FLEXIBILITY EXPLOITATION FOR RESIDENTIAL, TERTIARY AND INDUSTRIAL BUILDINGS

FLEXBIT PROJECT

CETPartnership Joint Call 2023

Deliverable 2.3 – Data Management and Security Protocols

Lead Author: Systems Research Institute Polish Academy of Sciences/Anastasiya Danilenka orcid: https://orcid.org/0000-0002-3080-0303

Contributed authors:
Szymon Smagowski orcid: https://orcid.org/0009-0003-1965-8099
Pio Alessandro Lombardi orcid: https://orcid.org/0000-0003-4598-9759
Gaetano Zizzo orcid: https://orcid.org/0000-0003-4413-4855
Mateusz Witkowski: https://orcid.org/0009-0005-4439-1458
Francesco Saverio Cannizzaro orcid: https://orcid.org/0009-0004-1528-1682
Lorenzo Bartolucci orcid: https://orcid.org/0000-0003-4258-4860
Finn Bennet Schröder
Tomasz Sikorski orcid: https://orcid.org/0000-0002-4423-7216
Christina Karatrantou orcid: https://orcid.org/0009-0008-9818-048X
Pavlos Tyrologou orcid: https://orcid.org/0000-0001-7706-1774
Nikolaos Koukouzas orcid: https://orcid.org/0000-0002-7094-3712
Muhammad Mahad Malik orcid: https://orcid.org/0009-0001-6425-8954
Alexander Micallef orcid: https://orcid.org/0000-0002-9497-5604
Alessandro Polimeni orcid: https://orcid.org/0009-0005-9359-8746
Lukasz Michalec orcid: https://orcid.org/0000-0002-2192-833X
Jacek Rezmer orcid: https://orcid.org/0000-0002-2822-7595

DOI: **10.5281/zenodo.20506409**

This research was funded by CETPartnership, the Clean Energy Transition Partnership under the 2023 joint call for research proposals, co-funded by the European Commission (GA N°101069750) and with the funding organisations detailed on https://cetpartnership.eu/funding-agencies-and-call-modules.

Co-funded by the European Union

Supported by:

Federal Ministry for Economic Affairs and Climate Action

on the basis of a decision by the German Bundestag

Ministero delle Imprese e del Made in Italy

GENERAL SECRETARIAT FOR RESEARCH AND INNOVATION

FlexBIT Consortium

Table of Contents

1 Executive Summary

This deliverable specifies the data management and security protocols adopted in the FlexBIT project. It builds upon the platform architecture defined in D1.1 and incorporates regulatory and compliance requirements set out in D1.2, such as data-access constraints, cybersecurity expectations, and country-specific regulatory differences relevant to demonstrators' integration. The deliverable defines the operational approach to data management, privacy, security, and interoperability in the FlexBIT platform. It further details the current platform integration and deployment status per demonstrator.

The scope of the data covered by this deliverable spans data produced for use by the platform and supporting development activities (e.g., AI model development and training), as well as platform-generated logs, monitoring records, and audit-related data, and data collected outside of the platform, such as proprietary demonstrator data, energy community survey results, and relevant metadata. This deliverable also specifies how data is acquired, accessed, stored, shared, and handled after project completion.

The objectives of this deliverable are as follows:

- describe data sources and data flows, including demonstrator-generated, platform-generated, and external data sources;
- establish the data governance model;
- specify the approach to data privacy and security controls;
- describe the data model, protocols, interfaces, and their current implementation within the FlexBIT platform;
- describe the current status of deployment across FlexBIT demonstrators.

As part of the WP2 deliverables set, this deliverable complements D2.1 and D2.2. As D2.1 lays down the energy exchange mechanisms, D2.3 provides the necessary context for their operation, detailing the adopted interoperability and security tools of the platform. Compared to D2.2, which focuses on forecasting and supporting datasets, D2.3 targets a broader data management perspective, covering both internal and external data sources, their lifecycle, and accompanying privacy and governance considerations.

This deliverable is structured in sections, covering the system context, data governance, privacy, security, digital twin technologies, interoperability/data exchange, and current deployment status across demonstrators.

2 System Context

2.1 Platform architecture overview

The FlexBIT platform is a modular, layered system that connects distributed energy resources across the project demonstrators in Germany, Italy, Poland, and Malta. Its full architecture is described in D1.1; this section recalls only the elements that are relevant for data management, privacy, security, and interoperability.

The platform is organised in layers. At the bottom sits the energy infrastructure, which includes PV systems, batteries, hydrogen storage, EV chargers, and controllable loads at each demonstrator site. Above it, the data acquisition and enforcement layer brings together smart meters, IoT gateways, edge controllers, and local SCADA or EMS systems that collect measurements and execute control actions. The data processing layer handles ingestion, storage, and analytics, while the energy business layer manages market-related and regulatory functions. Underneath everything, the digital infrastructure layer provides containerisation, orchestration, and deployment tooling, and a cross-cutting security and monitoring layer protects the platform end to end.

A key principle of the architecture is the separation between local autonomy and central coordination. Each demonstrator keeps full control of its physical assets through a fast local loop run by its own SCADA, EMS, or PLC systems. The FlexBIT central platform operates on a slower advisory loop, working on aggregated data and sending coordination signals back to the sites. This split shapes how data is managed in FlexBIT: local environments stay decoupled from central services through standardised adapters, and only normalised data crosses the boundary into the central platform. This architectural separation also enables proportional data exposure, controlled interoperability and more transparent governance boundaries between locally retained operational datasets and platform ingested information flows.

The central platform also includes a permissioned blockchain layer based on Hyperledger Besu with QBFT consensus. Detailed operational data stays off-chain, and only cryptographic fingerprints are recorded on-chain to provide tamper-evident proof of integrity.

2.2 Data producers and flows

The FlexBIT data acquisition approach is based on the separation between local site integration and central platform processing. Demonstrator-specific systems may differ in terms of devices, protocols, data structures, and local control environments, but data are exposed to the central platform through a common ingestion approach based on standardised adapters. These adapters act as the interface between heterogeneous field systems and the central FlexBIT services, reducing coupling between local implementations and the platform data layer. In this way, local technological differences remain encapsulated at adapter level, while the platform receives data through a harmonised access model. This approach reduces the fragmentation between demonstrators and supports long-term interoperability, traceability and extensibility of data services.

2.2.1 Data sources

At system level, the platform is designed to receive operational and contextual data originating from different categories of producers. These include:

- smart meters and site metering systems;
- IoT sensors and embedded monitoring devices;

- SCADA, EMS, PLC, or equivalent local supervisory systems;
- third-party aggregators or external data providers, where relevant;
- platform-generated operational data, such as logs, alarms, and processing results.

For the purposes of the central data stack, these sources are not integrated directly one by one into backend services. Instead, they are made available to the platform through standardised adapters, which provide a consistent interface for ingestion independently of the original source format or protocol. This is aligned with the broader FlexBIT architecture, in which local systems retain autonomy while central components operate on normalised and shareable data structures (the adapter standardised APIs are defined in Section 7). The list of relevant data sources per demonstrator is given in Table 2.1, specifying how data are produced.

Table 2.1: FlexBIT-relevant data sources and acquisition scope per demonstrator

Demonstrator	Data-producing assets or systems (what physically or digitally produces data)	Source signal families (what kinds of measurements are produced)	Local acquisition / adapter layer (how data are collected and exposed to the FlexBIT platform)
Smart Energy Parking in Palermo	• Hybrid PV-BESS inverter • EV charging station • V2G EV charging station	• AC, DC voltages • AC, DC currents • AC, DC power • Batteries SOC • Active Power injected to the grid • Power supplied to the loads • Power supplied to the emergency loads • Power requested from the grid • Power requested by the EV charging stations • Power injected to the grid from the V2G EV charging station	Data collected from the University LAN by a dedicated PC
Solar Cooling System in Pantelleria	Heat pump Thermal solar collectors	• Power requested by the heat pump • Power requested by the auxiliary devices • Water temperature • Air temperature	Data collected from the local energy management system Data collected from the University LAN by a dedicated PC
aRTE Möbel	• PV • BESS • charging station • Compressed Air System • all Machines used in the manufacturing processes (14) and office rooms	• AC, DC voltage and • AC, DC power for PV • AC power for all the loads included compressed air • SoC, SoH, voltage, current, active power • reactive power for BESS	Data collected from local energy management system and transmitted to Fraunhofer IFF. From Fraunhofer IFF are exposed to the FlexBIT platform
Aue Funeral	PV, BESS, Charging Stations (2) and Cooling aggregates (2)	• AC, DC voltage and • AC, DC power for PV • SoC, SoH, voltage, current, active power, reactive Power for BESS • AC power, AC current and charge mode for charging station • AC and temperature power for cooling aggregates.	Data collected from local energy management system and transmitted to Fraunhofer IFF. From Fraunhofer IFF are exposed to the FlexBIT platform

Malta Digital Twin	• MATLAB/Simulink residential microgrid model • Residential PV generation datasets • Residential load demand datasets • Battery Energy Storage System (BESS) model • Grid import/export simulation environment • Forecasting modules (LSTM, LightGBM, CatBoost, Hybrid models)	• PV generation power • Residential load demand • Grid import/export power • Battery State of Charge • Battery charging/discharging power • Active and reactive power measurements • Environmental variables (irradiance, temperature, humidity, rainfall) • Forecasted PV and load signals (1-min, 5-min, 15-min, and 30-min resolutions) • Time-series temporal and lag-based features	• Historical residential datasets imported into MATLAB/Simulink environment • Signal preprocessing and forecasting pipelines developed in Python/MATLAB • Simulated real-time signal exchange through the Simulink environment • Planned integration to the FlexBIT platform through FlexBIT API-based communication and platform-aligned interfaces • Ongoing preparation for Speedgoat HIL-compatible real-time data exchange
Tor Vergata Demonstrator	MATLAB/Simulink smart energy community model (local plant, prosumers, consumers).	• PV generation power • residential load demand • grid import/export power • battery SoC • battery charging/discharging power • stored hydrogen • hydrogen mass flow rate	Simulated real-time signal is exchanged with a Raspberry Pi, which communicates with the FlexBIT platform.
WUST Laboratory Demonstrator	• Microgrid hardware model using PV inverter with PV Array Simulator • AC Grid Simulator • battery energy storage (BESS) • controllable loads • measurements setup	• PV: AC, DC voltage, current, AC active, reactive, apparent • BESS: AC, DC voltage, current, SoC, AC active Power, reactive Power • Controllable load: AC voltage, current, active, reactive, apparent power • Measurement setup: power meters and power quality meters collecting AC power data, AC power quality parameters for particular devices. • Controller: control signals i.e. active power request for BESS, reactive power request for PV	Offline dataset: Characteristic of regulatory potential of microgrid devices, i.e. active-reactive power plane for PV, active power in relation to SOC for BESS. Simulated real-time scenarios of control demands, requests and responses of the microgrid devices Online dataset: Planned integration to the FlexBIT platform through FlexBIT API based on microcomputer, broker and protocols
Electrum/Alu Frost	• Industrial SCADA/EMS system (EMACS 2.0)	• Active/reactive power, energy import/export • Voltage, current, frequency • PV generation	Data collected from field devices via SCADA/EMS infrastructure using industrial protocols (MQTT, Modbus TCP, DNP3, REST

	• Energy Storage System (BESS) • PV installation • EV chargers • Industrial loads (production lines, auxiliary systems) • Heaters / HVAC systems • Sky Eye Station (all-sky camera) • Pyranometer • Meteorological station	• BESS SoC, charge/discharge power • EV charging power, status • Industrial load profiles, equipment states • Ambient temperature, humidity, wind speed • Solar irradiation • Sky images (all-sky) • Alarms, events, status signals	API). Raw data is normalised by the local API Agent and telemetry layer (RabbitMQ broker). Data is securely transmitted to the FlexBIT platform through HTTPS/TLS and Kafka Streams.

2.2.2 Stream processing

The ingestion backbone of the current platform stack is based on Redpanda, used as an Apache Kafka-compatible streaming layer. In this role, it supports asynchronous data intake from the adapter layer and decouples data producers from downstream consumers. This makes it possible to process incoming data flows in a scalable and modular way, without requiring tight synchronisation between acquisition and persistence services [13].

From an architectural point of view, Redpanda acts as the central event-streaming layer between the standardised adapters and the backend processing services. This supports:

- buffering of incoming data streams;
- asynchronous handoff to downstream services;
- separation between data production and persistence;
- easier extension of processing pipelines over time.

The topic authorisation is managed by the central FlexBIT platform through the API Keys system.

2.2.3 Persistence

Downstream of the ingestion layer, the current stack includes a dedicated data processor service. It can be described as a Kafka-to-QuestDB consumer (this component is designed to consume streamed data from the message backbone and forward or transform them toward persistence services without maintaining its own persistent state).

This processing pattern is important because it allows the data pipeline to remain modular:

- adapters expose and normalize data toward the stream layer;
- Redpanda handles ingestion and transport;
- the data processor consumes stream data and prepares them both for storage and for blockchain (attestation/ cryptographic proof);
- persistence is delegated to backend databases.

The persistence layer is split across multiple backend components, each serving a distinct role in the platform architecture:

- PostgreSQL is used for relational and application data, where structured records, platform entities, and general-purpose persistent information must be stored in a transactional database;
- QuestDB is used for time-series persistence, making it suitable for timestamped operational measurements and telemetry data;
- Redis is used for caching and low-latency temporary access, supporting fast retrieval of transient or frequently requested data.

2.3 Trust boundaries and security domains

The FlexBIT platform operates across a distributed environment involving multiple demonstrator sites, a central platform, and a blockchain attestation layer. From a system-context perspective, this means that data does not flow within a single homogeneous environment. Data crosses multiple trust boundaries, each characterised by different levels of exposure, different ownership, and different security responsibilities.

Four boundaries structure the FlexBIT platform end to end, consistent with the architecture described in D1.1 and illustrated in the security boundary diagram in Section 5.1. The summary of trust boundaries is given below and in Table 2.2.

2.3.1 Trust boundary 1: Local site

The first boundary encompasses the physical assets, local control system, and on-site acquisition infrastructure at each demonstrator. This includes PV inverters, BESS system, EV chargers, Heat Pumps, Electrolyze, Fuel Cell, Cooling Aggregators, Compressed Air System, smart meters, IoT sensors, PLCs, SCADA/EMS, and any edge devices involved in real-time monitoring and control. Data generated within this boundary consists of raw, high-frequency operational signals such as voltages, currents, power flows, state of charge, and temperature measurements. These data are metered and initially handled under the full responsibility of the demonstrator partner. Local control loops operate independently of the central platform, and physical asset safety remains entirely within this boundary. The FlexBIT central platform does not have direct access to field-level signals, it receives only what is explicitly exposed through the adapter layer.

