Context-driven Edge-based Data Sharing for Industrial IoT Data Spaces

An Ngoc Lam
SINTEF AS
Oslo, Norway
an.lam@sintef.no

Phu H. Nguyen
SINTEF AS
Oslo, Norway
phu.nguyen@sintef.no

Xiang Ma
SINTEF AS
Oslo, Norway
xiang.ma@sintef.no

Abstract

The increasing reliance on IoT ecosystems demands robust, secure, and context-aware data-sharing mechanisms that operate closer to data sources. Data spaces must leverage trusted edge-based system architectures for near real-time data processing, transformation, and enrichment while ensuring data privacy and security. However, the current International Data Spaces (IDS) Model lacks comprehensive support for Edge-based architectures and flexible, context-driven access control models essential for managing diverse applications within data space ecosystems.

To address these gaps, we propose IDS4Edge, an IDS-compliant approach that enables dynamic, context-driven, Edge-based IoT data sharing as a service. IDS4Edge integrates flexible access control policies on top of IDS connectors, tailored to specific IoT application contexts. These policies dynamically adapt in real-time to changes in IoT contexts and contractual agreements, ensuring secure and efficient data sharing at the Edge.

We validate our solution through a proof-of-concept implementation, demonstrating how IDS4Edge facilitates trusted, scalable, and real-time data sharing while maintaining compliance with IDS principles. This approach paves the way for enhanced (industrial) IoT applications and advanced data-sharing paradigms, such as Manufacturing-as-a-Service (MaaS).

CCS Concepts

• **Information systems** → **Data exchange**; • **Computer systems organization** → *Peer-to-peer architectures.*

Keywords

Access Control Policies, Context-Aware Data Sharing, Data Spaces, Digital Twins, Manufacturing as a Service (MaaS)

ACM Reference Format:
An Ngoc Lam, Phu H. Nguyen, and Xiang Ma. 2025. Context-driven Edge-based Data Sharing for Industrial IoT Data Spaces. In *The 40th ACM/SIGAPP Symposium on Applied Computing (SAC '25), March 31-April 4, 2025, Catania, Italy*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3672608.3707988

1 Introduction

The Internet of Things (IoT) [33] has revolutionized the integration of digital and physical entities, encompassing devices, individuals, and services. This integration has led to the formation of complex ecosystems that facilitate secure interactions across various domains. The Industrial Internet of Things (IIoT) extends IoT technologies into industrial production environments, enabling automation, predictive maintenance, and operational efficiency [15, 17]. As these ecosystems expand, they generate vast amounts of data, necessitating robust mechanisms for secure sharing and prudent utilization to uphold the confidentiality, integrity, and availability of sensitive information [18]. However, managing data sharing within the IoT framework presents multifaceted challenges.

Firstly, IoT systems are inherently complex, involving distributed sensing, actuation, and processing across multiple layers, including devices (Things), Edge, and Cloud resources [33]. Devices at the Thing and Edge levels operate in the physical world, collecting environmental data often including sensitive information. This data frequently crosses organizational boundaries and international borders, introducing a web of legal and compliance requirements.

Moreover, the diverse and interconnected nature of IoT ecosystems brings together various stakeholders, including device manufacturers, service providers, enterprises, and end-users, each with unique rights and responsibilities concerning data. This diversity enhances the value of IoT ecosystems by enabling seamless data sharing between enterprises, consumers, suppliers, and other stakeholders, thereby unlocking significant opportunities for collaboration and innovation [14, 22]. However, these benefits come with challenges, as ensuring security [31, 32], fostering trust among participants [2, 36, 40], and maintaining data quality [9, 21, 25, 26, 35] remain critical to sustainable IoT growth.

The International Data Spaces Reference Architecture Model (IDS RAM) [11], developed by the International Data Spaces Association (IDSA) [13], provides a standardized framework for establishing trust in data sharing. It addresses the critical aspects of data governance, including confidentiality, integrity, and interoperability among stakeholders. However, IDS RAM does not explicitly support Edge computing layers or dynamic application contexts, which are crucial for the real-time, context-driven nature of IoT applications. As IoT ecosystems grow, they increasingly rely on Edge computing to process and share data closer to the source, reducing latency and enabling near real-time decision-making. To facilitate Edge analytics and avoid the complexities associated with data security and privacy, data exchange must be performed at the Edge. This approach allows stakeholders from various sectors to utilize the data for context-specific business applications, unlocking significant opportunities within IoT ecosystems.

Despite the growing importance of secure IoT data sharing, Edge-focused data sharing remains under-explored and insufficiently supported by standardization bodies like IDSA. Existing studies have

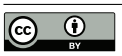

SAC '25, March 31-April 4, 2025, Catania, Italy

ACM ISBN 979-8-4007-0629-5/25/03
https://doi.org/10.1145/3672608.3707988

explored related topics such as data governance [1], blockchain solutions [20], and secure IoT data sharing frameworks [4], but the current IDS protocol still lacks native mechanisms for fully integrating Edge-based architectures. Furthermore, it does not address the need for flexible, context-driven access control models, which are essential for enabling dynamic policies tailored to decentralized and real-time IoT scenarios.

In this work, we propose **IDS4Edge**, an IDS-compliant solution for dynamic, context-driven access control and Edge-based Industrial IoT data sharing. Our approach focuses on enabling context-driven access control at the Edge layer while ensuring alignment with the IDS RAM. For future work, it should be possible to build on the foundation of our work in this paper to adopt more advanced access control models, such as delegation [24] or usage control models [29]. Our main contributions are:

- **Edge Integration with IDS:** We propose a solution to integrate IDS principles with Edge computing to enable real-time, secure IoT data sharing at the Edge.
- **Support for Manufacturing-as-a-Service (MaaS) Applications:** Our framework extends IDS capabilities to support real-time, context-aware operations in MaaS scenarios.
- **Proof-of-Concept Implementation:** We validate IDS4Edge through a prototype demonstrating real-time policy enforcement and secure stakeholder interactions.

