Enabling secure and trusted digital twin federations with data spaces

Cristian Martella
cristian.martella@unisalento.it
University of Salento
Lecce, Italy

Angelo Martella
angelo.martella@unisalento.it
University of Salento
Lecce, Italy

Antonella Longo
antonella.longo@unisalento.it
University of Salento
Lecce, Italy

Abstract

As digital twin ecosystems (DTEs) expand across various sectors, enabling secure and trusted collaboration among diverse instances remains a critical challenge. Current approaches often lack standardized mechanisms for data sharing, governance, and trust, hindering large-scale interoperability. This paper presents an approach for the secure federation of DTs using data spaces (DSs), grounded in NGSI-LD standards and policy-based governance models, including ODRL-based access control. This federation approach facilitates trusted data exchange among heterogeneous digital twin (DT) instances through a layered trust and governance approach, supporting decentralization and scalability. Moreover, the proposed approach addresses key barriers to DT federation and offers pathways to integrate participants from trusted federated context in complex DT environments securely, leveraging advantages of DS technology.

A case study documents an implementation of this federation approach, demonstrating how DSs enable secure context data sharing across different DT environments.

Keywords

data space, digital twin, federated digital twin, context data, data sovereignty

ACM Reference Format:
Cristian Martella, Angelo Martella, and Antonella Longo. 2025. Enabling secure and trusted digital twin federations with data spaces. In *The Twenty-sixth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '25), October 27–30, 2025, Houston, TX, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3704413.3765512

1 Introduction

The growing deployment of DT systems urgently needs mechanisms for seamless collaboration and interoperability between involved instances. As complexity and scale increase, individual DTs often operate in isolated environments, creating challenges such as data silos, incompatible standards, and security vulnerabilities.

Federation enables autonomous systems to collaborate and exchange information through standardized protocols and trust frameworks, while maintaining control over their data and operations. This approach promotes secure, interoperable, and scalable interactions to support decentralized cooperation without central authority. In the context of DT ecosystems, the federation integrates and coordinates multiple DTs representing different physical assets or systems, even across organizations, enabling real-time collaboration while preserving data sovereignty and security. DSs can enable federation by providing a trusted and standardized environment. This environment allows multiple entities to share and access data and services securely across organizational boundaries. Furthermore, using common APIs, data models, and governance frameworks, it becomes possible to ensure interoperability, data sovereignty, and trust.

Despite progress, DSs face challenges in supporting federation that hinder a unified and interoperable federation framework. These challenges include limited operational deployment, interoperability gaps, and diversity in governance frameworks, data models, and technical standards.

According to the specifications proposed in the International Data Spaces Association (IDSA) Rulebook [16] and collected in the Data Spaces Business Alliance (DSBA) Technical Convergence Recommendations document [12], DSs can support centralized, decentralized, and federated architectures. These architectures are mainly applied to the context of the trust authority for the deployment modes, leaving concrete opportunities to context data federation unexplored. In this sense, DSBA recommends the adoption of Next Generation Service Interface for Linked Data (NGSI-LD) to enable interoperable, secure, and scalable context sharing across diverse systems, and proposes mechanisms to establish federation of context sources.

This work generalizes the DS architectures discussed in [12, 16], and proposes an NGSI-LD compliant paradigm to federate context data among participants within an DS instance. Specifically, the paper highlights how DSs serve as a secure and trusted infrastructure for data sharing across organizations, with a particular focus on DT federation. DT ecosystems require real-time, synchronized, and context-aware data exchanges among diverse assets, which pose challenges in maintaining data consistency, low latency, and data security. The work employs DSs not just as data brokers but as a trust-enabling framework to support seamless, secure, and scalable federation of DT instances, effectively addressing these specialized requirements. In this sense, a conceptual and practical architecture for federating DTs using DSs, based on NGSI-LD standards and policy governance is proposed. The main purpose of this work is not to propose a novel federation model, but rather to present a federated architecture that leverages data spaces technology to enable secure, scalable, and interoperable digital twin ecosystems. Its main contribution is a proof-of-concept demonstrated through a case study,

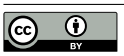

MobiHoc '25, Houston, TX, USA

ACM ISBN 979-8-4007-1353-8/2025/10
https://doi.org/10.1145/3704413.3765512

without presenting formal algorithms or quantitative performance results, to showcase the approach's feasibility and potential. To this end, a case study is documented corresponding to a working federation of contexts in DT environments, implementing the proposed model. The case study illustrates the practical feasibility of the proposed architecture, serving as a proof-of-concept that highlights key functionalities and integration points, leaving as possible future work its performance evaluation in terms of comparative metrics or optimization strategies.

After this introduction, the discussion of the paper follows by including Section 2 where the research landscape that frames this work is proposed and Section 3 describing the proposed approach. Section 4 documents an essential implementation of this approach in the form of a case study, while the corresponding results are analyzed and discussed in Section 5. Finally, Section 6 concludes the paper.

2 Background

2.1 Digital twin ecosystems

A DT offers a real-time representation of its physical analogue, facilitating simulation, testing, and performance optimization [24]. Unlike digital models and digital shadows, a DT enables bidirectional data exchange between the physical asset and its digital replica. Consequently, physical asset modifications automatically update the digital representation, and, conversely, virtual replica evolution may necessitate physical asset interventions [37].

The concept of DTE has been introduced to enhance product and service development, but also to identify novel product-service systems [36]. The DTE is a digital platform that leverages DTs to support product design and lifecycle management through a value network of twin-driven products and services. As defined in [36], the DTE comprises "an interconnected multiple instances of a digital twin or different digital twins that have been arranged into value networks using the different enabling technologies for digital twins". This concept enables industries to achieve real-time prediction and continuous optimization of system parameters, providing intelligent optimization instructions [21, 42].

DTEs provide various benefits, such as the integration of various data sources, real-time monitoring and prediction, and the simulation of policy scenarios. Existing research highlights gaps, including the need for privacy-preserving techniques, real-world case studies, and integrating DTs with decision support systems [2].

2.2 Federated digital twins

Federated DTs can be considered as a possible evolution of DTs aimed at enabling multiple local DT models to collaborate, sharing aggregated information without exchanging raw data [18]. Local models can be changed independently, while periodic updates are sent to the server for global model aggregation, enhancing accuracy and fidelity of collective representation by leveraging diverse local data and experiences. Proceeding in this way, a more comprehensive DT can be obtained for improved simulation and prediction of the physical system's behavior.

However, federated DTs pose security, privacy and trust challenges [43], including the limitations of passive defense mechanisms and data privacy risks due to decentralized architecture and sensitive data transmission. Proactive threat countermeasures and robust privacy-preserving schemes become crucial.

