E2E Security Architecture of Industrial Dataspaces

Manufacturing-X

End2End Security Topic Group

Preface **5**

About the Manufacturing-X E2E Security Topic Group **6**

1 Introduction and Context - Industrial Dataspace **10**

2 Ecosystem and roles **11**

2.1 Intended Use and Supported Use Cases *12*

3 Threat Landscape and Regulatory Drivers **13**

3.1 Evolving threats and vulnerabilities *13*

3.2 Core Regulatory Requirements (Must-Have) *14*

4 Challenges to creating a secure dataspace **15**

4.1 Main Challenges *15*

4.2 Challenges related to the development of secure e2e dataspace components *16*

4.3 Challenges related to enabling Scalability and large Rollouts *16*

4.4 Expectations and questions to solve *17*

5 Operation and Governance Framework **19**

5.1 Abstract dataspace definition and required aspects to create secure and trusted dataspaces *19*

5.2 Current Dataspace Architecture *22*

5.3 Emergency Containment and Major Incident Response *22*

6 Three-Layer Dataspace Security Architecture **24**

6.1 Associated Security Controls in the Gateway *24*

6.2 Data Provider Security Architecture *26*

6.3 OT/IT Integration and Shopfloor Considerations *29*

7 Data-exchange Gateway and Orchestrator **30**

7.1 Industry Core Hub (One open source implementation of the Data Exchange Orchestrator) *31*

7.2 Industrial Data Hub Implementation – Topic Group extension of the Gateway *32*

8 Use-Case-Driven Security via AAS Submodels **35**

9 Risk-based Cybersecurity Management **37**

9.1 Core Design Principles of a secure dataspace *37*

9.2 Trust Infrastructure and Identity Management *37*

9.3 Security building blocks of the Eclipse Dataspace Connector *39*

9.4 Technical implementation recommendations *40*

9.5 Secure Development Lifecycle for OSS and related artefacts *41*

10 Conclusions and Recommendations **44**

10.1 Recommendations *44*

10.2 Next Steps and Priority Actions *45*

10.3 Call to Action *46*

Appendix A: Relevant standards of design, operation, and governance **47**

A.1 Dataspace-Specific & Industrial Data Exchange Standards *48*
A.2 Industrial / OT Security Standards (Highly Relevant) *49*
A.3 Cross-Cutting IT Security & Risk Management *49*
A.4 Sector-Specific Standards *49*
A.5 Emerging New Topics *49*

Appendix B: Risk Management Template Table **50**

Appendix C: Attack Vectors and Baseline Security Controls **51**
C.1 Priority Attack Vectors for the Manufacturing-X Dataspace *51*
C.2 Baseline Security Controls and Requirements *51*

Appendix D: Current Security Measures Implemented in Eclipse Tractus-X **55**

IMPRINT

Publisher

Manufacturing-X Guidance Board

Project Office Manufacturing-X (to be deleted for internal use)

Open Industry 4.0 Alliance Implementation GmbH

Willy-Brandt-Platz 6

81829 Munich, Germany

Authors

Haris Avgoustinos	EXPO21XX GmbH
Harald Gossner	Intel
Michael Jochem	Robert Bosch GmbH/Plattform Industrie 4.0 AG „Sicherheit kollaborative Ökosysteme“
Dominik Goergen	Cofinity-X GmbH
Mathias Brunkow Moser	Catena-X Automotive Network e.V.
Johannes Diemer	ARENA2036 e.V.
Per-Henrick Addicks	SAP SE
Marius Pohl	Cofinity-X GmbH
Hans Kittelmann	Hochschule für Technik und Wirtschaft Dresden
Ingolf Gehrhardt	Hochschule für Technik und Wirtschaft Dresden
Paul Patolla	Hochschule für Technik und Wirtschaft Dresden
Niranjana Subrahmanyam	Siemens AG
Stefan Pyka	Siemens AG
Tobias Kutzler	Fraunhofer IFF
Monika Kamhuber	Fraunhofer AISEC
Gerd Brost	Fraunhofer AISEC
Hendrik Meyer zum Felde	Fraunhofer AISEC
Paul Georg Wagner	Fraunhofer IOSB
Pascal Birnstill	Fraunhofer IOSB
Christian Haas	Fraunhofer IOSB
Felix Hanke	Deutsches Zentrum für Luft- und Raumfahrt e. V. (DLR)
Claus Beyer	Deutsches Zentrum für Luft- und Raumfahrt e. V. (DLR)
Robin Thrift	EXPO21XX GmbH
Franz Wenninger	Fraunhofer EMFT
Florian Bieck	EXPO21XX GmbH

Mission Statement

The mission is to propose security concepts towards a secure industrial dataspace, which is a critical enabler for MX dataspaces to achieve acceptance of the interoperable dataspace concept in commercial workflows. This includes the perspective that all dataspaces share a common ground, including data and process interoperability, starting with, commonly trusted identities, common data exchange protocols and semantics, interoperable processes, and common security components and understanding in order to enable common interoperable use cases (see Figure 1).

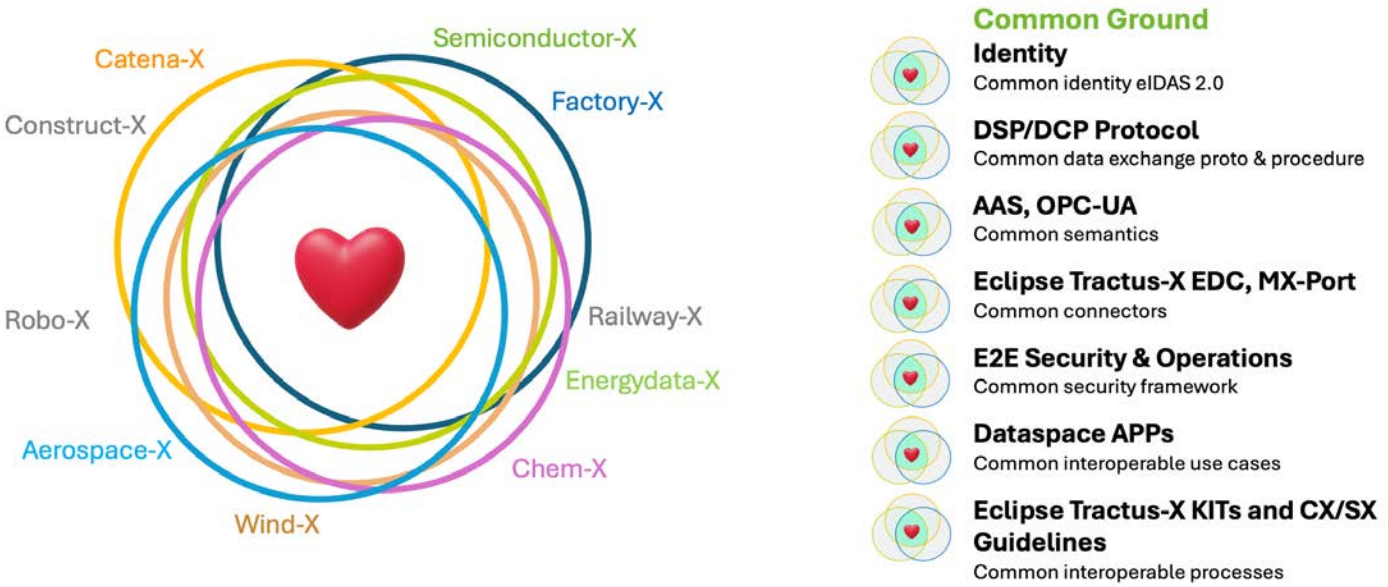


Figure 1: Providing a high security common ground and toolchain for all dataspaces

Preface

Only a **secure dataspace has the potential to become a successful dataspace** — its trustworthiness is essential for broad adoption, industrial collaboration, and long-term competitiveness. As such, robust security is not only a technical necessity but a key pillar for safeguarding **European economic strength and digital sovereignty,** as well as the resiliency of international Supply Chains.

This document provides strategic guidance for MX leadership, decision-makers in government and industry, the Tractus-X community, participants, operators, and application providers to initiate a dedicated coordinated set of activities on end-to-end (E2E) security for operational industrial dataspaces. The activities should define clear recommendations and requirements, design the new multi-layer dataspace architecture and the holistic security approach proposed in this document, including concepts and tools for auditability, resilience, and adaptive defense mechanisms. Together with the Tractus-X Kits, the proposed data sensitivity/criticality levels, and the MX-Connectors (including the EDC), these elements will enable **both secure and user-friendly access,** while the main target is to **reduce the entry barrier** on one side and **increase security for mission-critical use cases** on the other. The initiative should also lay the foundation for future technical and commercial advancements aligned with these goals, supporting dataspace participants & service providers on achieving **a secure and trusted architecture** in their productive business use case-oriented data exchange context

The intention is not only to provide guidance but to coordinate efforts toward a secure and continuously adaptive cyber-resilient dataspace for Europe's industrial future. To scale the E2E security concept, a **risk management process** including a risk-based security concept, security by design for components/services has to be implemented **in the open-source communities,** and effective tools need to be applied, which can be made available as commercial solutions. A close interaction with leading cyber security solution providers and a clear concept for secure operation, threat intelligence, and incident reaction teams is inevitable to secure the industrial dataspace in the long term. We encourage international exchange, taking regional cybersecurity directives into consideration. The impact of fast-developing technologies like AI and quantum computing has to be taken into account. This opens new threats and new capabilities of fast scaling zero-day attacks on the one side, while offering new horizons for better and more efficient AI-supported cyber defense. To ensure coordinated actions, we clearly recommend the continuous review of the security recommendations and mandatory

requirements of the industrial dataspace by an international council of leading experts with a clear mandate, operation model, and continuous financial support.

The paper places special focus on dataspace and connectivity supporting the DSP/DCP protocol including the Eclipse Tractus-X EDC and the MX-Port Hercules.

About the Manufacturing-X E2E Security Topic Group

The E2E Security Topic Group of Manufacturing-X was founded in February 2025 under the MX Guidance Board to address the extended security requirements in a wide-span industrial dataspace with system-critical partners and applications. The Topic Group has been driven by Members from MX projects. For the feature, we propose to continue the activities of the Group under the Umbrella of the Platform Industrie 4.0/AI in close collaboration with Tractus-X and to open it to leading security experts from different domain.

The Topic Group publishes this first Dataspace Security Whitepaper considering the Tractus-X and Manufacturing-X approach to support the dataspace participants and service providers to achieve a secure architecture, secure operations to protect the interests of participants joining & trusting the data space.

Disclaimers

The white paper does not constitute any legal or standard requirements. The recommendations are based on state-of-the-art expert knowledge and cannot guarantee E2E security liability.

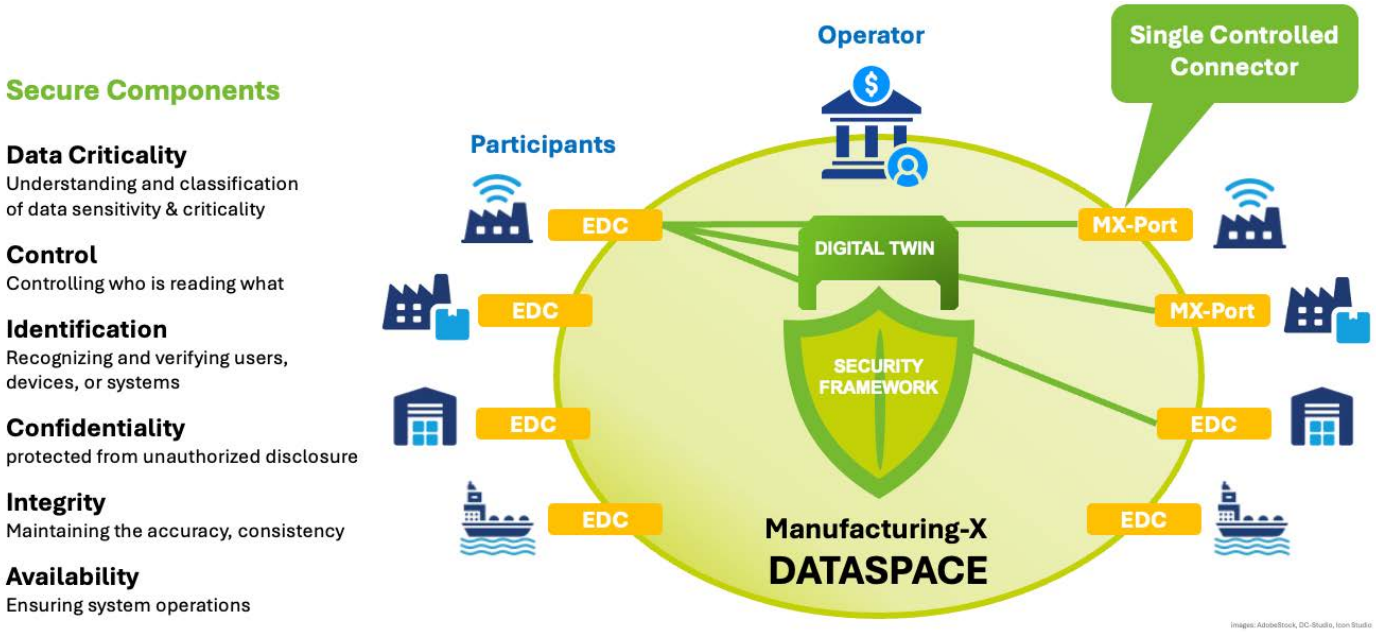
Executive Summary

E2E Security Architecture of Industrial Dataspaces

The case for secure dataspaces

Industrial data exchange today is highly fragmented. A typical company maintains hundreds of individual APIs, interfaces, and platform connections, each carrying its own security assumptions, its own understanding of what is sensitive data, and its own operational overhead. The result is a sprawling attack surface, inconsistent protection, and a daily risk of oversharing. This is not merely inefficient; it has become a material security risk. Cyber-attacks caused an estimated €202 billion in damages in Germany in 2025 alone, a large share of it traceable to legacy and fragmented IT structures.

Manufacturing-X dataspaces change this paradigm. Instead of operating a hundred bespoke channels, each participant connects through a single, controlled connector built on common identities, protocols, and semantics. This "one fits all" model promises a dramatic reduction in complexity, together with stronger security and genuine interoperability across value chains. The promise, however, holds under one condition: security must be an integral part of the architecture from the very beginning, not an afterthought. Both the Cyber Resilience Act (CRA) and the NIS2 Directive now make security-by-design and risk-based approaches a regulatory expectation rather than a best-effort aspiration.