2.3.2 Trust boundary 2: Integration layer

The second boundary covers the standardised adapters and the communication infrastructure between the local sites and the central platform. Adapters normalize local signals into the canonical FlexBIT data model and expose them through authenticated and encrypted channels towards the central services. The data crossing this boundary are already normalised and filtered: only the signals required for platform operation are transmitted, in line with the data minimisation principle. Communication at this boundary relies on HTTPS and Kafka Streams with TLS protection, as described in Section 7.2. API keys provisioned by the platform control which adapters are authorised to push data. Responsibility for correctly implementing and operating the adapter belongs to the demonstrator partner, while the platform operator defines and enforces the authentication of an authorisation model.

2.3.3 Trust boundary 3: Central platform

The third boundary encompasses the central FlexBIT services: the API and application layer, the data processing pipeline, and the storage components (PostgreSQL, QuestDB, Redis). Data entering this boundary have already been normalised and authenticated at the integration layer. Within this boundary, the platform is responsible for ingestion, persistence, access control, and processing of shared data from all demonstrators. Role-based access control limits visibility to authorised consortium partners according to their project roles and organisational affiliation, as detailed in Section 5.2. The platform operator holds responsibility for security, availability, and integrity of all services within this boundary.

2.3.4 Trust boundary 4: Blockchain

The fourth boundary is the permissioned Hyperledger Besu network, which provides distributed integrity assurance across the consortium. Data crossing into this boundary are not raw operational records, but cryptographic fingerprints: Merkle roots derived from aggregated off-chain measurements. Only these proofs are registered on-chain, while detailed telemetry remains in off-chain storage. This boundary involves multiple parties (platform operator, consortium partner nodes, validators) and its governance is distributed across the permissioned network. Responsibility for node operation is shared among the participating partners, while the overall attestation design and smart contract layer are managed by the platform operator.

Table 2.2: Summary of the four trust boundaries

Trust boundary	Component covered	Data crossing the boundary	Responsible party
1 – Local site	Physical assets, PLCs, SCADA/EMS, IoT sensors, smart meters, edge devices	Raw operational signals (voltages, current, power, temperatures, SoC, etc) internal to site	Demonstrator partner
2 – Integration layer	Standardised adapters, communication channels (HTTPS, Kafka), API authentication	Normalised and filtered canonical measurements; authentication credentials	Demonstrator partner (adapter operation); platform operator
3 – Central Platform	API/ application services, data processing, PostgreSQL, QuestDB, Redis	Authenticated canonical data from all demonstrators; platform-generated logs and records	Platform operator
4 - Blockchain	Hyperledger Besu nodes, validators, partner nodes, smart contracts	Cryptographic fingerprints (Merkle roots, hashes), attestation metadata	Platform operator (design), consortium partners (node operation)

3 Data Governance

This section defines ownership, access, sharing, and publication rules for FlexBIT data assets. The governance approach adopted in FlexBIT aligns with the overarching EU guidelines, emphasising transparency, accountability and stakeholder engagement throughout data lifecycle. Despite the progressive harmonisation introduced by General Data Protection Regulation (EU) 2016/679 (Art. 5, 24, 25, 32, 89), the Data Governance Act (EU) 2022/868, the Data Act (EU) 2023/2854 and the Electricity Directive (EU) 2019/944 (Art. 23-24), significant operational fragmentation still exists across European energy-data ecosystems regarding smart-meter accessibility, interoperability maturity, metadata standardisation, cybersecurity governance, and FAIR [28] implementation practices. These deviations are particularly relevant in cross-border demonstrator environments such as FlexBIT, where heterogeneous technical infrastructures and national governance practices coexist within a common interoperable platform.

For example, Germany has established comparatively mature smart-meter and energy-data governance structures through the Federal Data Protection Act (BDSG) and advanced digital-energy infrastructures, while Italy combines GDPR implementation with additional sector-specific provisions under Legislative Decree 196/2003 and Legislative Decree 101/2018. Malta and Poland, despite full GDPR applicability through the Data Protection Act (Cap. 586) and the Act of 10 May 2018 on the Protection of Personal Data respectively, still present varying levels of digital-energy ecosystem maturity, digital integration maturity and structured energy-data governance implementation. In this framework, demonstrator partners retain control and responsibility for the datasets generated locally. Simultaneously, interoperability, metadata standardisation, controlled data exchange and system-level processing are managed through collaborative governance and established technical interfaces.

To mitigate these cross-jurisdictional governance inconsistencies, the FlexBIT platform does not rely on country-specific operational assumptions or nationally bounded data-sharing mechanisms. Instead, interoperability and governance consistency are achieved through a federated governance architecture combining standard data representations, harmonised metadata structures, ontology-aligned semantic harmonisation practices, role-based access control, controlled API-mediated exchange and demonstrator-level stewardship responsibilities.

Data generated at demonstrator level remains under the responsibility of the respective demonstrator partner, which acts as the owner of the locally generated source data and the operator of the corresponding pilot systems. In line with the FlexBIT architecture, each demonstrator maintains local operational control, while the central FlexBIT platform receives and processes only the data shared for project purposes through the agreed interfaces and communication mechanisms.

Internal exchange of data, software components, documentation, and other intellectual-property-relevant materials is governed by the FlexBIT Cooperation Agreement. This section specifies the operational data-governance, FAIR, and publication practices applied within those contractual boundaries. The Cooperation Agreement therefore functions as the contract management layer defining partner responsibilities, access conditions, confidentiality obligations and permitted reuse conditions across the distributed FlexBIT data ecosystem.

Partners operating central platform components are responsible for the secure reception, storage, processing, and access management of the data made available to those components. Access is limited to authorised project partners according to project tasks and implemented access-control measures. Data provenance, traceability, and accountability are preserved throughout the platform lifecycle through harmonised metadata structures, timestamp consistency, source attribution, controlled ingestion

pipelines, and auditable access-management procedures across the FlexBIT architecture. Within this governance structure, metadata management and FAIR-oriented publication practices are important mechanisms that enable controlled interoperability, traceability and long-term reuse of project-generated data assets.

Robust metadata frameworks, aligned with EU standards, enhance data discoverability and interoperability, thus supporting the long-term sustainability of the FlexBIT project's data infrastructure.

Public release of datasets and metadata follows the principle "as open as possible, as closed as necessary" and is subject to GDPR, confidentiality, intellectual-property, and project-level release decisions. Before any dataset is released publicly, the responsible partner assesses privacy risk, confidentiality/IP constraints, aggregation/anonymisation needs, licence compatibility, and metadata completeness. Privacy aspects are addressed in Section 4. Datasets selected for dissemination are additionally curated to preserve methodological transparency, semantic consistency, provenance information, and reproducibility potential, thereby supporting long-term interpretability and reuse beyond the immediate project context. This governance-focused publication process ensures proportional openness while maintaining legal, commercial and cybersecurity limitations specific to each demonstrator. The publication-oriented governance approach adopted within FlexBIT is additionally aligned with the principles established under the Open Data Directive (Directive (EU) 2019/1024), particularly Articles 5, 8 and 10 concerning machine-readable formats, practical reusability arrangements, and the open availability of research data generated through publicly funded research activities. Especially, Article 10 encourages the dissemination and reuse of research data under the principle "as open as possible, as closed as necessary", while recognising that legitimate constraints related to privacy protection, confidentiality, intellectual-property rights, cybersecurity, commercial sensitivity, and third-party obligations may justify controlled-access conditions.

Within the FlexBIT demonstrators, these principles are operationalised through the selective publication of aggregated and non-sensitive datasets, the use of widely adopted non-proprietary and machine-readable formats, governance-aligned metadata structures, repository-based dissemination through Zenodo, GitHub and OSF, and demonstrator-level approval procedures prior to public release. Retrieval is ensured by digital object identifiers. This governance-oriented publication model additionally supports interoperability, provenance preservation, traceability, and long-term reusability of demonstrator-generated datasets across heterogeneous digital-energy environments. Once the data becomes openly available, it will remain open under the CC0 or CC-BY 4.0 licenses. Use of descriptive metadata and keywords will enable current and future users to find relevant data quickly.

The practical implementation of these principles may nevertheless differ across demonstrator countries due to variations in digital-energy ecosystem maturity, operational interoperability, metadata standardisation practices, and national open-data governance implementation. Consequently, the FlexBIT platform adopts a harmonised project-level governance approach that combines conventional data structures, controlled publication procedures, federated stewardship responsibilities, and governance-aware FAIR implementation mechanisms preserving cross-border interoperability and coordinated reuse without requiring full centralisation of operational datasets.

The FlexBIT approach therefore treats FAIR implementation and Open Data compliance as governance-managed processes rather than unconditional publication obligations.

The data sensitivity and accessibility categories are project-level governance categories, aligned with the need to account for privacy, confidentiality, intellectual-property rights, commercial interests, security considerations, and Cooperation Agreement obligations and listed in Table 3.1 and Table 3.2 for sensitivity

and accessibility accordingly. These categories do not constitute formal EU classified-information markings.

Table 3.1: Project-level sensitivity categories for FlexBIT data assets

Category	Meaning
Public	Data that are already public, such as weather data collected from open or governmental sources, or data that can be made public without privacy, confidentiality, intellectual-property, contractual, or security restrictions.
Internal	Project data intended for use within the consortium under FlexBIT cooperation agreement.
Personal data risk possible	Data that may become personal data if linked to identifiable persons, households, behaviour, etc.
Third-party restricted	Data provided by third-party under contractual, licence, or NDA conditions.
Anonymised/aggregated	Data processed before sharing or publication to reduce personal-data, confidentiality, or commercial-sensitivity risks. Raw source data may remain internal or restricted.
Restricted	Partner-owned proprietary, NDA-covered, IP-sensitive, commercially sensitive, or otherwise access-limited data.

Table 3.2: Project-level accessibility categories for FlexBIT data assets

Category	Meaning
Public	Data are publicly accessible or approved for open publication through a repository or other public source, under the applicable licence.
Public where approved	Data are intended for public release, subject to preparation and approval by the responsible partner
Normalised public, raw consortium only	Normalised data are intended for public release, while raw data remain restricted to consortium use under the general cooperation agreement.
Consortium only	Data are accessible only to authorised project partners for agreed project tasks.
Restricted	Access is controlled by the responsible partner, data provider, third-party licence, NDA, confidentiality obligation, or other contractual restriction.

The technical inventory in Section 2.2.1 identifies the main data-producing assets, signal families, and local acquisition mechanisms at the demonstrator level. Not all data generated by these systems is necessarily transmitted to the FlexBIT platform. Some data remain under local demonstrator control and may be used only for local operation, validation, model development, benchmarking, or reporting. The governance table below therefore covers FlexBIT-relevant data assets at demonstrator level, including platform-ingested data, locally retained project data, and selected derived or aggregated outputs, according to partner approval and project needs. In this table, “responsible partner” refers to the partner responsible for collecting, curating, providing, or managing the data asset within the FlexBIT project. For datasets derived from public or third-party sources, this does not imply ownership of the original source data and applicable source licences, attribution requirements, and access conditions remain valid. The distinction between locally retained datasets and platform-ingested datasets is essential for preserving demonstrator autonomy, limiting unnecessary data exposure, and enabling proportional and risk-aware data-sharing practices across heterogeneous operational environments.

Table 3.3: Governance view of FlexBIT-relevant data assets

Data asset	Responsible partner	Purpose	Sensitivity	Storage	Accessibility
BESS operational data (charge level, capacity, status signals)	Electrum	Monitoring and optimisation of energy storage operation	Internal	FlexBIT platform (cloud)	Consortium only
Energy storage performance data (available capacity, total capacity)	Electrum	Performance analysis and efficiency optimisation	Internal	FlexBIT platform	Consortium only
Calculated energy indicators (e.g. TOTAL_OWN_NEEDS_P)	Electrum	Energy balance analysis and decision support	Internal	FlexBIT platform	Consortium only
SCADA signals from BMS (status, measurements)	Electrum	Real-time monitoring and system integration	Internal	FlexBIT platform	Consortium only
Aggregated and anonymised outputs from demonstration data	Electrum	Cross-demonstrator benchmarking and research	Anonymised/aggregated	Cloud repository	Public where approved
FlexBIT-relevant data from the Smart Energy Parking in Palermo	UniPA	Monitoring the performance of the demonstrator and applying control action	Internal	FlexBIT platform	Piublic
FlexBIT-relevant data from the Solar Cooling System in Pantelleria	UniPA	Monitoring the performance of the demonstrator and applying control action	Internal	FlexBIT platform	Public where approved
Environmental data (solar power prediction)	aRTE Möbel based on external public sources	Prediction of power generation	Public	Local data storage at site level, Fraunhofer IFF and FlexBIT platform	Normalized public, raw public where approved
Load consumption data	aRTE Möbel	Real-time monitoring, Prediction of power generation and consumption, activation of batteries, compressed air, and charging station, passive control of industrial loads	Internal	Local data storage at site level, Fraunhofer IFF and FlexBIT platform	Normalised public, raw consortium only

FlexBIT-relevant data from Aue Funeral Facility	Aue Funeral	Real-time monitoring, Prediction of power generation and consumption, activation of batteries, cooling aggregators, and charging stations	Internal	Local data storage at site level, Fraunhofer IFF and FlexBIT platform	Public where approved
Environmental data (solar power prediction)	Aue Funeral based on external public sources	Prediction of power generation	Public	Local data storage at site level, Fraunhofer IFF and FlexBIT platform	Public
Residential load, PV generation, and grid import/export data	Malta Demonstrator	Energy system modelling, forecasting, and digital twin simulation	Third-party restricted	UOM Data Storage cloud/local simulation environment	Restricted
Environmental data (irradiance, temperature, humidity, rainfall)	Malta Demonstrator/External sources	Support PV modelling and forecasting accuracy	Public	FlexBIT platform/external datasets	Public
FlexBIT-relevant data from Tor Vergata Demonstrator	UniTOV	Real-time monitoring, control action suggestion.	Internal	FlexBIT platform, local storage	Consortium only
Test data from WUST Laboratory Demonstrator (static and dynamic response of regulation)	WUST	Characteristic of regulatory potential of microgrid devices, responses of the devices on active and reactive power control	Internal	FlexBIT platform/external datasets	Public where approved

FAIR-oriented governance practices within FlexBIT support the controlled discoverability, accessibility, interoperability and long-term reuse of project relevant datasets across diverse demonstrator environments. FlexBIT applies the FAIR principles according to the access level, ownership constraints, privacy risk, and publication status of each data asset listed in Table 3.3. For datasets chosen for publication into open access, FAIR principles are applied as follows:

1) Findable: all data assets selected for public release will be accompanied with relevant metadata, to provide necessary context and integrity for long-term use. Collected metadata must include, but not limited to: title, description, provider/creator, contact point, date, version, provenance, format, access conditions, license, persistent identifier, and relevant methodological context. Metadata specifications are intended to support discoverability, provenance tracking, dataset versioning, interoperability, and future integration with external research infrastructures and emerging European energy-data ecosystems.
2) Accessible: open data repositories will be used to share data with public access. The main repositories to be used are Open Science Framework and Zenodo. To support future ease of locating relevant data, persistent identifiers will be assigned to publicly shared data assets. The FlexBIT project communities are present in Zenodo [18] and OSF [19].
3) Interoperable: to facilitate seamless exchange of data, metadata will be used to consistently describe the structure of data exchanged. To ensure interoperability of shared and published datasets across software and devices, common non-proprietary data formats will be used, such as CSV, RDF, XML, JSON, where appropriate. This reduces dependence on proprietary software and supports reuse by external users. Semantic interoperability is further supported through alignment with the standard FlexBIT data model, shared metadata conventions, and recognised ontological references described in Section 7. This reduces fragmentation between demonstrators and promotes future integration with emerging European energy-data spaces and cross-domain digital infrastructures.
4) Reusable: to promote data reuse upon project completion, publicly available data assets will be licensed under permissive Creative Commons licences, giving preference to CC-BY 4.0 [27] where possible and approved by the dataset-submitting partner. Reuse conditions additionally depend on demonstrator-specific confidentiality constraints, cybersecurity considerations, third-party licensing conditions, and the sensitivity and accessibility categories assigned to each dataset.

