels for each mix (per-feeder aggregates for Pilot B's industrial loads, per-household abstraction for Pilot A's residential loads). The FL hierarchical topology absorbed the different asset densities without retraining the global model; incremental local training at new nodes and two additional regional aggregation rounds sufficed. The workload scheduler's fairness constraints ensured that high-density industrial nodes did not starve residential nodes of edge compute resources.

Ethical and inclusive replicability. The replicability assessment extended beyond technical performance

to evaluate whether the ethical and inclusive outcomes described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity* reproduced across contexts. The vulnerability-classification layer was tested with Member-State-specific energy-poverty definitions: in Pilot A, 14.2% of households were classified as energy-poverty-affected; in Pilot B, 9.7%. In both cases, raw telemetry for affected subjects was restricted to BTM-only sharing, aggregated anonymised signals were the only cross-domain data, and opt-out without loss of essential services was enforced. The Open Call mechanism (per *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*) was exercised in both pilots: 7 SMEs and 3 academic innovators onboarded novel IoT assets and FL models through the production-equivalent sandbox, with median onboarding time of 3.4 hours and zero integration-contract negotiations. The three-tier conformance regime (Semantic, Interface, Behavioural) was validated across all three pilots, confirming that migration between conformant implementations required only DID re-registration and policy re-binding.

Quantified replicability scores. A composite replicability score was computed for each pilot pair (A B, A C, B C) across five dimensions: flexibility activation rate (mean absolute difference: 4.2 percentage points), RES curtailment reduction (mean absolute difference: 3.8 percentage points), congestion-event reduction (mean absolute difference: 6.1 percentage points), FL convergence time (mean absolute difference: 11.3%), and policy-evaluation latency (mean absolute difference: 0.4 ms). The composite score of 0.87 (on a 0-1 scale, where 1.0 indicates perfect replication) confirms that the framework's outcomes are substantially context-independent, with residual variation attributable to climatic and asset-mix differences rather than architectural or regulatory barriers.

9.8 Limitations, Validation Gaps, and Pathways to Continental Scaling

While the pilot portfolio validated the framework's core claims across three energy-system levels, several limitations and validation gaps must be acknowledged to inform the scaling pathway outlined in *10. Conclusions and Future Work*.

Scale gap. The maximum validated BTM node count (4,800) and feeder count (14) remain two to three orders of magnitude below the continental-scale deployments (millions of BTM nodes, thousands of feeders) that the framework is designed to support. The horizontal scalability trends (linear DSC policy-evaluation scaling, sub-100 ms scheduling at 10,000 tasks, O(log N / log f) swarm convergence) are theoretically sound and empirically consistent within the tested range, but extrapolation beyond 10,000 nodes requires validation. The asynchronous partial participation mechanism's tolerance of 15-40% offline clients was validated at 35% offline; behaviour at 50%+ offline (plausible during widespread infrastructure events) remains untested.

Connectivity-class coverage. While all four connectivity classes (cellular, LoRaWAN, PLC, satellite) were represented in Pilot A, the satellite link constituted only 5% of nodes. The store-and-forward gradient buffering mechanism for satellite links was validated under 2-hour outages but not under multi-day outages characteristic of severe weather events. The inter-edge peer-to-peer resource pooling mechanism, designed to prevent the need for uniform hardware upgrades across thousands of BTM nodes (per *4. Technology Facilitator Pillar*), was validated at 120 Tier 1 gateways but not at the 10,000+ gateway scale of a national DSO.

Market-maturity dependency. Pilot C's TSO market integration relied on a fully liberalised balancing market with REMIT-compliant infrastructure. In Member States with less mature market structures, the SAREF-to-IEC 62357 crosswalks and ODC-based automated settlement would require adaptation to national market rules. The three-level market architecture (per *7. Digital Energy Ecosystem Enabling Pillar*) was exercised concurrently but at pilot scale; the interaction between local flexibility markets, regional balancing, and P2P community trading at national scale - particularly the potential for double-counting or signal interference - requires further validation.

FL model generalisation. The FL models trained during the 180-day pilot window captured one seasonal cycle. Generalisation to multi-year operational horizons, including extreme weather events (heatwaves, cold snaps, storms) and long-term trends (EV fleet growth, PV cost reduction, electrification of heating), was not assessed. The co-evolutionary FL-swarm loop's long-term stability - whether the feedback between predictive accuracy and operational responsiveness converges or oscillates over multi-year horizons - remains an open question.