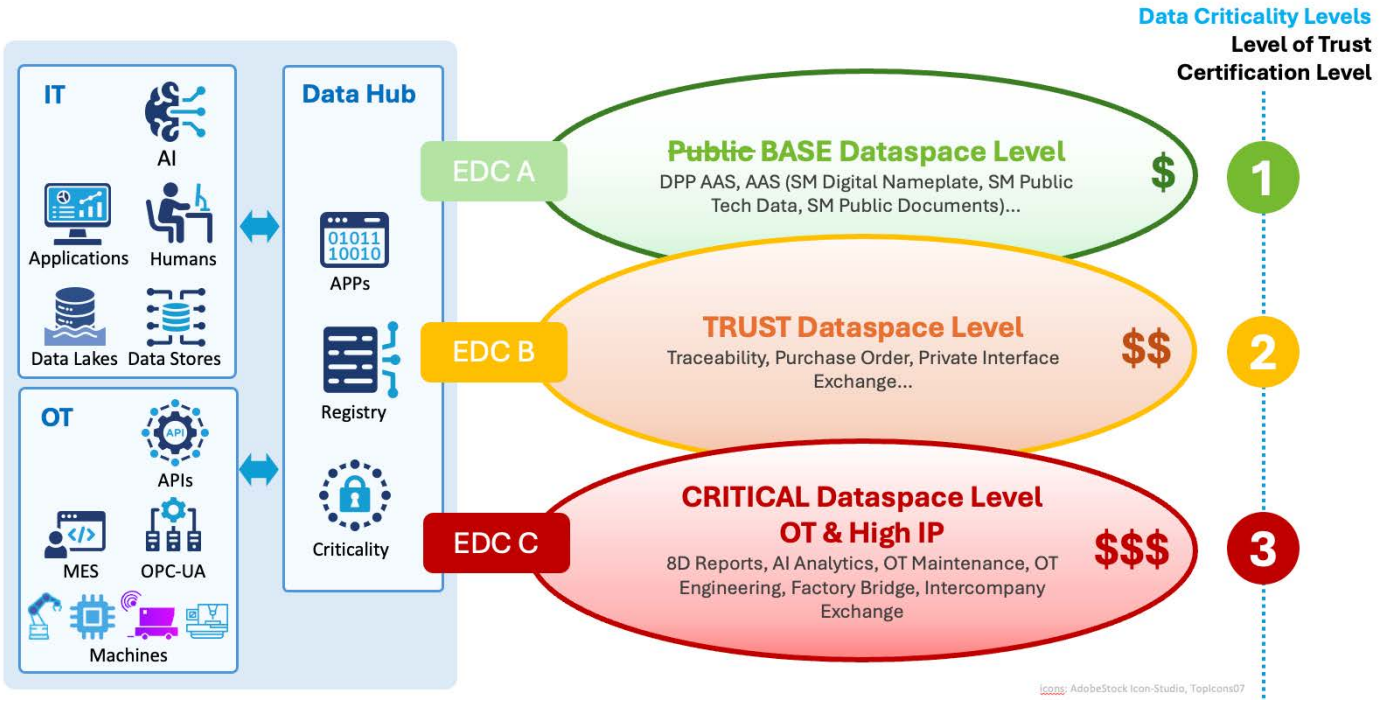
The central difficulty is that complexity is itself the primary adoption barrier, particularly for SMEs. Where authorization models, access policies, and usage contracts become too intricate, participants make configuration mistakes, apply overly broad permissions, or simply ignore policy detail — and sensitive data is exposed. Add to this the integration of OT and shopfloor systems, the need for coordination across different dataspaces, and an emerging threat landscape shaped by AI-driven attacks and the "harvest now, decrypt later" logic of quantum computing, and it becomes clear that a coherent, end-to-end security architecture is required rather than a loose collection of point measures. End-to-end security here goes well beyond protecting data in transit: it spans IT and OT components, connectors and data apps, open-source components, identity and access management, usage policies, lifecycle governance, and the trust mechanisms that hold a federated environment together.

A proposal built on five foundational elements

The proposed E2E Security Architecture extends the existing Eclipse Tractus-X and Manufacturing-X framework rather than replacing it. It rests on five foundational elements that, taken together, make strong security both achievable and economically viable.

First, a three-layer security model — Basic, Trust, and Critical — is directly aligned with data criticality and use-case sensitivity. Each layer carries its own security controls, technical requirements, deployment patterns, and running costs. The *Basic* level offers lightweight, cost-efficient protection that dramatically lowers the entry barrier for SMEs and low-risk use cases such as digital product passports and product carbon footprint data. The *Trust* level adds usage-control enforcement, data encryption, integrity verification, and auditability for confidential business data. The *Critical* level delivers the highest assurance — hardened nodes, technical policy enforcement, secure tunnels, hardware-backed attestation, and continuous monitoring — for intellectual property, OT telemetry, and mission-critical applications.

Because all three layers share a common data fabric and remain fully interoperable, participants implement only the level they actually need and can scale upward later as their use cases demand.



Second, the Industrial Data Hub and Industry Core Hub, positioned at the core of the dataspace gateway as the intelligent, AI-augmented interface between a participant's internal IT/OT systems and the external dataspace. All data entering or leaving the participant passes through the Hub, where it performs real-time criticality assessment, validates intended use, checks for corruption or anomalies, interprets and enforces contracts, and dispatches each data flow to the appropriate connector instance and security level. Crucially, this assessment happens before data leaves the company's security perimeter, so usage control is enforced as close as possible to the data source. Conceived primarily for SMEs and brownfield environments, the Industrial Data Hub enables the fast enrolment of existing applications without requiring deep in-house dataspace expertise.

Third, native, use-case-driven security through Asset Administration Shell (AAS) submodels. Traditional approaches that classify and protect data at the level of individual attributes quickly become unsustainable: they demand manual, per-field decisions, fine-grained policy authoring, and continuous maintenance, and they carry a high risk of misconfiguration and "policy drift" across participants. The AAS changes this equation. Because AAS submodels are purpose-built for concrete business scenarios, and because each Eclipse Tractus-X KIT maps directly to one or more standardized submodels, the path from use case to data criticality is already implicit in the data structure. The correlation Use Case → Tractus-X KIT → AAS Submodel → Data Criticality → Security Level lets the dataspace software assign security levels systematically and with AI support, rather than configuring thousands of attributes by hand. Classification becomes far more consistent and significantly less error-prone than attribute-level alternatives.

Fourth, systematic risk-based cybersecurity management applied across the full lifecycle — both in the open-source development of dataspace components within the Eclipse Tractus-X community and in operational environments. This encompasses threat modelling for each KIT, tiered control catalogues aligned with the three security levels, coordinated vulnerability disclosure and SBOM-based vulnerability management, and structured incident response mechanisms. Security is treated as a process rather than a product — embedded in planning, design, implementation, testing, deployment, and operation alike, and supported by centralized security intelligence, continuous monitoring, and dedicated response teams.

Fifth, a security operations and governance framework – since security of a gateway cannot exist on its own, a concept for the secure operation of a data space needs to be provided, taking into account management of security incidents and a formalization of security levels for components. This is closely attached to the risk-based cybersecurity management process, directly shaping how gateways are operated in a secure manner and how identities are managed.

What the architecture delivers

Taken together, these elements finally deliver the long-promised "one fits all" connector model while making security realistic, practical, and economically viable for SMEs and large enterprises alike. The "one fits all" concept includes

variations of the MX-Port developed in Factory-X to support AAS and OPC-UA data exchange. The strength of controls, the implementation effort, and the operating cost are scaled precisely to the levels and use cases each participant chooses to support. By combining Security by Design, Zero Trust principles, Defense-in-Depth, and a strict separation of Control Plane and Data Plane, the proposal addresses head-on the most pressing challenges identified throughout this white paper: complexity as the dominant adoption barrier, misconfiguration and oversharing risk, the secure integration of OT and shopfloor systems, cross-dataspace interoperability, and the emerging threats posed by AI and quantum computing.

It should be stressed that the connector itself is not a turnkey secure product. The open-source EDC and the wider Eclipse Tractus-X stack will always require ongoing hardening, correct deployment, monitoring, and clear operator responsibilities. The value of this architecture lies in making those obligations explicit, structured, and proportionate to the criticality of the data being exchanged.

Conclusions and recommendations

The Manufacturing-X initiative builds on the proven open-source foundation of Catena-X and extends it across broader ecosystems such as Factory-X, Semiconductor-X, and Aerospace-X, with the Eclipse Tractus-X KITs providing the core technical building blocks — the Connector KIT, the Data Trust & Security KIT, and the supporting governance, digital-twin, and industry-core kits. The E2E security proposal is designed to slot into this framework and to remain regulator-friendly, auditable, and implementable.

To translate the architecture into practice, the Topic Group recommends a focused set of priority actions. The most urgent is to introduce systematic threat modelling for every Eclipse Tractus-X KIT and to develop tiered, catalogue-based cybersecurity concepts and baseline controls, structured according to the three Gateway Security Levels and aligned with VDE-AR-E 2802-20. In parallel, a joint risk management and secure development lifecycle should be established together with the Eclipse Tractus-X community, covering coordinated vulnerability disclosure, SBOM management, and post-market surveillance in line with the CRA. A certification and assessment framework — including self-assessment toolkits for participants and conformity schemes for connectors and Data Hub implementations, differentiated by security level — is needed to give CIOs and decision-makers the credible, testable evidence they require to adopt the model with confidence.

Beyond these, the reference implementation of the Industrial Data Hub, with built-in AI criticality assessment and AAS submodel awareness, should be prioritized as the key enabler for both security and broad adoption, especially in brownfield scenarios. Finally, as industrial dataspaces grow into part of Europe's critical infrastructure, the appropriate role of state and governmental actors — bodies such as BSI and ENISA — must be actively discussed, alongside concrete mechanisms for cross-dataspace security intelligence sharing, joint incident response, and consistent policy enforcement across European and international initiatives.

Call to action

The concepts presented here offer a coherent and implementation-oriented path toward secure and trusted industrial dataspaces, but success now depends on coordinated execution across the Manufacturing-X projects, the Eclipse Tractus-X community, and industry stakeholders. The Topic Group recommends that the Manufacturing-X Guidance Board, in close collaboration with Platform Industrie 4.0/AI and leading security experts, establish a dedicated implementation roadmap and working group to drive these priorities, and that it initiates a focused working stream on state actor integration and strategic resilience. Only a secure dataspace can become a successful dataspace — and its trustworthiness remains the decisive condition for broad adoption, industrial collaboration, and Europe's long-term digital sovereignty.

1 Introduction and Context - Industrial Dataspace

Today's data exchange in industry is highly fragmented. Companies typically maintain hundreds of individual APIs, interfaces, and platform connections. This creates a complex, hard-to-manage network with numerous attack surfaces, inconsistent security and data sensitivity understanding, and high operational effort. This fragmented architecture is not only inefficient — it has become a significant security risk. Cyber-attacks caused approximately €202 billion in damages in Germany in 2025[1], with a large share resulting from legacy and fragmented IT structures.

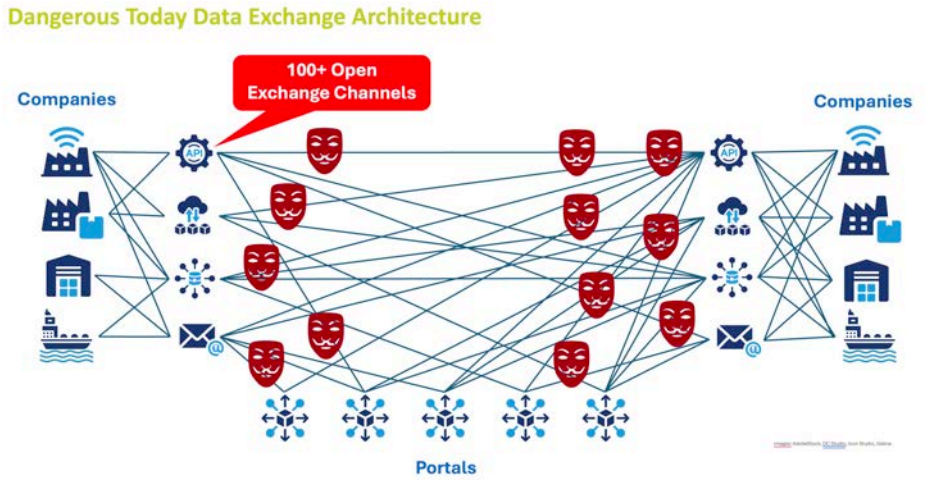


Figure 2: Today's data exchange involves hundreds of custom connections

Dataspaces fundamentally change this paradigm. Instead of managing 100+ individual exchange channels, they introduce a universal, standardized, and secure approach to data exchange, often based on a single controlled connector per participant. This "one fits all" model promises to significantly reduce complexity while improving security and interoperability. To address specific industrial requirements, Factory-X extended this concept by introducing the MX-Port offering three use case specific connector versions. The MX-Port versions Hercules and Orion support the standard DSP/DCP protocols and offers Asset Administration Shell (AAS) and OPC-UA based data exchange, providing a second implementation option while remaining part of the overall Manufacturing-X connector strategy.