General every day data, draft reports, reports, presentations among others will be stored and shared in the private Collaborative Platform with restricted access (username + password) to authorised users. As an initial step, only the Consortium Partners will have access to the cloud storage where dataset and metadata are filed. Sensitive data will be protected against unwanted disclosure. Sensitive data may include: personal, confidential, a combination of different datasets, biological, personal and sensitive metadata. In addition, the General Data Protection Regulation will apply where research data that contains personal data with which a living person can, directly or indirectly, be identified. This concerns both textual data and image and sound data. For personal data, fully informed consent will be requested for collecting, processing and storing data.

Risks and mitigation: The main risks relevant to data governance in FlexBIT are limited interoperability between heterogeneous infrastructures, privacy risks when handling sensitive multi-stakeholder data, and compliance challenges arising from evolving legal frameworks. These risks are mitigated through agreed interfaces and early integration testing, controlled data handling and encryption measures, and periodic legal and procedural review. The privacy analysis for the data present in the project is done in Section 4, while precise data interoperability and exchange within the platform is described in Section 7.

The governance framework is intentionally designed to support federated data stewardship, diverse operational environments and progressive interoperability maturity, while remaining adaptable toward future interoperable and cross-domain European energy-data ecosystems.

Data retention is determined by the responsible partner or platform component owner, subject to project needs, confidentiality obligations, legal requirements, and GDPR storage-limitation principles where applicable. Data are kept only as long as needed for the relevant project, validation, reporting, security, legal, or reproducibility purposes. Data in certain cases may be subject to embargo periods for up to 5 years before any public release is considered.

In distributed digital-energy platforms such as FlexBIT, data governance cannot be treated independently from cybersecurity governance. The increasing convergence between operational technology (OT), IoT infrastructures, cloud-based analytics, and cross-platform interoperability introduces cross-domain governance dependencies that extend beyond classical personal-data protection frameworks. For this reason, the FlexBIT governance approach additionally incorporates principles aligned with Articles 21 and 23 of the NIS2 Directive (EU) 2022/2555 concerning cybersecurity risk-management measures, supply-chain security, access-control governance, and operational resilience obligations. Similarly, where demonstrator data are used to develop or operate AI-based forecasting, optimisation, or control-support components, the applicability of EU AI Act [23] should also be assessed in relation to the intended use, particularly with regard to data quality, documentation, transparency, traceability, and human oversight.

A further challenge across European research and energy ecosystems concerns the uneven operational implementation of FAIR principles, particularly regarding provenance preservation, semantic interoperability, reusable metadata structures, and machine-readable governance information. In practice, datasets may formally satisfy open-access requirements while remaining insufficiently interoperable or reusable across heterogeneous digital-energy environments.

The FlexBIT governance framework addresses these limitations through harmonised metadata conventions, provenance-oriented dataset management, approved semantic representations, and governance-aware semantic interoperability mechanisms capable of supporting controlled reuse across heterogeneous demonstrator environments.

Consequently, the FlexBIT governance framework is intentionally designed as an adaptive and semantically interoperable federated architecture capable of accommodating heterogeneous demonstrator environments, evolving European energy-data regulations, differentiated national governance practices, and future cross-domain digital-energy ecosystems without requiring full centralisation of operational datasets or governance authority.

4 Privacy

The project FlexBIT handles data from heterogeneous sources across multiple European countries and partner institutions. The privacy framework addresses the applicable regulatory obligations, the categories of data processed within the FlexBIT platform, and the specific arrangements governing the two main data collection streams:

1. technical operational data shared among partners,
2. survey data collected for the Social Acceptance study (WP 4).

4.1 Regulatory framework

Regulation (EU) 2016/679 (GDPR) [1] is the primary regulatory instrument governing personal data processing in FlexBIT. It establishes a harmonised legal framework applicable across all EU member states. The GDPR defines personal data as any information relating to an identified or identifiable natural person and sets out the key principles guiding lawful processing: lawfulness, fairness and transparency, purpose limitation, data minimisation, storage limitation, and integrity and confidentiality.

In addition to the GDPR, the FlexBIT platform operates within a broader EU digital data governance landscape. The Data Governance Act (EU) 2022/868 applies to the sharing of demonstrator datasets for research and innovation purposes. The Data Act (EU) 2023/2854 is relevant to data generated by connected devices, including IoT sensors and smart meters deployed at demonstrator sites. The NIS2 Directive establishes the cybersecurity framework applicable to the FlexBIT digital infrastructure and is addressed in Section 5.

Within the FlexBIT project, the regulatory requirements identified in deliverable D1.2 are implemented through the technical and organisational design of the platform architecture. Specifically, the FlexBIT framework applies privacy-by-design and data-minimisation principles across the full data lifecycle, from local acquisition and pre-processing at demonstrator level to secure transmission, storage and controlled access within the central platform infrastructure.

The platform architecture is designed to preserve demonstrator autonomy while limiting the exchange of unnecessary or directly identifiable information. Although only normalised and function-specific datasets are exposed to the central FlexBIT services through authenticated interfaces and standardised adapters, raw operational data remain under the control of demonstrator partners. Access to shared data is restricted through role-based access control mechanisms, while encrypted communication channels and controlled storage environments support the confidentiality and integrity requirements established under GDPR [1] and related EU digital legislation.

In alignment with the Data Act [7] and the Data Governance Act [5], the FlexBIT platform also promotes controlled and interoperable data sharing between diverse systems while preserving security, transparency and data ownership responsibilities across consortium partners.

Given the multinational composition of the consortium, national implementing legislation applies in parallel with EU-level instruments.

Table 4.1 summarizes the core EU regulations and their national counterparts across the jurisdictions covered by the FlexBIT demonstrators.

Table 4.1: Relevant EU and national legislation applicable to FlexBIT project

EU Regulation/Directive	Objective	Relevance to FlexBIT	Relevant National Legislation
GDPR (EU) 2016/679 [1]	Protection of personal data	Processing of energy consumption data, EV charging data and stakeholder datasets	**Germany**: BDSG [2] **Italy**: Legislative Decree 196/2003 [3] **Malta**: Data Protection Act [4] **Poland**: Personal Data Protection Act [5]
Data Governance Act (EU) 2022/868 [5]	Secure data sharing across sectors	Sharing of demonstrator datasets for research and innovation	Implemented through national authorities
Data Act (EU) 2023/2854 [7]	Regulation of data generated by connected devices	Relevant for IoT devices, smart meters and digital energy infrastructure	National digital governance frameworks
NIS2 Directive (EU) 2022/2555 [8]	Cybersecurity framework	Security of FlexBIT digital infrastructure and data flows	**Germany**: IT Security Act 2.0 [9] **Italy**: Legislative Decree 65/2018 [10] **Malta**: CAP 460 [11] **Poland**: National Cybersecurity System Act [12]

Additional legislation of relevance is:

1. Open Data Directive [20], that allows for reuse of public sector data;
2. INSPIRE Directive [21];
3. The Digital Operational Resilience Act (DORA) [22];
4. The Artificial Intelligence Act [23];
5. ePrivacy Directive [24];
6. Electricity Directive [25].

4.2 Privacy-by-design and data minimisation

The FlexBIT platform adopts a privacy-by-design approach in accordance with Article 25 of the GDPR [1]. Privacy and data-protection considerations are integrated into the system architecture from the earliest stages of platform design, rather than being treated as external compliance requirements. This approach is particularly important due to the distributed and heterogeneous nature of the FlexBIT demonstrators, which involve multiple countries, infrastructures, and operational environments.

Data minimisation principles are applied throughout the platform by limiting data exchange to the measurements and metadata strictly required for platform operation, interoperability, forecasting, flexibility management, and research activities. Local demonstrator systems retain control over raw operational datasets, while only filtered, aggregated, or normalised information is transmitted through the FlexBIT integration layer. High-frequency field-level signals and directly identifiable information are avoided wherever technically and operationally feasible.

Technical safeguards supporting this approach include authenticated API access, encrypted communication channels (TLS/HTTPS), role-based access control (RBAC), controlled storage environments and logical separation between local demonstrator infrastructures and central platform services. In addition, the blockchain layer implemented within the FlexBIT architecture stores only cryptographic

certifications and integrity proofs, rather than raw operational measurements. This approach reduces unnecessary exposure of sensitive information while preserving traceability and data integrity.

Figure 4.1 shows the privacy-protecting data flow implemented within the FlexBIT platform architecture.

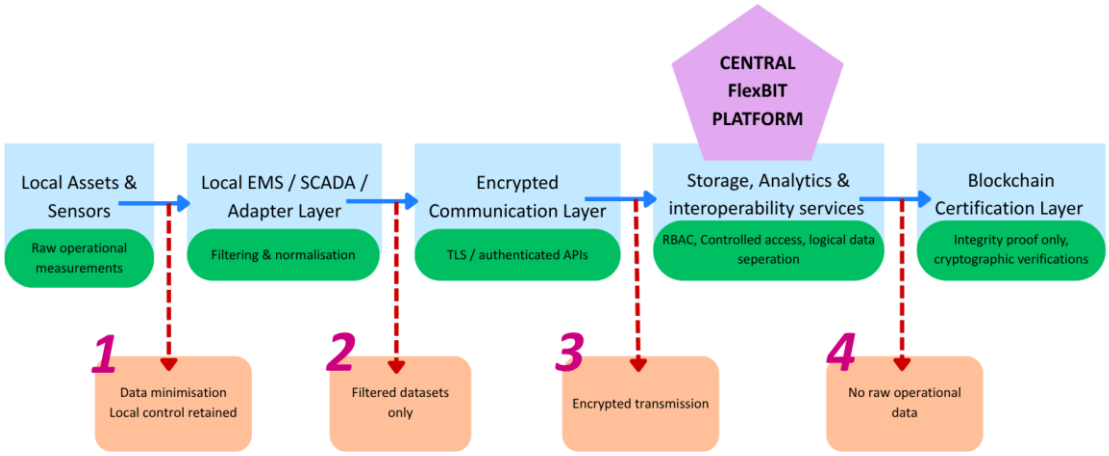

Figure 4.1: Privacy-preserving operational data flow and privacy controls implemented within the FlexBIT platform architecture. The blue-green boxes indicate the operational flow whereas the orange ones the privacy control

4.3 Data categories and personal data risk assessment

The FlexBIT platform integrates data from a range of sources, including smart meters, IoT sensors, SCADA systems, electric vehicles (EV) charging infrastructure, and third-party aggregators. While the primary purpose of these datasets is operational (supporting energy forecasting, optimisation, and flexibility exploitation), certain categories of data may be associated with identified or identifiable natural persons and therefore qualify as personal data under Article 4 (1) GDPR. In particular, electricity consumption patterns derived from smart meter data may reveal information about household behaviour and occupancy. EV charging session logs, including timestamps and location data, may reveal individual mobility patterns. IoT sensor data carries a lower risk in isolation but may become personal data when linked to individuals through other identifiers. Forecasting and aggregated datasets present minimal risk when properly anonymised prior to use. Although most operational datasets processed within the FlexBIT platform are technical in nature, certain categories of measurements may become indirectly linked to distinct individuals or behavioural traits, particularly in residential or small-scale community environments. As a result, the project evaluates privacy risks based on explicit personal identifiers and the potential for indirect identification, behavioural inference, occupancy detection and user profiling from energy consumption patterns and device interactions.

Table 4.2 summarizes the main data categories and the corresponding regulatory considerations.

Table 4.2: Data categories and applicable legal framework

Data Category	Description	Personal Data Risk	Applicable Regulation	Mitigation Measures
Smart meter data	Energy consumption measurements	May reveal household behaviour	GDPR; Data Act, ePrivacy Directive, Electricity Directive	Aggregation, data minimisation, controlled access, restricted sharing of high-frequency measurements
EV charging data	Charging sessions and timestamps	May reveal mobility behaviour	GDPR; Data Act, ePrivacy Directive, AI Act if AI optimisation used based on a personal profile	Tokenisation, controlled access, avoidance of direct user identifiers
IoT sensor data	Operational data from energy assets	Low risk unless linked to individuals	GDPR; NIS2, AI Act [23] if an autonomous decision system is used	Purpose limitation, filtered transmission, separation from recognisable datasets
Forecasting datasets	Aggregated datasets for energy prediction	Low risk when anonymised	GDPR; Data Governance Act, Open Data Directive, Data Act, AI Act	Data consolidation, identity masking, use of anonymous training datasets
Survey and questionnaire data	Stakeholder responses, social acceptance datasets	Direct or indirect participant identification	GDPR, Data Governance Act	Informed consent, anonymisation, aggregation before publication
Platform logs, authentication records	System access logs, operational monitoring metadata	User traceability, activity monitoring	GDPR, NIS2	RBAC, retention policies, access logging, restricted administration rights

4.4 Technical and operational data

The technical operational data shared among FlexBIT partners from platform operation include:

- load profiles;
- generation profiles;
- consumption measurements;
- SCADA output;
- related metadata;
- laboratory test results.

Within the FlexBIT architecture, technical and operational datasets are primarily processed for platform operation, interoperability, forecasting, flexibility management, and demonstrator validation activities. Although these datasets are generally technical in nature, the project recognises that certain operational measurements may indirectly relate to identifiable individuals or behavioural patterns, particularly within

residential or community-scale demonstrators. This applies especially to smart-meter measurements, EV charging records, and timestamped operational datasets that could enable indirect occupancy inference or activity profiling when combined with other contextual information.

These data are in the majority of cases purely of a technical nature and do not directly identify natural persons. However, certain data streams (from smart meters and EV charging data) may constitute personal data in contexts where they can be attributed to an identifiable individual.