Adversarial and cyber-physical resilience. The cyber-resilience tests (99.6% rejection of spoofed telemetry, BFT consensus under 33% Byzantine nodes) were conducted under controlled injection scenarios. Real-world adversarial campaigns involving coordinated, multi-vector attacks spanning BTM, DSO, and TSO Data Spaces simultaneously were not simulated. The interaction between cyber-attacks and physical grid events (e.g., a coordinated attack during a frequency-deviation event) represents a validation gap requiring dedicated red-team exercises.

Pathways to continental scaling. The validation gaps identified above define the research and deploy-

ment agenda for the next phase. Specifically: (i) a 100,000-node BTM pilot across three Member States to validate horizontal scalability and cross-border Data Space federation; (ii) a 200-feeder DSO pilot to validate swarm-computing convergence at national distribution-network scale; (iii) a multi-year FL model-stability study incorporating extreme-weather scenarios; (iv) a coordinated cyber-physical red-team exercise spanning all four Data Space domains; and (v) a market-maturity sensitivity study across Member States with varying balancing-market structures. The Open Call mechanism (per *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*) and the Standards Working Group's RFC process (per *6. Standardisation Pillar*) provide the governance infrastructure for these next-phase activities, ensuring that scaling proceeds through open, vendor-neutral, and ethically governed pathways consistent with the framework's design invariants.

10. Conclusions and Future Work

10.1 Summary of Contributions

The HEDGE-IoT Framework addresses a structural deficiency in the digitalisation of energy systems: the absence of a unified, multi-dimensional reference architecture capable of transforming heterogeneous IoT assets - spanning behind-the-meter devices, distribution-network infrastructure, and transmission-system platforms - into coordinated, privacy-preserving, market-ready flexibility resources. As established in *1. Introduction and Motivation*, five interlocking barriers (protocol/semantic heterogeneity, non-linear scalability demands, divergent latency requirements, data-sovereignty constraints, and persistent vendor lock-in) have rendered single-layer and single-axis interventions structurally inadequate. The framework's core proposition - that a four-pillar, one-continuum architecture unified under a horizontal computational orchestration layer is necessary and sufficient to resolve these barriers simultaneously - has been validated through architectural design, formal specification, and empirical multi-level pilot deployment.

The principal contributions relative to prior work, as catalogued in *2. State of the Art and Related Work*, are fivefold. First, HEDGE-IoT is the first framework to jointly optimise technology facilitation, interoperability, standardisation, and ecosystem enablement within a single coherent envelope, rather than treating connectivity, data management, analytics, or market access in isolation. Second, it introduces energy-specific federated learning and swarm-computing mechanisms with formal convergence guarantees under highly non-IID data distributions and intermittent heterogeneous connectivity - conditions that existing edge/fog paradigms (OpenFog, ETSI MEC, SENSE4US) do not accommodate. Third, it provides a deployable, energy-specific Data Space blueprint spanning BTM IoT telemetry, DSO operational data, and TSO market platforms under a unified yet federated governance model, closing the gap where IDSA/Gaia-X and SAREF/WoT efforts existed as parallel but unintegrated initiatives. Fourth, it elevates standardisation from a retrofitted afterthought to a first-class architectural pillar with binding influence on every interface, data artefact, and governance contract. Fifth, it supplies pilot-validated evidence of impact on flexibility activation, RES hosting, and grid resilience, together with a replicability assessment confirming substantial context-independence across regulatory, climatic, and asset-mix environments.

The six design invariants articulated in *3. HEDGE-IoT Framework Architecture* - edge-autonomy by default; data never leaving its sovereign boundary unprocessed; open standards as the default interface; graceful degradation under connectivity loss; ethical and inclusive design as architectural constraints; and empirical validation at every tier - constrain all implementation choices and ensure the framework is deployable, verifiable, and extensible rather than merely conceptual. The cross-pillar integration produces emergent capabilities absent from any single pillar, as demonstrated through the three representative end-to-end flows (privacy-preserving flexibility discovery, federated congestion forecasting, new-asset onboarding) and confirmed empirically across the multi-level pilot portfolio described in *9. Scalability, Replicability, and Validation.*

10.2 Demonstrated Impact on Energy System Intelligence and Flexibility

The empirical validation campaign, conducted over 180-day observation windows across three coordinated pilots (BTM/Residential at 1,200→4,800 nodes, DSO/Feeder-Level across 14 MV feeders serving 8,400 customers, and TSO/Market-Level at bidding-zone balancing), confirms that the framework's value is emergent from cross-pillar coupling rather than achievable through any single-pillar intervention. The quantitative outcomes substantiate the framework's central economic and operational propositions.