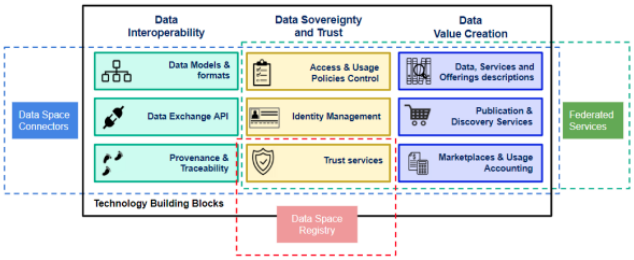

Figure 1: DS technical pillars proposed by OpenDEI

The prominent approaches to federate DTs focus on integrating advanced technologies to enhance collaboration, security, and efficiency. These approaches include federated learning (FL) frameworks to enable distributed model training across multiple DTs while preserving data privacy [29, 40]. Such frameworks are often enriched with blockchain technology to ensure secure data management and improve trust [1]. To optimize performance in heterogeneous networks, hierarchical FL frameworks have been studied that adopt mobile edge computing [34, 35]. Moreover, some approaches employ adaptive aggregation techniques, such as trust-based aggregation and adaptive thresholding, to improve model accuracy and detect anomalies [38]. These enhancements aim to address the challenges of federating DTs, enabling secure, efficient, and collaborative operations in industrial and IoT environments.

Federated open data models are also being developed to improve interoperability across scales and domains [29]. Innovative DT-based Federated Learning (DT-FL) frameworks are also emerging to enable secure communication and virtual monitoring of remote clients [34].

As a result, current research is mainly focused on specific technical challenges, with limited work on comprehensive frameworks for a secure, trusted, and privacy-preserving federation of context sources in DTEs. This paper tries to fill this gap by proposing a novel approach to federation that adopts DS technology, to address data sovereignty and privacy preservation requirements.

2.3 Data spaces

DSs are trusted ecosystems that enable secure, organized, and reliable data exchange between organizations, where participants offer or consume data and/or digital services. Recent DS initiatives have evolved independently, prompting efforts to establish standard technical specifications [23].

The European Interoperability Framework[1] addresses technical and semantic interoperability requirements by introducing specific technical pillars. In particular, the OpenDEI initiative proposes three technical pillars for the development of DSs, as shown in Figure 1. These pillars are documented in [26] and discussed in the following by introducing the corresponding building blocks for each pillar as DS requirements.

[1] https://op.europa.eu/en/publication-detail/-/publication/bca40dde-deee-11e7-9749-01aa75ed71a1/language-en

- Data Interoperability: DSs require efficient data exchange among participants, achieved through common APIs, data models, and traceability and provenance mechanisms.
- Data Sovereignty and Trust: DSs must implement standardized participant identity management, truthfulness verification, and policy enforcement to ensure trust and data sovereignty among participants.
- Data Value Creation: DSs have to enable multi-sided markets by implementing: (1) specifications for terms and conditions, (2) data service publication and discovery, and (3) contract lifecycle accountability for data access and usage.

Data Space Connectors (DSCs) [12] are novel technologies that serve as agents of DS participants implementing a selection of the building blocks. As shown in Figure 1, the Federated Services of a DS instance insist on pillars related to both the Data Value Creation and Data Sovereignty and Trust. In this case, the focus lies on the Trust Services building block that provides the Data Space Registry, managed by a Data Space Authority, as detailed in Section 2.3.2.

DSs enable real-time guarantees and edge resource management, facilitating efficient data sharing and processing. They also support FL, allowing AI algorithms to maintain privacy and explainability. In effect, they enable collaboration between multiple entities without disclosing raw data, fostering context-aware decision-making and orchestration. In addition, DSs promote security and resilience in distributed computing, particularly in cyber-physical systems, through secure data exchange protocols [9, 19, 33].

DSs can improve the scalability of swarm intelligence frameworks by integrating various data sources, benefiting smart cities and industrial automation [9]. Moreover, DSs can promote trust in distributed systems through transparent data governance and accountability, supporting effective monitoring and resource coordination [27].

The standardization of data formats enables interoperability across environments, supporting decentralized data pipelines for learning and analytics. In this regard, the EU promotes ETSI's NGSI-LD API for Context Information Management (CIM) [4] as a standard for semantic interoperability. NGSI-LD is an open standard for context information management, facilitating data exchange and integration across systems and applications. Using JSON-LD (JavaScript Object Notation for Linked Data) as a data format, data can be represented as a graph of interconnected entities and concepts. The Context Broker (CB) is a software component that is responsible for supporting CIM by implementing the corresponding NGSI-LD interfaces. The NGSI-LD API can natively support various architectural settings, including centralized, distributed, and federated configurations.

Current endeavors are mainly devoted to develop more efficient implementations of decentralized identity management [20, 22]. In particular, the use of decentralized identities (DIDs) is widely perceived as a key enabler for self-sovereign identity management and privacy-preserving data sharing. The available implementations can support key requirements in terms of user control over personal data, privacy enhancement, and secure digital authentication [22].

Recent research is investigating how to integrate DTs and DSs with the aim to implement trusted federation of DTEs that are capable of facilitating secure and sovereign data sharing in industrial ecosystems [28]. One first result in this sense is reported in [37], where the Digital Twin Space is introduced to address complex data management challenges in distributed DT scenarios, incorporating DSs concepts. Another contribution that leverages this kind of integration is documented in [11]. In this case, the application of the proposed approach enables flexible data sharing and seamless interaction in urban DTs, promoting interoperability and autonomous operation. The same approach is also applied in the energy domain, where DTs of power systems integrate data from various sources to improve performance and efficiency [39].

In summary, current literature investigating the use of DS technology to support DTEs is particularly limited due to the low-level maturity of the available implementations [24, 37]. Furthermore, companies' reluctance to fully perceive the value of open standard architectures is also contributing to this limitation, although data sharing is even more recognized as a key factor in various industries [17].

2.3.1 European initiatives and technical specifications. The European Commission is currently supporting IDSA and Gaia-X initiatives, which provide technical specifications for the development of DSs, focusing on data sovereignty and trust between participants. Other initiatives, such as the Big Data Value Association (BDVA) and the Data Spaces Support Centre (DSSC), focus on data management and provide shared blueprints for DS design. Furthermore, international standardization bodies, such as the DSs Business Alliance (DSBA), promote technical convergence and uniform rule sets to streamline seamless and inter-DS data sharing and promote the adoption of common standards [25].