The governance of partner-shared data within the project is primarily regulated by the FlexBIT Cooperation Agreement (CA), which has been signed by all consortium partners, while demonstrator partners retain ownership and operational control over locally generated raw datasets. Where technical data streams carry a residual personal data risk, the FlexBIT platform adopts the following protective measures in line with the privacy-by-design and privacy-by-default principles required under Article 25 GDPR:

- pseudonymisation or anonymisation of datasets before integration into analytical and forecasting workflows;
- application of the data minimisation principle, limiting data collection to what is strictly necessary for operational purposes;
- role-based access control restricting data access to authorised project partners according to their project tasks;
- encrypted data transmission between demonstrator sites and the central platform;
- secure storage procedures ensuring that sensitive information remains protected at rest.

No separate data processing agreement is required for the handling of purely technical operational data within the project context, given that the CA provides the applicable governance framework. Where personal data processing cannot be excluded, the responsible partner acts as data controller and must ensure compliance with GDPR obligations applicable in their jurisdiction.

4.5 Survey data and social acceptance studies

The FlexBIT project includes Social Acceptance studies (WP4) involving questionnaires administered to stakeholders and members of the public, as committed to the project proposal. The survey results may be provided in open-data formats where appropriate, designed to facilitate broad stakeholder engagement while preserving anonymity and preventing respondent re-identification. These surveys are conducted within WP4 activities under the coordination of CERTH, which supports the design, administration, and governance-related management of the survey activities and associated data-handling procedures within the project framework.

The questionnaires are designed to collect responses in fully anonymous form from the point of collection. No directly or indirectly identifying information, such as names, contact details, IP addresses, or unique identifiers, is collected at any stage and no profiling or automated individual-level assessment is performed on survey participants. Sociodemographic and socioeconomic data, e.g. age range, gender, professional sector, are collected at a categorical or aggregated level that does not allow the re-identification of individual respondents. Based on the current survey design and data-collection methodology, the questionnaires are intended to operate on an anonymous basis, without collecting directly identifiable personal information. The collected responses are structured to minimise re-identification risks and are analysed only in aggregated form for research and dissemination purposes. Consequently, the activity is expected to fall outside the scope of identifiable personal-data processing under Article 4(1) GDPR.

The aggregated and anonymised survey results will be made publicly available through Open Science Frameworks (OSF) and Zenodo, in accordance with the project's open science commitments. Published datasets will be accompanied by appropriate metadata and licensed under a Creative Commons license (CC BY or CC0) where compatible with the nature and sensitivity of the published datasets, following the principles "as open as possible, as closed as necessary". Where necessary, additional aggregation or disclosure-control measures may be applied prior to publication in order to preserve anonymity and ensure compliance with applicable ethical and regulatory requirements.

5 Security

5.1 Goals, context, boundaries

The security framework of the FlexBIT platform is intended to protect data, services, and interactions across a distributed multi-partner environment. Its purpose is to provide a coherent security model for the exchange of operational data between demonstrators, the centralised platform, and the supporting digital infrastructure.

The main security goals are:

- confidentiality, to protect sensitive operational and partner data from unauthorised access;
- integrity, to ensure that exchanged and stored data are not modified in an unauthorised way;
- availability, to support reliable operation of platform services and interfaces;
- traceability, to preserve evidence of relevant operations and data handling;
- controlled access, to ensure that users, partners, and software components can only access authorised functions and resources.

From a system-boundary perspective, FlexBIT includes the local demonstrator environments, the adapter and communication layer, the central platform services, the data-processing and storage components, and the blockchain attestation layer (Figure 5.1). Within these boundaries, local sites retain operational control of physical assets, while the centralised platform manages shared services, data exchange, and supervisory coordination.

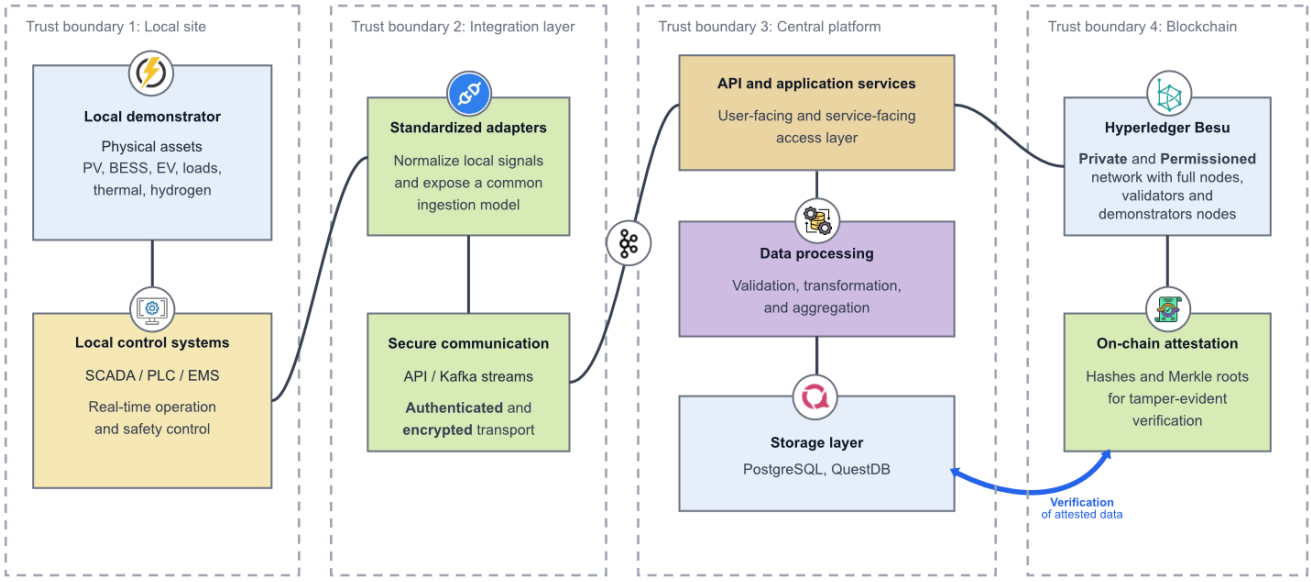

Figure 5.1: Trust boundaries diagram

The blockchain layer adds a distributed trust component by supporting [14][15]:

- data attestation, through on-chain registration of cryptographic proofs;
- transparency and immutability, through tamper-evident anchoring of selected platform data;
- privacy preservation, since detailed operational datasets remain off-chain and only their cryptographic fingerprints are registered on-chain [17].

Overall, FlexBIT combines centralised service management with distributed integrity assurance, providing a security model suited to heterogeneous demonstrators and consortium-level data exchange.

5.2 Access control and authentication

Access control in the FlexBIT platform combines user authentication, role-based administration, organisation-based separation, and API-key-based machine access. This provides one common security model for both interactive use of the platform and automated data exchange from demonstrator integrations. The underlying authentication framework supports email/password login, administrative roles, organisations, session handling, and API keys.

In the current implementation, access is structured as follows in Table 5.1.

Table 5.1: Access layers

Layer	Current approach	Main purpose
User authentication	Email and password with session-based login	Interactive access to the platform
Roles	User and admin roles	Administrative control over users, invitations, and integration credentials
Organisations	Users grouped inside the same partner/demonstrator organisation	Shared access to the same data and configuration space
API keys	Platform-generated credentials for programmatic access	Adapters and modules sending metrics or interacting automatically
Streaming auth	One Redpanda user provisioned per API key	Alignment between API credentials and Kafka authentication

The organisation model is particularly important for demonstrator operation. It allows multiple people belonging to the same partner or demonstrator to work inside the same configuration and data scope, instead of duplicating access per user. Invitation-based onboarding is used to add members to the same organisation, while administrators retain control over organisation-level settings and credentials.

For machine-to-machine integration, the platform generates API keys that are used by adapters, gateways, and other backend components. In the current model, each API key is associated with a partner account context and can be used to push metrics for any site or device registered under that partner. The specific source site or asset is then identified in the submitted payload through metadata. This means that authorisation is currently enforced at partner scope, not yet at single-device scope.

The same credential lifecycle is also used for streaming integration. Every time an API key is created on the platform, a corresponding Redpanda user is provisioned so that the API key identifier and secret can be reused in the Kafka authentication flow. This creates a one-to-one correspondence between platform API credentials and streaming credentials, simplifying demonstrator integration and keeping HTTP/API and Kafka access aligned under the same administrative process.

Enforcement and tracking are handled centrally through the authentication layer and related backend records.

5.3 Cryptography

Cryptographic protection in FlexBIT covers both secure communication and data integrity assurance. The first aspect protects data while they are exchanged between adapters, platform services, streaming components, and blockchain infrastructure. The second aspect provides verifiable proof that attested data have not been altered after registration.

For communication security, transport encryption is applied to the main integration layers of the platform [16]. In the streaming layer, Redpanda supports TLS protection for the Kafka API and for internal service interfaces, including the internal RPC server used by the broker. In all environments (testing, staging and production) these interfaces will be protected with TLS in addition to authenticated access.

At blockchain layer, FlexBIT relies on a permissioned Hyperledger Besu network organised around nodes, validators, and partner nodes. TLS is supported for Besu client-facing JSON-RPC interfaces, and the consortium blockchain deployment should protect node and validator communications within this permissioned environment through encrypted transport and controlled peer connectivity. This is consistent with the network role separation already defined for the FlexBIT blockchain infrastructure (D1.1).

The cryptographic design therefore addresses three complementary objectives:

- protection of API and stream communication channels;
- protection of blockchain service and node-level communications;
- integrity verification of platform data through hashing and attestation.

For integrity assurance, FlexBIT uses a private Hyperledger Besu network with QBFT consensus and an attestation flow based on hashing and Merkle-tree anchoring. Daily measurements are hashed and their cryptographic fingerprints are recorded on-chain. This allows partners to verify that the corresponding off-chain data have not been modified after attestation, while avoiding the overhead of storing raw operational datasets directly on the blockchain.

In practical terms, this means that the platform combines:

- TLS-protected transport, for secure exchange of operational data and service traffic;
- authenticated and authorised messaging, for controlled access to Kafka-based integrations;
- hash-based attestation, for tamper-evident verification of platform measurements and records.

This combined approach ensures that FlexBIT does not rely only on access control, but also on cryptographic mechanisms that protect data in transit and support long-term immutability and auditability of attested information.

5.4 Monitoring

At the current stage of the project, a complete monitoring stack is not yet fully deployed in operational form. Its implementation is planned as part of the WP3 platform maturation activities, together with the broader transition from the current testing environment toward a more structured pre-production and demonstrator-integration setup. For this reason, the present section describes the target monitoring approach and the main monitoring capabilities that are planned for the FlexBIT platform.

The monitoring framework is intended to provide visibility over both the digital infrastructure and the application and blockchain services composing the FlexBIT environment. In particular, the objective is to

support operational supervision, fault detection, and auditability across the different layers of the platform.

The planned monitoring scope includes at least the following categories:

- machine and host metrics, such as CPU, memory, disk, and network usage;
- Kubernetes metrics, covering pod status, resource consumption, service health, and workload lifecycle;
- application and platform metrics, related to the behaviour of platform services and data-processing components;
- blockchain metrics, in particular metrics exposed by Hyperledger Besu nodes and validators;
- logs and events, collected across the main software modules and infrastructure components.

The intended implementation is based on a Prometheus + Grafana stack. In this architecture, Prometheus will act as the main metrics collection component, scraping exposed endpoints from infrastructure and application services, while Grafana will provide dashboards and operational views for visualisation and analysis. This setup is particularly suitable for FlexBIT because it allows the platform to combine infrastructure-level supervision with service-specific and blockchain-specific observability in a single environment.

A relevant part of the planned setup is the monitoring of the Hyperledger Besu layer. Since Besu exposes Prometheus-compatible metrics, blockchain nodes and validators can be integrated directly into the common monitoring framework. This will support observation of node health, peer connectivity, synchronisation status, and block progression.

In parallel, the platform is expected to integrate a centralised logging system across the main modules and infrastructure elements. Its role is to collect and retain logs from platform services, adapters, processing components, and blockchain-related elements, supporting troubleshooting and operational traceability. In this sense:

- metrics support continuous supervision of system state and performance;
- logs support event traceability and root-cause analysis.

5.5 Blockchain

The FlexBIT blockchain layer provides an additional security mechanism for data integrity, traceability, and consortium-level trust. It is not intended to store operational datasets directly on-chain. Instead, detailed telemetry and platform records remain in off-chain storage, while only cryptographic proofs of selected data are registered on the permissioned blockchain. This approach preserves confidentiality and scalability while enabling later verification of data integrity. The current FlexBIT design is based on a private Hyperledger Besu network with QBFT consensus, where nodes, validators, and partner nodes support controlled participation and distributed verification across the consortium.

The core security mechanism is hash-based attestation. Operational data are collected through the platform ingestion pipeline, processed and aggregated off-chain, and then converted into cryptographic fingerprints. This allows a partner to verify that a specific off-chain dataset, or a single metric included in it, has not been modified after attestation, without exposing the full dataset on the blockchain.

Smart contracts are used as trusted registries for attested data fingerprints and related metadata. The information stored on-chain should be limited to what is necessary for verification, such as the Merkle root, dataset identifier, time window, source or organisation identifier, version, and timestamp of

attestation. Raw measurements, commercially sensitive operational data should remain off-chain. This is consistent with the FlexBIT security principle that blockchain provides tamper-evident proof of integrity rather than a general-purpose data storage layer. More detailed information on the smart-contract layer, its role in platform verification, and its connection with energy-sharing and energy-exchange mechanisms is provided in deliverable D2.1.

Communication between partner nodes in the permissioned distributed ledger is protected through TLS, ensuring authenticated and encrypted peer-to-peer channels for validator and participant-node interaction. This reduces the risk of interception, manipulation, or unauthorised participation in the blockchain network and is aligned with the broader FlexBIT security approach based on encrypted transport, authenticated access, and controlled peer connectivity.

As the platform evolves, additional privacy-preserving access-control techniques may be evaluated for cases where multiple partners need to compute, verify, or access information without exposing unnecessary underlying data. Two relevant families are Secure Multi-Party Computation (SMC) and Attribute-Based Encryption (ABE). SMC may be considered where several demonstrators or partners need to compute aggregate flexibility, settlement, benchmarking, or performance indicators while keeping their local inputs confidential. Candidate criteria for SMC adoption include the number of participating parties, sensitivity of the input data, acceptable computation latency, required precision, auditability of the result, and compatibility with the platform’s existing off-chain processing pipeline.

6 Digital Twin Technologies for Real-Time Monitoring and Simulation

6.1 Digital twin concept in FlexBIT

The Malta Digital Twin is being developed as a flexibility-oriented residential energy management and validation environment within the FlexBIT project. Its purpose is to provide a predictive simulation framework for analysing how residential prosumers can contribute to flexibility services through coordinated operation of photovoltaic generation, load demand, battery storage, and grid import/export behaviour.