Flexibility activation. The four-stage conversion pipeline described in *7. Digital Energy Ecosystem Enabling Pillar: RES Integration, Resilience, and Market Opportunities* - Discovery and Classification, Aggregation and Progressive Abstraction, Predictive Forecasting and Swarm-Based Execution, and Market Integration and Settlement - increased activated flexibility from 8.3% to 61.7% of technically available capacity within 90 days. Swarm-based dispatch achieved 4.7% MAPE against FL forecasts, with 94.2% of events completing within the 500 ms consensus window. The co-evolutionary FL-swarm loop, formalised in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*, reduced dispatch MAPE by 31% over 180 days. Activated flexibility absorbed 22.4% of peak RES export without physical network reinforcement, deferring an estimated €2.1M in DSO capital expenditure over a 10-year horizon. The self-reinforcing flexibility flywheel was observed empirically: 34% asset-registration growth correlated with 19% FL accuracy improvement and 12% revenue increase per household.

RES integration. FL-derived probabilistic forecasts (Brier score 0.087) enabled proactive swarm dispatch rather than reactive curtailment, reducing RES curtailment by 59.2% and congestion events by 76.6%. BFT-consensus swarm protocols delivered 1.2 MW synthetic inertia in <380 ms, providing TSO-grade frequency control without synchronous generation while tolerating $(N-1)/3$ faulty agents. Seasonal robustness was

confirmed across a full autumn-to-spring transition, with FL accuracy degrading only 8.3% during the transition and recovering within 14 days via the co-evolutionary retraining loop.

Grid resilience. Resilience was validated as a multi-dimensional, distributed property: during a 4-hour backhaul outage affecting 35% of BTM nodes, delta-synchronisation bounded accuracy loss to 1.4% MAPE; a simulated DSC compromise was isolated within 4.8 seconds through policy-as-code enforcement; full system recovery after a 12-hour DSO outage was achieved in 47 minutes; and spoofed-telemetry injection was rejected at 99.6% by schema-level payload inspection. The layered-sovereignty Data Space architecture prevented cascade failures, and the Federated Application Broker's stateless cross-Data-Space coordination eliminated single points of failure.

Scalability and replicability. DSC policy-evaluation latency remained <5 ms at 4,800 BTM nodes; the workload scheduler sustained sub-100 ms placement decisions at 10,000 concurrent tasks; and new-asset onboarding time remained stable at 2.3-4.1 hours regardless of node count, confirming the "months to days" reduction. The composite replicability score of 0.87 (0-1 scale) across regulatory, climatic, and asset-mix environments confirms substantial context-independence, with residual variation attributable to physical and climatic differences rather than architectural or regulatory barriers.

Ecosystem and inclusivity outcomes. The Open Call mechanism achieved median onboarding of 3.4 hours with zero integration contracts; energy-poverty protections, three-tier conformance, and algorithmic transparency requirements reproduced across all three pilots. The progressive abstraction protocol converted GW-scale latent flexibility into a coordinated, market-ready resource without requiring any single entity to possess a complete centralised view, fulfilling the framework's core proposition stated in *1. Introduction and Motivation*.

10.3 Limitations and Open Challenges

The validation campaign, while comprehensive within its scope, identified specific boundaries beyond which the framework's behaviour remains untested. The maximum validated scale (4,800 BTM nodes, 14 feeders) remains two to three orders of magnitude below national deployments. Satellite-link multi-day outages, FL behaviour at 50%+ offline-client rates, multi-year model stability under extreme weather events, coordinated multi-vector cyber-physical attacks, and market-maturity sensitivity across less-liberalised Member States constitute open empirical questions. The 0.87 replicability score, while encouraging, leaves 13% of variance unexplained - primarily attributable to climatic and asset-mix heterogeneity that may compound at continental scale.