The remainder of this paper is organized as follows: Section 2 provides key background concepts, followed by a motivational example in Section 3. Section 4 presents our IDS4Edge approach, while Section 5 describes its proof-of-concept implementation. Related work is discussed in Section 6, and conclusions along with potential future work are presented in Section 7.

2 Background

2.1 IDSA and IDS Connector

A significant organization in the data exchange domain is the International Data Spaces Association (IDSA) [13], which aims to innovate the future of data exchange in Europe and beyond by establishing crucial technical standards. IDSA has established a framework for enabling secure and trusted data sharing across organizations while maintaining data sovereignty. At the core of this framework is the Reference Architecture Model (RAM) [11], which defines the key components, roles, and interactions necessary for establishing trusted data ecosystems.

Figure 1 illustrates the interaction between key components within the IDS framework, emphasizing the flow of data and metadata between participants in a data space. At the core of this architecture are the *IDS Connectors*, which serve as interfaces for secure data exchange between the *Data Provider* and *Data Consumer*. These connectors enforce security policies, manage access control, and ensure compliance with usage restrictions defined in contractual agreements. A comprehensive list of IDS Connector implementations is available in [8]. In this work, we utilized the Eclipse Data Space (EDC) Framework[1], an open-source solution that incorporates the latest IDS protocol [12].

[1]https://eclipse-edc.github.io/

In addition to the IDS Connector, the IDS RAM includes various supporting components such as the *Identity Provider*, which authenticates participants and establishes trust relationships, and the *Clearing House*, which logs transactions to ensure accountability and traceability. The *Metadata Broker* facilitates the discovery of available data by indexing and providing metadata descriptions, while the *Vocabulary Hub* ensures semantic interoperability by standardizing terminologies and data formats used in the data space. The *App Store* provides a secure platform to distribute *IDS Data Apps*, supporting operations such as app registration, publication, and provisioning to the IDS Connectors within the data space.

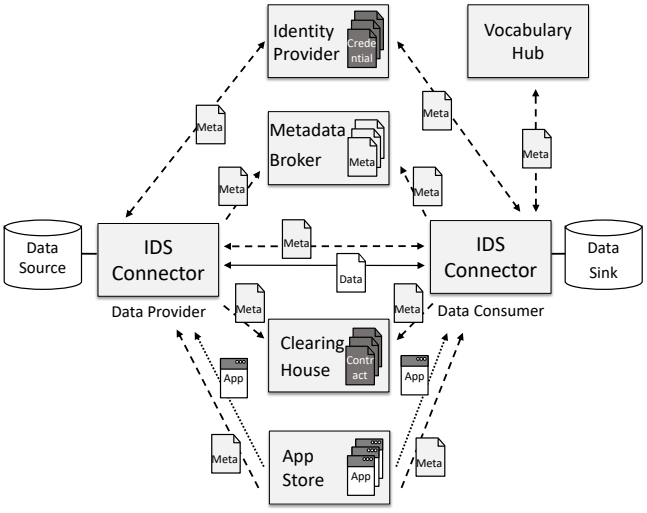

Figure 1: Interaction of technical components in IDSA RAM (Adapted from [11]).

2.2 GAIA-X

GAIA-X [5] is a European initiative designed to establish a federated and secure data infrastructure, promoting data sovereignty, innovation, and interoperability across industries. Rooted in European values of trust, transparency, and self-determination, GAIA-X provides a common framework for seamless and secure data exchange while maintaining control over data access and usage.

At its core, GAIA-X emphasizes building trust through transparency, ensuring that participants maintain full control over their data while enabling secure and reliable interactions. A fundamental aspect of this effort is the GAIA-X Trust Framework [6], which outlines the policies and requirements for establishing a secure and verifiable environment within the data ecosystem. The IDSA has adopted this framework, integrating Decentralized Identifiers (DIDs) [37] to strengthen identity and access management. DIDs provide a self-sovereign and secure method for identifying entities, aligning with the principles of data sovereignty and trust central to both GAIA-X and IDSA.

2.3 IoT architecture and Edge Computing as basis for data sharing

The IoT architecture is commonly structured using a seven-layer model as defined by the IoT World Forum Reference Model [3]. These layers include:

- **Physical Devices & Controllers:** Includes all the "Things" in IoT, such as machines, sensors, and devices.
- **Connectivity:** Responsible for communication and processing units.
- **Edge Computing:** Handles data analysis and transformation at the edge of the network.
- **Data Accumulation:** Manages data storage for subsequent processing.
- **Data Abstraction:** Facilitates data aggregation and access.
- **Application:** Supports reporting, analytics, and control functionalities.
- **Collaboration & Processes:** Involves people and business processes to enable decision-making and collaboration.

Edge Computing layer plays a pivotal role in enabling real-time processing and transformation of data close to its source, bridging the gap between physical devices and higher-level applications [23]. By operating at the edge of the network, it reduces latency, improves efficiency, and enhances privacy [30]. This approach provides lightweight solutions for local, small-scale data storage and processing, forming the foundation for Edge-based data sharing-as-a-service models. These capabilities are particularly valuable when integrated into standardized data spaces, enabling more efficient and secure data ecosystems.

One of the key challenges in Edge-based data sharing is ensuring *real-time, context-sensitive* access control for diverse stakeholders. Each stakeholder operates under distinct contractual agreements, which may depend on the dynamic context of the data. Since these services are often automated in real-time, access control mechanisms must be equally dynamic, adapting instantaneously to changes in context and contractual conditions to maintain secure and efficient operations.

2.4 Manufacturing as a service (MaaS)

MaaS is an emerging business model that leverages digital technologies and cloud-based platforms to provide manufacturing resources and capabilities on demand. Instead of requiring companies to own and maintain physical production facilities, MaaS allows them to access advanced manufacturing processes as a service, enabling cost-effective, flexible, and scalable operations [39].