At the current stage, 15-minute load and PV forecasting models have been implemented and validated using residential time-series datasets. In parallel, additional forecasting models operating at 1-minute, 5-minute, and 30-minute temporal resolutions are currently under development to analyse how data granularity affects forecasting accuracy, operational stability, and flexibility-oriented control performance. Evaluating multiple time resolutions is important because different operational scenarios require different response times and control behaviours. Higher-resolution forecasting supports more dynamic and fast-response energy management scenarios, while lower-resolution forecasting enables longer-horizon operational optimisation and stability analysis.

The digital twin is currently implemented as a MATLAB/Simulink-based virtual environment using residential datasets that include PV generation, load demand, grid import/export power, and storage modelling. The Simulink model is operational and is being continuously refined to support realistic energy balancing, supervisory control development, and future integration of forecasting-driven energy management strategies.

In parallel, the system components are being redesigned and adapted for deployment on the Speedgoat Hardware-in-the-Loop (HIL) platform. This process requires synchronisation between the Simulink model architecture and the Speedgoat execution environment to ensure compatibility for real-time operation. As a result, each subsystem, including battery control, PV modelling, grid interaction, and supervisory EMS logic, is being modelled in a modular and hardware-compatible manner.

The modelling approach is important because it provides a controlled environment for testing residential energy management and flexibility-oriented control strategies before real-world deployment. This enables repeatable scenario analysis, validation of supervisory control approaches, and safer transition toward future real-time HIL-based implementation within the FlexBIT framework.

6.2 Data inputs and system representation

The Malta Digital Twin operates using multivariate residential time-series datasets representing both electrical and environmental behaviour of the residential energy system. The current implementation utilises historical residential measurements collected at multiple temporal resolutions, with the validated operational baseline currently focused on 15-minute data. Additional datasets at 1-minute, 5-minute, and 30-minute resolutions are under evaluation to support future flexibility-oriented scenario analysis and forecasting studies.

The primary electrical datasets include:

- residential load demand;
- PV generation;
- grid import/export power;

- battery operational variables;
- derived energy-balancing indicators.

In addition to electrical measurements, environmental datasets are used during forecasting-model development and are planned for integration into future real-time digital twin operation.

These include:

- solar irradiance;
- ambient temperature;
- humidity;
- rainfall measurements.

The data engineering pipeline includes timestamp alignment, missing-data handling, normalisation, rolling statistical computation, and temporal feature extraction. Time-derived features such as hour-of-day, day-of-week, and cyclical temporal encodings are generated to capture repetitive daily and seasonal energy patterns that influence residential consumption and PV production behaviour.

For forecasting-oriented operation, lag-based autoregressive features and rolling-window statistics are extracted from historical PV and load signals. These engineered features are used by machine learning and deep learning forecasting models integrated into the digital twin environment. The forecasting framework is being prepared for integration with the supervisory EMS layer to support future predictive control evaluation.

Within the Simulink environment, all datasets are synchronised into a unified timestep representation to ensure consistent interaction between the forecasting layer, battery dispatch logic, PV generation model, and grid exchange model. This synchronised representation enables realistic simulation of residential energy balancing under varying operating and environmental conditions.

6.3 Simulink-based system modelling

The Malta Digital Twin is implemented in MATLAB/Simulink as a modular residential energy system simulation environment. The modelling idea is based on representing the interaction between distributed energy resources, forecasting algorithms, and supervisory control logic within a unified timestep-based framework. The objective is to reproduce realistic residential energy behaviour while enabling evaluation of flexibility-oriented control strategies under varying operating conditions.

The simulation operates using discrete timestep execution synchronised with the temporal resolution of the forecasting datasets. The current validated implementation is based on 15-minute operation, while additional high-resolution scenarios (1-minute and 5-minute) are under development for fast-response flexibility analysis and real-time control evaluation. At each timestep, the system updates PV generation, load demand, battery state, and grid interaction variables.

The Simulink architecture is organised hierarchically into multiple interacting layers:

- data and forecasting layer;
- supervisory EMS layer;
- component-level energy system models;
- grid interaction and balancing layer.

The forecasting layer is under development to provide predicted PV generation and load demand values. These forecasts will be transferred to the supervisory EMS, which will evaluate system conditions and determine operational control actions. Within the proposed architecture, the EMS acts as the central decision-making layer responsible for coordinating battery dispatch, power balancing, and grid interaction.

An energy balancing strategy is currently under development to prioritize local energy utilisation before external grid exchange. The planned supervisory EMS operation will first allocate available PV generation to residential demand and then evaluate battery charging availability based on state-of-charge (SoC) limits and operational constraints. Under excess generation conditions, additional energy will be directed toward battery charging and subsequently toward grid export if storage limits are reached. During insufficient PV production, the future EMS strategy is intended to prioritize battery discharge within predefined operational windows before importing additional energy from the grid.

The BESS model incorporates operational constraints including:

- minimum and maximum SoC limits;
- charging and discharging power limits;
- round-trip efficiency;
- charge/discharge transition logic.

Battery dispatch decisions are currently governed by real-time simulated system conditions, while integration of forecasting-driven predictive dispatch is under development. The supervisory control framework is being designed to maintain predefined SoC operational windows in order to protect battery lifetime and preserve reserve availability during peak-demand or low-generation periods.

Grid import/export arbitration is also planned as part of the future supervisory balancing logic. The intended control strategy will evaluate the power mismatch between generation and demand and determine whether energy should be imported from or exported to the grid.

This will support future assessment of:

- self-consumption optimisation;
- peak shaving operation;
- reduction of grid dependency;
- flexibility-oriented energy exchange behaviour.

The modular Simulink implementation is additionally being prepared for future integration with Hardware-in-the-Loop operation on the Speedgoat platform. For this reason, subsystem models are being designed using synchronised and hardware-compatible structures to support consistency between offline simulation and future real-time execution environments.

6.4 Integration of forecasting algorithms

The Malta Digital Twin integrates multiple forecasting models to support predictive energy management and to evaluate the suitability of different forecasting approaches under varying residential operating conditions. The forecasting framework currently includes deep learning and machine learning approaches, specifically Long Short-Term Memory (LSTM) networks and gradient boosting methods such as LightGBM, CatBoost, and hybrid residual-learning architectures.

The use of multiple forecasting models is methodologically important because residential energy behaviour exhibits both temporal dependencies and nonlinear feature interactions. PV generation and residential load profiles contain sequential patterns linked to daily and seasonal cycles, while also being influenced by rapidly changing environmental and operational conditions. Different forecasting models therefore capture different characteristics of the system behaviour.

LSTM models are used because they are effective in learning sequential temporal dependencies from historical time-series data, making them suitable for capturing repetitive consumption and PV generation patterns over time. In contrast, gradient boosting approaches such as LightGBM and CatBoost operate on engineered feature spaces and are more effective in modelling nonlinear relationships, short-term fluctuations, and complex interactions between temporal and environmental variables.

For residential load forecasting, the hybrid framework combines LightGBM and XGBoost using an ensemble-learning strategy. Both models are trained independently using the same engineered temporal, lag-based, and statistical features. The final load prediction is obtained through a weighted combination of the individual model outputs. This approach is motivated by the complementary strengths of the two gradient boosting methods: LightGBM provides efficient learning over large feature spaces, while XGBoost offers strong regularisation and improved handling of complex nonlinear interactions. The ensemble structure therefore improves generalisation performance and reduces model-specific forecasting errors under varying residential demand conditions.

For PV forecasting, a different hybrid strategy is employed based on residual learning. In this framework, the LSTM model first captures the sequential temporal dynamics of PV generation using historical time-series observations. The residual error between the actual PV value and the LSTM prediction is then modelled using LightGBM with engineered temporal and lag-based features. The final prediction combines the LSTM output with the residual correction generated by LightGBM. This two-stage architecture is particularly suitable for PV forecasting because solar generation exhibits both strong temporal patterns and highly nonlinear weather-dependent fluctuations.

The forecasting models are currently validated independently from the real-time EMS environment. Integration of forecasting outputs into the supervisory EMS layer is under development and is intended to support future predictive energy balancing and battery dispatch evaluation within the digital twin framework. This future integration will enable evaluation of forecast-driven operational strategies, including predictive charging/discharging control, peak-load reduction, self-consumption optimisation, and grid import/export management.

The use of multiple forecasting approaches also supports comparative analysis across different temporal resolutions and operational scenarios. This is important within the FlexBIT framework, where different flexibility services may require different forecasting characteristics depending on response speed, optimisation horizon, and operational objectives.

6.5 Energy management and control logic

The Malta Digital Twin is being developed to support supervisory Energy Management System (EMS) evaluation for flexibility-oriented residential energy management scenarios. The current Simulink environment enables standalone simulation of residential PV generation, load demand, battery storage, and grid interaction behaviour under different operating conditions.

At the present stage, forecasting models and the Simulink-based energy system are being evaluated independently. Integration of forecasting outputs into the supervisory EMS framework is currently under development to support future predictive energy management and flexibility-oriented control strategies.

Two operational control approaches are being considered within the digital twin framework:

- rule-based EMS operation;
- predictive forecasting-driven EMS operation.

The current implementation primarily focuses on rule-based operational behaviour using predefined control conditions and threshold logic. In this approach, battery charging is prioritised during periods of excess PV generation, while battery discharge is activated during periods of high residential demand or reduced PV production. Grid import reduction strategies are also evaluated whenever local generation or stored energy is available. This provides a baseline operational strategy for comparison with future predictive EMS approaches.

The predictive EMS framework is currently under development and is intended to incorporate forecasting outputs from the load and PV forecasting models. The objective is to enable future forecast-driven battery dispatch, energy balancing, and flexibility-oriented operational optimisation within the digital twin environment. This will enable forward-looking battery dispatch and improved energy balancing decisions.

The optimisation horizon depends on the operational scenario and timestep resolution currently under evaluation. The validated 15-minute forecasting implementation currently supports standalone short-term forecasting analysis, while higher-resolution forecasting scenarios (1-minute and 5-minute) are being investigated for future fast-response flexibility services and dynamic control applications.

The EMS dispatch framework under development follows a prioritised operational hierarchy:

1. satisfy residential load demand using locally generated PV power;
2. utilize battery charging during periods of excess PV generation;
3. discharge battery storage during peak-demand or low-generation periods;
4. minimize unnecessary grid import/export exchanges;
5. maintain battery operation within safe SoC and power limits.

The battery dispatch strategy incorporates operational constraints including minimum and maximum SoC windows, charging/discharging power limits, and energy balancing requirements. Future EMS integration is intended to evaluate both forecasted and real-time system conditions in order to determine charging, discharging, import, and export priorities at each timestep.

Grid import/export arbitration is also being considered as part of the future supervisory control logic. Under the intended operational strategy, excess generation conditions will be evaluated to determine whether energy should be stored locally or exported to the grid depending on battery availability and operational objectives. During supply deficits, battery discharge strategies will be evaluated prior to additional grid import. This framework is intended to support future assessment of residential flexibility behaviour under varying operational scenarios.

The primary flexibility-oriented objectives currently under evaluation within the digital twin include:

- self-consumption maximisation;
- peak shaving and demand reduction;
- reduction of grid dependency;
- improved utilisation of residential battery storage;
- forecasting-driven energy scheduling;
- evaluation of flexibility service readiness.

Overall, the ongoing EMS and forecasting integration activities are intended to enable the digital twin to evolve from a standalone simulation environment toward a predictive validation platform for residential flexibility strategies within the broader FlexBIT architecture.

6.6 Hardware-in-the-Loop (HiL) deployment

Following validation in the Simulink environment, the digital twin will be deployed on a Speedgoat Hardware-in-the-Loop (HiL) platform (Figure 6.1).

This enables real-time execution of the model, allowing:

- testing under real-time constraints;
- validation of control algorithms in a hardware environment;
- safe experimentation before real-world deployment.

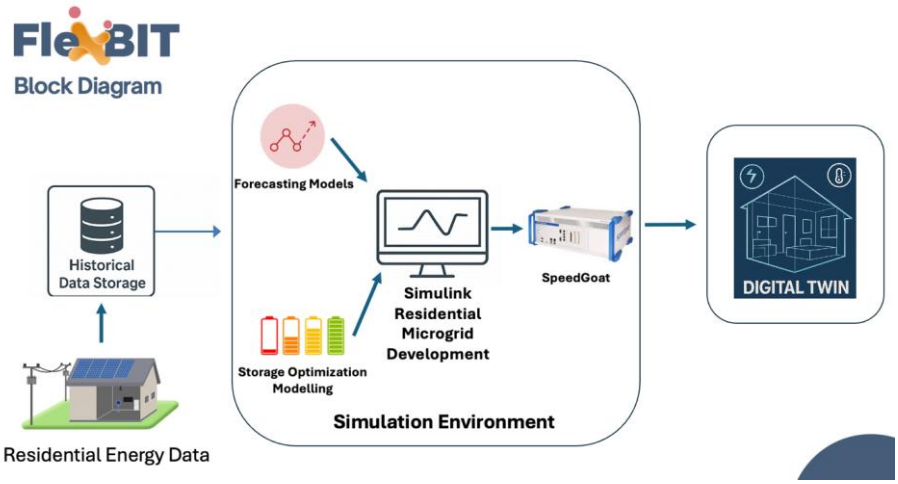

Figure 6.1: HiL environment diagram

The HIL setup acts as a bridge between simulation and physical implementation, improving reliability and reducing deployment risks.

7 Interoperability and Data Exchange

Interoperability within the FlexBIT platform is approached not only as a technical communication problem, but also as a governance, semantic-consistency, and controlled data-exchange challenge across heterogeneous digital-energy environments. The interoperability architecture therefore combines standardised communication interfaces, standard data representations, governance-aware exchange mechanisms, and federated integration principles ensuring demonstrator autonomy while enabling coordinated platform-level operation.

Interoperability in FlexBIT is addressed at two complementary levels. First, each demonstrator may retain its local devices, control systems, data structures, and communication protocols, reflecting the diversity of the project's residential, tertiary, industrial, and laboratory environments. Second, once data crosses the demonstrator boundary and enters the FlexBIT platform, they are represented through a common data model, standardised metadata, and controlled exchange interfaces. This approach allows local autonomy to be preserved while enabling cross-demonstrator monitoring, forecasting, optimisation, flexibility coordination, and later data reuse.

This section, therefore, focuses on the practical interoperability layer of the FlexBIT platform. The section covers the canonical data model, the protocol stack used between field systems, gateways, and the platform, the API and streaming interfaces, and the rules for sharing curated datasets within the consortium.