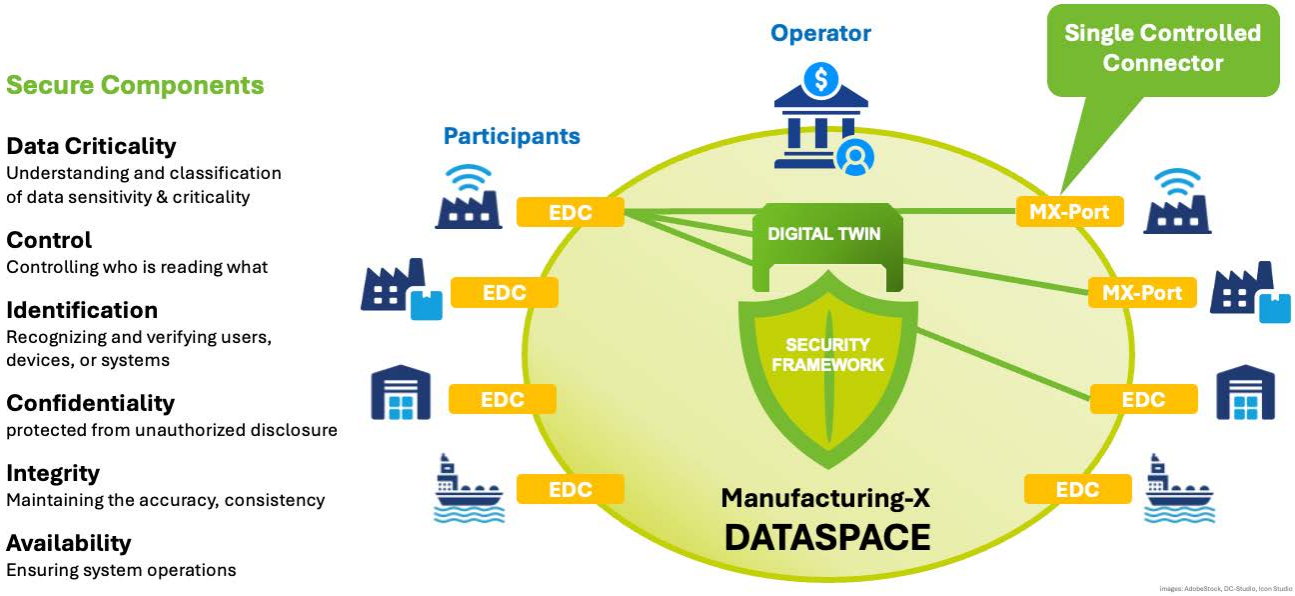


Figure 3: Moving toward a "one fits all", highly secure connection

Manufacturing-X, therefore, aims to establish robust industrial dataspaces across value chains. For this vision to succeed at scale, security must be an integral part of the architecture from the very beginning — not an afterthought. Both the Cyber Resilience Act (CRA) and NIS2 mandate security-by-design and systematic risk management. Only if these requirements are met can industrial dataspaces fulfill the promise of acting as a "turbo" for secure and trusted digitalization.

[1] Bitkom "Wirtschaftsschutz 2025" study (published in September 2025).

2 Ecosystem and roles

In the new data ecosystems promoted by the European Union and the BMWE, interacting actors are able to exchange data through dataspaces to implement digital data supply chains and enable digital business cases, and leverage the implementation of new domain-specific AI applications. Currently, a vast number of different dataspaces exist, emerging from various initiatives. We describe the characteristics of a data ecosystem that has emerged from current implementations in Manufacturing-X (e.g., Semiconductor-X, Catena-X) and that can serve as a blueprint for other projects. The current use cases rely on the Eclipse Tractus-X EDC (Eclipse Dataspace Components Connector), the dataspace connector component developed in the Catena-X and Manufacturing-X projects. Eclipse Tractus-X software components are available as open-source software and are supported by Tractus-X implementation guidelines as well as interoperable process and use-case templates (Eclipse Tractus-X KITs). Eclipse Tractus-X is a community-driven open-source project and is to be understood as continuously evolving, which increases its impact on the one hand while presenting new challenges in implementing security and risk management, which will be discussed in this paper.

The Manufacturing-X dataspace involves a well-defined set of actors with clear roles and responsibilities:

Primary actors in the dataspace:

- **Data Provider** – A dataspace participant owning and sharing data while retaining full sovereignty and defining access and usage policies.
- **Data Consumer** – A dataspace participant who can only discover and request data that is explicitly designated for them through access policies. Once access is granted, the consumer uses the data strictly according to the agreed-upon contracts and usage policies.
- **Data Space Core Service Operator** – Operates technical dataspace components, ensuring security, availability, and policy enforcement.

Federation-level actors:

- **Dataspace Authority** – Defines governance rules, participation conditions, and compliance requirements for the dataspace.
- **Dataspace Operator** – Operates and coordinates core federation services (onboarding, identity registry, discovery/catalog service, monitoring). The operator needs to orchestrate the Risk Management in Operation with Security Intelligence, Vulnerability management, Security compliance (updates), and PSIRT teams.
- **Identity Provider / Trust Anchor** – Issues and manages verifiable credentials and attestations. Alignment with European identity frameworks such as eIDAS and the upcoming European Digital Identity (EUDI) Wallet is strategically planned to strengthen cross-dataspace interoperability and trust establishment.
- **Monitoring / Cyber Response Entities** – Continuous security monitoring, anomaly detection, and coordinated incident response. While a common understanding of the importance of this role exists among all players and is discussed in this paper, the exact allocation to responsible entities must be further evaluated and implemented in the near future.
- **State / Governmental Actors** (e.g., BSI, ENISA, national cybersecurity authorities) – Provide threat intelligence on geopolitical and state-sponsored attacks, support the protection of critical infrastructure, and enforce regulatory compliance. In the long run, when the dataspaces reach a critical level of usage, the discussion regarding a possible deeper involvement of state actors and a direct role in the security concept of the dataspace should be started. This could include supplying real-time attack information, defining secure network operation requirements, coordinating national-level incident response, and integrating governmental security intelligence into the overall dataspace security architecture.

Supporting actors:

- **Auditors / Conformity Assessment Bodies** – Independent verification of compliance and security posture (e.g., TÜV or equivalent).
- **Reference Implementation Provider** - Develops and maintains standard-compliant reference implementations for dataspace components (Eclipse Tractus-X EDC, Eclipse Tractus-X Identity Hub, ...). Development adheres to security-by-design principles and established secure development lifecycle (SDLC) practices.
- **Participants' internal systems** – IT/OT landscapes, brownfield systems, MES, AAS Digital Twins, and local data sources that connect to the dataspace via the EDC/MX-Port.

- **Participant Administrator** – Manages onboarding, certificates, wallets, and connector configuration for the participating organization, and acts as the central point of contact for governance-related topics

A clear definition of roles and responsibilities is essential to avoid gaps in accountability and to enable effective governance.

2.1 Intended Use and Supported Use Cases

The primary intended use of the Manufacturing-X dataspace is **sovereign, interoperable, and secure data exchange across industrial value chains** while maintaining full data sovereignty for each participant.

Use cases are derived from Eclipse Tractus-X KITs and the various Manufacturing-X projects (Factory-X, Semiconductor-X, Aerospace-X, etc.) and include:

- **Supply Chain data exchange** – traceability, Product Carbon Footprint (PCF), Digital Product Passport (DPP), quality data, capacity information.
- **Federated Applications** – cross-company applications that consume and process shared data without moving raw data unnecessarily.
- **Federated AI / Analytics workloads** – collaborative model training, predictive analytics, and federated learning while protecting IP and sensitive data.
- **OT / Shopfloor integration** – machine data, real-time telemetry, predictive maintenance, digital twins, and modular production data exchange.
- **Manufacturing as a Service (MaaS)** and **Modular Production** – dynamic, on-demand manufacturing scenarios.
- **Circularity and sustainability use cases** – end-of-life data, recycling information, and environmental impact data.

3 Threat Landscape and Regulatory Drivers

In the following, we highlight the threat landscape and industry- and country-specific regulations that influence the design of an industrial dataspace.

3.1 Evolving threats and vulnerabilities

Industrial dataspaces are high-value targets for sophisticated attacks. Unlike traditional systems with clear perimeters, they operate in a highly federated, perimeter-less environment with participants of varying security maturity.

Key threat categories include:

- **Supply chain** threats— weaknesses or compromises anywhere in the software, hardware, and service supply chain of dataspace components (EDC, Eclipse Tractus-X, extensions), commercial providers, managed services, and protocol variants (e.g., MX-Port with OPC-UA) undermine the confidentiality, integrity, and availability of the entire dataspace.
- **Involuntary data disclosure due to complexity in understanding data criticality and implementing access rules** — misclassification of data criticality and mis-specified access and usage control rules in a highly complex policy landscape lead to unintended disclosure, accidental public exposure, or over-sharing of sensitive data, even by legitimate participants. This is aggravated by the limited granularity of access control at the EDC level, which cannot enforce endpoint-specific restrictions.
- **Identity and trust fragmentation** — inconsistent identity management, trust anchors, and revocation and certification practices across participants lead to misidentification of parties, unauthorized access, and loss of confidence in dataspace-wide trust decisions.
- **Data poisoning and manipulation** — asset manipulations, including the manipulation of catalogs and payment accounts, or the manipulation of distributed supply chain information like PCF, traceability, undermining trust and accountability. This also includes unauthorized data modification or deletion through, e.g., missing HTTP method enforcement.
- **Layer misalignment and lack of interoperability** — misalignment between legal agreements, organizational policies, and technical enforcement mechanisms, aggravated by missing standardized policy templates for Eclipse Tractus-X KIT use cases, results in ineffective protection and uneven enforcement of data usage constraints across participants.
- **Cross-Dataspace coordination gaps** — Inconsistent policy enforcement, identity concepts, and incident response when different dataspaces interact can threaten the confidentiality of the data.
- **Changing understanding of security perimeter enforcement and zero-tolerance policy failure** — extension of dataspaces into more critical use cases (e.g., OT integration, federated AI, shopfloor data) without a robust redefinition of the security perimeter and strict, reliable enforcement of access and usage policies leads to systemic weaknesses and high-impact failures.
- **AI-driven threats** — autonomous exploitation and advanced attack generation by new AI systems, and prompt injection and model poisoning in AI systems implementing federated use cases.
- **Quantum computing risks** — “Harvest now, decrypt later” future quantum computing will render current cryptographic protections ineffective, enabling retrospective decryption of today’s long-lived encrypted industrial data.
- **State-sponsored cyber operations** — highly capable, state-supported threat actors with effectively unrestricted resources actively target industrial dataspaces. These threats are amplified by the dynamic nature of participants, heterogeneous trust levels, and the increasing integration of OT systems and shopfloor data.
- **Broken access control and privilege escalation at the connector boundary** — insufficient authorization enforcement in backend applications behind the EDC connector (e.g., missing OAuth2, API key, or object-level authorization) allows attackers to access, modify, or abuse API endpoints and object properties beyond their intended permissions. Since the EDC was not designed to enforce fine-grained authorization for every backend endpoint, the entire API surface behind a connector may be exposed, enabling systemic abuse of trust boundaries, tenant isolation failures, and unrestricted resource consumption.

- **Internal secret and credential exposure via dataplane proxy** — authorization-related endpoints (e.g., /jwks.json, /oauth/token) of backend applications, originally restricted to internal networks and potentially implemented with lower security standards, may become reachable through the dataplane proxy, enabling attackers to extract secrets, tokens, or cryptographic keys and compromise the authentication infrastructure.

3.2 Core Regulatory Requirements (Must-Have)

As the industrial dataspace shall operate across countries and thus legislative borders as well as in differing industries, we highlight and discuss regulations and requirements that the dataspace must consider.

Two EU cybersecurity regulations significantly influence how dataspaces have to be developed and operated: the NIS2 Directive and the Cyber Resilience Act (CRA). NIS2 focuses on organizational cybersecurity and operations, thus affecting operators of central services and of connectors-as-a-service, whereas the Cyber Resilience Act focuses on the security of digital products and software, thus affecting suppliers of connectors.

Europe combines a broad operator-security regime with a horizontal product-security law. Other jurisdictions often split this into critical-infrastructure laws, data/security laws, incident-reporting rules, IoT labels, sector standards, and procurement requirements rather than one direct equivalent. China's framework is built around the Cybersecurity Law, Data Security Law, PIPL, CII regulation, cybersecurity reviews, and MLPS. The US has CIRCIA for critical-infrastructure cyber incident reporting, while product security is more fragmented across IoT, NIST, procurement, and sector rules. Japan and Singapore are notable for IoT security labelling schemes such as JC-STAR and CLS.

NIS2 introduces mandatory cybersecurity risk management and incident reporting for organizations providing important digital services. Operators of central dataspace services – such as identity providers, registries, and onboarding services, are likely to fall within the scope of NIS2. Managed connector providers offering connector-as-a-service may also fall under the directive. The directive requires systematic cybersecurity risk management, technical safeguards, operational monitoring, incident response capabilities, and management accountability.

The Cyber Resilience Act (CRA) regulates software products placed on the European market. It primarily affects companies that distribute connectors or related software commercially. It mandates a secure software development lifecycle, systematic vulnerability management, SBOM, security updates, and post-market surveillance.

For more specific frameworks and guidelines, see Appendix A.

Implications for Manufacturing-X

For industrial dataspaces, the central compliance challenge is managing shared security responsibilities across multiple organizations. Primary objectives for operators should include secure operations of data space, security of software supply chains, and monitoring, and potentially contribute to the evaluation of the security of open-source components used in their services, while suppliers distributing connectors must implement secure development and vulnerability management processes. At the same time, operators of central dataspace services must ensure that central services like identity providers are operated with strong security monitoring, incident response, and resilience measures.

As a result, governance mechanisms such as connector certification programs, coordinated vulnerability disclosure processes, and security requirements for ecosystem participants have to be introduced or revised. NIS2 and the CRA push dataspaces toward stronger security governance, transparent software supply chains, and resilient shared infrastructure.

4 Challenges to creating a secure dataspace

Typically, security objectives can be classified into three major categories:

- **Confidentiality** – Prevent unauthorized disclosure of information. Protected through access control, usage policies, and (for critical data) end-to-end encryption at the point of origin.
- **Integrity** – Prevent unauthorized modification or tampering. Ensured by cryptographic signatures, attestations, and tamper-evident logging.
- **Availability** – Ensure that information and services can be accessed when required. Supported by resilient deployment patterns, monitoring, and business continuity measures.

However, a secure dataspace requires an additional dimension of protection:

- **Trustworthiness** – Ensure that the system will behave as expected. Includes authenticity, non-repudiation, provenance, and verifiability through VCs, DIDs, and attestations.

A common misconception is that securing industrial dataspaces is primarily about protecting data in transit. In the Manufacturing-X context, **End-to-End Security** goes far beyond transport layers and must be addressed holistically across the entire ecosystem. It encompasses the security of IT and OT components, dataspace services (connectors, data apps), Open-Source Components, identity management, access control, data sovereignty, usage policies, lifecycle governance, as well as trust mechanisms and certifications. **Security must be consistently orchestrated across all participants,** including dataspace members, software providers, and operators, to ensure reliable trust in a federated environment.

Finally, the practical use of a dataspace also requires that data sensitivity and criticality can be labelled, understood, and managed at a level that is usable for common use cases and everyday participants. This is essential to reduce the daily risk of oversharing. If authorization models, access policies, and ODR-based usage contracts become too complex, participants are likely to make configuration mistakes, ignore policy details, or apply overly broad permissions. The result is unwanted data exposure, a risk that is already common in today's fragmented, non-dataspace data ecosystems.