Architecturally, the current orchestration layer's multi-objective placement optimisation (latency, energy, communication, fairness) employs genetic-algorithm search with constraint filtering, which, while achieving sub-100 ms decisions at 10,000 tasks, may face combinatorial challenges at 10 -10 concurrent workloads characteristic of national-scale deployments. The FL convergence guarantees, validated under 15-40% offline-client conditions, require extension to the more adversarial regimes (50%+ offline, Byzantine clients, adversarial gradient injection) anticipated in open, multi-stakeholder environments. The semantic mediator's runtime ontology mapping, while effective for the current TSTO/FTO/CDCO stack, has not been stress-tested against the full combinatorial explosion of SAREF extensions, IEC 61850 profiles, and emerging WoT Thing Description variants at continental scale.

10.4 Future Research Directions

The identified limitations and the framework's inherent extensibility define a structured research agenda organised along four principal axes.

10.4.1 Advanced Orchestration Algorithms

The computational orchestration layer, described in *3. HEDGE-IoT Framework Architecture* and operationalised in *4. Technology Facilitator Pillar: AI/ML Federated Learning and Swarm Computing at the Grid Edge*, requires algorithmic evolution to sustain sub-100 ms placement decisions and cross-Data-Space coordination at national and continental scale.

Reinforcement-learning-augmented placement. The current genetic-algorithm search, while effective at 10,000 concurrent tasks, exhibits diminishing returns as the task-resource matrix grows super-linearly. Future work will develop multi-agent reinforcement learning (MARL) schedulers that learn placement policies from historical workload traces, incorporating grid-topology-aware reward shaping (voltage sensitivity, congestion proximity, RES-penetration gradients) and constraint-aware action masking to guarantee hard policy-as-code compliance. The objective is to maintain <100 ms decision latency at 10 concurrent tasks while improving energy-efficiency of placement by 15-25% relative to the current GA baseline.

Predictive orchestration and proactive resource provisioning. The current reactive placement model will be extended with a predictive layer that anticipates workload demand from FL training cycles, swarm consensus rounds, and market-clearing windows. Temporal graph neural networks (TGNNs) trained on historical orchestration telemetry will forecast resource contention 5-15 minutes ahead, enabling proactive DSC scaling, gradient-buffer pre-allocation, and swarm-participant pre-warming. This predictive orchestration is particularly critical for the store-and-forward gradient buffering mechanism under satellite-link multi-day outages, where buffer sizing must be optimised a priori rather than reactively.

Cross-Data-Space orchestration under adversarial conditions. The Federated Application Broker's stateless coordination protocol must be hardened against Byzantine orchestration nodes, Sybil attacks on DID-based identity, and coordinated policy-evasion attempts. Future research will formalise a Byzantine-fault-tolerant orchestration consensus (extending the BFT swarm protocols to the orchestration plane) with formal guarantees under up to $(N-1)/3$ adversarial orchestrators, and develop zero-knowledge proof mechanisms enabling policy-compliance verification without revealing underlying data or model parameters to the broker.

Multi-objective optimisation under dynamic regulatory constraints. The living traceability matrix mapping the standards stack against GDPR, DGA, NIS2, AI Act, and grid-code requirements introduces a time-varying constraint surface. Future orchestration algorithms must incorporate regulatory-change prediction (e.g., anticipated DGA implementing acts, AI Act conformity-assessment deadlines) as soft constraints with configurable relaxation, enabling graceful policy adaptation without service interruption.

10.4.2 Expanded Standardisation Scope

The Standardisation Pillar, detailed in *6. Standardisation Pillar: SAREF and Open Standards for Energy IoT*, establishes SAREF-anchored ontologies, OpenAPI/AsyncAPI specification surfaces, and a three-tier conformance regime. Future standardisation work must extend this foundation to address emerging domains, deepen cross-sectoral alignment, and ensure long-term resilience across 10-15-year asset lifecycles.

SAREF extension for emerging asset classes. The current FTO taxonomy covers residential storage, V2G, thermal loads, and curtailable industrial processes. Future extensions must incorporate: green-hydrogen electrolyser flexibility (with ramp-rate and minimum-run-time constraints); grid-forming inverter capabilities (synthetic inertia, virtual synchronous machine parameters); multi-energy coupling (electricity-heat-gas-mobility); and AI-native energy services (model-as-a-service, data-as-a-service, orchestration-as-a-service). These extensions will be submitted as candidate SAREF 4.x/5.x modules through the Standards Working Group's RFC process, with SHACL shapes and backward-compatible mapping adapters maintained for at least two major versions.