The Dataspace Protocol https://docs.internationaldataspaces.org/ids-knowledgebase/dataspace-protocol is a set of specifications aimed at publishing, negotiating, and accessing data within a DS, allowing entities to participate in a federation of technical systems.

IDSA and Gaia-X are key contributors to the development of the DS technical specifications. In particular, IDSA established a reference architecture, an information model, and specific tools for secure data exchange, including DSCs for usage control policies, metadata brokerage, and identity management through trusted authorities. IDS-RAM is the IDSA's framework for implementing data access, sharing, and usage rule enforcement, complying with both the General Data Protection Regulation (GDPR) [5] and FAIR (Findable, Accessible, Interoperable, and Reusable) [41] principles. In line with IDS standards, Gaia-X aims to develop a federated data infrastructure with a special focus on interoperability, decentralization, and sovereignty.

As documented in [12], a decentralized Identity and Access Management (IAM) framework is available among the Data Space Connector (DSC) components implemented by the FIWARE Foundation [7]. This framework supports Gaia-X trust frameworks and is compliant with both the OpenID Connect (OIDC) and Self-Sovereign Identity standards from the W3C consortium (decentralized identity, verifiable credentials).

Concerning the federation of data and services offered by participants in a DS, most approaches remain largely theoretical or in the early stages of development [3, 32]. The IDS architecture represents a prominent conceptual framework, which uses connectors

and brokers, while Gaia-X provides blueprint-level federation services [30]. Again in this regard, most federation implementations are still in the early to intermediate stages or exist only as pilots and prototypes [31], highlighting a significant gap between conceptual frameworks and real-world operational deployments [8].

2.3.2 Governance architectures. As introduced in Section 2.3, the key components of a DS federation are the DS Governance Authority and the DS Registry. The DS Governance Authority defines rules for a DS instance, which are also used by the DS Registry. Common governance and rules can be adopted to enable cross-DS interoperability, using a shared meta-registry such as the Gaia-X Registry. The DS Registry may adopt various mechanisms to identify trusted participants that can be either public or private. In the DSBA Reference Technology Framework [12], identification is based on Verifiable Credentials (VC) issued by Trusted Issuers that can be in turn accredited or registered on the DS Registry.

For developing trusted DSs, IDSA proposes mandatory and optional functional requirements. Three design approaches can be followed for developing the DS Registry according to centralized, decentralized, or federated models. While supporting each of the aforementioned approaches, IDS-RAM currently adopts the centralized approach, while Gaia-X is planning to integrate a wider support [12]. In this regard, IDS-RAM v5.0 is expected to include decentralized and federated approaches as well, e.g. in Trust Framework [14].

This paper extends existing concepts to enable secure and seamless context data federation between DTEs. Its innovative aspect lies in treating DTEs as participants within a DS instance, leveraging DS security and trust frameworks to ensure sensitive data privacy protection and trustworthiness among DT units in the federation.

3 Context federation in Data Spaces for connected digital twins

3.1 Data Spaces in Digital Twins

Participants in DSs to offer or consume data and/or digital services must be characterized before starting any secure data exchange. In the context of DTs, various approaches can be considered to identify participants, based on well-defined organizational boundaries. For instance, the organizational unit that manages the cyber-physical assets may differ from the organizational unit(s) that offers digital twinning data services. In this sense, both would need to be registered as participants of a DS instance.

In general, a provider offers data or services to other trusted participants within the same DS via a trusted contracting interface. Instead, a consumer uses the data or services offered by a provider within the same DS. Anyway, scenarios where participants can act simultaneously as a consumer and a provider are also possible.

The data exchange process within distributed DSs is characterized by a series of well-defined steps that ensure secure and trusted data sharing among participants. These steps include (1) onboarding, where participants are registered and their identities verified through trusted registries using VCs; (2) authentication, which confirms the identities of the entities involved; (3) authorization, where access policies are enforced to regulate data usage based on predefined rules; and finally, (4) data exchange itself, which is conducted through standardized interfaces and protocols that facilitate efficient, traceable, and privacy-preserving data transfer.

Federation plays a crucial role in managing context data within federated data services by enabling seamless and secure interoperability among multiple autonomous systems and organizations. Using standardized protocols and governance mechanisms, federation establishes a trusted environment where diverse data sources can share, access, and exchange context data without compromising their individual control or sovereignty. This interconnected framework supports the dynamic collaboration of various stakeholders, ensuring that context-specific information remains consistent, reliable, and accessible across different domains.

In this regard, the proposed approach extends DS technology towards a dedicated DT federation framework by integrating context-aware trust management and dynamic access policies tailored for physical assets' digital representations. This integration enables DT environments that are not only data-connected but also resilient and trustworthy, supporting real-time decision-making, synchronized operations. This fosters the development of multi-stakeholder ecosystems where data can be exchanged efficiently, unlocking new insights and value creation.

3.2 Reference federation architecture

Federated deployments consist of distributed deployments that transparently offer access to context information made available by federated players. In particular, such context information is managed via CB instances associated with the corresponding federated participants. A federated scenario is not technically different from other distributed scenarios, with two main key differences: (1) it goes over administrative boundaries; and (2) it does not adopt a centralized control approach.

In federated scenarios, the focus typically is on accessing (and possibly aggregating) context information from multiple CBs, while the management of the information (create, update, delete) takes place locally in each domain.

According to NGSI-LD specifications, the CB-to-CB approach is the default federation implementation. Based on the requirements, other possible distributed deployment scenarios may be: (1) support for actuation, (2) support for "lazy" attributes, (3) explicit distribution of context information, (4) backup of context sources, and (5) implementation of complex data sharing scenarios.

Such distributed deployment scenarios require different subsets of operations with respect to the default federation implementation, which means that a full CB is often not required. Anyway, for some scenarios, security concerns may explicitly deny specific operations. In other scenarios, some entities or attributes are exclusively available on a single specific context source (CS), e.g. device actuation, while in any others the same context could be augmented referring to multiple sources with limited availability.

This work generalizes the federated architecture introduced in the DSBA Technical Convergence Recommendations [12] and further detailed in the IDSA Rulebook [16].

The general model consists of a multi-level hierarchical structure for federated CSs where each participant aggregates contributions offered by downstream context units (i.e. federated providers) and

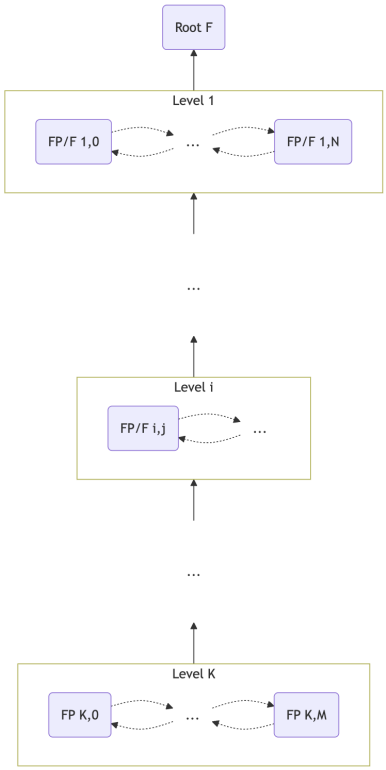

Figure 2: Federated context sources general architecture

potentially becomes a federated provider for its upstream federators. This scenario is illustrated in Figure 2.