4.1 Main Challenges

Several fundamental challenges make the implementation of robust End-to-End Security particularly demanding:

- **Complexity as the primary adoption barrier:** The biggest obstacle for broad participation, especially among SMEs, is not a lack of willingness but the high technical and operational complexity of integration and onboarding.
- **Fragmented implementations and protocol extensions:** Different initiatives have extended the core connector concept with the **MX-Port**, adding support for both DSP/DCP and OPC-UA protocols. While these extensions are necessary for real industrial requirements (e.g., shopfloor integration), they increase risks around interoperability and consistent security.
- **Geostrategic and regulatory complexity:** Differing continental approaches to data sovereignty combined with overlapping regulations such as NIS2 (EU), the Cyber Resilience Act (CRA, EU), NIST (US), CIRCIA (US), JC_STAR (Japan), and other regional Cybersecurity Basic Acts and Laws (Japan, China) create additional requirements for risk management, incident reporting, and security-by-design.
- **Architectural trade-offs:** Centralized architectural models simplify governance but create single points of failure, while decentralized models enhance sovereignty but add complexity. Hybrid approaches are required to ensure common defense strategies, dataspace security intelligence, and coordinated incident response, but are difficult to balance.
- **Open-source ecosystem risks:** Securing community-driven components such as the EDC and Eclipse Tractus-X components demands strong DevSecOps practices and clear responsibilities across the ecosystem.
- **Cross-Dataspace Security Coordination:** Manufacturing-X dataspaces do not exist in isolation. Different European dataspaces (e.g., Catena-X, Factory-X, Semiconductor-X, Aerospace-X, Health-X) must securely interact with each other. This requirement extends to international dataspaces and platforms worldwide, such as those emerging in North America, Asia (e.g., China's national industrial dataspaces), and other regions. Effective cross-

dataspace collaboration requires coordinated security activities, consistent policy enforcement, interoperable identity and trust mechanisms, as well as joint incident response capabilities across organizational and national boundaries.

Many aspects have been tackled in some level of detail across various dataspace initiatives and projects, may it be IDSA, Gaia-X, Manufacturing-X, and so on. The focus should be on harmonizing existing standards rather than creating new ones. To overcome the challenges listed above, we need to collect and integrate what is out there to create a consistent security architecture for dataspaces. Proven principles such as **Threat Modeling**, **Zero Trust, Defense-in-Depth, and Penetration Tests** must be applied consistently both while developing and operating the dataspaces.

To **enable broad adoption**, dataspaces must reduce both component complexity and onboarding complexity through modular architectures, clear layering concepts, and reusable building blocks. Most importantly, **security must be intrinsic and built-in by design**—enabling seamless, near **"one-click" onboarding with automated, adaptive, and AI-supported security mechanisms** that ensure both simplicity and robustness without requiring deep expertise.

4.2 Challenges related to the development of secure e2e dataspace components

The development of critical dataspace components (EDC, Identity Hub, Wallets) currently lacks a formalized security process. This results in several challenges that must be addressed when trying to create a secure dataspace:

- **Component Integration:** The integration of components emerging from various projects to create a secure e2e dataspace poses a serious challenge. The various projects use different architecture models of secure dataspaces, development concepts, and technologies to create the components, which can result, e.g., in components that are based on competing assumptions and incompatible interfaces.
- **Risk Management:** When using components to form a secure e2e dataspace, it is often unclear which risks their use might entail, as no clear risk management process was used during component development. Making the risks clear and transparent, and establishing a CRA (Cyber Resilience Act) risk management process during component development, poses a challenge for the community-driven development of e2e components. This includes clear definitions of intended use and the recognition and mitigation of threats.
- **Secure Development:** Similarly, base security requirements and controls are often unclear. Establishing an integrated DevSecOps approach to enable secure coding, CI/CD pipeline integration, and continuous vulnerability detection and patching poses a challenge for the community-driven development of e2e components.
- **Centralized Threat Model:** It is a challenge to create a centralized threat model and vulnerability database that are needed to identify, prioritize, assess risks throughout the lifecycle, and define mitigation strategies.
- **Lifecycle Management:** Component development should include lifecycle management and post-market surveillance (including data space services incidence response (DSIRT) and product incidence response (PSIRT) processes, security updates, patches, and forward compatibility).
- **Secure Deployment:** Components should offer guidance on securing the overall system, including the platforms, the components are deployed onto and all connected systems.
- **Adaptive Threat Modelling:** In an ecosystem with an increasing number of use cases and dataspace applications, it is challenging to establish adaptive Threat Modelling when Baseline Security Controls and concepts are missing.

In Chapter 9.5, we will describe how a secure development lifecycle of dataspace components can be defined.

4.3 Challenges related to enabling Scalability and large Rollouts

The rollout of the first operating company Cofinity-X, and the experience of several MX-Projects highlight several critical pain points that must be addressed to ensure secure, scalable, and interoperable data exchange across diverse dataspaces and enable a large rollout of M-X technologies:

- **Lack of Unified Security Understanding:** Because there is no globally shared understanding of the implemented security concepts, certifications, and standards, the initiative is encountering an overload of fragmented security definitions and dictionaries, leading to confusion and misalignment between stakeholders.

- **Missing Structured Onboarding:** To enable secure and consistent onboarding, participants need a standardized self-assessment and evaluation toolkit to validate their readiness and best practice guidelines for secure integration with the dataspace. This includes installation and configuration guidelines for EDC, identity components, and the underlying platforms. Security certifications should be required for participation. At least one organization should be designated to conduct the security certification/tests.
- **Secure and Interoperable Data Exchange:** Interoperability is challenged by differing security protocols, tools, and expectations between ecosystems. The establishment of Long-Term Confidentiality, Integrity, and Availability (CIA), including post-quantum encrypted data sharing across various dataspaces, poses a serious challenge.
- **"Crossboarding":** Onboarding new participants from other supply chains (worldwide) is complex due to diverse industry-specific security standards (e.g., different ISO frameworks) and varying maturity levels and security practices among participants.
- **Undefined Security Requirements:** It is unclear what minimum security requirements a participant must meet to join a dataspace. These requirements likely differ by industry sector (e.g., medical, aerospace, banking), use cases, and technologies (e.g., the different MX-Ports and connector concepts: Leo, Hercules, and Orion).
- **Unclear Acceptance Criteria:** CIOs and decision-makers need clear, credible arguments to trust and adopt the security concept. Current concepts lack transparency and validation through industry-accepted metrics or test results.
- **Measurable Security Assurance:** Currently, there is a lack of such measures, which could also help to present decision-makers with clear arguments for the use of dataspaces. Such measures would include the definition of security metrics (e.g., performance under attack, compliance levels), the evaluation of EDC & Eclipse Tractus-X components for robustness across a range of test scenarios, and the communication of what happens when components are used outside their intended design.
- **Defending Measures and Post-Market Surveillance:** For large rollouts, it is vital to define the post-market surveillance process and dedicated responsible bodies for all dataspace components employed. A Security Incident Response Team and protection mechanics must be defined to ensure that the processes are clear for when a dataspace component is compromised.

4.4 Expectations and questions to solve

The Topic Group has identified the following key expectations that must be addressed for broad acceptance of the Manufacturing-X security concept:

- What minimum and enhanced security requirements must participants meet, differentiated by **industry sector**, **role**, and **dataspace application security level (Basic, Trust, Critical)?**
- What clear, credible arguments and evidence do **CIOs and decision-makers need to trust the security concept**?
- How can we provide **measurable security assurance** (security levels of components, behavior under test cases, robustness in intended and unintended use, and incident response capabilities)?
- **How to implement adequate risk management processes** for operators, participants, use-cases, and the open-source community (EDC, Eclipse Tractus-X)?
- How to establish effective **cross-dataspace security coordination** (security intelligence sharing, joint incident response, consistent policy enforcement)?
- How can we **keep the dataspace secure while significantly reducing entry barriers** for all participants, especially SMEs?
- **How do we respond to emerging AI-driven threats**, such as advanced models from Anthropic (Mythos, Fable of the Claude Series) and other new AI systems, in terms of threat modeling, detection, and mitigation?
- **How to establish effective incident response and secure operations mechanics**, clearly defining the roles and responsibilities of participants, the OSS community, software developers, operators, **and state actors** (including eIDAS 2.0)?

- **What governance, auditing, and certification mechanisms** (including verifiable credentials and independent bodies) are required?
- What role will a Dataspace-AI play in the near future? Can AI take the role of a governance, auditing, and certification body in the future?
- How do we deal with emerging cybersecurity threats? How can we deal with a very real threat of cyber-attacks from rogue countries? A potential cyber war is becoming a possibility.

Addressing these questions transparently with concrete deliverables is essential for the successful scaling of Manufacturing-X dataspaces.

5 Operation and Governance Framework

We propose to fundamentally align the E2E security architecture for Manufacturing-X with data sensitivity and criticality by introducing a three-layer architecture, to which all Eclipse Tractus-X KIT use cases are systematically assigned according to the criticality and protection needs of the data. This graduated model enables participants to implement only the security layers they require, dramatically lowering the access barrier with lightweight Basic deployments while allowing seamless scaling to higher layers.

In addition, we propose to enhance the EDC with an “Industrial Data Hub” (an extended implementation of the Eclipse Tractus-X Data Exchange Orchestrator) that acts as the secure interface between internal IT/OT systems and the external dataspace, ensuring interoperability based on the AAS and automatically assessing criticality, validating intended use, with a full understanding of data criticality. While the “Industrial Data Hub” is mainly targeted at SMEs, it can ensure a fast enrolment of brownfield applications.

A central element of this holistic approach is enforcing **Security by Design,** combined with systematic **Risk Management** at all levels:

- **Risk management for open-source components**, OSS Risk Management for Eclipse Tractus-X EDC and Eclipse Tractus-X IdentityHub, etc.
- **Threat modelling of concrete use cases**, supply chain data exchange, federated applications, federated AI, OT integration
- **Operational risk management**, including centralized security intelligence, continuous monitoring, and **dedicated incident response teams and mechanisms**

Only through this comprehensive, **lifecycle-oriented security approach** can Manufacturing-X dataspaces achieve the required level of trustworthiness and scalability.

5.1 Abstract dataspace definition and required aspects to create secure and trusted dataspaces

Most of the current implementations can be condensed to an abstract view of a dataspace. Most participants in a dataspace use data sources or sinks located either on the IT or OT side. One or more gateways are used to exchange data with other participants.

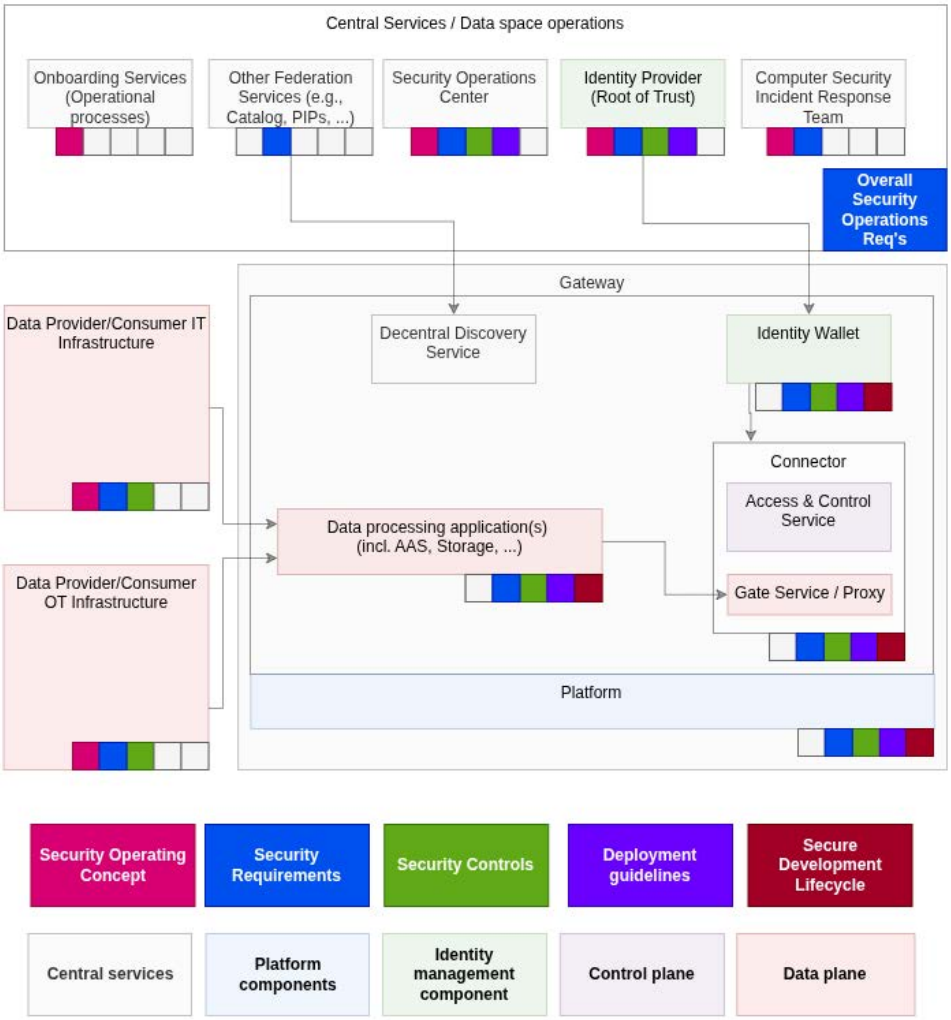


Figure 4 – Artefacts and Aspects needed for secure dataspace

To enable a dataspace that is secure and trustworthy, different aspects need to be tackled, and components need to be developed, as is depicted in Figure 4 and explained in the following sections.

5.1.1 Central Services / Dataspace Operations

Onboarding Services: These include organizational measures to let participants join the dataspace and technical measures to equip participants with necessary digital certificates or other authentication tokens provided by the Identity Provider to link given cryptographic material to participating technical entities.

Federation Services: These include components such as the Catalog, a central instance keeping track of available products and participating vendors of data sets, or Policy Information Points (PIPs), which provide information concerning policies.

Security Operation Center: To keep track of security incidents within the dataspace realm, all components must provide detected security alerts to the centralized Security Operation Center (SOC) of the dataspace. In the SOC, the general security and technical health status of the dataspace is surveilled. This way, the analysis of large-scale attacks on the dataspace can be measured, and cyber threat intelligence (CTI) can be generated and provided to all dataspace participants.