7.1 Data model

The FlexBIT data model acts as a canonical interoperability layer abstracting heterogeneous local demonstrator structures into governance-aligned and semantically consistent platform representations. This approach reduces dependency on demonstrator-specific schemas and supports scalable integration of heterogeneous cyber-physical energy systems across distributed operational environments.

The FlexBIT data model is designed to harmonize heterogeneous demonstrator data within a common platform representation. Since the participating sites use different devices, local naming conventions, and protocol-specific signal structures, harmonisation was carried out through a canonical mapping process in which equivalent measurements are translated into shared field identifiers before being exposed to platform services and APIs. As a result, the same physical quantity is represented in a uniform way across all demonstrators, independently of the original source system. A typical example is the state of charge of a battery, which may originate from different local tags at different sites but is normalised into a single canonical field used consistently by the platform. The detailed canonical field catalogue is maintained as part of the FlexBIT API documentation. For each canonical field, the catalogue specifies the field identifier, field data type, unit, description, semantic/ontological reference where applicable, and demonstrator usage.

The canonical model is based on a structured field naming convention in the form [Domain]_[Equipment]_[Property]_[Qualifier]. This organisation separates the source context, the relevant subsystem, the measured or controlled quantity, and any additional qualifier such as phase, direction, or type. In practice, this provides a stable internal representation for platform services and ensures that local signal diversity is resolved at the mapping stage rather than propagated to higher-level applications.

Examples for field segments are given in Table 7.1.

Table 7.1: Field ID Structure

Segment	Rule	Examples
Prefix (Domain)	Where is the data coming from?	grid, bess, ev, pv, therm, h2, caes, meta
Middle (Object)	What physical part is it?	bus, inv, bat, chg, storage, fc
Core (Property)	What is the property being measured?	voltage, current, power, energy, temp
Suffix (Context)	What phase or specific type?	l1, l2, l3, dc, ac, active, max, min

For instance, this way, grid bus voltage on phase L1 is represented as grid_bus_voltage_l1.

The canonical model is complemented by semantic alignment with established energy and IoT ontologies. This means that local demonstrator tags are not only renamed, but also mapped to a shared meaning. For this purpose, SAREF Core [29] ontology and its extensions (SAREF4ENER [30], SAREF4BLDG [31], SAREF4GRID [32]) are used, with QUDT [33] acting as a units ontology. This is further supported by aligning data with existing industry standards, such as CIM/CGMES [34], IEC 61850-7-420 [35], OCPP 2.0.1 [36], and SunSpec information model [37].

The ontology alignment mainly serves three purposes. First, it ensures that equivalent quantities from different demonstrators are interpreted consistently by platform services. Second, it supports machine-readable documentation of field meaning, units, data types, and quality information. Third, it provides a path for future integration with external energy-data spaces, digital twins, and other tools.

The harmonised model has three main characteristics:

- it groups data by common energy-resource categories such as BESS, EV charger, PV production, thermal storage, hydrogen, community controller, demand and loads, CAES, and grid metering;
- it uses shared metadata fields, including timestamp, site identifier, asset identifier, data-quality indicator, and source-system reference;
- it supports consistent use of normalised fields across APIs, SDKs, storage components, and monitoring tools.

The harmonisation process therefore went beyond simple renaming. It established a common organisation of asset classes and a shared metadata layer that is used throughout the platform. Records include common fields such as *meta_timestamp_utc*, *meta_site_id*, *meta_asset_id*, *meta_data_quality*, and *meta_source*, which support traceability, provenance control, and quality assessment across sites. This common metadata structure is particularly important in a project where data originate from different technical environments and may have different acquisition characteristics.

More broadly, the data model provides the basis for interoperability across the FlexBIT ecosystem. It enables API-based access and stream-based exchange over a common semantic structure, while remaining aligned with external semantic references and energy-domain conventions. In practical terms, heterogeneous demonstrator data are first normalised into the canonical model, and only then consumed by platform components and applications. This makes the harmonisation step a completed and central part of the current FlexBIT platform design.

7.2 Protocols

The coexistence of industrial protocols, IoT communication layers, streaming infrastructures, and simulation-oriented interfaces within FlexBIT reflects the heterogeneous nature of modern digital-energy

ecosystems. The interoperability strategy therefore prioritises protocol abstraction, asynchronous integration, semantic harmonisation, and modular extensibility over tight protocol coupling.

The FlexBIT protocol strategy reflects the heterogeneous nature of the demonstrators and does not require local installations to expose the same field-level protocol, as each site already operates with its own combination of meters, controllers, SCADA/EMS systems, IoT gateways, and supervisory software. Instead, interoperability is enforced at the adapter and platform boundary. Local protocols remain site-specific, while data exchanged with the FlexBIT platform follows a harmonised platform-level communication model.

At the local level, demonstrators may use protocols and interfaces such as Modbus TCP, MQTT, HTTP, OPC UA, IEC 61850, DNP3, or proprietary SCADA/EMS interfaces, depending on the available equipment and existing operational environment. At the platform level, the current baseline is based on HTTPS APIs and Kafka-compatible streaming over TLS [26][16].

As shown in Figure 7.1, demonstrators' heterogeneity is encapsulated at the IoT adapter level. The adapter translates demonstrator-specific signals into the FlexBIT canonical data model and exposes them to the central platform using HTTPS APIs and Kafka-compatible streams protected by TLS. This design preserves local autonomy while enabling cross-demonstrator monitoring, forecasting, optimisation, and secure data exchange.

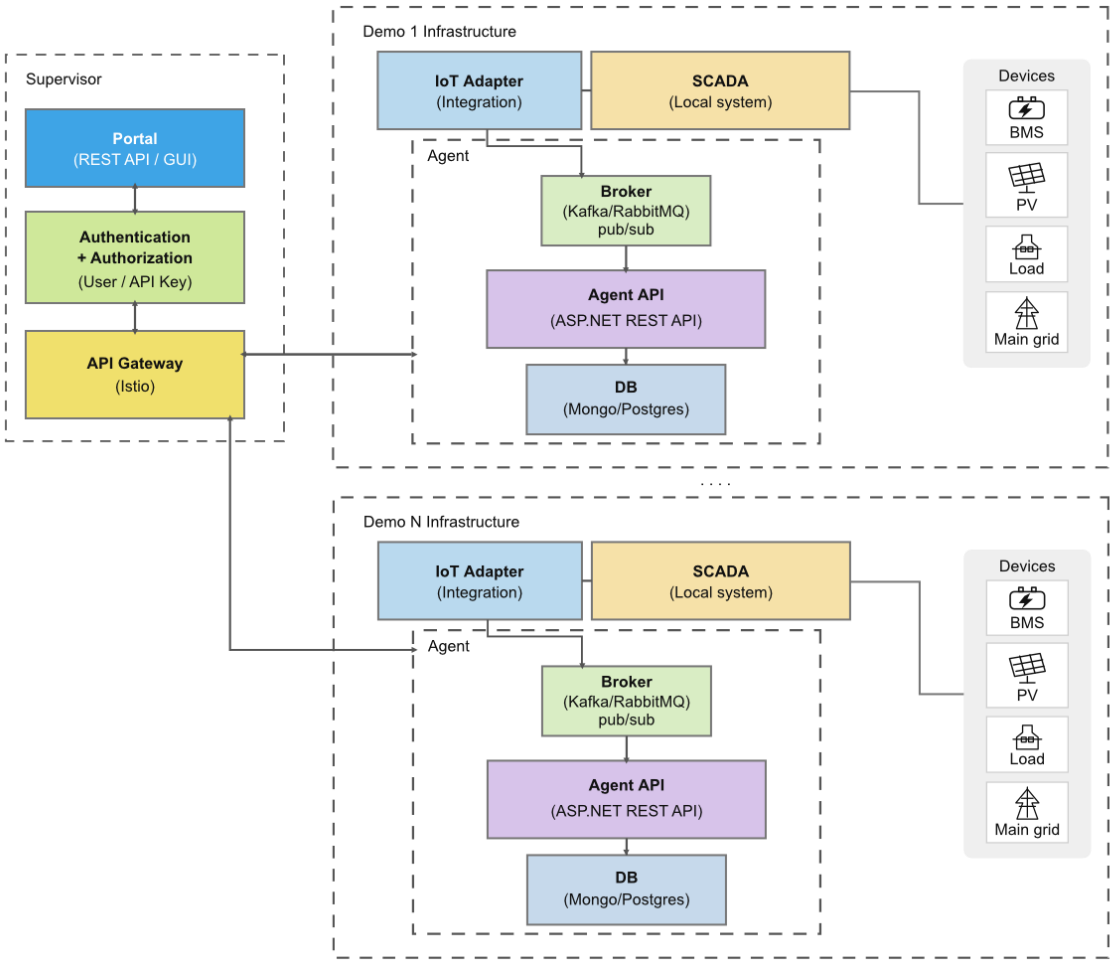

Figure 7.1: Demonstrator integration diagram

In practice, the interoperability workflow distinguishes between local field-level communication within each demonstrator environment and the standardised gateway-to-platform communication layer used for coordinated platform operation. This separation allows local infrastructures to preserve their operational autonomy, internal technological configurations, and demonstrator-specific protocols, while still enabling harmonised cross-platform integration, controlled data exchange, and scalable interoperability across the FlexBIT ecosystem.

Table 7.2 summarizes the current protocol and gateway configuration per demonstrator. The table distinguishes local field-to-gateway communication from gateway-to-platform communication, highlighting this central distinction in the FlexBIT interoperability model.

Table 7.2: Gateway table per demonstrator

Site	Gateway HW/SW	Protocols (field to gateway)	Polling/Push rate	Notes
German (IFF)	Local PCs with Data collector and Forwarder (1 Aue Funeral, 1 aRTE Möbel) → Mosquitto Broker (IFF) → SCADA System (IFF)	HTTP (aRTE Möbel PV) MQTT (aRTE Möbel Loads, Aue Funeral Cooling Loads) MODBUS TCP (aRTE Möbel Compressor, aRTE Möbel Battery, Aue Funeral Battery, Aue Funeral PV, Aue Funeral Cooling)	aRTE Möbel: Loads, Compressor and Battery = 1s PV = 10s Aue Funeral: Loads = 10s to 15s Battery, PV = 3s Cooling = 3s	N/A
UNITOV HiL	Local PC	MQTT, TCP	15 min	N/A
Smart energy parking in Palermo	Local PC	HTTP	15 min	N/A
Solar Cooling System in Pantelleria	Local EMS	HTTP	15 min	N/A
Malta digital twin	MATLAB/Simulink Environment	Simulated data streams	5 min, 15 min, 30 min	Virtual environment using residential PV, load, and grid interaction data
Electrum/Alu-Frost	IoT Adapter, SCADA MQTT Modbus TCP DNP3	Broker-RabbitMQ Agent API-REST API	15 min	N/A
WUST	Local PC, local microcomputer	Modbus, MQTT, VE. Bus, HTTP	1 min, 15 min	LAB test bench of microgrid with real PV inverters and BESS, controllable load

This layered approach allows each demonstrator to preserve its local control architecture while ensuring that the central platform receives harmonised data suitable for monitoring, forecasting, optimisation, and energy exchange mechanisms.

7.3 API layer and integration interfaces

The API layer additionally functions as a controlled interoperability boundary enforcing authentication, role-based access control, protocol normalisation, and governance-aware data exchange policies between local demonstrator infrastructures and central platform services.

The FlexBIT API provides a unified application interface for accessing and exchanging harmonised data across the project demonstrators. Its role is to expose platform data through a consistent structure, so that external integrations, partner applications, and internal software modules can interact with the same canonical model independently of local site-specific implementations. In this way, the API operationalizes the interoperability layer of the platform and makes the harmonised data model available in a structured and controlled form.

The API is organised by asset category, reflecting the main classes of resources and systems managed in FlexBIT. The current structure includes dedicated ingestion paths for BESS, EV chargers, PV systems, thermal systems, hydrogen, community controllers, demand, CAES, and grid metering. This category-based organisation keeps the interface aligned with the platform data model and ensures that each family of assets is handled through a stable and recognizable API structure. The same interface layer also supports asset-oriented retrieval patterns, so that the status of a specific asset can be queried through standardised site and asset identifiers.

The main API characteristics are summarised in Table 7.2:

Table 7.3: API characteristics

Aspect	Description
Interface model	Unified API layer for harmonised cross-demonstrator data access
Organisation	Category-based structure by asset family
Main ingestion areas	BESS, EV charger, PV, thermal, hydrogen, community, demand, CAES, grid metering
Retrieval logic	Asset-oriented status access through standardised site and asset identifiers
Payload format	JSON payloads combining common metadata and category-specific canonical fields
Security scheme	Bearer-token access for API endpoints; platform-generated API credentials for integration
API key management	Credentials are generated on the platform and used to authenticate data submission and integration services
Current sending scope	At present, credentials are partner-scoped rather than device-scoped: any registered device of that partner can send metrics and self-identify in the payload
SDK support	Native client SDKs in TypeScript/JavaScript, Python, and C#
SDK integration model	Built-in Kafka integration for demonstrator adapters and internal modules
Stream authentication	Kafka SASL/SCRAM-SHA-256 using platform-generated API ID and API secret
Rate limits	To be defined; expected to be set according to demonstrator data frequency and operational requirements

Each API payload combines canonical telemetry or status values with a small common metadata set, so that platform services always operate on normalised and traceable information rather than on local raw tags. In practice, the payload structure always includes:

- canonical asset signals or status values;
- core metadata such as timestamp, site ID, and asset ID;
- optional quality and source information for traceability.

Security is based on platform-generated API credentials. These credentials are created from the platform and used to authorize integrations that send metrics on behalf of partner assets. In the current model, authorisation is effectively managed at partner scope rather than at single-device scope: one credential set can be used to publish metrics for any device registered under that partner, while the specific source device is identified in the payload through metadata fields. A finer device-level restriction can be introduced later if needed.

To simplify implementation, the API layer is complemented by official client SDKs in TypeScript/JavaScript, Python, and C#. These libraries are based on Kafka interaction between local adapter and remote module (for ingestion/control flow). The SDKs support two main integration roles:

- device-side integration, for adapters, gateways, and third-party monitoring systems that send telemetry and receive commands;
- module-side integration, for EMS services, schedulers, and optimisation engines that subscribe to telemetry and publish control actions.

Regarding access control, the current implementation supports authenticated access and partner-level credential management, while a more granular scope model can be further developed as the platform evolves. Explicit rate limits are not yet fixed and can be changed in future, to be defined according to demonstrator data frequency, adapter behaviour, and overall platform operational constraints.

7.4 Data sharing

The interoperability framework preserves demonstrator-level data sovereignty by ensuring that local infrastructures retain ownership, operational control, and governance responsibility over raw operational datasets, while only governance-approved and operationally relevant information is exchanged with the central platform.