The diagram in Figure 2 shows the general tree-based topology obtained by properly organizing the participants in Context Federator (F) and Federated Context Provider (FP). In particular, for a generic node $FP/F_{i,j}$, let $i \in [0, K]$ be the level of the node in the hierarchy starting from the root, where K corresponds to the depth of the tree. Similarly, let $j \in [0, Q]$ be the index of the node, where Q corresponds to the maximum degree of its parent nodes. The node $FP/F_{i,j}$ can participate in the federation of multiple upstream and/or peer Fs, while combining the contributions provided by its downstream FPs. Hence, in such a tree representation the root is an F node, while the leaves are FPs.

It is possible to combine multiple tree-based structures to achieve advanced federation schemas. For instance, multiple F nodes can be defined, one for any given set of downstream FPs. In this work, the CB-to-CB federation approach is adopted, which can be implemented through Context Source Registrations (CSRs). A detailed discussion of this approach is proposed in the next section.

3.3 Context Source Registration to enable federation of contexts

CB architectures assume that entity data do not need to be centralized within a single CB instance. However, when querying context information, retrieval of entity data can be considered as a unitary operation, masking the fact that each registered CB is receiving a separate distributed Context Consumption request [4]. To process each Context Consumption request efficiently, it is necessary for the CB to initially make a broad request addressed to each registered CS whose registration is matching the request.

When an CS is registered, an operation mode is selected. This defines the basis for distributed operations and also defines whether or not the CB is allowed to hold context data about the entities and attributes locally itself. Furthermore, it is possible to limit a registered CS to operate exclusively on one or more specific entity types, but also to indicate its availability to serve just a limited subset of API operations.

CBs must take into account these situations to avoid sending additional distributed operation requests that will fail anyway. In alternative, the access to some CS endpoints (such as updates) may be granted only to authorized users and not accessible to the other CBs. Limited access is likely to be the case in extended data sharing scenarios, where a registered CS may belong to an external third party, along with its corresponding data.

Four types of CSR operations exist, organized into Additive and Proxied categories (as illustrated in Figure 3).

In case of Additive Registrations, a CB is allowed to hold context data about entities and attributes locally, while also obtaining data from (possibly multiple) external sources.

Additive CSRs can be differentiated into Inclusive and Auxiliary CSRs. Inclusive CSR is the default mode of operation. It allows a CB instance to consider any registered CSs as equals in distributing operations to them. In this case, if relevant data are directly available within the CB itself, any result will be integrated in the final response. On the other hand, an Auxiliary CSR never overrides any data held directly available within a CB. Auxiliary distributed operations are limited to context information consumption: context data from Auxiliary CSs are included exclusively when they provide a supplementary contribution to the directly available one.

Concerning Proxied Registrations, instead, a CB is not allowed to hold context data locally itself. In effect, all context data are obtained from external registered sources.

Proxied Registrations can be differentiated into Exclusive and Redirect CSR. In an Exclusive CSR, any registered context data is held in a single location which is external to the CB. An Exclusive Registration must be fully defined and it always relates to specific attributes present in a single entity. The CB itself holds any data locally about the registered attributes on the entity. Given these characteristics, Exclusive CSRs are specifically suitable for actuations. In Redirect CSRs, registered context data is still held externally, but potentially multiple distinct Redirect registrations can be enabled at the same time.

4 Federation of contexts: a case study

This section documents the assessment of the proposed approach through a case study that showcases the practical use of DSC to implement the federation of contexts. In the following, the concept of context is meant to describe the set of entities, relationships, and vocabulary managed by a DTE. The DSC implementation adopted to develop this case study is the FIWARE Data Space Connector (FDSC). The reasons behind this choice are the support for NGSI-LD APIs and the relative maturity level of the implementation.

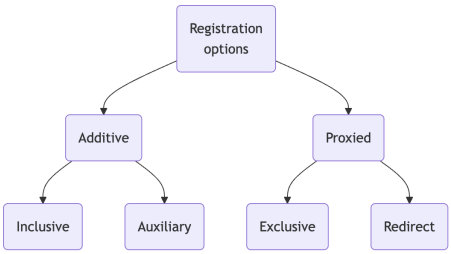

Figure 3: Context Registration operations

The FDSC is a modular and flexible connector that enables secure, policy-compliant, and semantically interoperable data exchange among diverse participants in DSs. It adheres to the DSBA Technical Convergence recommendations [12] and supports integration with established protocols such as the IDSA Dataspace Protocol [13] and the Gaia-X Trust Framework [10]. FDSC offers essential DS features including catalog exploration, data transfer management, and contract negotiation. It further promotes the adoption of NGSI-LD APIs for CIM and provides a library of NGSI-LD compliant data models. FDSC integrates with the Gaia-X Trust Framework, leveraging its Digital Clearing House (GXDCH) as a Trust Anchor to establish secure and trustworthy identities. It also supports policies and mechanisms for data usage control, incorporating contractual terms and operational policies into data sharing workflows. It can also enforce usage rights through standardized policy frameworks such as the Open Digital Rights Language (ODRL).

In the following, a description of the architecture that implements the case study is reported along with the definition of the contexts and the corresponding managed data entities. Subsequently, a detailed discussion of both the adopted policies of data sharing and the federation workflow is also reported.

4.1 Architecture

The proposed case study discusses the development of an architecture consisting of one F node and two FP nodes, namely FP-A and FP-B. FPs are responsible for managing the context data related to photovoltaic (PV) measurements that are provided by the corresponding PV devices. The two PV devices considered in this scenario are named PV-Device-A and PV-Device-B.