Identity Provider: This system serves as a central point for acquiring authenticating materials within the dataspace, such as digital certificates or authentication tokens. All connection setups between technical entities within the dataspace involve the PKI hosted by this entity.

Computer Security Incident Response Team (CSIRT): This organizational structure manages immediate actions in case of security breaches and may decide to block traffic to certain contaminated or hacked entities within the data space to minimize structural damage from attacks.

5.1.2 Technical Components of the Gateway

Identity Wallet: This technical entity holds the cryptographic material that must be used by the gateway to authenticate itself in the name of the participant. The key material should be stored in hardware-protected storage (e.g., TPM). The corresponding certificates or authentication tokens are provided by the Identity Provider.

Connector: This central component manages connection setups to other connectors within the dataspace and acts as a connection point for data transmission.

Data Processing Applications: These applications may store, analyze, secure and transfer data from the participant's OT or IT infrastructure, perform data operations, and pass on the data further to the connector for data transmission.

Platform: This component covers hardware, firmware, the hypervisor, and the operating system hosting the other technical components of the gateway, which make up the software stack of the gateway. In case confidential computing is required for a gateway, this platform must support a Trusted Execution Environment (TEE) and remote attestation capabilities for additional security guarantees.

5.1.3 Aspects of a Dataspace

Secure Operating Concept: All components interacting with the gateway of the dataspace must comply with a secure operation concept, which must cover all areas from organizational measures concerning the onboarding process, to operation and management of running software artifacts, and interacting third-party services.

Security Requirements: Before security concepts can be developed, the fundamental prerequisite is to gather the technical and organizational requirements that must be maintained for security. This aspect also covers almost all areas of the dataspace and must be done first.

Security Controls: The best concept and carefully executed initialization is useless unless it can be security-wise surveilled and checked during the execution. Therefore, all technical components taking part in data transmission must provide logs, analyses, and security reports.

Deployment Guidelines: All components that span direct communication within the dataspace must be deployed correctly to ensure security. However, the deployment aspect only covers the dataspace gateways and connectors and their underlying platform, the dataspace's identity provider, and the security operation center. Already deployed areas in the OT and IT do not need to be covered in this guideline.

Secure Development Lifecycles: Security is a process, not a product. Therefore, the gateway and all its internal components must be regularly improved and security-wise checked through a secure development lifecycle, including frequent code-vetting to address new features being demanded in the data processing applications. In parallel, the underlying platform must be continuously patched and maintained as part of secure operations. The gateway's permanent security is paramount for the dataspace.

Figure 5 shows an abstract implementation reflecting the components described above. The identification of the participants is implemented, facilitating a DID document. The trust base between companies is established based on the framework of agreements and membership conditions imposed by the operating company. The EDC serves as a proxy for data access by Data Consumers.

Trust is built between the two companies using Verifiable Identities and agreements, which allow both companies to enable access to internal business applications and storage systems. This trust is maintained by a third trust anchor, which owns the "root of trust", which is operated by a third party known and trusted by both parties. And of course, to answer the question, "But why would I trust the root of trust?", the answer is "Governance", which is currently provided in dataspaces like Catena-X, via an Association, which nominates the companies acting as "root of trust" and defines the operation rules of the game (Operating Model).

5.2 Current Dataspace Architecture

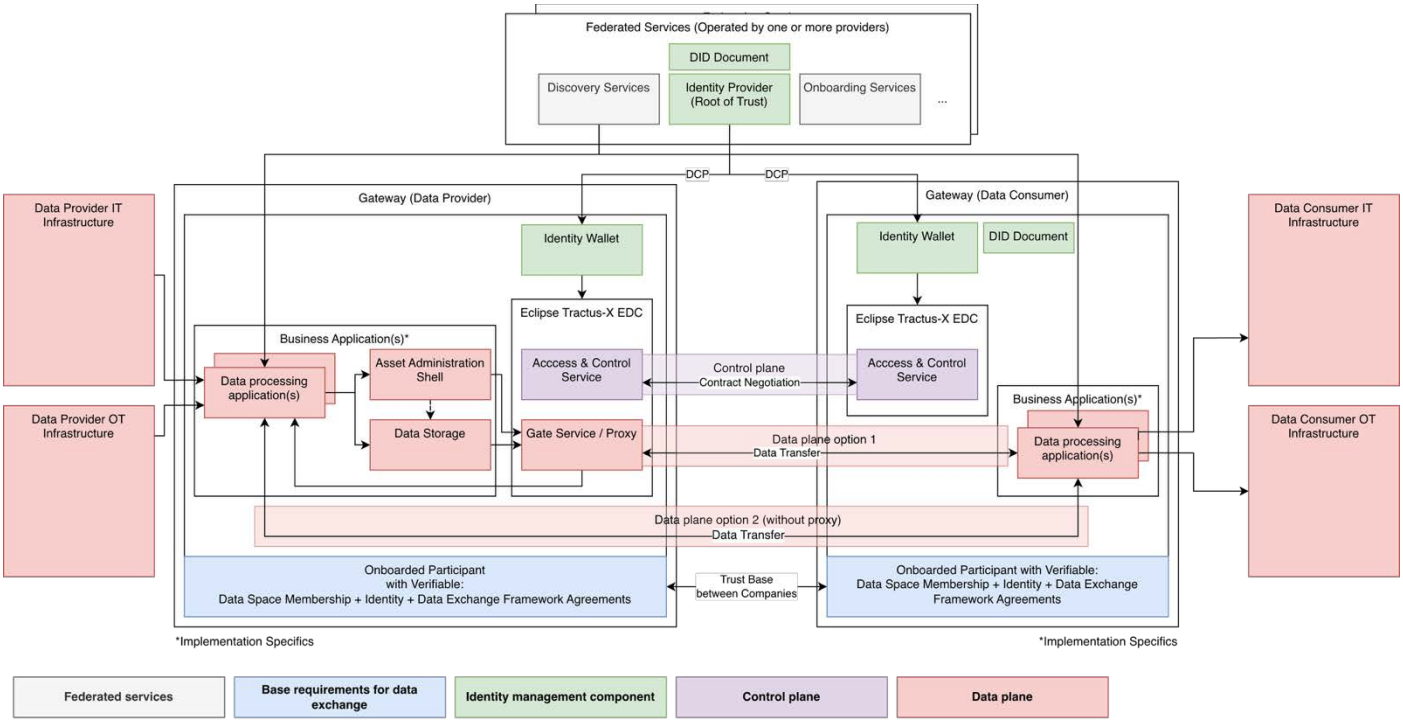


Figure 5 - Current abstract example implementation (e.g., in Catena-X)

5.3 Emergency Containment and Major Incident Response

Federated dataspaces are built on the principles of decentralization, interoperability, and data sovereignty. Such decentralized architectures create specific challenges when responding to cyber attacks or critical vulnerabilities. In such situations, the ecosystem must react quickly and consistently without undermining the autonomy of its participants.

This chapter provides an initial overview of containment measures. Additional detailed analysis and development are required to fully address the topic.

When an ongoing cyber attack is detected, the primary objective is to contain malicious activity while maintaining trust and operational continuity across the ecosystem. Unlike centralized cloud platforms, a federated dataspace cannot simply be “switched off” through a single technical mechanism. Instead, response strategies focus on the trust and interoperability layers. Compromised participants, connectors, or identities can be isolated through temporary credential revocation, suspension of trust relationships, or restrictions on data exchange. In practice, this functions as a controlled containment mechanism that prevents affected entities from continuing to interact with other participants until the incident has been resolved. Roles and actors will have to be fully addressed with further detailed work.

Such measures require clearly defined governance structures and emergency procedures. The ecosystem operator or trust authority must be able to coordinate incident response, distribute trusted security advisories, and define escalation paths for high-severity situations. At the same time, transparency and accountability are essential to preserve confidence in the federation. Emergency actions should therefore be formally documented, time-limited, and auditable. The goal is not to centralize operational control, but to ensure that the ecosystem can react coherently during crises.

In addition to governance measures, technical resilience plays a major role. Dataspaces should support segmentation and graceful degradation, allowing parts of the ecosystem to continue operating even when others are affected. Operators may temporarily disable selected services, protocols, or interoperability functions while maintaining essential operations. Coordinated threat intelligence sharing is equally important. Participants should rapidly exchange indicators of compromise, attack patterns, affected software versions, and mitigation recommendations in order to reduce reaction times and prevent lateral spread across the federation.

A similar approach applies when a critical vulnerability is discovered in a shared component or service. In this scenario, the challenge is usually less about immediate containment and more about coordinated remediation before exploitation occurs. Effective response requires a trusted vulnerability management process capable of identifying affected participants, distributing mitigation guidance, and defining remediation timelines. Software transparency mechanisms

such as Software Bills of Materials (SBOMs) and dependency tracking become particularly valuable because they allow organizations to quickly determine whether they are exposed to the vulnerability.

Federated dataspaces should therefore define staged operational modes that can be activated depending on risk severity. Under normal conditions, participants operate with full interoperability. During elevated risk situations, additional monitoring, stricter policy enforcement, or temporary interoperability restrictions may be introduced. In extreme cases, vulnerable services or connectors may be temporarily excluded until patches or mitigations have been applied.

Ultimately, the resilience of a federated dataspace depends on balancing decentralization with coordinated trust governance. The most practical and sustainable strategy combines decentralized operations with transparent, auditable, and narrowly scoped emergency coordination mechanisms that enable rapid containment, recovery, and restoration of trust across the ecosystem.

6 Three-Layer Dataspace Security Architecture

Following the core principle of scalability of security mechanisms, our security architecture must provide a selection of appropriate security measures for given use cases. To reduce complexity and simplify onboarding, we base our security architecture on a three-tiered data criticality classification. Each data criticality is associated with a security level, which defines the security measures necessary to protect data assets of this criticality.

Data Criticality	Use Case Examples	Security Level	Security Characteristics
Semi-public	Eco Pass Kit, Digital product passports (metadata), Public IDLink, Product specifications, Product Carbon Footprint Kit	**Basic**	Lightweight protection, high interoperability
Confidential	Circularity Kit, Demand and Capacity Management Kit, Logistics Kit, Purchase Order Kit, Requirements Kit, Commercial Documents Kit, Material Accounting Kit	**Trust**	Policy agreements, data encryption, system integrity verification, auditability
Strictly confidential	IP (blueprints, recipes, algorithms), Data Driven Quality Kit, Online Simulation Kit, Model Based Production Kit, ...	**Critical**	Hardened nodes, technical policy enforcement, secure tunnels, hardware-backed attestation, continuous monitoring

Based on the security by design principle, we propose a three-layer security architecture for Manufacturing-X that uses the security levels corresponding to the data criticality levels defined in the previous section. This layered approach balances security effort with functionality, performance, cost, and ease of use, depending on data sensitivity and use-case criticality.

It defines three distinct yet fully interoperable dataspace layers, corresponding to the three security levels: Basic, Trust, and Critical. Each layer is equipped with security controls, technical requirements, deployment models, and running costs, which are tailored according to the security level that this layer implements:

- **The Basic dataspace layer** provides lightweight, cost-efficient protection with minimal overhead — dramatically lowering the entry barrier for SMEs and low-risk use cases.
- **The Trust dataspace layer** adds enhanced policy enforcement, auditability, and commercial-grade safeguards for confidential business data.
- **The Critical dataspace layer** delivers the highest assurance for intellectual property, OT telemetry, and mission-critical applications.

All layers share a common data fabric but apply differentiated security stacks. Cross-layer communication is governed by strict, predefined rules and secure channels. Participants can choose to implement the security layer(s) that best match their needs and capabilities. The dataspace gateways and network infrastructure must be designed and implemented to meet the layer's security requirements.

6.1 Associated Security Controls in the Gateway

To ensure adequate protection of exchanged data, each criticality level and dataspace layer implies the necessity of certain standard **security controls**. Security controls are technical, organizational, or physical measures that reduce cybersecurity risk to IT/OT systems and data. As such, they lie at the heart of every cybersecurity concept.

To achieve a secure dataspace infrastructure, appropriate security controls must be identified and implemented according to the relevant IT/OT security standards and guidelines. Typically, a use-case-based threat modelling and risk assessment process is used to identify the security requirements of systems/components and select appropriate security controls depending on the individual protection needs. The effectiveness of the implemented security controls must be continuously (re-) evaluated as part of the established cybersecurity program (ISMS/CSMS). This common methodology is closely linked to the relative industry sectors and their respective Regulations and Guidelines (e.g., ISO 81001-5 for medical devices).

Gateway Security Levels

Crucially, our security architecture must identify which technical security controls are necessary at the dataspace gateways to adequately protect data assets of different criticalities. To integrate with a risk-based cybersecurity approach

and preserve flexibility, we refrain from obligating concrete technical solutions and configurations. Instead, we define which security measures gateway implementers must consider as countermeasures during the detailed risk assessment. The resulting control catalogue offers a time-saving approach to achieve a common understanding of gateway security over different use cases. However, it is not intended to fully replace traditional use-case-based risk assessment. Additional requirements arising from the use cases, which have not been defined in the per-level common risk assessment, still necessitate individual threat modelling.

		Gateway Security Levels		
Category	**Common Security Controls**	**Basic**	**Trust**	**Critical**
Authentication & Access Control	EDC, AAA services, OAuth, …	✓ Company level	✓ Department level	✓ Individual level
Usage Control Enforcement	EDC, ODRL, Policy-aware applications	○	✓ Connector level	✓ Application level
Data protection in transit	Secure communication protocols (TLS, IPSec, …), Secure tunnels (VPN)	✓	✓	✓ Use hybrid PQC
Data protection at rest	Symmetric encryption	○	✓	✓
Data protection in use	Confidential Computing, Encrypted VMs	○	○	✓
Platform integrity protection	Secure Boot, Trusted Boot, TPMs	○	✓	✓
Proof of integrity / enforcement	Remote Attestation	○	○	✓
Secure hardware	Secure elements, HSMs	○	(✓)	✓
Availability	DoS protection, Rate limiting, Load balancing	(✓)	✓	✓

○ Usually not required. (✓) Should be considered. ✓ Must be considered.