This section covers the practical sharing and reuse of curated datasets within the consortium for development and research activities, such as AI/ML model training, data analysis, validation, benchmarking, and other tasks requiring collaboration between industry and research partners. It does not apply to every raw data stream, operational metric, alarm, or API payload exchanged with the FlexBIT platform during runtime. Runtime data exchange is addressed through the governance, access-control, security, and interoperability mechanisms described in Sections 3, 5, and 7.1–7.3.

Datasets produced by demonstrators and used in WP2 and related project activities are shared within the consortium under specific access conditions. For controlled internal data sharing, the project SharePoint environment is used. Upon agreement, the requested data are prepared by the data-owning partner and made available to authorised consortium partners. Such datasets may include cleaned historical time series, forecasting datasets, model-validation datasets, aggregated demonstrator datasets, and other curated data extracts needed for project development activities.

Within the FlexBIT interoperability framework, data exchange between demonstrators and platform services is implemented according to controlled-access and governance-aware interoperability principles. The exchange architecture is intentionally designed to preserve demonstrator-level operational autonomy and data sovereignty while enabling coordinated cross-platform communication, semantic consistency, and scalable integration across heterogeneous digital-energy environments.

Consequently, only operationally relevant and governance-approved information is exchanged through the interoperability layer, using standardised interfaces, protocol abstraction mechanisms, and canonical

data representations that support secure, traceable, and semantically harmonised cross-demonstrator interaction.

Post-project reuse of data is governed by the same dataset-level conditions recorded in Section 3. Datasets marked as public may remain available through the selected repository under the assigned license. Datasets marked as internal or restricted to consortium use will not be reused or redistributed beyond the agreed project or partner-specific conditions unless further permission is granted by the responsible partner.

The interoperability architecture adopted within FlexBIT is intentionally designed to remain compatible with emerging European energy-data spaces, cross-domain digital infrastructures, and future federated data-sharing ecosystems requiring semantic interoperability, controlled data exchange, and distributed governance coordination.

8 Deployment Status Across Demonstrators

The FlexBIT platform and its supporting infrastructure are currently being deployed in a testing environment to enable the first interactions with the demonstrators. At this stage, the objective is to validate core services, interfaces, and data-exchange mechanisms under controlled conditions. This phase supports the initial onboarding of pilot sites and the verification of adapters, APIs, and data flows. The current state of central platform implementation is given in D2.1 Section 5. Moreover, description of current demonstrator implementation is presented in D2.1 Section 7. The continuation and consolidation of demonstrator integration is planned in WP3.

Thus, this section summarises the current deployment status of demonstrators and extends it towards future deployment activities and validation, by providing validation targets in the form of demonstrator-specific KPIs. KPIs reported in this section are deployment-to-validation indicators that will be measured and refined once demonstrator data flows and platform-facing interfaces are consolidated.

8.1 German demonstrators (IFF)

8.1.1 Current deployment status

The German demonstrator consists of two main components: a Data Collection and Forwarding module, which is installed at the local sites of the two industrial partners, and a central SCADA system, which runs at the IFF. Communication between the two modules is handled through a MQTT broker instance. The Data Collection and Forwarding module supports HTTP, MQTT, and Modbus TCP connections to various devices at the Aue Funeral and aRTE Möbel facilities. These connections can be created and configured through a UI component for optimal scalability. Collected signals from all connections are stored in a local QuestDB instance for backup purposes and are simultaneously forwarded to the central SCADA system. Figure 8.1 shows the architecture of demonstrators' integration with IFF facilities and FlexBIT platform.

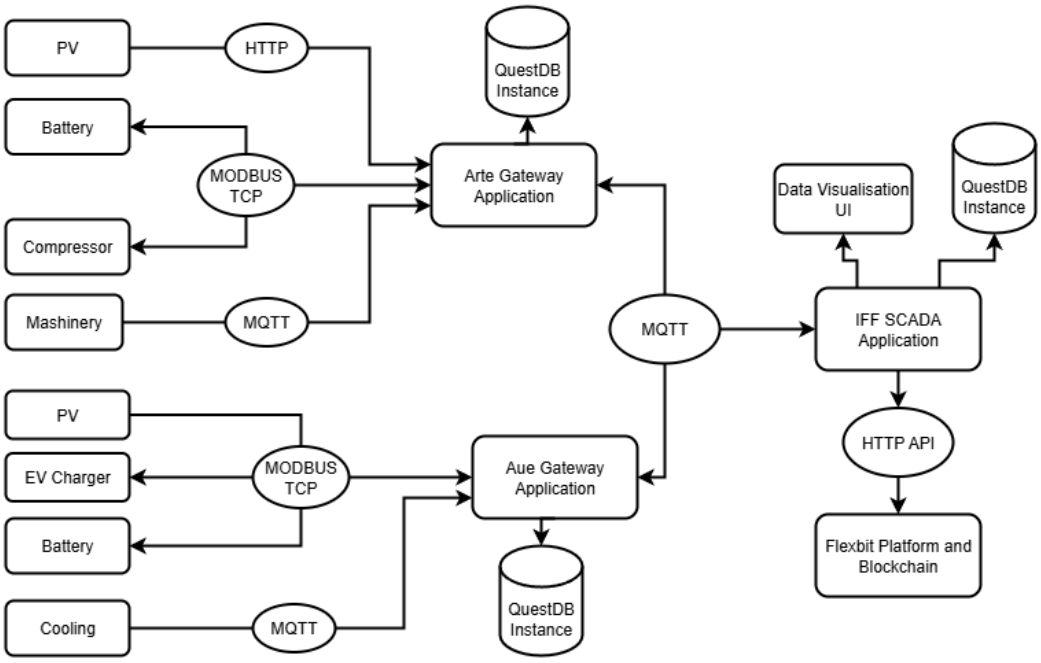

Figure 8.1: IFF demonstrator integration diagram

8.1.2 Validation and KPIs

The key performance indicators against which the demonstrator is evaluated are aligned with the FlexBIT project objectives and listed in Table 8.1.

Table 8.1: German demonstrator (IFF) KPIs

KPI	Description	Target
RES Integration	Percentage increase in renewable energy source integration enabled by flexible device control	>10%
Self-Consumption	Share of total energy generation from solar production consumed within the two demonstrators (aRTE Möbel and Aue Funeral)	Up to 100% of daily energy consumption
Energy Sharing	Amount of energy which is shared among partners	10 kWh – 20 kWh within one hour
System Responsiveness	Average site and visualisation load times of the SCADA Application	<5 seconds
Flexibility Index	Quantification of the system's ability to modulate load/generation in response to FlexBIT advisory signals	Benchmarked against non-flexible baseline
Cooling Flexibility Contribution	Quantification of system’s ability to modulate load by changing cooling aggregators temperature	Benchmarked against non-flexible baseline

8.2 UNITOV HiL demonstrator

8.2.1 Current deployment status

The UNITOV demonstrator is a Hardware-in-the-Loop (HiL) laboratory facility located at the SCERG Lab (Smart and Microgrid Laboratory) of the University of Rome Tor Vergata. The demonstrator operates in the context of the FlexBIT architecture as an Italian laboratory-scale pilot, representing a residential energy community use case. It complements the broader FlexBIT platform by providing a testbed for advanced Energy Management System (EMS) algorithms, Virtual Power Plant (VPP) strategies, and simulated ancillary services, without the constraints of a live grid deployment.

The HiL setup operates under a dual-loop control paradigm: a fast local control loop for real-time hardware operation and safety, and a slow advisory loop through which the central FlexBIT platform receives aggregated data and returns optimisation recommendations via MQTT/REST API interfaces. A smart meter bridges the physical laboratory setup to the FlexBIT platform, enabling real-time online monitoring and data-driven advisory control via MQTT/REST APIs over a secured VPN channel.

The demonstrator is also planned for integration with the University of Palermo (UNIPA) demonstrator (smart energy parking), forming a hybrid microgrid node for cross-site coordinated optimisation involving PV generation, hydrogen production, and EV charging.

8.2.2 Validation and KPIs

The UNITOV HiL demonstrator is designed to validate the effectiveness of multi-vector energy management strategies and the role of hydrogen in providing flexibility services. Validation is carried out through the HiL environment, which allows controlled, repeatable testing of EMS algorithms against realistic load and generation profiles.

The key performance indicators against which the demonstrator is evaluated are listed in Table 8.2.

Table 8.2: UNITOV HiL KPIs

KPI	Description	Target
RES Integration	Percentage increase in renewable energy source integration enabled by flexible storage and control	>60%
Self-Consumption	Share of total energy demand met by local flexible resources	>20% of maximum daily energy consumption
Flexibility Index	Quantification of the system's ability to modulate load/generation in response to FlexBIT advisory signals	Benchmarked against non-flexible baseline
Prediction Accuracy	Improvement in forecasting accuracy for energy demand and supply	>10%
Hydrogen Flexibility Contribution	Assessment of HT-PEM fuel cell and metal hydride system contribution to ancillary service provision and VPP operations	Validated through HiL scenarios

Beyond individual indicators, the demonstrator will produce evidence on the comparative impact of multi-vector storage options (electrochemical batteries vs. hydrogen vs. combined) on the above KPIs, supporting design guidance for future residential energy community deployments.

8.3 UNIPA Smart energy parking in Palermo

8.3.1 Current deployment status

The smart energy parking is a demonstrator built in 2023 at the Campus of the University of Palermo with the aim of testing the participation of this type of end-users in various Demand Response actions or in the REC framework, both through the use of V2G technology and through the control of a 46 kWh lithium stationary battery storage system. The system is equipped with a 40 kW photovoltaic system and a parking light control system and communicates with a central controller located in the Smart and Microgrid laboratory of the Department of Engineering. The demonstrator can be used for testing both frequency regulation services via the V2G charging point and load shifting actions.

Control and communication infrastructure includes: smart meter in the electrical switchboard, LAN of the University Campus, and dedicated PC in the smart and microgrids Lab of the Engineering Department.

8.3.2 Validation and KPIs

The key performance indicators against which the demonstrator is evaluated are aligned with the FlexBIT project objectives (Table 8.3).

Table 8.3: UNIPA smart energy parking KPIs

KPI	Description	Target
RES Integration	Percentage increase in renewable energy source integration enabled by flexible storage and control	>35%
Self-Consumption	Share of total energy demand met by local flexible resources	>40% of maximum daily energy consumption
Flexibility Index	Quantification of the system's ability to modulate load/generation in response to FlexBIT advisory signals	Benchmarked against non-flexible baseline
BESS/V2G Flexibility Contribution	Assessment of BESS/V2G contribution to ancillary service provision and VPP operations	Validated through field tests

8.4 Solar cooling system in Pantelleria

8.4.1 Current deployment status

The demonstrator is an innovative air-conditioning system (Freescoo 3.0 VMC) operating with low temperature heat (solar energy, heat pump, district heating, or waste heat) designed for controlled mechanical ventilation applications in the residential and tertiary sectors. The system is based on an original adsorption air treatment cycle capable of ensuring the control of air temperature and humidity, and guaranteeing adequate air exchange in the building. The main advantages are: operation based on the use of low-enthalpy heat (T>60 °C). For this reason, Freescoo can be integrated with any type of heating distribution system; low electricity consumption due to the movement of circulating air and water.

The system will be used to test the possibility of its use in load shifting actions by exploiting thermal inertia. Current control and communication infrastructure includes: local EMS and LAN of the building.

The current configuration identifies the main thermal and electrical components required for monitoring and later integration with the FlexBIT platform:

- Solar collectors: 3 south-southwest-facing flat-plate solar collectors with a total gross surface area of 7.1 m² (azimuth angle of approximately 10°);
- Cooling system: Freescoo 3.0 device, with the characteristics listed in Table 8.4Table 8.5;
- Hot water storage: 80-liter storage tank with a single coil and auxiliary heating element for water heating.

Table 8.4: Pantelleria solar cooling system technical characteristics

Cooling	The system cools the air drawn in from outside
Dehumidification	YES
Air exchange	YES
Heating	YES
Recupero di calore	YES
Sizes	900 x 1900 x 497 mm
Weight	120 kg
Cooling and dehumidification performance	
Total cooling capacity	3,7 kW
Cooling capacity supplied to the building	1,5 kW
Air inlet temperature (ambient air)	26 °C
Absolute humidity of the incoming air	9 g/kg
Maximum airflow	350 m^3/h
Vapor removed from the building	1,9 l/h
Total removed vapor	3,8 l/h
Air change rate	100%
Heat required for regeneration	4,3 kW
Inlet water temperature under design conditions	70 °C
Outlet water temperature under design conditions	60 °C
Power consumption	0,22 kW
Water consumption	5,2 l/h
EER	17.6

8.4.2 Validation and KPIs

The key performance indicators against which the demonstrator is evaluated are aligned with the FlexBIT project objectives (Table 8.5).

Table 8.5: Pantelleria solar cooling system KPIs

KPI	Description	Target
RES Integration	Percentage increase in renewable energy source integration enabled by Freescoo control	>35%
Self-Consumption	Share of total energy demand for heating and cooling met by local thermal solar collector	>30% of maximum daily energy consumption
Flexibility Index	Quantification of the system's ability to modulate load/generation in response to FlexBIT advisory signals	Benchmarked against non-flexible baseline

8.5 Malta digital twin

8.5.1 Current deployment status

The Malta demonstrator is based on a digital twin environment developed in MATLAB/Simulink for modelling and validation of residential energy management and flexibility-oriented control strategies. The demonstrator represents a residential microgrid architecture integrating photovoltaic (PV) generation, residential load demand, grid interaction (import/export), and Battery Energy Storage System (BESS) operation. The digital twin follows a timestep-based simulation architecture implemented in Simulink. At each simulation interval, the environment updates energy generation, demand, battery state, and grid exchange variables to reproduce residential operational behaviour under varying conditions. Development of the supervisory EMS framework and forecasting-driven control integration is currently ongoing.

The current validated implementation is based on standalone 15-minute forecasting analysis and Simulink-based energy system simulation, while additional 1-minute, 5-minute, and 30-minute forecasting scenarios are under development to analyse the impact of temporal granularity on flexibility-oriented operation and forecasting performance. Future work will focus on integration of forecasting outputs into the supervisory EMS framework for predictive energy management evaluation.

The forecasting framework is also being prepared for coupling with the supervisory EMS to support future predictive operational control. The Simulink implementation is additionally being prepared for deployment on a Speedgoat Hardware-in-the-Loop (HIL) platform. API-based interface is planned for future FlexBIT platform integration, with data exchange structures aligned with the FlexBIT canonical data model.

8.5.2 Validation and KPIs

The Malta demonstrator focuses on validating forecasting-driven residential flexibility management strategies within a controlled simulation and future HIL-based environment.

Validation scope includes:

- accuracy and robustness of load and PV forecasting models;
- performance of predictive battery dispatch strategies;
- impact of forecasting uncertainty on EMS operation;
- evaluation of grid import/export balancing strategies;

- behaviour of flexibility-oriented control under varying environmental conditions;
- transition readiness from offline simulation to real-time HIL execution.