Cross-sectoral semantic alignment. The TSTO, FTO, and CDCO ontologies must be formally aligned with: IEC 61850-7-420 (DER modelling), IEC 62357 (market information exchange), ISO 15118 (V2G communication), ETSI TS 103 621 (smart metering), and the EU's forthcoming Energy Data Space reference architecture. Bidirectional crosswalk mappings, validated through conformance testing against reference implementations from CEN/CENELEC CLC/TC 57X and IEC TC 57, will ensure the framework's internal SAREF-anchored representation remains the single source of truth while external systems consume their native formats without loss of semantic fidelity.

AI Act and algorithmic-transparency standardisation. The algorithmic transparency requirements described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity* - Verifiable Credentials recording model lineage, feature-selection rationale, and differential-privacy budgets - must be formalised into a standardised AI-Act conformity-assessment artefact set. This includes: standardised model cards for FL models (training-data provenance, non-IID characterisation, convergence guarantees, fairness metrics stratified by CDCO consumer-archetype metadata); standardised swarm-policy documentation (BFT parameters, convergence bounds, fault-tolerance guarantees); and standardised orchestration-audit records (placement decisions, policy evaluations, SLA compliance). These artefacts will be contributed to the AIOTI working groups and the European Commission's AI Act implementing-act drafting process.

Temporal and spatial standardisation for distributed consensus. The IEEE 1588 PTP, NTPv4, and hybrid logical clock mechanisms providing temporal semantics for voltage regulation, market clearing, and FL gradient ordering require extension to multi-timezone, multi-jurisdiction deployments. Future standardisation work will define a cross-border temporal-synchronisation profile for swarm consensus and FL gradient ordering, incorporating leap-second handling, daylight-saving transitions, and regulatory settlement-period boundaries across Member States.

Conformance-regime evolution for AI-native components. The three-tier conformance regime (Semantic, Interface, Behavioural) must be extended to accommodate AI-native components whose behaviour is non-deterministic and data-dependent. This includes: standardised model-performance envelopes (accuracy

bounds under specified data distributions); standardised drift-detection and retraining triggers; standardised adversarial-robustness certification for FL models and swarm policies; and standardised explainability interfaces enabling DSOs, TSOs, and regulatory authorities to verify compliance without access to raw training data.

10.4.3 Deeper AI/ML Integration

The Technology Facilitator Pillar establishes hierarchical FL, swarm computing, and their co-evolutionary coupling. Future AI/ML research must deepen integration across all four pillars, extend model capabilities, and address the open convergence and robustness challenges identified in *9. Scalability, Replicability, and Validation.*

Foundation models for energy-system intelligence. The current FL architecture trains task-specific models (load forecasting, congestion prediction, voltage regulation). Future work will explore energy-specific foundation models - large-scale FL-trained transformers pre-trained on multi-domain energy telemetry (BTM consumption, DSO feeder loading, TSO market data, weather, grid topology) - fine-tuned via federated transfer learning for downstream tasks. The hierarchical FL topology (local → regional → global) provides a natural curriculum for progressive pre-training. Key challenges include: communication-efficient fine-tuning under intermittent connectivity (extending the 60-80% federated feature-selection reduction to LoRA/QLoRA-style parameter-efficient adaptation); non-IID pre-training across heterogeneous consumer archetypes; and privacy-preserving model-architecture sharing under differential-privacy budgets (1.0 for residential data).

Multi-modal and multi-task federated learning. Current FL jobs are single-task. Future research will develop multi-task FL frameworks where a single training round simultaneously optimises load forecasting, congestion prediction, voltage-regulation policy, and flexibility-availability estimation - exploiting shared representations across tasks while maintaining task-specific heads. Multi-modal FL will integrate heterogeneous data types (time-series telemetry, grid-topology graphs, weather imagery, market-price sequences) within a unified federated training loop, with the semantic mediator (SAREF ontologies) providing the cross-modal alignment substrate.

Causal and physics-informed ML for grid operations. Purely data-driven FL models, while effective for prediction, lack the structural guarantees required for safety-critical grid operations. Future work will integrate physics-informed neural networks (PINNs) encoding power-flow equations, Kirchhoff's laws, and thermal constraints into FL model architectures, ensuring that swarm control policies derived from FL predictions respect physical feasibility by construction. Causal-inference methods (structural causal models, do-calculus) will be integrated into the FL pipeline to distinguish correlation from causation in flexibility-activation signals, preventing spurious market signals that could destabilise grid operations.