A gateway at **Security Level Basic** fulfills the minimal requirements for a secure dataspace infrastructure, including participant authentication and secure data exchange. A gateway satisfying **Security Level Trust** offers additional security controls to protect data at rest and to restrict data processing. The latter is achieved through usage control policies that define under which circumstances data may be transferred and processed on the data provider side. The level also introduces a secure platform that ensures the gateways' integrity. To adequately protect even strictly confidential data, gateways of **Security Level Critical** include a secure platform that ensures and attests the gateways' integrity and capability to enforce adherence to the defined usage control policies in the connector and all utilized data applications. This includes authenticating and verifying the hardware-backed platform identities and software states of gateways.

Note that security controls applied to the gateway alone are not sufficient to cover all data criticalities. For semi-public data, access control enforced by the gateway may be adequate. However, for strictly confidential data, it must be guaranteed that only the intended receiver can decrypt the data. This requires end-to-end encryption performed directly at the point of origin (where the data is generated), not solely at the gateway level.

Control Catalogues

Typical Risk Management is difficult to apply in dataspaces where new use cases emerge continuously and form a kind of moving target. "G*lobal ecosystems require comprehensive security architectures for all participants, regardless of regional and national structures. This necessitates* ***interoperable cybersecurity guidelines*** *that are recognized by all parties involved, as well as a shared regulatory framework*" (VDE-AR-E 2802-20), *and can be commonly used by the different use cases*. Therefore, to assist dataspace participants and developers in identifying necessary security controls for their scenarios during a cybersecurity risk assessment, we propose Manufacturing-X to adopt a procedural model in accordance with VDE-AR-E 2802-20 (Horizontal cybersecurity concept for Industry 4.0), providing **catalogue-based cybersecurity concepts and controls as baseline requirements**. The catalogues of concepts and controls are tiered into three Gateway Security Levels, each corresponding to one of the presented data criticality levels.

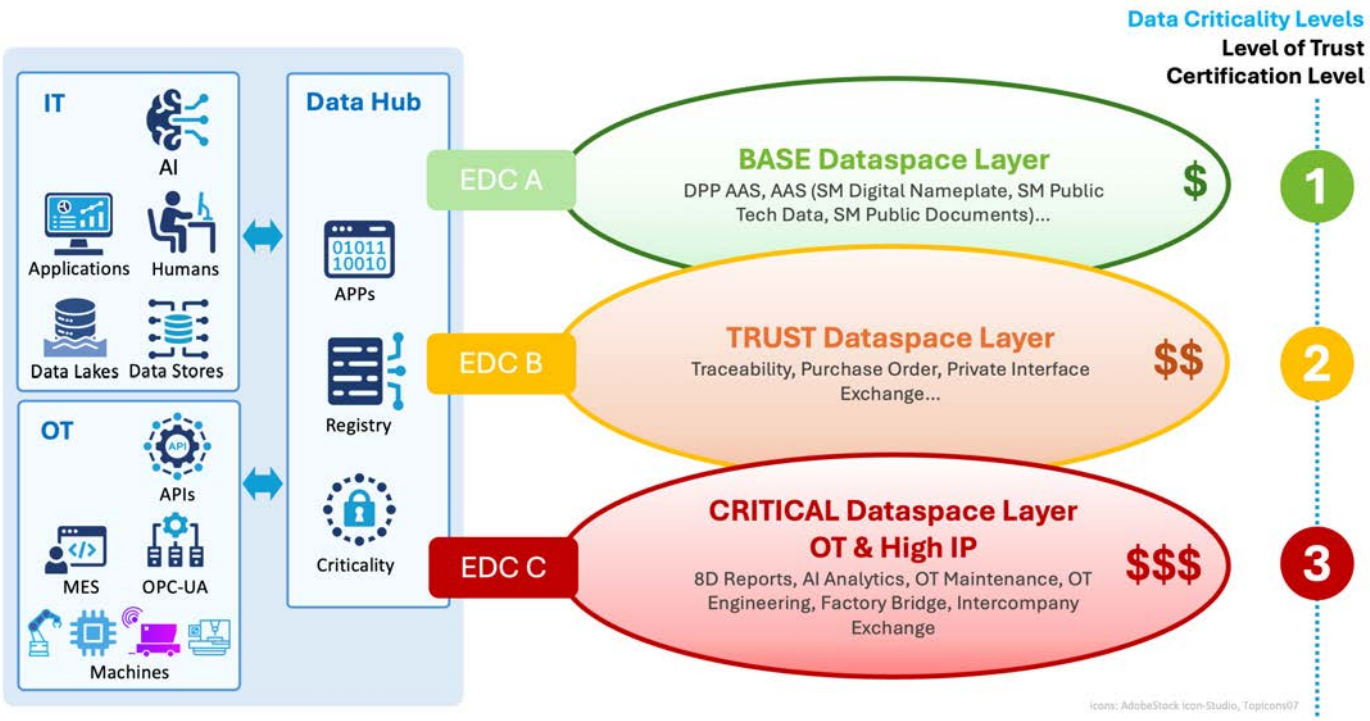


Figure 6: A three-layer staggered, use-case-oriented security architecture balances security effort with functionality, performance, cost, and ease of use, depending on the sensitivity of the data and the criticality of the use case.

This graduated model finally delivers the long-promised "one fits all" solution: every participant connects via a single controlled connector (EDC described in the framework of the MX-Port Herkules and Orion) of a specific security level. The security strength, controls, implementation effort, and operating costs are scaled precisely to the levels and use cases the participant has chosen to support.

By enabling participants to select the right dataspace layer for their specific business needs, the architecture achieves an optimal balance between strong security, usability, scalability, performance, and cost efficiency. It enables broad adoption across the entire manufacturing ecosystem while ensuring that higher-criticality data and use cases receive the robust protection they require. Participants can start at the Basic layer and upgrade to the Trust or Critical layer at a later stage, if required by the implemented Eclipse Tractus-X Use Cases (KITs).

6.2 Data Provider Security Architecture

By using or implementing a dataspace infrastructure based on the best practices from a security perspective, the main security strategy should be the **Zero Trust approach,** combined with a **Layered Architecture** scope. A high level of security always implies complexity, additional checks, and redundancy. Applying the three-layer security principle to the data provider securely protects data of different data criticality:

- **Application Access & Usage Sovereignty:** The outer sovereignty layer handles access control and secure exchange of data with other dataspace participants (Semi-public data → Security level Basic).
- **Data Access & Usage Sovereignty**: Internal applications use more comprehensive submodels, which are valuable and must be protected accordingly (Confidential data → Security level Trust).
- **Data Source Sovereignty**: Data that is sourced directly in the IT/OT infrastructure contains mission-critical information and IP. Hence, these data assets require the highest possible degree of security and isolation (Strictly confidential data → Security level Critical).

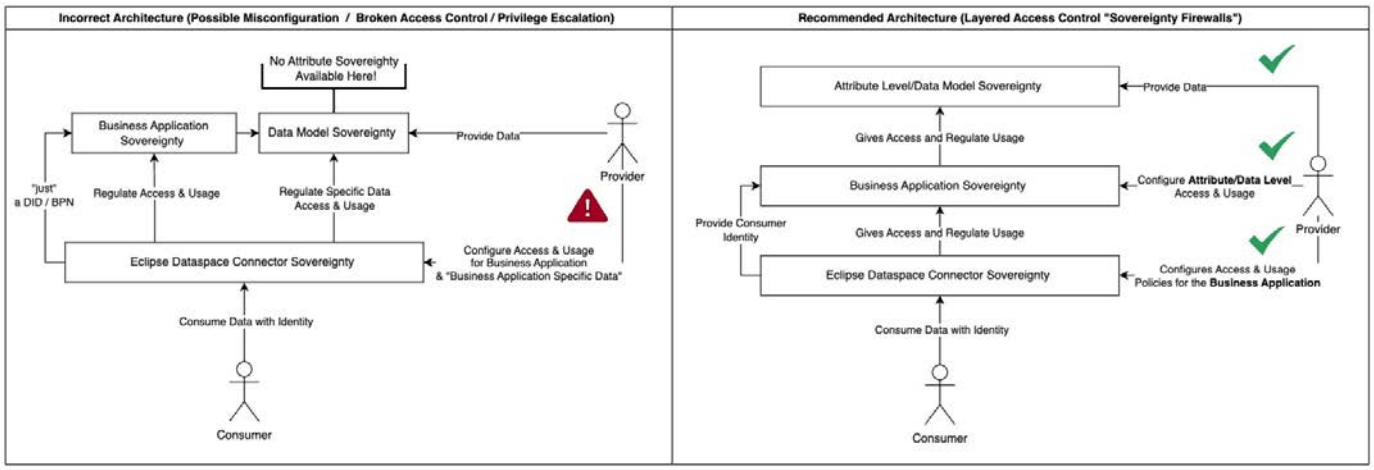


Figure 7 Data sharing architecture recommendations and risks

It is important to note that the way the underlying authorization wiring is made can influence the overall robustness of the complete internal dataspace setup, keeping its Data Sovereignty flexible enough and granularly secured. Figure 7 shows two ways to configure data access and usage. On the left side, the access permissions are configured only in the Eclipse Dataspace Connector, both for specific data endpoints (directly the data APIs) and business application systems (AAS, PCF App, etc). On the right side, data access and usage are configured using a layered architecture for more robustness.

There exists the misconception that every single piece of data from a company is “directly” configured in the Eclipse connector. However, that is not the case for what the EDC Connector was designed for. This misuse of the EDC poses a real risk for Dataspace Participants, who do not have access to the information on how to build their dataspace setup internally in a secure way, which may lead to several vulnerabilities risks (see Section 3.1).

The configuration on the right side of Figure 7 uses the EDC Connector to regulate the “Usage” of specific data endpoints/systems, mostly on “Type” level data endpoints. Access control at the EDC Connector is mostly only related to the “Membership” / “Identity” / “Framework Contracts” validation of the participants. It is not meant to duplicate the configuration of access permissions for every single piece of data in a company. Each business application/system which is providing the APIs configured in the EDC with attached Usage Policies and Access limitations, has the necessary Business Domain Knowledge to restrict the best the access to each data endpoint, even if general access was already granted in the EDC, so to have a complete secure data sharing configuration, additional layers of granular security permissions must be also defined in the applications behind the EDC.

By following this three-layered approach, the following vulnerabilities (see Section 3.1) can be prevented:

- Involuntary data disclosure due to complexity in understanding data criticality and implementing access rules
- Data poisoning and manipulation
- Broken access control and privilege escalation at the connector boundary
- Internal secret and credential exposure via dataplane proxy

Recommended Layered Security Patterns

To prevent this EDC from being a “single point of configuration” insecure configuration, there are two recommended security patterns/approaches to set up the data provider side:

1. **Layered Architecture with Verifiable Identities + Internal Identity Provider (see Section 6.2.1)**
    a. This is an option broadly used by applications offered in the productive Catena-X Dataspace, where application providers set up a second layer of security when it comes to accessing APIs via the dataplane.
    b. It is useful since most of the internal company applications that already exist already contain some authorization schemes that use OAuth2 or OpenID Connect technologies.
2. **Layered Architecture with only Verifiable Identities**
    a. This option is currently not yet available in reference implementations. However, it would be optimal for securing access to internal resources, thanks to the reusability of the verifiable credential identity provided during the contract exchange for the EDC Connector. In this way, not only is the DID identity from the participant proxied, but also its identity credentials, provided in the first security layer.

6.2.1 Layered Architecture protected with Verifiable Identities & IDP

Securing data is a very important task when it comes to sharing it via a Dataspace. And for that, considering reusing synergies available at a company's infrastructure can be beneficial when designing your dataspace setup. Most companies already have some sort of Identity Provider in their company infrastructure. Therefore, to secure their dataspace applications, they can reuse existing internal identity providers and configure them at the EDC Control Plane for specific "client ids" which would be configured for a specific "asset" (an EDC Management Api proxy endpoint securing concept). As observed in the diagram below, the Dataspace Application would still have the responsibility of determining which identity profile should have access to which specific data granularity. Since the Dataspace Application has much more business domain knowledge, for limiting the scope of the authentication to its own data endpoints (APIs).

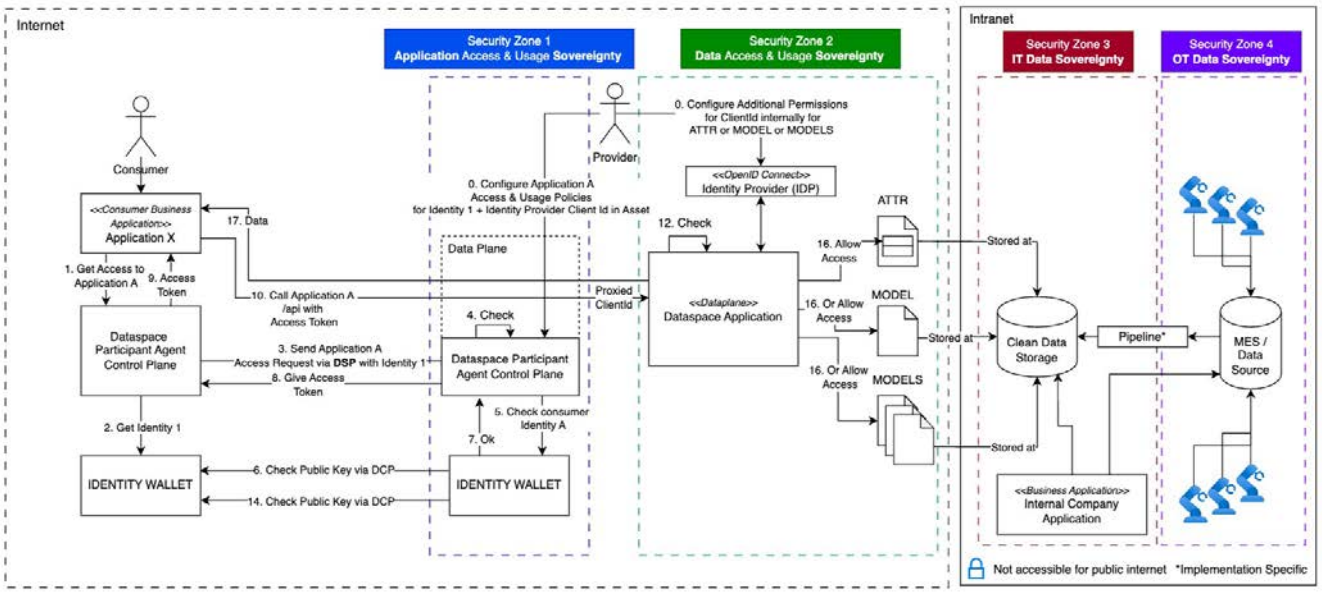


Figure 8 Dataspace layers Authorization model using a combination of existing Identity Providers and Verifiable Identities.