FP-A and FP-B hold context data and metadata related to the PV-Device-A-001 and PV-Device-B, respectively. In particular, managed metadata also include data about the current state of the monitored device. Listings 1 and 2 report the structure of data entities managed in the FP-A context. A similar structure is used for the data entities involved in the FP-B context.

```
{
  "id": "urn:ngsi-ld:PhotovoltaicDevice:
      PhotovoltaicDevice:PV-Device-A",
  "type": "PhotovoltaicDevice",
  "name": "PV-Device-A",
  "location": {
    "type": "Point",
    "coordinates": [
      50.774783,
      6.084345
    ]
  },
  "NominalPower": 20,
  "MaximumSystemVoltage": 24,
  "@context": [
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-
        context.jsonld",
    "https://raw.githubusercontent.com/smart-data-
        models/dataModel.GreenEnergy/master/
        context.jsonld"
  ]
}
```

Listing 1: PV-Device-A PhotovoltaicDevice data entity detail.

```
{
  "id": "urn:ngsi-ld:PhotovoltaicMeasurement:PV-
      Measurement-A-2025072510000000Z",
  "type": "PhotovoltaicMeasurement",
  "dateObserved": "2025-07-25:10:00.000Z",
  "location": {
    "type": "Point",
    "coordinates": [
      50.774783,
      6.084345
    ]
  },
  "nominalPeakPowerGeneration": 18,
  "temperature": 23.4,
  "refPhotovoltaicDevice": "urn:ngsi-ld:
      PhotovoltaicDevice:PV-Device-A",
  "@context": [
    "https://raw.githubusercontent.com/smart-data-
        models/dataModel.GreenEnergy/master/
        context.jsonld"
  ]
}
```

Listing 2: PV-Measurement-A PhotovoltaicMeasurement data entity detail.

4.1.1 Context Federator. In a FDSC-enabled Inclusive federated context, the F instance keeps track of any CSR associated with the corresponding FP node. Essentially, for a given entity, it is possible to combine different bits of information provided by multiple local CSs and output a comprehensive representation of the same entity that ultimately integrates such contributes.

To this end, the F instance must include a data service, namely a CB, to manage any CSR to the corresponding FP node. FDSC proposes Scorpio [6] as a CB implementation because it natively supports federation through CSRs, in compliance with the NGSI-LD API.

The F node holds data entities and attributes provided by trusted and authorized providers. Hence, the F node recipe must include the Authentication and Authorization service stack. In particular, Authentication services are registered and connected to the Verifiable Data Registry in order to identify trusted DS participants and verify the corresponding credentials against their roles and authorized operations.

In addition, services for data marketplace and contracting are also offered to enable the definition of the digital product catalog, including product specifications and offerings. In this sense, the digital product offered by an FP node enables the definition of a CSR for sharing context data entities.

For simplicity, the present scenario does not include marketplace operations and makes it possible for any authenticated user to accomplish authorized operations using the corresponding VCs. Anyway, advanced operations and policies corresponding to the configuration of contract management can be implemented by following a few extra configuration steps.

For the purposes of this simple demo scenario, a user that operates on behalf of a federated data provider can register a CSR to the F node and/or interact with this node to fetch entities of a given type. The F node is also responsible for enforcing proper ODRL policies to grant or restrict access to specific data entities included in the federated contexts and allowing operations on the F data service itself. The relevant policies for this case study are detailed in Section 4.2.

4.1.2 Federated Context Provider. An FP node offers local context data and services to an upstream F node. The F node aggregates and shares data provided by the FP nodes in favor of any other authorized DS participants, as discussed in Section 4.1.1.

For managing and mirroring local context to the F node, each FP node refers to its FIWARE's Scorpio CB instance.

In this scenario, the FP node authenticates on the F node and acquires the necessary authorization. This authorization allows the FP node to consume its data service and create a CSR, consequently. The FP node can be considered as a hybrid participant between Provider and Consumer roles. It consists of a CB instance and a trusted VC issuance service, as well. The latter service is responsible for issuing VCs that can be used to authenticate on the F node.

4.2 Data sharing policies

For the implementation of the case study, two data sharing policies were defined and registered through the Policy Access Point (PAP) interface. The first policy allows the creation of a CSR instance for owners of user VCs associated with FPs. The second policy, instead, grants access to "PhotovoltaicMeasurement" data entities, permitting only consultation of this entity type.

By default, the implementation denies access to entities without a targeted policy, preventing the F node from sharing context data related to PV devices due to the absence of a policy allowing consultation of "PhotovoltaicDevice" entities.

Listings 4 and 5 represent the code snippets corresponding to the ODRL target refinements for the two policies, respectively. They also include information related to the policy assignees and the corresponding permitted actions. In particular, for each policy refinement, the target entity type is specified along with the action allowed to the designated set of assignees.

For simplicity, both policies are assigned to any VC owner. However, more complex scenarios can be developed by chaining further refinement policies that specify ranges of property values, time, or geospatial constraints, but also by assigning them to different classes of VC owners or to specific subjects.

```
...
  "odrl:target": {
    "@type": "odrl:AssetCollection",
    "odrl:source": "urn:asset",
    "odrl:refinement": [
      {
        "@type": "odrl:Constraint",
        "odrl:leftOperand": "ngsi-ld:entityType",
        "odrl:operator": { "@id": "odrl:eq" },
        "odrl:rightOperand": "
            ContextSourceRegistration"
      }
    ]
  },
  "odrl:assignee": { "@id":"vc:any" },
  "odrl:action": { "@id":"odrl:use" },
...
```

Listing 3: ODRL policy: target refinement detail to allow the creation of CSRs.

```
...
  "odrl:target": {
    "@type":"odrl:AssetCollection",
    "odrl:source":"urn:asset",
    "odrl:refinement": [
      {
        "@type":"odrl:Constraint",
        "odrl:leftOperand":"ngsi-ld:entityType",
        "odrl:operator": { "@id":"odrl:eq" },
        "odrl:rightOperand":"
            PhotovoltaicMeasurement"
      }
    ]
  }
  "odrl:assignee": { "@id":"vc:any" },
  "odrl:action": { "@id":"odrl:read" },
...
```

Listing 4: ODRL policy: target refinement detail to grant access to PhotovoltaicMeasurement entities.

4.3 Federation workflow

The first step to implement the architecture proposed as a case study consists of onboarding FP and F nodes as participants in the DS-enabled federation. To this end, participating organizations must be registered in the global trust registry before issuing the corresponding VCs for each node. In particular, trusted issuance authorities are delegated to issue VCs, which are then used to check a node's membership in the same DS instance and its corresponding permissions. It is worth noting that the onboarding process is not documented in the following discussion essentially because it is outside the scope of this work.

Assuming that the CB instances corresponding to the two FP nodes are populated with context data entities and that the proper ODRL policies are in place, as documented in Section 4.2, the procedure to create a federation of participants consists of the following three steps: (1) token issuance, (2) creation of a CSR, and (3) federated context consultation. A detailed description of these steps is proposed below.