MaaS heavily relies on data-driven technologies like Digital Twins (DTs), IoT devices, and Edge computing to streamline operations and facilitate real-time decision-making. In such systems, data from sensors, machines, and other connected devices is collected and analyzed to monitor production progress, ensure quality, and improve operational efficiency [28]. For instance, IoT-enabled factories can integrate their physical and digital assets to provide customers with up-to-date information about production progress, material status, and delivery schedules.

3 Motivational Example

In this section, we introduce a training factory as a representative example of the MaaS model. We utilize the Fischertechnik Training Factory Industry 4.0 24V[2], which simulates an end-to-end process encompassing ordering, manufacturing, and delivery. This simulation incorporates various pieces of equipment, including the following:

- **Highbay Warehouse (HBW)**: Serves as a storage facility within the factory, housing raw materials of types red, blue, and white.
- **Robot (VGR)**: Represents the factory's automation system, handling tasks such as material handling and product assembly.
- **Sensor Unit (SSC)**: Consists of sensors distributed throughout the factory, collecting real-time data on parameters such as cameras, temperature, pressure, and machine statuses.
- **Delivery and Pickup (DPS)**: Manages the logistics, including the reception of raw materials from suppliers and the pickup of finished products for customers.
- **Multi-Processing Station (MPO)**: Simulated furnace for handling the production process of raw materials.
- **Sorting Line (SLD)**: Handles the sorting and categorization of finished products based on its color.

The Training Factory generates numerous time-series outputs from its physical assets, which communicate using OPC-UA [16] or MQTT through two corresponding gateways. A detailed list of the 37 time-series values is provided in Table 1.

Table 1: Sensor data collected from the Factory.

Assets	Sensor Values	Protocol
SSC	14 values for temperature, humidity, air quality and pressure, camera positions and its image	MQTT
MPO	2 values for MPO status	OPC-UA
SLD	2 values for SLD status	OPC-UA
HBW	6 values for HBW status and positions	MQTT
DPS	4 values for DPS status	OPC-UA
VGR	9 values for VGR status and positions	OPC-UA

The factory is integrated with a Digital Twin (DT) implemented as an instance of the DT platform known as SINDIT [43]. SINDIT synchronizes in real time with the physical assets, storing all relevant data in a time-series database and offering various REST APIs for real-time data access. Details about the implementation of the Fischertechnik Factory's DT are discussed in [19]. Listing 1 and 2 provide examples of the data returned by the SINDIT APIs for the inventory status of raw materials of type blue and the temperature sensor respectively.

The highly data-driven processes and increasingly interconnected physical assets present new opportunities for enabling data and asset sharing through the integration of DT and Edge-Cloud computing technologies. These technologies serve as the foundation for implementing the MaaS model. To illustrate this concept, we present a simplified example of how IDS-compliant, Edge-based IoT data sharing can support MaaS. In this scenario, the DT developed by SINDIT serves as an intermediate layer between the Physical Twin and the IDS, providing its services to external users via the IDS protocol. This enables an external customer, based on a

[2]https://www.fischertechnik.de

predefined contract, to access real-time data or trigger production processes within the factory.

The MaaS platform facilitates the sharing of real-time updates on production progress, process quality, and the production footprint. Customers can place orders and initiate production while ensuring that production-related data is shared exclusively with them. Sensor and camera data are made available only to customers whose orders are currently being processed. Similarly, suppliers are granted access solely to inventory information related to the specific type of materials they provide (e.g., red, blue, or white).

Listing 1: Data Value for Blue Raw Material.

```
{
    "class_uri": "urn:samm:sindit.sintef.no:1.0.0#
         StreamingProperty",
    "uri": "http://sindit.sintef.no/2.0#blueRawMaterial",
    "label": "Blue Raw Material",
    "propertyDescription": "Stock information of blue raw
         materials in storage",
    "propertyDataType": {
        "uri": "http://www.w3.org/2001/XMLSchema#integer"
    },
    "propertyValue": 1,
    "propertyValueTimestamp": "2024-11-25T12:39:39.946Z",
    "propertyConnection": {
        "uri": "http://sindit.sintef.no/2.0#mqtt-connection"
    },
    "streamingTopic": "inventory/rawMaterials",
    "streamingPath": "data['blue']"
}
```

Listing 2: Data Value for Temperature sensor.

```
{
    "class_uri": "urn:samm:sindit.sintef.no:1.0.0#
         StreamingProperty",
    "uri": "http://sindit.sintef.no/2.0#temperature",
    "label": "Temperature",
    "propertyUnit": {
      "uri": "urn:samm:org.eclipse.esmf.samm:unit:2.1.0#
           degreeCelsius"
    },
    "propertyDescription": "Temperature data from the sensor",
    "propertyDataType": {
      "uri": "http://www.w3.org/2001/XMLSchema#float"
    },
    "propertyValue": 21.9,
    "propertyValueTimestamp": "2024-10-18T14:35:37.307751+02:00"
         ,
    "propertyConnection": {
      "uri": "http://sindit.sintef.no/2.0#mqtt-connection"
    },
    "streamingTopic": "i/bme680",
    "streamingPath": "data['t']"
}
```

Despite its immense potential, several key challenges must be addressed to fully realize MaaS's capabilities:

- **Trustworthy Data Sharing:** Establishing reliable and secure data sharing mechanisms among diverse stakeholders, including MaaS owners, customers, and suppliers.
- **Security and Privacy:** Ensuring robust security and privacy across the platforms involved in the MaaS ecosystem.
- **Real-Time Process Support:** Supporting real-time MaaS operations through flexible, context-aware access control, seamless data sharing, and effective data management and governance across the shop-floor (OT-IT) and Edge-Cloud computing continuum.

In this work, we address these critical challenges by enhancing IDS technological components with the advanced functionalities provided by the DT.