Adversarial robustness and Byzantine resilience in FL. The validation campaign confirmed convergence under 15-40% offline clients but did not test adversarial gradient injection, model-poisoning attacks, or coordinated Byzantine client behaviour. Future research will develop: Byzantine-robust aggregation rules (trimmed mean, Krum, multi-Krum) with formal guarantees under up to (N−1)/3 adversarial clients; differential-privacy budget monitoring with automatic training-round suspension upon -exceedance; and adversarial-robustness certification for deployed FL models, integrated into the Verifiable Credential model-card framework.

Swarm-computing extensions: hierarchical multi-agent reinforcement learning. The current stigmergic coordination and gossip-based consensus protocols achieve sub-500 ms voltage-regulation equilibria for 500-node feeders. Future work will extend swarm computing to hierarchical multi-agent reinforcement learning (H-MARL), where micro-level swarm agents (Tier 0-1 devices) learn local control policies through FL-informed reward shaping, meso-level agents (Tier 1-2 gateways) coordinate feeder-level equilibria, and macro-level agents (Tier 3) optimise bidding-zone dispatch. The co-evolutionary FL-swarm loop will be formalised as a joint optimisation problem with convergence guarantees under the hierarchical topology, extending the current empirical 31% MAPE reduction to theoretically bounded improvement rates.

Digital-twin integration for predictive orchestration. The distributed observability plane provides per-task telemetry and anomaly detection. Future work will develop energy-system digital twins - physics-based simulation models continuously calibrated by FL-derived parameters and swarm-execution outcomes - enabling what-if analysis, predictive maintenance, and pre-deployment validation of new FL models and swarm policies before production instantiation. The digital twin will be instantiated within the BTM and DSO Data Spaces, respecting the "data never leaves its sovereign boundary unprocessed" invariant.

10.4.4 Broader European Deployment Pathways

The replicability assessment (composite score 0.87) and the identified validation gaps define a structured continental-scaling pathway. Future deployment research must address the two-to-three-order-of-magnitude

scale gap, cross-border regulatory heterogeneity, market-maturity variation, and the socio-technical dimensions of multi-stakeholder adoption.

100,000-node cross-border BTM pilots. The next-phase validation will instantiate BTM Data Spaces across 100,000+ nodes spanning at least four Member States with divergent regulatory frameworks (e.g., fully liberalised vs. partially regulated markets), climatic zones (Mediterranean, Continental, Nordic), and asset mixes (high PV penetration, high wind penetration, mixed). The Federated Application Broker will mediate cross-border FL jobs and swarm-computing cycles, testing the progressive abstraction protocol's ability to maintain per-household data sovereignty while enabling cross-border flexibility aggregation. Satellite-link multi-day outage scenarios and 50%+ offline-client FL behaviour will be explicitly tested.

200-feeder DSO studies and multi-year FL stability analyses. DSO-level validation will expand from 14 to 200+ MV feeders across multiple distribution operators, testing the workload scheduler at 10 -10 concurrent tasks and the DSC policy engine at 10 + BTM nodes per DSO domain. Multi-year FL stability analyses (3-5 years) will assess model drift under climate-change-driven load-pattern shifts, asset-base evolution (EV fleet growth, heat-pump adoption, battery-storage proliferation), and regulatory-change impacts on feature distributions. The co-evolutionary retraining loop's long-term convergence properties will be formally characterised.

Coordinated multi-vector cyber-physical attack testing. Red-team exercises will simulate coordinated attacks spanning: DSC compromise with policy-evasion attempts; spoofed-telemetry injection at scale (>10 simultaneous injections); FL model-poisoning via adversarial gradient manipulation; swarm-consensus disruption via Sybil node injection; and orchestration-layer denial-of-service. The framework's layered-sovereignty isolation, BFT consensus, schema-level payload inspection, and policy-as-code enforcement will be stress-tested against multi-vector, multi-stage attack chains. Formal security proofs for the cross-Data-Space federation protocol under adversarial conditions will be developed.

Market-maturity sensitivity studies. The three-level market architecture (local flexibility, regional/national balancing, P2P/community) will be evaluated across Member States with varying degrees of market liberalisation, aggregator regulation, and consumer-protection frameworks. The SAREF-to-IEC 62357/REMIT crosswalks will be validated against national market-clearing platforms (e.g., EPEX SPOT, Nord Pool, MIBEL) to confirm that flexibility signals produced by the progressive abstraction protocol are directly consumable by existing market infrastructure without bespoke integration. Market-maturity sensitivity will inform the Open Call selection criteria and the Standards Working Group's regulatory-alignment roadmap.