4.3.1 Token issuance. This simple scenario aims to demonstrate how to set-up an essential federation of participants, implementing a basic workflow where any owner of user VCs is authorized to obtain an access token.

The interaction with the F node's data service is protected by an authentication and authorization layer. Therefore, any participant who needs to interact with it must first obtain an access token. In this case, the candidate FP node can request an access token for a user by presenting the corresponding issued VCs. In particular, the authentication and authorization process starts with the request for the issuance of VCs for a user. These VCs must be issued by the candidate FP's VC issuer and used to obtain the corresponding user access token. The F authentication layer is responsible for assigning user access tokens. Eventually, the issued token must be embedded in any request sent to the F data service, including the creation and management of CSRs.

4.3.2 Creation of a Context Source Registration. The creation of a CSR allows the federation of participants to take place.

```
curl -s -X POST 'http://mp-data-service-federator.127.0.0.1.nip.io:8080/ngsi-ld/v1/csourceRegistrations' \
  -H "Authorization: Bearer $USER_TOKEN" \
  -H 'Content-Type: application/json' \
  -d '{
    "id": "urn:ngsi-ld:ContextSourceRegistration:csr-federated-participant-a",
    "type": "ContextSourceRegistration",
    "information": [
      {
        "entities": [
          { "type": "PhotovoltaicMeasurement" }
        ]
      }
    ],
    "endpoint": "http://scorpio-federated-participant-a.127.0.0.1.nip.io:8080",
    "@context": ["https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.7.jsonld"]
  }'
```

Listing 5: CSR creation POST request detail.

As illustrated in Listing 5, the payload of the request contains information on entities that have to be mirrored to the F data service ("PhotovoltaicMeasurement" entities, in this case) along with the target FP data service endpoint.

More advanced configurations are available to enable fine-grained federation rules, as documented in clauses 5.9 and 5.2.9 of the ETSI NGSI-LD API specification [4] corresponding to the CSR Definition and API Operations, respectively. It is worth highlighting that in this case study only entities belonging to the "PhotovoltaicMeasurement" type are mirrored to the F node and not those having type corresponding to "PhotovoltaicDevice".

4.3.3 Federated context consultation. The last step consists in checking whether the contexts exposed by the federated FP nodes can be accessed. In other words, the goal is to verify the access to the FPs' data entities on the F node through to the established CSRs.

As discussed in Section 4.1.1, by registering specific ODRL policies to the F's PAP, it becomes possible to mirror on the F node entities with "PhotovoltaicMeasurement" type held in the FPs' contexts.

```
[
  {
    "id":"urn:ngsi-ld:PhotovoltaicMeasurement:PV-Measurement-A-202507251000000Z",
    "type":"PhotovoltaicMeasurement",
    "dateObserved":"2025-07-25:10:00.000Z",
    "location": {
      "type":"Point",
      "coordinates": [
        50.774783,
        6.084345
      ]
    },
    "nominalPeakPowerGeneration":18,
    "temperature":23.4,
    "refPhotovoltaicDevice":"urn:ngsi-ld:PhotovoltaicDevice:PV-Device-A-001",
    "@context":[
      "https://raw.githubusercontent.com/smart-data-models/dataModel.GreenEnergy/master/context.jsonld"
    ]
  },
  {
    "id":"urn:ngsi-ld:PhotovoltaicMeasurement:PV-Measurement-B-202507251030000Z",
    "type":"PhotovoltaicMeasurement",
    "dateObserved":"2025-07-25:10:30.000Z",
    "location": {
      "type":"Point",
      "coordinates": [
        50.789466,
        6.049514
      ]
    },
    "nominalPeakPowerGeneration":15,
    "temperature":22.8,
    "refPhotovoltaicDevice":"urn:ngsi-ld:PhotovoltaicDevice:PV-Device-B-001",
    "@context":[
      "https://raw.githubusercontent.com/smart-data-models/dataModel.GreenEnergy/master/context.jsonld"
    ]
  }
]
```

Listing 6: List of federated PhotovoltaicMeasurement data entities detail.

According to established CSRs, when a local context is updated by the corresponding FP node, the change is subsequently reflected in the aggregated F node consultation output.

Eventual requests requiring details on the PV devices that provide the measurements yield an unauthorized access error, according to the configured policies (see Section 4.2 for more details).

5 Final remarks and discussion

Following the detailed discussion of the case study, some considerations can be formulated and discussed. Firstly, the main contribution provided by the case study consists in demonstrating the robust data governance of the proposed essential federation architecture by ensuring data sovereignty and privacy through policy enforcement mechanisms. At the same time, seamless interoperability can be achieved using compliance with NGSI-LD interfaces.

In effect, the developed federated architecture designates the F node as a passive context data aggregator, characterized by read-only access to entities managed by FP nodes, as discussed in Section 3.2. This approach prioritizes decentralization and data sovereignty, preventing the F node from attempts to data-changing operations on the entities, and thus upholding the autonomy of local context sources.

By complying with EU specifications, this architecture supports interoperability and secure federation among DS participants. The adoption of uniform standards fosters seamless integration of diverse participants, supporting a heterogeneous environment that promotes collaboration and secure data sharing. This flexibility is crucial for scaling DTEs and promoting widespread adoption.

A key requirement for achieving functional compatibility is support for the NGSI-LD API, which enables CIM and data exchanges. Its adoption is mandatory for mechanisms like CSRs, otherwise serious limitations may occur in terms of integration, trust, and policy enforcement within the federation.

Federated DS-based implementations contribute to secure DTE development by enabling multiple DTs and data sources to collaborate while maintaining decentralized data control. This approach can improve data sovereignty, allowing organizations to retain ownership and control over sensitive information, thereby reducing cyberattack risks associated with centralized repositories.

Furthermore, federated DS frameworks also integrate advanced identity and access management to enable secure interactions, including DIDs and trust frameworks. These trust frameworks ensure transparency, accountability, and regulatory compliance, such as GDPR, and govern sensitive data exchanges with protective policies. They also incorporate security measures such as cryptography and policy enforcement to prevent malicious activities and data tampering. As NGSI-LD compliant enablers, CSRs standardize data source description and discovery, enabling seamless context data integration and interoperability across federated DTEs. These aspects are critical for effective operation and informed decision-making in interconnected DTs, while also ensuring adaptability and manageability across diverse data sources and applications.

ODRL policies complement CSRs by specifying usage rights and access controls for context data, protecting against unauthorized access. When a CSR registers a CS, associated ODRL policies define access conditions, permissions, and constraints for data consumers. These policies are enforced by access control mechanisms to guarantee that only authorized entities can access sensitive context information. This approach integrates policy-based governance with registration, maintaining data sovereignty and privacy while enabling secure, policy-compliant federation of context data across DTs ecosystems.