4 Approach

Figure 2 shows the overall IDS4Edge architecture. To develop the data space, we employed the EDC Framework, which is an open-source Java framework for implementing the IDS RAM [11]. Our solution is based on a decentralized architecture built around Edge Hubs, which act as local entities managing data collection, processing, and sharing at the edge of the network. These Edge Hubs correspond to participants of the data space (e.g., customers, suppliers, and factories), and each is equipped with an **Identity Hub**[3], a **IDS Connector**[4], and different **Data Assets**. The Identity Hub handles identity management, ensuring that data exchange within the system is secure and that only authorized entities participate in the data space. The IDS Connector ensures that all data transactions follow the IDSA standards for security, privacy, and data sovereignty. Each IDS Connector provides services for the Data Provider to publish its Data Assets to the **Local Catalog** and for the Data Consumer to discover and utilize these Data Assets. The **Federated Catalog**[5] serves as an aggregated catalog for all participants within the data space. It periodically collects the registries from the Local Catalogs and consolidates them in a local cache, facilitating the discovery of Data Assets across the entire data space.

The participants of the data space (i.e., customer, supplier or factory) host their own Data Assets in the backend and make their endpoints available for sharing data via the IDS Connector. Details of data description will be discussed in the Section 5. The **Factory** participant is the main component where MaaS is realized. As discussed earlier, the factory offers its production capabilities as a service to external consumers. These consumers can access real-time data on the manufacturing process, track order progress, and monitor quality through the **SINDIT Digital Twin**. The Factory participant ensures that only relevant data is shared with authorized users, facilitated by the IDS Connector, which follows strict access control policies. For example, customers can access data related to their specific orders, including sensor data, camera feeds, and status updates, while suppliers may only access inventory data for the materials they provide. The edge-based architecture enables near real-time data processing and sharing, ensuring that data is available when needed without compromising security or privacy. In order to validate these real-time data sharing policy, the Factory participant is extended with **IDS4Edge Extension**. This extension consists of the **Local Context-Based Policy Enforcer** and the **Contextual Information Database**.

The Contextual Information Database collects real-time contextual data from the factory's operational environment through the DT. This data includes production status, order status, and machine conditions, and other relevant operational details. It is utilized to enable context-aware decision-making for data access. For example, access to specific production data can be granted or restricted based on the current order status or the role of the requesting participant

[3]https://github.com/eclipse-edc/IdentityHub
[4]https://github.com/eclipse-edc/Connector
[5]https://github.com/eclipse-edc/FederatedCatalog

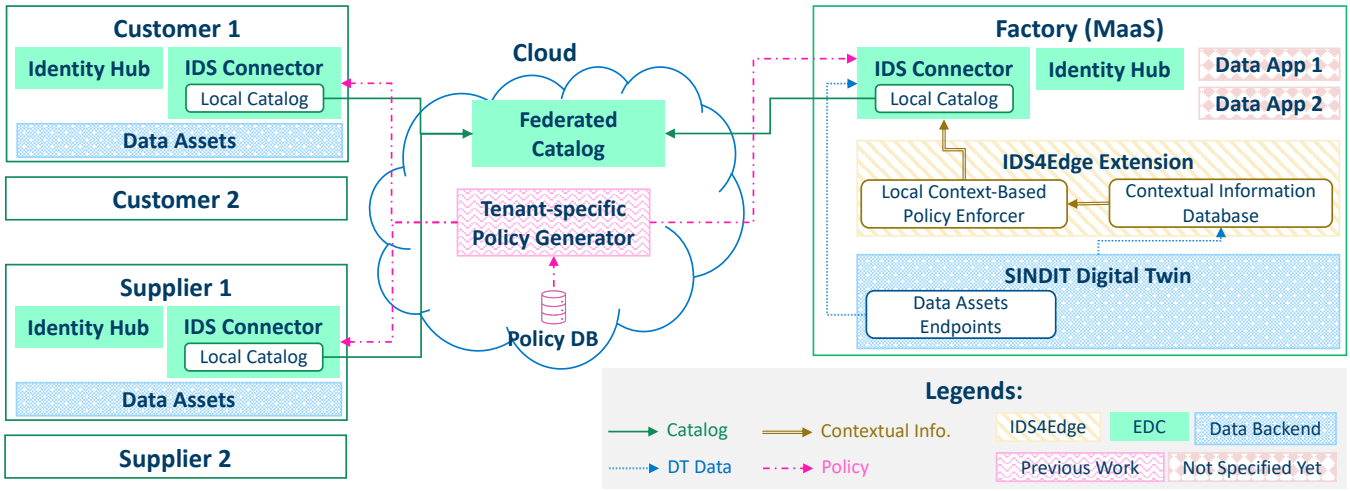

Figure 2: IDS4Edge Architecture.

(e.g., customer, supplier). By integrating context into access control decisions, the system ensures that only relevant data is shared with the appropriate stakeholders at the right time.

The Local Context-Based Policy Enforcer is responsible for validating dynamic access control rules based on the current context, ensuring that data is shared in compliance with specific agreements and business rules. It works by enforcing policies directly at the edge, where the data is generated, allowing for real-time access as conditions evolve.

The **Policy Database** and the **Tenant-specific Policy Generator** are the components adapted from previous work [22]. The Policy Database maintains all access control rules in a secure and scalable manner, while the Tenant-specific Policy Generator dynamically generates policies tailored to specific users and contexts, ensuring compliance with contractual obligations and privacy regulations. Together with the Federated Catalog, these components will be deployed in the cloud to provide centralized management of policies and data discovery services. The cloud-based deployment ensures that these components can handle the scalability requirements of large, multi-tenant industrial IoT environments, while also enabling seamless integration with various edge nodes.

5 Proof of Concept

This section presents three proof-of-concept scenarios that demonstrate the real-time, context-driven data-sharing capabilities of the proposed IDS4Edge framework. These scenarios highlight interactions between suppliers, customers, and the factory, illustrating how real-time data sharing and access control policies are implemented. The framework is developed based on the EDC Minimum Viable Dataspace (MVD)[6] and is available in our GitHub repository[7].