The list of KPIs for the demonstrator is given in Table 8.6.

Table 8.6: Malta Digital Twin KPIs

KPI	Description	Target
Forecasting Accuracy (Load)	Accuracy of residential load forecasting models (LSTM, LightGBM, CatBoost, Hybrid) evaluated using MAE/MAPE/RMSE	Hybrid MAPE <10%
Forecasting Accuracy (PV)	Accuracy of PV generation forecasting under varying environmental conditions	Hybrid MAPE <5%
Multi-Resolution Forecasting Capability	Evaluation of forecasting performance across 1-min, 5-min and 15-min.	Stable forecasting across all resolutions
Simulink Digital Twin Availability	Operational availability of the MATLAB/Simulink residential microgrid environment	It is in operating mode
BESS Model Operational Stability	Stable execution of battery charging/discharging simulations within operational limits	Stable operation without convergence issues
EMS Integration Readiness	Readiness of forecasting modules for future EMS integration.	Successful forecasting-to-EMS interface preparation

8.6 Electrum/Alu-Frost

8.6.1 Current deployment status

The Electrum / Alu-Frost demonstrator represents an industrial-scale energy system located at the Alu-Frost manufacturing facility near Białystok, Poland. The site integrates multiple distributed energy resources (DERs), including photovoltaic (PV) generation, a Battery Energy Storage System (BESS), electric vehicle (EV) charging infrastructure, and industrial loads. Demonstrator also employs smart meters and power quality measurement devices and a local weather station.

The demonstrator operates within the FlexBIT architecture, combining local operational autonomy (fast control loop) with centralised optimisation and advisory control (slow loop). Local control is handled by SCADA/EMS systems (EMACS), ensuring real-time reliability and safety, while the FlexBIT platform provides predictive analytics, optimisation strategies, and coordination signals. Communication is handled by IoT gateways and communication modules (MQTT, Modbus TCP, REST API interfaces), enforcing secure communication via TLS and VPN. For data analytics and model training, an HPC data center is available, while for the ML inference, an edge device with AI accelerator was established. Analytical and AI artifacts currently comprise: ESS Digital Twin for predictive control, PV forecasting models (based on time-series and sky imagery), load forecasting models, and optimisation engine (energy scheduling and flexibility management).

8.6.2 Validation and KPIs

The Electrum / Alu-Frost demonstrator focuses on validating technical, operational, and economic performance of advanced energy management strategies.

Validation scope includes:

- AI-based optimisation of BESS operation;
- PV generation forecasting accuracy;
- load prediction and demand response capability;
- integration of industrial processes into flexibility schemes;
- interoperability with FlexBIT platform and other demonstrators.

The KPIs for the demonstrator are listed in Table 8.7.

Table 8.7: Alufrost KPIs

KPI	Description	Target
BESS Availability	Operational availability of the Energy Storage Station integrated with SCADA/EMS	>95%
BESS SoC Accuracy	Accuracy of State of Charge monitoring and synchronisation	Error <5%
PV Self-Consumption	Share of locally generated PV energy consumed within the demonstrator	>70%
Sky Imaging Availability	Availability of Sky Eye Station image acquisition infrastructure	>95%
Pyranometer Data Accuracy	Availability and consistency of solar irradiation measurements	>98% valid measurements
SCADA/EMS Interoperability	Successful integration of Energy Storage, PV, EV Charger, Heater and sensors through unified telemetry layer	MQTT / REST API / Modbus operational
Real-Time Monitoring Availability	Average response time from FlexBIT advisory signal to flexibility activation	>99%
Forecasting Infrastructure Availability	Availability of AI forecasting infrastructure (Storage VM, Tesla A100, Jetson Orin)	>95%

8.7 WUST demonstrator

8.7.1 Current deployment status

The WUST (Wroclaw University of Science and Technology) power-system microgrid demonstrator includes a DC photovoltaic system hardware simulator, a PV inverter, an AC grid simulator, an energy storage charger, an energy storage system, and a controllable AC load. In this subsection, “demonstrator” denotes a controlled laboratory setup used to emulate selected microgrid components and validate monitoring, control, and data-exchange functions. The setup forms a laboratory microgrid equipped with a monitoring and control system for the microgrid devices (Figure 8.2).

The monitoring and control system uses the available device communication protocols: MQTT, VE.Bus, MODBUS, and HTTP. Data in the internal computer network are converted to MQTT and published collectively to the broker on the main campus network server. A microcomputer acts as an intermediary, forwarding selected MQTT data to the external broker of the FlexBIT platform. The campus-network

microcomputer also subscribes to topics published by the external FlexBIT platform broker and forwards them to the main server.

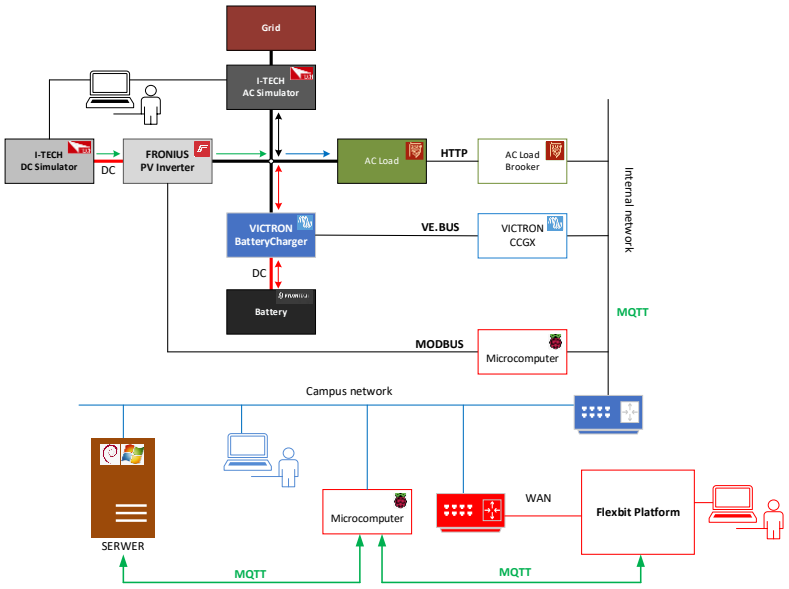

Figure 8.2: Schematic diagram of the WUST power-system microgrid demonstrator and device communication protocols

The main server for the monitoring and control system is a virtual machine running a Unix operating system. It hosts services deployed in the IoT stack convention (a container-based IoT service stack), enabling a flexible system for data storage, development and testing of monitoring and control applications, data processing, data subscription and publication, and visualisation.

8.7.2 Validation and KPIs

The WUST demonstrator is designed to test at laboratory scale the control strategies for activation of demand side flexibility.

The key performance indicators against which the demonstrator is evaluated are aligned with the FlexBIT project objectives and include:

Table 8.8: WUST LAB Test Bench KPIs

KPI	Description	Target
Flexibility Contribution	Quantification of the system's ability to modulate load/generation in response to FlexBIT advisory signals	Benchmarked against non-flexible baseline
PV/BESS Flexibility provision features	Assessment of PV/BESS contribution to ancillary service provision	Validated through laboratory tests
Innovation in Energy Technologies and Tools	Number of new technologies and tools implemented	1 tool

9 Conclusions

This deliverable established the data management and security framework for the FlexBIT platform. It translated the architecture and the regulatory context into practical principles and mechanisms for data governance, privacy, secure exchange, interoperability, and presented the current deployment status of the demonstrators. The document defined the main data categories handled in FlexBIT, the governance and FAIR principles applied to them, and the security controls used to protect platform communication and access.

The report also consolidated the technical basis for secure and interoperable data exchange across heterogeneous demonstrator environments. In particular, it described the trust-boundary model, the adapter-based integration approach, the use of canonical data structures, APIs and streaming mechanisms, the role of platform storage components, and the blockchain-based attestation mechanism used to support data integrity without exposing raw operational datasets.

The deployment overview shows that FlexBIT integration is progressing through a demonstrators onboarding process. At the current stage, the platform and supporting infrastructure are being used to establish initial demonstrator interactions, align local acquisition mechanisms with the common interfaces, and prepare for further validation activities in WP3.

Further work will focus on completing onboarding of demonstrators, strengthening platform monitoring, supporting platform maturation, and validating demonstrator-specific data flows. Public datasets and metadata will be prepared for publication where approved by the responsible partners and where compatible with privacy, confidentiality, cybersecurity, intellectual-property, and project-governance constraints.

10 References

[1] Regulation (EU) 2016/679 of the European Parliament and of the Council on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), in 2016/679. 2016, L 119/1, 4.5.2016

[2] Gesetz über den Schutz personenbezogener Daten bei der Datenverarbeitung (Bundesdatenschutzgesetz – BDSG), in BGBl. I p. 1858; 2022 I p. 1045. 2017: Germany.

[3] Italian Personal Data Protection Code (Codice in materia di protezione dei dati personali), in Legislative Decree No. 196 of 30 June 2003, as amended by Legislative Decree No. 101/2018 implementing Regulation (EU) 2016/679 (GDPR). 2018, Italian Republic: Rome, Italy.

[4] Data Protection Act, in CAP. 586. 2018: Malta.

[5] ustawa z dnia 10 maja 2018 r. o ochronie danych osobowych, Dz.U. z 2018 r. poz. 1000.

[6] Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Act), in PE/85/2021/REV/1. 2022, OJ L 152, 3.6.2022.

[7] Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data and amending Regulation (EU) 2017/2394 and Directive (EU) 2020/1828 (Data Act), in 2023/2854. 2023, L. 22.12.2023.

[8] Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive), in PE/32/2022/REV/2. 2022, OJ L 333, 27.12.2022,. p. 80–152.

[9] Zweites Gesetz zur Erhöhung der Sicherheit informationstechnischer Systeme (IT-Sicherheitsgesetz 2.0), in BGBl. I 2021 S. 1122. 2021: Germany.

[10] Attuazione della direttiva (UE) 2016/1148 del Parlamento europeo e del Consiglio, del 6 luglio 2016, recante misure per un livello comune elevato di sicurezza delle reti e dei sistemi informativi nell'Unione. (18G00092), in D.Lgs. n. 65/2018. 2018, Gazzetta Ufficiale / Normattiva.

[11] Measures For A High Common Level Of Cybersecurity Across The European Union (Malta) Order, 2024, in CAP. 460. 2024: Malta.

[12] Act of 5 July 2018 on the National Cybersecurity System, in Journal of Laws 2018, item 1560, as amended (including amendments implementing Directive (EU) 2022/2555 - NIS2). 2018, Republic of Poland: Warsaw, Poland.

[13] Redpanda. "Redpanda Documentation." *Redpanda Docs*. Available at: https://docs.redpanda.com/current/ Accessed: 23 May 2026.

[14] LF Decentralised Trust, Besu. 2026, LF Decentralised Trust.

[15] Hyperledger Besu, Configure QBFT Consensus. 2026, Hyperledger Besu Documentation.

[16] Rescorla, E., The Transport Layer Security (TLS) Protocol Version 1.3, in RFC 8446. 2018, Internet Engineering Task Force.

[17] Merkle, R.C., A Digital Signature Based on a Conventional Encryption Function, in Advances in Cryptology — CRYPTO '87, Lecture Notes in Computer Science, vol. 293. 1988, Springer: Berlin, Heidelberg. p. 369–378.

[18] Zenodo. "CETP-FlexBIT-Flexibility Exploitation for residential, tertiary and industrial buildings." *Zenodo*. Available at: https://zenodo.org/communities/cetp-flexbit/ Accessed: 23 May 2026.

[19] Open Science Framework. "FlexBIT." OSF. Available at: https://osf.io/b4z78/overview Accessed: 23 May 2026.

[20] Directive (EU) 2019/1024 of the European Parliament and of the Council of 20 June 2019 on open data and the re-use of public sector information (recast). *OJ L 172, 26.6.2019, pp. 56–83*.

[21] Directive 2007/2/EC of the European Parliament and of the Council of 14 March 2007 establishing an Infrastructure for Spatial Information in the European Community (INSPIRE). *OJ L 108, 25.4.2007, pp. 1–14*.
[22] Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011. *OJ L 333, 27.12.2022, pp. 1–79*.
[23] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828. *OJ L, 2024/1689, 12.7.2024*.
[24] Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector. *OJ L 201, 31.7.2002, pp. 37–47*.
[25] Directive (EU) 2019/944 of the European Parliament and of the Council of 5 June 2019 on common rules for the internal market for electricity and amending Directive 2012/27/EU. *OJ L 158, 14.6.2019, pp. 125–199*.
[26] Apache Software Foundation, "Apache Kafka Documentation". Available at: https://kafka.apache.org/documentation/ 2026, Accessed: 23 May 2026.
[27] Creative Commons. "Attribution 4.0 International (CC BY 4.0)." *Creative Commons*. Available at: https://creativecommons.org/licenses/by/4.0/ Accessed: 23 May 2026.
[28] Wilkinson, M., Dumontier, M., Aalbersberg, I. *et al.* The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* **3**, 160018 (2016). https://doi.org/10.1038/sdata.2016.18
[29] ETSI TS 103 264 (V4.1.1): "SmartM2M; Smart Applications; Reference Ontology and oneM2M Mapping".
[30] ETSI TS 103 410-1 (V2.1.1): "SmartM2M;; Extension to SAREF; Part 1: Energy Domain".
[31] ETSI TS 103 410-3 (V2.1.1): "SmartM2M;; Extension to SAREF; Part 3: Building Domain".
[32] ETSI TS 103 410-12 (V2.1.1): "SmartM2M;; Extension to SAREF; Part 12: Smart Grid Domain".
[33] FAIRsharing.org. "QUDT: Quantities, Units, Dimensions and Types." DOI: 10.25504/FAIRsharing.d3pqw7. Accessed: 23 May 2026.
[34] International Electrotechnical Commission. "IEC 61970-452:2021: Energy management system application program interface (EMS-API) - Part 452: CIM static transmission network model profiles." *IEC Webstore*, 2021. Available at: https://webstore.iec.ch/en/publication/64844 Accessed: 23 May 2026.
[35] International Electrotechnical Commission. "IEC 61850-7-420:2021: Communication networks and systems for power utility automation - Part 7-420: Basic communication structure - Distributed energy resources and distribution automation logical nodes." *IEC Webstore*, 2021. Available at: https://webstore.iec.ch/en/publication/34384 Accessed: 23 May 2026.
[36] Open Charge Alliance. "Open Charge Point Protocol." Open Charge Alliance. Available at: https://openchargealliance.org/protocols/open-charge-point-protocol/ Accessed: 23 May 2026.
[37] SunSpec Alliance. "SunSpec Information Model Reference." *SunSpec Alliance*. Available at: https://sunspec.org/sunspec-information-model-reference-sunspec-alliance/ Accessed: 23 May 2026.