The case study discussed in this work highlights the data sharing process that enables and supports real-time scenarios such as anomaly detection, real time system monitoring, and performance optimization, facilitating secure and trustworthy exchanges of contextual information for enhanced operational effectiveness.

In summary, the case study demonstrates how secure, scalable, and interoperable data sharing in federated DTEs can be enabled according to the proposed approach, preserving data sovereignty and privacy through ODRL-based policies. The approach can be extended to other disparate domains, such as smart cities and industrial automation, enabling trust and interoperability among different DT environments. Again in these domains, compliance with NGSI-LD and the adoption of decentralized IAM frameworks are the key enablers for the federation.

6 Conclusion

The growing complexity of DT architectures requires the urgent need for mechanisms to facilitate secure and trustworthy federation between DTEs. This paper examines the role of DSs in enabling and implementing this kind of mechanisms. In effect, DSs can significantly contribute to implement collaboration across organizational boundaries and uphold data sovereignty and security through advanced policy and trust mechanisms. This contribution can also be meant in terms of standardized frameworks for data exchange and governance.

This work generalizes the federated model adopted in decentralized identity authority and extends it to context data management. In this way, the resulting architecture enables secure, trust-based, and interoperable sharing of contextual information across federated DTEs. This result becomes possible by complying with the NGSI-LD's CSR to enable the federation of CSs.

The case study details a federated data exchange architecture within DTEs by implementing both the NGSI-LD standards and the ODRL-based policy governance. The context federation lifecycle involves multiple CSRs, with access policies governing data sharing, including creation grants for CSR owners and data-specific inclusion and exclusion restrictions. Federated context providers (FP nodes) authenticate and establish secure interactions with federation nodes (F nodes), which enforce access control policies to facilitate trustworthy and real-time data exchange.

In future work, it is first planned to investigate the implications of integrating the IDS's Transfer Process Protocol specifications [15] in order to further control the data stream between federated participants. Finally, some studies regarding the evaluation of the federation efficiency are in the work-plan aimed to explore quantitative assessment, scalability analysis, and the development of specific algorithms.

7 Acknowledgments

This research was partially supported by grant from the Italian Research Center on High Performance Computing, Big Data and Quantum Computing (ICSC) funded by EU-NextGenerationEU (PNRR-HPC, CUP: C83C22000560007), SCIAME project (CUP: F89J22003510004), FIWARE Foundation and D.M. n. 117/2023 (CUP: 9311).

References

[1] I. K. B. Ababio, J. Bieniek, M. Rahouti, and et al. 2025. A Blockchain-Assisted Federated Learning Framework for Secure and Self-Optimizing Digital Twins in Industrial IoT. *Future Internet* 17, 1 (Jan. 2025), 13. https://doi.org/10.3390/fi17010013

[2] M. Aghaabbasi and S. Sabri. 2025. Potentials of digital twin system for analyzing travel behavior decisions. *Travel Behaviour and Society* 38 (Jan. 2025), 100902. https://doi.org/10.1016/j.tbs.2024.100902

[3] T. Dam, L. Daniel Klausner, S. Neumaier, and T. Priebe. 2023. A Survey of Dataspace Connector Implementations. (2023). https://doi.org/10.48550/ARXIV.2309.11282

[4] ETSI Group Specification. 2025. Context Information Management (CIM); NGSI-LD API Version 1.9.1. https://www.etsi.org/deliver/etsi_gs/CIM/001_099/009/01.09.01_60/gs_CIM009v010901p.pdf. [Accessed 24-07-2025].

[5] European Parliament. 2016. General Data Protection Regulation (GDPR). https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679 Accessed: 2025-07-08.

[6] FIWARE Foundation e.V. 2020. ScorpioBroker. https://scorpio.readthedocs.io/en/latest/introduction.html. [Accessed 25-07-2025].

[7] FIWARE Foundation e.V. 2025. FIWARE Data Space Connector. https://github.com/FIWARE/data-space-connector. [Accessed 28-07-2025].

[8] Noardo F., R. Atkinson, L. Bastin, and et al. 2024. Standards for Data Space Building Blocks. *Remote Sensing* 16, 20 (Oct. 2024), 3824. https://doi.org/10.3390/rs16203824

[9] B. Farahani and A. K. Monsefi. 2023. Smart and collaborative industrial IoT: A federated learning and data space approach. *Digital Communications and Networks* 9, 2 (2023), 436–447.

[10] Gaia-X. 2025. Gaia-X Trust Framework - Gaia-X Trust Framework - Version fb420580. https://gaia-x.gitlab.io/policy-rules-committee/trust-framework/gaia-x_trust_framework/. [Accessed 25-07-2025].

[11] J. Gil, D. Petrova-Antonova, and G. JL Kemp. 2024. Redefining urban digital twins for the federated data spaces ecosystem: A perspective. *Environment and Planning B: Urban Analytics and City Science* (Dec. 2024). https://doi.org/10.1177/23998083241302578

[12] P. Gronlier, J. Hierro, and S. Steinbuss. 2023. Data Spaces Business Alliance Technical Convergence - Discussion Document Version 2.0. https://data-spaces-business-alliance.eu/wp-content/uploads/dlm_uploads/Data-Spaces-Business-Alliance-Technical-Convergence-V2.pdf. [Accessed 24-07-2025].

[13] International Data Spaces Association. 2024. Dataspace Protocol. https://docs.internationaldataspaces.org/ids-knowledgebase/dataspace-protocol. [Accessed 25-07-2025].

[14] International Data Spaces Association. 2025. IDS-RAM v5.0 working draft. https://docs.internationaldataspaces.org/ids-ram-5-working-draft. [Accessed 31-07-2025].

[15] International Data Spaces Association. 2025. IDS Transfer Process Protocol Specification. https://docs.internationaldataspaces.org/ids-knowledgebase/dataspace-protocol/transfer-process/transfer.process.protocol. [Accessed 01-08-2025].

[16] International Data Spaces Association. 2025. IDSA Rulebook - Creating a Data Space. https://docs.internationaldataspaces.org/ids-knowledgebase/idsa-rulebook/idsa-rulebook/3.-functional_requirements/3.5-creating_a_data_space. [Accessed 25-07-2025].