In this IDS4Edge data space, each participant is issued two types of *Verifiable Credentials (VCs)* [38]:

- *Membership Credential:* Contains information about the holder's membership in the data space. A valid membership credential is required for the participant to join the data space and to query the catalog.
- *Actor Credential:* Specifies the participant's actual role (e.g., customer, supplier) within the data space. This credential is necessary for negotiating contracts and transferring data.

These credentials are secure certificates digitally signed by a trusted authority. Listing 3 shows an excerpt of an Actor Credential issued to a supplier. These credentials are stored in the participants' Identity Hubs and are used for authentication and authorization in the data space, as outlined in DID architecture [37].

Listing 3: Excerpt of a Verifiable Credential of a Supplier.

```
{
  "id": "40e24588-b510-41ca-966c-c1e0f57d1b16",
  "participantId": "did:web:localhost%3A7083",
  "verifiableCredential": {
    "credential": {
      "id": "http://org.yourdataspace.com/credentials/3456",
      "type": ["VerifiableCredential", "ActorCredential"],
      "issuer": { "id": "did:example:dataspace-issuer" },
      "credentialSubject": [
        {
          "id": "did:web:localhost%3A7083",
          "claims":
          {
            "actorId": "Supplier12345",
            "actorName": "Supplier Inc.",
            "actorType": "Supplier"
          }
        }
      ]
    }
  }
}
```

The Contextual Information Database[8] is implemented as a simple properties file. To support this, we introduced an extension module to the EDC MVD that provides an interface for the SINDIT

[6] https://github.com/eclipse-edc/MinimumViableDataspace
[7] https://github.com/SINTEF-9012/IDS4Edge-MVD
[8] https://github.com/SINTEF-9012/IDS4Edge-MVD/tree/main/extensions/ids4edge-context

Digital Twin to write contextual information to the file, and for the Local Context-Based Policy Enforcer[9] to read the necessary data for policy validation. However, leveraging the flexibility of the EDC framework, other database backends can easily replace the current implementation as long as they adhere to the same API interface.

5.1 Scenario 1: Supplier Integration

This scenario represents a regular data-sharing scenario involving suppliers and the Fischertechnik Factory. Suppliers are integrated into the factory's system, allowing them to monitor the inventory status of the materials they produce in real time.

Suppliers connect to the factory's system to access real-time data specific to the materials they provide. For instance, a supplier delivering raw materials of type "blue" can only access the data endpoint related to the inventory status of blue materials. As shown in Listing 1, the shared data primarily includes the real-time inventory status of specific materials, such as the available quantity and stock levels for different colors (e.g., red, blue, white). This data-sharing mechanism allows suppliers to proactively monitor inventory levels, enabling them to plan restocking without waiting for manual requests from the factory. The automated and seamless flow of information reduces delays and minimizes the risk of material shortages that could disrupt production.

Listing 4: Policy for inventory status of blue materials.

```
{
  "@context": [
    "https://w3id.org/edc/connector/management/v0.0.1"
  ],
  "@type": "PolicyDefinition",
  "@id": "require-blue-supplier",
  "policy": {
    "@type": "Set",
    "permission": [
      {
        "action": "use",
        "constraint": {
          "leftOperand": "ActorCredential.actorType",
          "operator": "eq",
          "rightOperand": "Supplier"
        }
      }
    ],
    "obligation": [
      {
        "action": "use",
        "constraint": {
          "leftOperand": "SupplierType",
          "operator": "eq",
          "rightOperand": "Blue"
        }
      }
    ]
  }
}
```

As an example, Listing 4 defines a policy that restricts inventory data for blue materials exclusively to the suppliers. Access is granted only to participants with an Actor Credential where the *actorType* is set to "Supplier" (this corresponds to the first condition of the policy). To validate this condition, the Policy Enforcer verifies the Actor Credential (example shown in Listing 3) retrieved from the requesting participant's Identity Hub to determine the *actorType* value. Additionally, the supplier must satisfy the second condition, which requires their *supplierType* to be "Blue." To validate this condition, the Policy Enforcer retrieves contextual information about the supplier type from the Contextual Information Database. These validations ensure that only authorized suppliers who meet both conditions can access the inventory data.

[9]https://github.com/SINTEF-9012/IDS4Edge-MVD/tree/main/extensions/ids4edge-policy-impl

5.2 Scenario 2: Real-Time Order Tracking

This scenario describes a contextual-based data-sharing scenario where customers, such as other manufacturers, interact with the Factory to track the real-time production of their orders. Customers are granted access to live production data based on predefined contextual policies. For instance, they can only view the relevant sensor and production data while their specific order is being processed. This ensures that access is tightly controlled and limited to the necessary time frame, enhancing data security and privacy while maintaining transparency in the production process.

Listing 5: Policy for real-time production data.

```
{
  "@context": [
    "https://w3id.org/edc/connector/management/v0.0.1"
  ],
  "@type": "PolicyDefinition",
  "@id": "require-actor-supplier",
  "policy": {
    "@type": "Set",
    "permission": [
      {
        "action": "use",
        "constraint": {
          "leftOperand": "ActorCredential.actorType",
          "operator": "eq",
          "rightOperand": "Customer"
        }
      }
    ],
    "obligation": [
      {
        "action": "use",
        "constraint": {
          "and": [
            {
              "leftOperand": "ProductionStatus",
              "operator": "eq",
              "rightOperand": "Active"
            },
            {
              "leftOperand": "ParticipantID",
              "operator": "eq",
              "rightOperand": "CustomerOrderID"
            }
          ]
        }
      }
    ]
  }
}
```

Customers can access various production details, such as order initiation time, production progress, and estimated delivery time. Contextual policies play a significant role here—customers only have access to this data during the active production of their orders. If a failure or issue arises during production, additional data, such as live camera feeds and sensor readings (e.g., temperature, humidity, air pressure), may be made available to provide deeper insights into the problem. This conditional access ensures that customers

are informed when necessary, while preventing overexposure of factory operations during normal production times.