Two good examples of reference implementations that follow this pattern are the Eclipse Tractus-X Digital Twin Registry and the Eclipse Tractus-X Industry Core Hub.

6.2.2 Layered Architecture with only Verifiable Identities

Another alternative to the combined usage of IDP and Verifiable Identities would be to do the authentication at the Dataspace Application level based only on the Self Sovereign Identities (SSI) concept. In this way, the SSI mechanisms of authentication (Identity Wallet) can be reused by business applications to scope authorization in a more granular way for resources. The Dataspace Application knows best its own interfaces, data models, and APIs to "restrict" the access permissions additionally to the first security zone.

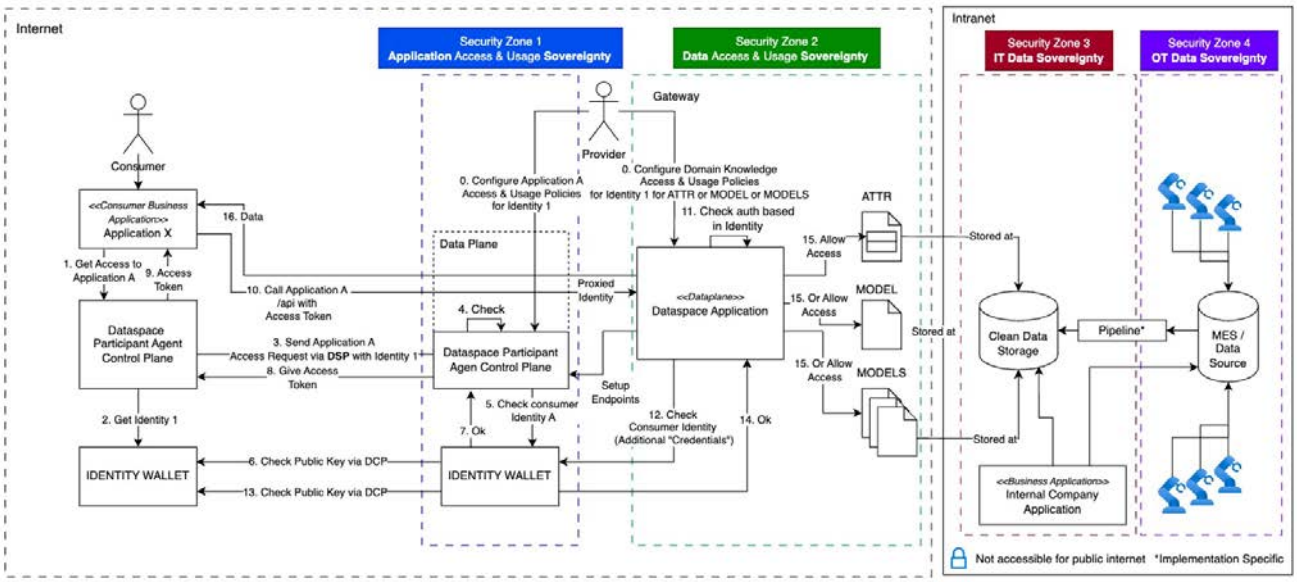


Figure 9 Dataspace layers Authorization model using only Verifiable Identities to secure dataspace applications + EDC

In the diagram above, a third level of security is described. This third level is a very critical infrastructure of a company, usually only accessible via VPNs and internal intranet connections. It is not recommended to connect your Dataspace infrastructure/apps directly to these systems, for example, by skipping security zone 2 and embedding/configuring an EDC Connector directly at the machine software on your shop floor. This would be very insecure and offer a risk of operation from your production systems. So, the recommendation here is, in this case, to keep a separation of concerns, extract the data from your internal systems/machines, collect it, and clean it at a safe data storage (which can then be accessed by your dataspace applications and other internal company applications). In this way, you keep your data formatted for using it at a Dataspace Use Case Data Exchange, for internal Industrial Artificial Intelligence usage, data dashboards, specific business applications, and many more!

6.3 OT/IT Integration and Shopfloor Considerations

Manufacturing-X applications and use cases may source data from internal IT/OT systems. As a result, clear boundaries are required between the dataspace and internal systems. The MX-Port concept is used to securely bridge OT data while maintaining zero-trust principles at the edge.

In addition, security requirements on the level of the dataspace must be extended into the OT domain, guided by OT security frameworks such as IEC 62443. Since the underlying OT infrastructure acts as a data source (and potentially sink) in Semiconductor-X and Manufacturing-X applications, its resilience against cybersecurity threats must also be considered for an overall security concept. To achieve end-to-end security in industrial dataspaces, each dataspace participant must establish a suitable Cybersecurity Program for their IACS systems based on risk assessments. Furthermore, IACS-specific security goals, threats, vulnerabilities, and consequences should be integrated into the security risk assessments on the level of the dataspace, in order to gain an understanding of which IACS security controls and countermeasures are required to achieve the dataspace security requirements. For example, if strictly confidential data, which may be used on the OT level, is shared with a dataspace participant, the target security level (SL-T) for Foundational Requirements FR2 (Use Control) and FR4 (Data Confidentiality) should be chosen accordingly by the data receiver. Similarly, if the integrity of input data sourced from a remote OT network is critical for a dataspace application, the data provider should choose an appropriately high target security level for FR3 (System Integrity) and implement the resulting IACS security controls and countermeasures. This way, a risk-based approach to cybersecurity allows for translating security requirements originating in the dataspace down to the OT level. In the future, Manufacturing-X should provide additional guidance for this translation, e.g., by mapping dataspace requirements and security levels to concrete security requirements for IACS systems and components (IEC 62443 parts 3-3 and 4-2).

7 Data-exchange Gateway and Orchestrator

The proposed Manufacturing-X architecture introduces the Dataspace Gateway (see Figure 9) as the central technical boundary between a participant's internal IT/OT environment and the external federated dataspace. The gateway is composed of several tightly integrated zones that work together to enforce security, data sovereignty, and usage control while ensuring seamless interoperability.

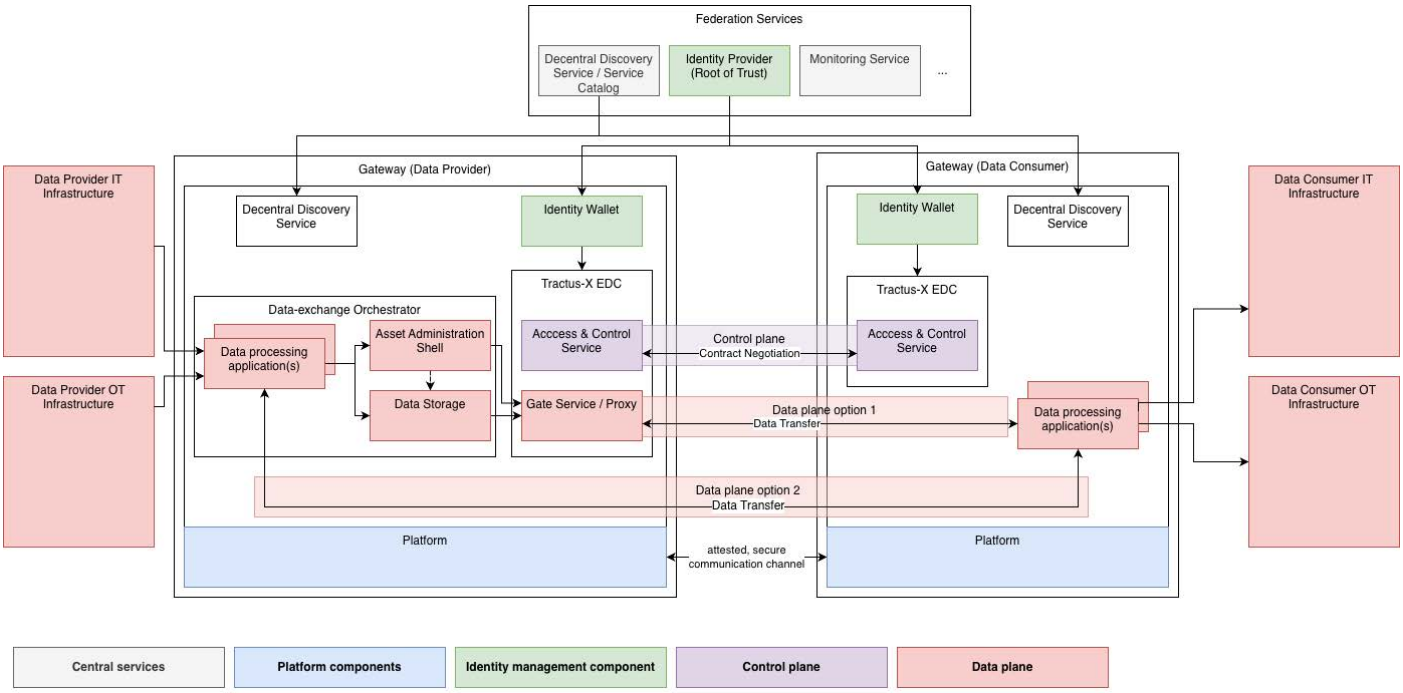


Figure 10 Reference Secure Implementation for data exchange in Manufacturing-X dataspaces

Gateways serve as main access points to the dataspace. We view them as holistic systems, which integrate all components that are essential for achieving the data processing mission while protecting the security of the exchanged data assets. Hence, gateways include (but are not to be confused with) standardized connector components such as the EDC, as well as the data processing applications themselves. Gateways are responsible for managing the control and data plane, negotiating contracts, initiating the transfer of data, and implementing access and (connector-level) usage control. The data plane may either be integrated into the connector, effectively making it a proxy application that is directly involved with the data transfer, or it can run outside of the connector. For both options, it must be ensured that the data transfer on the data plane is only performed when a data sharing agreement exists, both communication partners are authenticated, and defined prerequisites (e.g., necessary security requirements/restrictions defined by usage control policies/contracts) are fulfilled. The gateways may also be connected to the company's internal IT and OT infrastructure, serving as a data source/sink. Depending on the required security level, the gateway must also provide a secure platform that allows for local integrity protection, confidential computing, and remote attestation.

The components of the gateway are mapped into security zones in accordance with the concept of IEC 62443 (Industrial Automation and Control Systems Security), where security is built on two fundamental constructs: The **Security Zone,** a logical or physical grouping of assets (systems, applications, devices, or data) that share common security requirements and the **Conduit,** a controlled communication path (logical or physical) that connects two or more zones. Conduits enforce security policies, inspect traffic, and ensure that only authorized data flows between zones. They act as the "security boundaries" between zones.

The diagram adapts the IEC 62443 zone-and-conduit model to a dataspace participant architecture (combining IT/OT systems with Eclipse Tractus-X / Catena-X connectivity). The zones are organized by trust level and exposure:

Zone	Name in Diagram	Description	Trust / Exposure*	Security Focus
Zone 4	Industrial OT Applications	Shopfloor systems: MES, OPC-UA servers, machines, APIs, controllers	Highest internal trust	Highest protection (OT integrity & availability)

Zone 3	IT Applications	Enterprise IT: Applications, AI systems, humans, data lakes, data stores	High internal trust	Strong protection for business/IT data
Zone 2	Data-exchange Orchestrator (Industry Core Hub, Industrial Data Hub)	Central middleware: Orchestrator, AI Criticality Proof, Apps, Data Fabric/Registry	Medium trust (internal gateway)	Dynamic criticality assessment + policy enforcement
Zone 1	Dataspace Edge & Connector / DMZ	Exposed edge: EDC (Control & Data Plane), WAFI, Identity Wallet	Low trust (dataspace-facing)	Strong perimeter controls, authentication, monitoring
Zone 0	External / Internet	External participants, remote identities, internet services	No trust	Public / untrusted network

**Important Note on Trust Levels: The trust levels assigned to the zones (Zone 0 – Zone 4) are intended solely to create a common understanding of potential attack vectors and the relative exposure of different parts of the architecture. They do not imply a reduction of the Zero Trust concept.*

Figure 10 highlights the actual internal component structure of a secure dataspace gateway – including the Data-exchange Orchestrator. This abstract architecture is component-based rather than data-oriented. It can serve as a blueprint for upcoming projects and future implementations. The next sections will describe the actual Eclipse Tractus-X implementation of the Data-exchange Orchestrator, the "Industrie Core Hub", and the security-enhanced SME reference architecture, "Industrial Data Hub", with AI-supported security components proposed by the topic group and under implementation in Semiconductor-X. Both perspectives are valid and interoperable. They demonstrate that the graduated security model can be realized through different technical realizations while preserving the core principles of data criticality, sovereignty, and federation. The following subsections describe the reference implementation chosen for Manufacturing-X.

7.1 Industry Core Hub (One open source implementation of the Data Exchange Orchestrator)

The Industry Core Hub is the open-source reference implementation of the Eclipse Tractus-X Industry Core KIT. It serves as a middleware layer that connects your business applications and legacy systems to the Tractus-X dataspace — hiding the complexity of its individual components.