[17] M. Jurmu, I. Niskanen, A. Kinnula, and et al. 2023. Exploring the Role of Federated Data Spaces in Implementing Twin Transition within Manufacturing Ecosystems. *Sensors* 23, 9 (April 2023), 4315. https://doi.org/10.3390/s23094315

[18] Y. Li, T. Li, K. Xu, and et al. 2025. Federated Digital Twin-Empowered Online Control and Optimization for Cyber-Physical Systems. *IEEE Transactions on Industrial Cyber-Physical Systems* (2025), 1–11. https://doi.org/10.1109/ticps.2025.3586992

[19] S. Liu, J. Z. Huang, P. P.C. Lee, and G. C. Polyzos. 2025. Incentivized Federated Learning in Data Spaces. In *2025 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*. IEEE Computer Society, 410–415.

[20] Y. Liu, B. Zhao, Z. Zhao, and et al. 2024. SS-DID: A Secure and Scalable Web3 Decentralized Identity Utilizing Multilayer Sharding Blockchain. *IEEE Internet of Things Journal* 11, 15 (Aug. 2024), 25694–25705. https://doi.org/10.1109/jiot.2024.3380068

[21] R. Liyanage, N. Tripathi, T. Päivärinta, and Y. Xu. 2022. Digital twin ecosystems: Potential stakeholders and their requirements. In *International Conference on Software Business*. Springer, 19–34.

[22] D. Maram, H. Malvai, F. Zhang, and et al. 2021. CanDID: Can-Do Decentralized Identity with Legacy Compatibility, Sybil-Resistance, and Accountability. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1348–1366. https://doi.org/10.1109/sp40001.2021.00038

[23] A. Martella, C. Martella, and A. Longo. 2025. Designing Data Spaces: Navigating the European Initiatives Along Technical Specifications. *arXiv preprint arXiv:2503.15993* (2025).

[24] A. Martella, A.I.H.A Ramadan, C. Martella, and et al. 2023. *State of the Art of Urban Digital Twin Platforms*. Springer Nature Switzerland, 299–317. https://doi.org/10.1007/978-3-031-43401-3_20

[25] C. Martella, A. Martella, and A. Longo. 2024. European data spaces for urban digital twins: user-and implementation-driven recommendations. In *2024 IEEE International Conference on Big Data (BigData)*. IEEE, 5496–5505. https://doi.org/10.1109/bigdata62323.2024.10826100

[26] C. Martella, A. Martella, and A. I. H. A. Ramadan. 2023. Identifying key factors in designing data spaces for urban digital twin platforms: a data driven approach. In *2023 IEEE International Conference on Big Data (BigData)*. IEEE, 3985–3994.

[27] A. Mitrovska, M. F. Rosas, P. Safari, and et al. 2024. Federated Learning Governance using Eclipse Dataspace Components Connectors. In *2024 IEEE International Conference on Big Data (BigData)*. IEEE, 7946–7954.

[28] F. Möller, I. Jussen, V. Springer, and et al. 2024. Industrial data ecosystems and data spaces. *Electronic Markets* 34, 1 (Aug. 2024). https://doi.org/10.1007/s12525-024-00724-0

[29] N. Moretti, X. Xie, J. Merino Garcia, J. Chang, and A. Kumar Parlikad. 2023. Federated Data Modeling for Built Environment Digital Twins. *Journal of Computing in Civil Engineering* 37, 4 (July 2023). https://doi.org/10.1061/jccee5.cpeng-4859

[30] B. Otto, M. T. Hompel, and S. Wrobel. 2019. *International Data Spaces: Reference architecture for the digitization of industries*. Springer Berlin Heidelberg, 109–128. https://doi.org/10.1007/978-3-662-58134-6_8

[31] J. Pampus, B. F. Jahnke, and R. Quensel. 2022. *Evolving Data Space Technologies: Lessons Learned from an IDS Connector Reference Implementation*. Springer Nature Switzerland, 366–381. https://doi.org/10.1007/978-3-031-19762-8_27

[32] H. Pettenpohl, M. Spiekermann, and J. R. Both. 2022. *International Data Spaces in a Nutshell*. Springer International Publishing, 29–40. https://doi.org/10.1007/978-3-030-93975-5_3

[33] B. Potteiger, F. Cai, Z. Zhang, and X. Koutsoukos. 2022. Data space randomization for securing cyber-physical systems. *International Journal of Information Security* 21, 3 (2022), 597–610.

[34] A. Rizwan, R. Ahmad, A. N. Khan, and et al. 2023. Intelligent digital twin for federated learning in AIoT networks. *Internet of Things* 22 (July 2023), 100698. https://doi.org/10.1016/j.iot.2023.100698

[35] M. M. Salim, D. Camacho, and J. H. Park. 2024. Digital Twin and federated learning enabled cyberthreat detection system for IoT networks. *Future Generation Computer Systems* 161 (Dec. 2024), 701–713. https://doi.org/10.1016/j.future.2024.07.017

[36] H. Diogo Silva, M. Azevedo, and A. L. Soares. 2021. A Vision for a Platform-based Digital-Twin Ecosystem. *IFAC-PapersOnLine* 54, 1 (2021), 761–766. https://doi.org/10.1016/j.ifacol.2021.08.088

[37] A. Somma, A. De Benedictis, M. Zappatore, and et al. 2023. Digital Twin Space: The Integration of Digital Twins and Data Spaces. *2023 IEEE International Conference on Big Data (BigData)*, 4017–4025. https://doi.org/10.1109/BigData59044.2023.10386737

[38] Q. Song, S. Lei, W. Sun, and Y. Zhang. 2021. Adaptive Federated Learning for Digital Twin Driven Industrial Internet of Things. In *2021 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 1–6. https://doi.org/10.1109/wcnc49053.2021.9417370

[39] M. M. Thwe, A. Ştefanov, V. S. Rajkumar, and P. Palensky. 2025. Digital Twins for Power Systems: Review of Current Practices, Requirements, Enabling Technologies, Data Federation, and Challenges. *IEEE Access* 13 (2025), 105517–105540. https://doi.org/10.1109/access.2025.3580055

[40] Y. Wang, Z. Su, S. Guo, and et al. 2023. A Survey on Digital Twins: Architecture, Enabling Technologies, Security and Privacy, and Future Prospects. *IEEE Internet of Things Journal* 10, 17 (Sept. 2023), 14965–14987. https://doi.org/10.1109/jiot.2023.3263909

[41] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, and et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data* 3, 1 (2016), 1–9.

[42] X. Yan. 2021. Construction of digital twin ecosystem for coal-fired generating units. *Journal of Physics: Conference Series* 1748, 5 (Jan. 2021), 052037. https://doi.org/10.1088/1742-6596/1748/5/052037

[43] L. Zhang, Z. Wu, H. Xu, and et al. 2025. Digital Twin-Driven Federated Learning for Converged Computing and Networking at the Edge. *IEEE Network* 39, 2 (March 2025), 20–28. https://doi.org/10.1109/mnet.2024.3504520