Listing 5 defines an example policy for controlled access to real-time production data by customers in the data space. First, the requester must possess an Actor Credential that specifies their *actorType* as "Customer," ensuring that only customers can initiate the action. As previously discussed, this condition is validated by the Policy Enforcer by retrieving and verifying the Actor Credential from the participant's Identity Hub. Additionally, the policy includes an obligation requiring two conditions to be met: the *ProductionStatus* must be "Active," ensuring access is granted only during ongoing production, and the *ParticipantID* (i.e., the requester ID) must match the *CustomerOrderID*, confirming that the requester is authorized to access resources associated with their specific order. The Policy Enforcer validates these obligations by retrieving contextual information, such as the *ProductionStatus* and *CustomerOrderID*, from the Contextual Information Database. This policy enables fine-grained, context-aware access control, ensuring that customers can access production data only under conditions relevant to their role and active orders.

Upon the successful completion of the manufacturing process, customers receive a comprehensive report of their order. This includes product lifecycle data, material specifications, and quality assurance metrics, all of which are governed by the same contextual access control policies. The ability to access additional information in the event of an issue—while restricting access during regular production—provides customers with transparency and trust while safeguarding the factory's sensitive operations. This approach also exemplifies the factory's commitment to maintaining a secure and efficient MaaS environment.

5.3 Scenario 3: Digital Twin Optimization

This scenario illustrates an edge-based application deployment in which AI/ML model developers collaborate with the Factory DT to leverage edge-based computation for optimizing digital twins and enabling predictive maintenance. The key focus here is that AI developers have restricted access to only sampled data, ensuring the privacy and security of the factory's full dataset. However, the AI models themselves are trained, evaluated and deployed on the complete dataset, allowing for robust predictive capabilities while maintaining strict data access control. This approach balances the need for data protection with the requirement for comprehensive model training to ensure accuracy and reliability in the AI models.

The data involved includes real-time sampled sensor data, which is accessible to the AI developers for initial model development, alongside historical error logs used for enhancing predictive accuracy. The full data set—containing detailed sensor readings and machine status—is utilized by the models during training and evaluation, allowing for precise prediction of potential equipment failures. While this scenario relies on the Data App concept from the IDSA specification to deploy edge-based applications, it has not been fully implemented in the current proof-of-concept. The Data App structure presents a future direction for enabling flexible, edge-based AI applications that provide predictive maintenance without compromising the privacy and security of sensitive factory data.

6 Related Work

Surprisingly, very few primary studies have leveraged standardization efforts like the IDSA's reference architecture. GAIA-X has seen limited adoption, and only two studies [7, 34] have explored IoT data-sharing approaches using the IDSA architecture in the seaport sector. Gimenez et al. [7] demonstrate how their INTER-IoT solution facilitates secure and robust data exchange among stakeholders within the port community, positioning it as an economical and user-friendly option for both stakeholders and system integrators. Similarly, the authors of [34] introduce a seaport data space that enables secure and interoperable data sharing between stakeholders. Their approach employs IDS Connectors alongside the FIWARE IoT platform to process data (e.g., cleaning, filtering, aggregation) at the edge before sharing, ensuring that only refined data is exchanged rather than raw data.

In [41], the authors explore the role of DTs in data spaces and propose an extension for real-time synchronization of Asset Administration Shells (AAS) within the data space. In [10], the authors present a context-aware security framework that integrates context management with access control policies, such as XACML [27], to enable secure data-sharing decisions. Their approach introduces a mechanism for secure data sharing among groups of smart objects based on contextual data. However, it does not address Edge-based enforcement solutions for managing cross-sector data sharing driven by dynamic contexts.

In [22], the authors propose a context-driven, Edge-based IoT data sharing-as-a-service framework. This framework enables authentication and authorization of multiple tenants in IoT environments at the Edge level, allowing tenant applications to be deployed directly within the Edge Hubs. Despite its advanced capabilities, the framework is not built on the IDSA reference architecture.

7 Conclusions and Future Work

In this paper, we introduced **IDS4Edge**, an IDS-compliant approach, which supports dynamic, context-driven IoT data sharing at the Edge. IDS4Edge enables flexible, context-aware IoT data sharing by integrating access control policies on top of IDS connectors, customized for specific application-level IoT contexts. We have implemented a proof-of-concept to demonstrate the effectiveness of this solution. Our Edge-based IDS4Edge solution dynamically enforces participant-specific access control policies on shared data, adjusting in real-time to evolving IoT contexts and contractual agreements at the Edge.

For future work, we plan to extend IDS4Edge to support more diverse and complex use cases beyond the current factory testbed as well as more sophisticated access/usage control models. Another key focus will be integrating Digital Product Passports (DPPs) [42] to enable traceability, transparency, and secure sharing of product lifecycle data in supply chains, leveraging IDS and Edge computing to meet real-time and privacy-sensitive requirements.

Acknowledgments

This work is supported by the European Union's Horizon Europe research and innovation programme under the Grant Agreements 101058585 (Circular TwAIn), 101058477 (COGNIMAN), 101178405 (REED), and the SINTEF SEP project IDS4Edge.

References

[1] Majid Al-Ruithe, Elhadj Benkhelifa, and Khawar Hameed. 2019. A systematic literature review of data governance and cloud data governance. *Personal and Ubiquitous Computing* 23 (2019), 839–859.

[2] Vijay Atluri, Yuan Hong, and Soon Ae Chun. 2020. Security, Privacy and Trust for Responsible Innovations and Governance. In *The 21st Annual International Conference on Digital Government Research* (Seoul, Republic of Korea) *(dg.o '20)*. Association for Computing Machinery, New York, NY, USA, 365–366. https://doi.org/10.1145/3396956.3396978

[3] Cisco. 2024. Fast Innovation require Fast IT. https://www.cisco.com/c/dam/global/en_ph/assets/ciscoconnect/pdf/bigdata/jim_green_cisco_connect.pdf.

[4] Fabian De Prieëlle, Mark De Reuver, and Jafar Rezaei. 2020. The role of ecosystem data governance in adoption of data platforms by Internet-of-Things data providers: Case of Dutch horticulture industry. *IEEE Transactions on Engineering Management* 69, 4 (2020), 940–950.