SME and innovator ecosystem scaling. The Open Call mechanism, described in *8. Stakeholder Engagement, IoT Ethics, and Inclusivity*, will be scaled from pilot-level sandbox access to a permanent, EU-funded innovation pipeline. Future work will develop: automated conformance-certification tooling reducing SME onboarding from 3.4 hours to <1 hour; a marketplace for FL models, swarm policies, and Data Space connectors enabling peer-to-peer component exchange; and a structured mentorship programme pairing SMEs with DSO/TSO operational teams to ensure pilot-to-production transition. The two reserved Standards Working Group seats for SME/innovator representatives will be expanded to a dedicated SME/Innovator Advisory Board with formal voting rights on RFCs affecting integration costs and conformance requirements.

Cross-border Data Space federation and EU policy alignment. The four-domain Data Space topology will be extended to a pan-European federation spanning all 27 Member States, with the Cross-Domain Federation Space (Tier 4) operating under a multi-stakeholder governance charter aligned with the European Data Strategy, the Data Governance Act, the Data Act, and the forthcoming European Energy Data Space. The living traceability matrix will be maintained as a public, machine-readable artefact updated at each regulatory milestone, providing a single auditable compliance reference for DSOs, TSOs, aggregators, and regulatory authorities. Formal liaison with AIOTI, the Smart Grids Task Force, ENTSO-E, ACER, and the European Commission's DG ENER/DG CNECT will ensure the framework's evolution remains synchronised with EU digitalisation and energy-policy trajectories.

Replicability-to-reproducibility transition. The 0.87 replicability score must be elevated toward 0.95+ through: standardised deployment playbooks per Member State (regulatory configuration, grid-code parameters, market-platform interfaces); automated vulnerability-classification threshold calibration per national social-welfare frameworks; and a continental reference implementation (open-source, vendor-neutral) serving as the canonical instantiation against which national deployments are validated. The Standards Working Group's RFC process will govern the evolution of these playbooks, ensuring that local adaptations do not fragment the underlying architecture.

10.5 Closing Remarks

The HEDGE-IoT Framework demonstrates that the transformation of heterogeneous, fragmented IoT assets into a coordinated, privacy-preserving, market-ready flexibility resource is achievable through a multi-dimensional, open reference architecture that jointly optimises technology facilitation, interoperability, standardisation, and ecosystem enablement under a unified computational orchestration layer. The empirical validation across the full cloud/edge continuum - from 1,200 BTM nodes to bidding-zone market coordination - confirms that the framework's value is emergent from cross-pillar coupling: the 61.7% flexibility activation, 59.2% curtailment reduction, 76.6% congestion-event reduction, and sub-500 ms swarm consensus are not attributable to any single mechanism but to the architecturally enforced mutual reinforcement of all four pillars.

The framework's six design invariants - edge-autonomy, data sovereignty, open standards, graceful degradation, ethical inclusivity, and empirical validation - ensure that this transformation proceeds without centralised data accumulation, vendor lock-in, or the erosion of consumer rights. The embedding of ethical and inclusive safeguards as architectural defaults, rather than procedural overlays, positions HEDGE-IoT as a reference model for trustworthy digitalisation of critical infrastructure across the European energy system.

The future research agenda - advanced orchestration algorithms, expanded standardisation scope, deeper AI/ML integration, and broader European deployment pathways - defines a structured, multi-year programme to scale the framework from validated pilot to continental infrastructure. The identified limitations (two-to-three-order-of-magnitude scale gap, adversarial FL robustness, multi-year model stability, market-maturity sensitivity) are not architectural deficiencies but empirical boundaries of the current validation campaign, addressable through the next-phase pilot portfolio governed through the Open Call mechanism and Standards Working Group RFC process. The framework's openness, extensibility, and vendor-neutrality by construction ensure that this scaling proceeds through transparent, multi-stakeholder, and ethically governed pathways - consistent with the EU's digital sovereignty objectives, the European Green Deal's decarbonisation imperatives, and the fundamental right of every energy consumer to participate in, and benefit from, the intelligent energy system of the future.