[5] GAIA-X European Association for Data and Cloud. 2022. *Gaia-X Architecture Document*. Technical Report. Gaia-X European Association for Data and Cloud AISBL, Avenue des Arts 6-9, 1210 Brussels. https://docs.gaia-x.eu/technical-committee/architecture-document/22.10/ Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

[6] GAIA-X European Association for Data and Cloud. 2022. GAIA-X Trust Framework. https://docs.gaia-x.eu/policy-rules-committee/trust-framework/22.10/ Accessed: 2024-11-27.

[7] Pablo Gimenez, Miguel Llop, E Olivares, C Palau, M Montesinos, and M Llorente. 2020. Interoperability of IoT platforms in the port sector. In *Proceedings of 8th Transport Research Arena TRA*. TRAFICOM, Helsinki, Finland, 27–30.

[8] Giulia Giussani, Sebastian Steinbuss, Tobias Prasse, and Nora Gras. 2024. *Data Connector Report*. Technical Report. International Data Spaces Association, Dortmund, Germany. https://doi.org/10.5281/zenodo.13838396

[9] Arda Goknil, Phu Nguyen, Sagar Sen, Dimitra Politaki, Harris Niavis, Karl John Pedersen, Abdillah Suyuthi, Abhilash Anand, and Amina Ziegenbein. 2023. A Systematic Review of Data Quality in CPS and IoT for Industry 4.0. *ACM Comput. Surv.* 55, 14s, Article 327 (July 2023), 38 pages. https://doi.org/10.1145/3593043

[10] Jose L. Hernandez Ramos, Jorge Bernal Bernabe, and Antonio F. Skarmeta. 2015. Managing Context Information for Adaptive Security in IoT Environments. In *2015 IEEE 29th International Conference on Advanced Information Networking and Applications Workshops*. IEEE, Gwangju, Korea (South), 676–681.

[11] IDSA. 2022. *IDS RAM 4*. Technical Report. International Data Spaces Association, Anna-Louisa-Karsch-Str. 2, 10178 Berlin, Germany. https://docs.internationaldataspaces.org/ids-knowledgebase/ids-ram-4 Creative Commons Attribution 4.0 International License.

[12] IDSA. 2024. Dataspace Protocol. https://docs.internationaldataspaces.org/ids-knowledgebase/dataspace-protocol Accessed: 2024-11-27.

[13] IDSA. 2024. Innovating the future of daa exchange in Europe and beyond. https://internationaldataspaces.org/we/ Accessed: 03.05.2024.

[14] Stephanie Jernigan, David Kiron, and Sam Ransbotham. 2016. Data sharing and analytics are driving success with iot. *MIT Sloan Management Review* 58, 1 (2016), 1–18.

[15] An Ngoc Lam. 2021. *Dynamic Adaptation in Industrial IoT Systems*. Ph. D. Dissertation. Luleå University of Technology.

[16] An Ngoc Lam and Øystein Haugen. 2019. Implementing OPC-UA services for Industrial Cyber-Physical Systems in Service-Oriented Architecture. In *IECON 2019-45th Annual Conference of the IEEE Industrial Electronics Society*, Vol. 1. IEEE, Lisbon, Portugal, 5486–5492.

[17] An Ngoc Lam, Øystein Haugen, and Jerker Delsing. 2021. Interoperability for industrial internet of things based on service-oriented architecture. In *IECON 2021–47th Annual Conference of the IEEE Industrial Electronics Society*. IEEE, Toronto, ON, Canada, 1–6.

[18] An Ngoc Lam, Øystein Haugen, and Jerker Delsing. 2022. Dynamical orchestration and configuration services in industrial iot systems: An autonomic approach. *IEEE Open Journal of the Industrial Electronics Society* 3 (2022), 128–145.

[19] An Ngoc Lam, Gøran Brekke Svaland, Miguel Ángel Barcelona, Shane Keaveney, Wissam Mallouli, Luong Nguyen, Assia Belbachir, Xiang Ma, Akhilesh Kumar Srivastava, and Ahmed Nabil Belbachir. 2024. SINDIT: A Framework for Knowledge Graph-Based Digital Twins in Smart Manufacturing. In *Internet of Things. 7th IFIPIoT 2024 International IFIP WG 5.5 Workshops (IFIP Advances in Information and Communication Technology, Vol. 738)*, Gaëtan Rey, Jean-Yves Tigli, and Erwin Franquet (Eds.). Springer Cham, Nice - Sophia Antipolis, France, 1–18.

[20] Sin Kuang Lo, Yue Liu, Su Yen Chia, Xiwei Xu, Qinghua Lu, Liming Zhu, and Huansheng Ning. 2019. Analysis of Blockchain Solutions for IoT: A Systematic Literature Review. *IEEE Access* 7 (2019), 58822–58835. https://doi.org/10.1109/ACCESS.2019.2914675

[21] Tim Menzies, Bowen Xu, Hong Jin Kang, Jie M. Zhang, Jiri Gesi, Sagar Sen, Beatriz Cassoli, Nicolas Jourdan, Jieke Shi, Phu Nguyen, and Valentina Golendukhina. 2024. SEA4DQ 2024 Workshop Summary. *SIGSOFT Softw. Eng. Notes* 49, 4 (Oct. 2024), 29–30. https://doi.org/10.1145/3696117.3696125

[22] Huu-Ha Nguyen, Phu H. Phung, Phu H. Nguyen, and Hong-Linh Truong. 2022. Context-driven Policies Enforcement for Edge-based IoT Data Sharing-as-a-Service. In *2022 IEEE International Conference on Services Computing (SCC)*. IEEE, Barcelona, Spain, 221–230. https://doi.org/10.1109/SCC55611.2022.00041

[23] Phu Nguyen, Arda Goknil, Gencer Erdogan, Shukun Tokas, Nicolas Ferry, and Thanh Thao Thi Tran. 2024. Advances in Secure IoT Data Sharing. *Foundations and Trends® in Privacy and Security* 7, 1 (2024), 1–73. https://doi.org/10.1561/3300000042

[24] Phu H Nguyen, Gregory Nain, Jacques Klein, Tejeddine Mouelhi, and Yves Le Traon. 2014. Modularity and dynamic adaptation of flexibly secure systems: Model-driven adaptive delegation in access control management. *Transactions on Aspect-Oriented Software Development XI* 8400, 1 (2014), 109–144.

[25] Phu H. Nguyen, Sagar Sen, Beatriz Bretones-Cassoli, Nicolas Jourdan, and Maria Chiara Magnanini. 2023. Software Engineering and AI for Data Quality in Cyber-Physical Systems/Internet of Things - SEA4DQ'22 Report. *SIGSOFT Softw. Eng. Notes* 48, 1 (Jan. 2023), 108–111. https://doi.org/10.1145/3573074.3573103

[26] Phu H. Nguyen, Sagar Sen, Nicolas Jourdan, Beatriz Cassoli, Per Myrseth, Mikel Armendia, and Odd Myklebust. 2022. Software Engineering and AI for Data Quality in Cyber- Physical Systems - SEA4DQ'21 Workshop Report. *SIGSOFT Softw. Eng. Notes* 47, 1 (Jan. 2022), 26–29. https://doi.org/10.1145/3502771.3502781

[27] OASIS. 2013. eXtensible Access Control Markup Language (XACML) Version 3.0. OASIS Standard. http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html

[28] Deepak Pahwa and Binil Starly. 2021. Dynamic matching with deep reinforcement learning for a two-sided Manufacturing-as-a-Service (MaaS) marketplace. *Manufacturing Letters* 29 (2021), 11–14.

[29] Jaehong Park and Ravi Sandhu. 2004. The UCONABC usage control model. *ACM Trans. Inf. Syst. Secur.* 7, 1 (Feb. 2004), 128–174. https://doi.org/10.1145/984334.984339

[30] Gopika Premsankar, Mario Di Francesco, and Tarik Taleb. 2018. Edge Computing for the Internet of Things: A Case Study. *IEEE Internet of Things Journal* 5, 2 (2018), 1275–1284. https://doi.org/10.1109/JIOT.2018.2805263

[31] Tanusan Rajmohan, Phu H. Nguyen, and Nicolas Ferry. 2020. Research Landscape of Patterns and Architectures for IoT Security: A Systematic Review. In *2020 46th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. IEEE, Portoroz, Slovenia, 463–470. https://doi.org/10.1109/SEAA51224.2020.00079

[32] Tanusan Rajmohan, Phu H Nguyen, and Nicolas Ferry. 2022. A decade of research on patterns and architectures for IoT security. *Cybersecurity* 5, 1 (2022), 1–29.

[33] Karen Rose, Scott Eldridge, and Lyman Chapin. 2015. The internet of things: An overview. *The internet society (ISOC)* 80, 15 (2015), 1–53.

[34] David Sarabia-Jácome, Ignacio Lacalle, Carlos E. Palau, and Manuel Esteve. 2019. Enabling Industrial Data Space Architecture for Seaport Scenario. In *2019 IEEE 5th World Forum on Internet of Things (WF-IoT)*. IEEE, Limerick, Ireland, 101–106. https://doi.org/10.1109/WF-IoT.2019.8767216

[35] Sagar Sen, Erik Johannes Husom, Arda Goknil, Simeon Tverdal, Phu Nguyen, and Iker Mancisidor. 2022. Taming Data Quality in AI-Enabled Industrial Internet of Things. *IEEE Software* 39, 6 (2022), 35–42. https://doi.org/10.1109/MS.2022.3193975

[36] J. E. Siegel, S. Kumar, and S. E. Sarma. 2018. The Future Internet of Things: Secure, Efficient, and Model-Based. *IEEE Internet of Things Journal* 5, 4 (Aug 2018), 2386–2398. https://doi.org/10.1109/JIOT.2017.2755620

[37] Manu Sporny, Amy Guy, Markus Sabadello, Drummond Reed, Dave Longley, Orie Steele, and Christopher Allen. 2022. Decentralized Identifiers (DIDs) v1.0: Core architecture, data model, and representations. W3C Recommendation 19 July 2022. https://www.w3.org/TR/did-core/

[38] Manu Sporny, Ted Thibodeau Jr, Ivan Herman, Michael B. Jones, Gabe Cohen, Dave Longley, David Chadwick, and Orie Steele. 2024. Verifiable Credentials Data Model v2.0. W3C Candidate Recommendation Draft 19 October 2024. https://www.w3.org/TR/vc-data-model-2.0/

[39] Gianluca Tedaldi and Giovanni Miragliotta. 2023. Early adopters of Manufacturing-as-a-Service (MaaS): state-of-the-art and deployment models. *Journal of Manufacturing Technology Management* 34, 4 (2023), 580–600.

[40] Ikram Ullah, Gerard De Roode, Nirvana Meratnia, and Paul Havinga. 2021. Threat modeling—how to visualize attacks on IoTA? *Sensors* 21, 5 (2021), 1834.

[41] Friedrich Volz, Gerhard Sutschet, Ljiljana Stojanovic, and Thomas Usländer. 2023. On the role of digital twins in data spaces. *Sensors* 23, 17 (2023), 7601.

[42] Joerg Walden, Angelika Steinbrecher, and Maroye Marinkovic. 2021. Digital product passports as enabler of the circular economy. *Chemie Ingenieur Technik* 93, 11 (2021), 1717–1727.

[43] Maryna Waszak, An Ngoc Lam, Volker Hoffmann, Brian Elvesæter, Maria Flavia Mogos, and Dumitru Roman. 2022. Let the Asset Decide: Digital Twins with Knowledge Graphs. In *2022 IEEE 19th International Conference on Software Architecture Companion (ICSA-C)*. IEEE, Honolulu, HI, USA, 35–39. https://doi.org/10.1109/ICSA-C54293.2022.00014