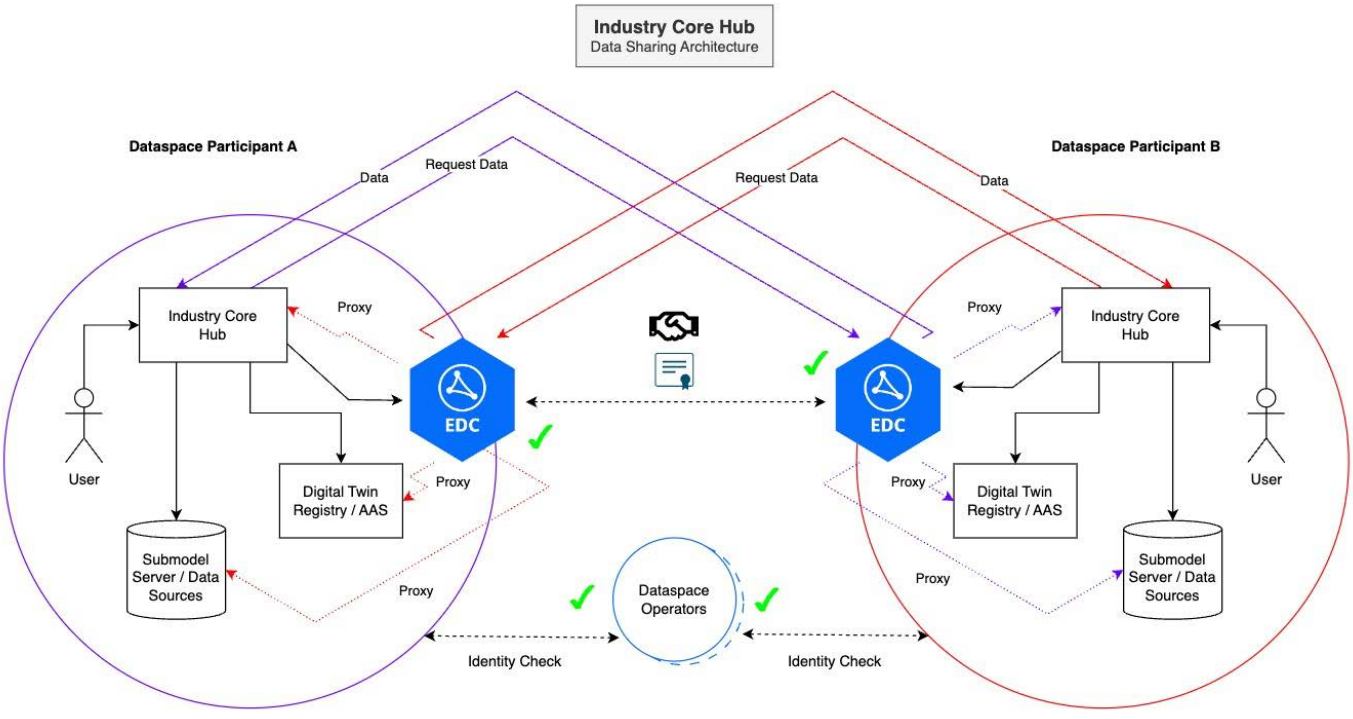


Figure 11 Secure Data Sharing Orchestration Architecture from the TX Industry Core Hub in an Operative Dataspace Setup.

Adopting these technologies typically requires integrating the Eclipse Dataspace Connector (EDC), Digital Twin Registry (DTR), and Discovery Services. Setting up policies, asset registrations, and data contracts correctly demands weeks or months of specialized engineering work that most organizations lack in-house.

For easing the configuration effort, the usage of an orchestration system for sharing and consuming data in a safe and secure manner is important. In the end, for every dataspace participant user, the complexity of data sharing should be very minimal. Reference implementations like the Eclipse Tractus-X Industry Core Hub provide several levels of abstraction for the interaction of a user in a dataspace, reducing the risks of a misconfiguration to a minimum. In the following diagram (See Figure 11), we can observe the architecture, where all the different security levels are configured by the Industry Core Hub.

Currently the Industry Core Hub is being developed in collaboration with Catena-X, Railway-X, Construct-X and Semiconductor-X.

7.2 Industrial Data Hub Implementation – Topic Group extension of the Gateway

To strengthen security, simplify the onboarding of new use cases, and introduce a semantic data layer for Asset Administration Shells in the dataspace, we propose the **Industrial Data Hub**.

Positioned at the core of the Dataspace Gateway, the Industrial Data Hub (Zone 2 – Dataspace Application Service) serves as the intelligent, data-centric orchestrator. It bridges internal IT and OT systems with the external dataspace while enabling dynamic data classification and policy enforcement (see Figure 10).

The Industrial Data Hub is divided into four main functional blocks:

- **Orchestrator –** Manages data routing pipelines, fractional AAS processing, and aggregated AAS repositories. It coordinates the flow of data between internal sources and the dataspace.
- **AI Criticality Proof –** The AI-augmented core component that performs real-time data criticality assessment, contract and request/transfer proof, contract interpretation/mapping, and evaluation of PII/privacy obligations. It automatically determines the required security level (1, 2, or 3) for each data flow.
- **Apps –** Provides reusable application services such as Converter, Projections, Aggregator, and Industry Core that transform and prepare data according to use-case requirements.
- **Data Fabric / Registry –** Serves as the central repository for AAS native storage and ECS storage, maintaining a consistent, queryable view of all assets and digital twins.

The Industrial Data Hub is positioned between the strictly controlled conduits coming from internal IT (Zone 3) and OT (Zone 4) systems on one side and the Dataspace Edge & Connector (Zone 1) on the other. All data entering or leaving the participant passes through the Industrial Data Hub, where AI-driven criticality assessment and policy enforcement are applied before the data reaches the EDC Control Plane and Data Plane.

This design ensures that:

- **Data criticality** is evaluated before data leaves the company's secure perimeter.
- **The correct security level** (Basic, Trust, or Critical) is **automatically** applied.
- **Usage control policies** are enforced as close as possible to the data source.
- **Interoperability** with the external dataspace is maintained while preserving **full data sovereignty**.

The combination of the Industrial **Data Hub** (Zone 2) and the **EDC / WAFI** (Zone 2) forms the complete **Dataspace Gateway**, providing layered, defense-in-depth protection that scales with the chosen security level while remaining practical and achievable for industrial participants.

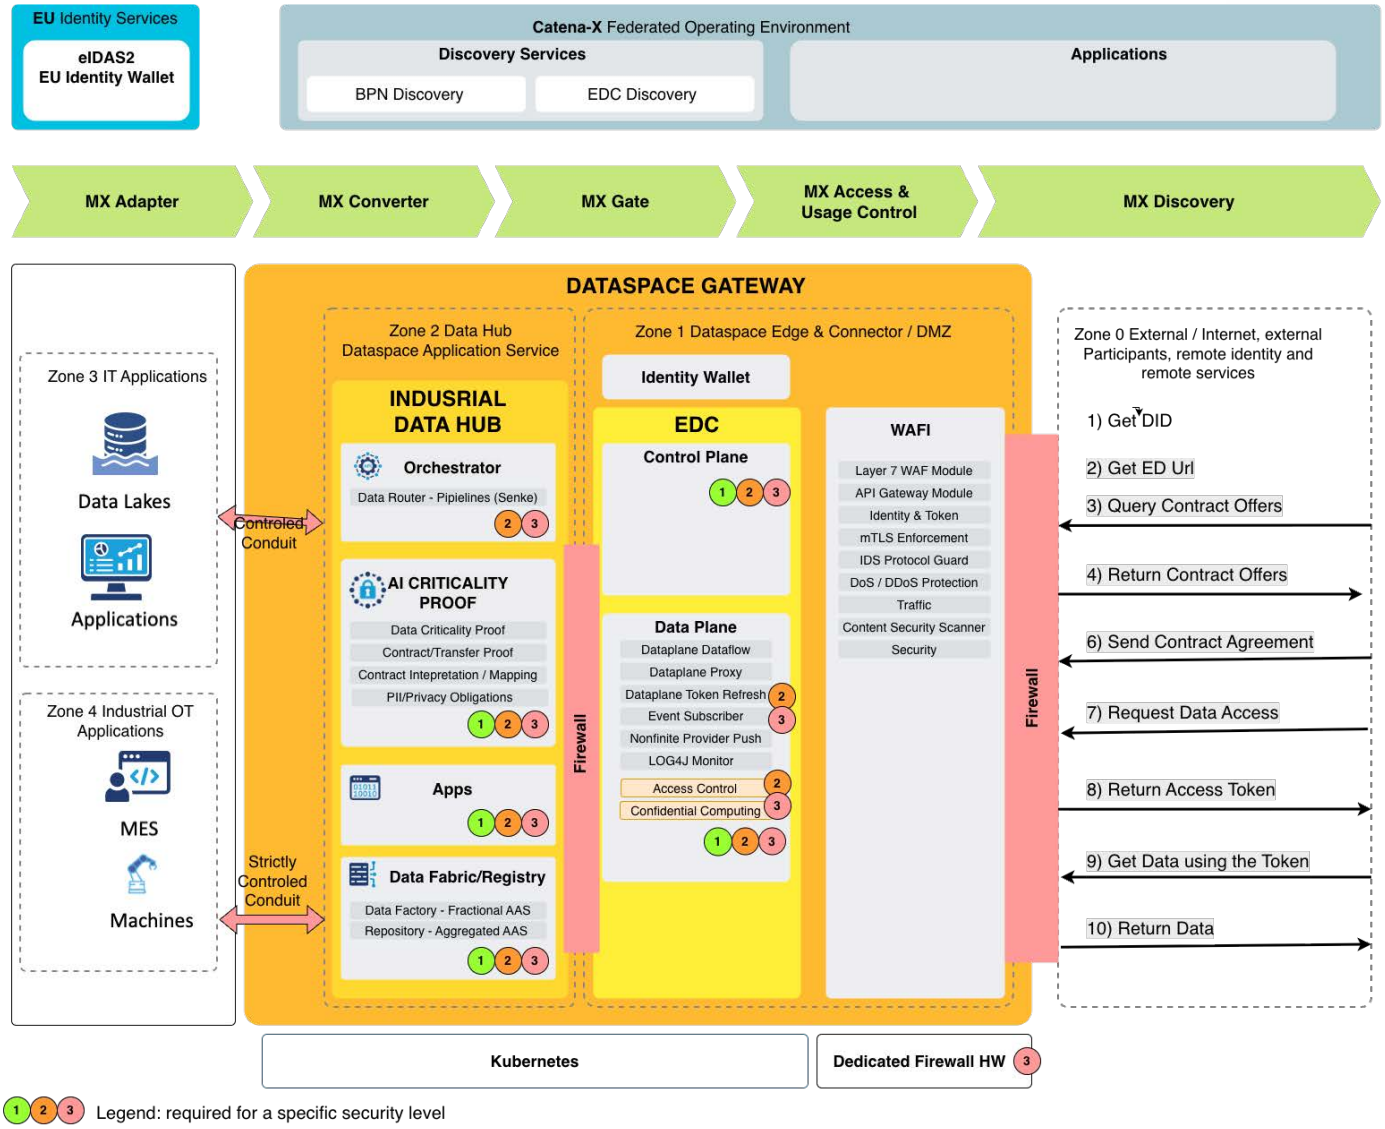


Figure 12 Manufacturing-X Dataspace Gateway – overall zoned architecture with Industrial Data Hub (Zone 3)

The **Industrial Data Hub** handles data classification, policy enforcement, transformation, and controlled AAS publication. It has the following components:

Orchestrator

The Orchestrator coordinates all data flows inside the Data Hub, ensuring alignment with contracts, criticality levels, and usage policies.

- **Data Router – Pipelines (Data Sink):** Manages configurable data routing and sink pipelines that receive, transform, and forward data streams based on metadata and real-time policy decisions.

AI Criticality Proof

The AI Criticality Proof module automatically assesses data assets and generates verifiable proofs for criticality classification, contract compliance, and privacy obligations.

- **Data Criticality Proof:** Uses AI to classify or validate the criticality level (Basic / Trust / Critical) of data assets based on AAS submodel semantics and context.
- **Contract and Request/Transfer Proof:** Provides auditable evidence that data requests and transfers comply with active contracts and usage policies.
- **Contract Interpretation / Mapping:** Translates contract terms (e.g., ODRL policies) into enforceable technical rules across the EDC and Data Hub components.
- **PII/Privacy Obligations:** Detects personal or sensitive data and enforces privacy requirements such as purpose limitation, consent, or anonymization.

Data Fabric/Registry

- **Data Factory – Fractional AAS:** Connects to the backend systems in IT and OT and dynamically generates partial ("fractional") Asset Administration Shells containing only the submodels required for a specific use case or consumer.
- **Data Repository – Aggregated AAS:** Stores and manages aggregated AAS views that combine data from multiple internal sources into unified, queryable representations.

Apps Layer

The Apps layer contains specialized microservices that perform data processing within the trusted Zone 3 environment under Orchestrator and AI Criticality Proof control.

WAFI

WAFI is the dedicated security gateway located in the Dataspace Edge & Connector zone (DMZ). It acts as the first line of defense and the primary entry point for all incoming traffic from external dataspace participants (Zone 1). WAFI sits in front of the Eclipse Dataspace Connector (EDC) and provides advanced Layer 7 security and API protection tailored to the requirements of industrial dataspaces. Its main responsibilities include:

- **Layer 7 WAF and API Gateway functions** — inspects and filters HTTP/HTTPS and dataspace protocol traffic (DSP/DCP) at the application layer.
- **Identity & Token Validation and mTLS Enforcement** — validates Decentralized Identifiers (DID), Verifiable Presentations, and enforces mutual TLS for all connections.
- **IDS Protocol Guard** — ensures that only valid, protocol-compliant dataspace messages (according to the International Data Spaces standards) are forwarded to the EDC.
- **DoS / DDoS Protection and Traffic Segmentation/Routing** — protects against volumetric and application-layer attacks and routes traffic according to security level and data criticality.
- **Content Security Scanner** — scans payloads for malicious content, policy violations, or unexpected data patterns before they reach the connector.
- **Security Monitoring & Logging** — provides centralized visibility and audit trails for all edge traffic, supporting incident detection and NIS2 reporting obligations.

By placing strong security controls in WAFI, the architecture reduces the attack surface of the more complex EDC Control and Data Plane components and enables differentiated protection levels (Basic / Trust / Critical) already at the network edge.

Data Plane

We propose to enhance the Data Plane with two additional components:

- **Access Control Enforcement:** Access Control Enforcement in the EDC Data Plane performs the final, fine-grained authorization at the moment of actual data transfer.
- **Confidential Computing:** For Strictly Confidential data, the Data Plane should support Confidential Computing. Data remains encrypted until it reaches a hardware-protected Trusted Execution Environment (TEE) on the consumer side. The provider can remotely attest the integrity of the consumer's TEE before releasing the decryption key.

8 Use-Case-Driven Security via AAS Submodels

In large-scale, multi-party data ecosystems, labeling data to determine the right protection level is complicated and requires deep use case and data management knowledge. Data sensitivity is highly context-dependent: a data element (e.g., a material specification or a production parameter) can range from semi-public to strictly confidential. Traditional approaches that attempt to classify and protect data at the individual attribute level quickly become unsustainable because they require:

- Manual, per-attribute, or per-field classification decisions.
- Complex, fine-grained policy authoring.
- Continuous maintenance as new use cases, partners, and data types are added.
- High risk of misconfiguration, inconsistent enforcement, and "policy drift" across participants.

The result is high onboarding delays, significant operational risk, and poor scalability — exactly the opposite of what an open, growing ecosystem needs.

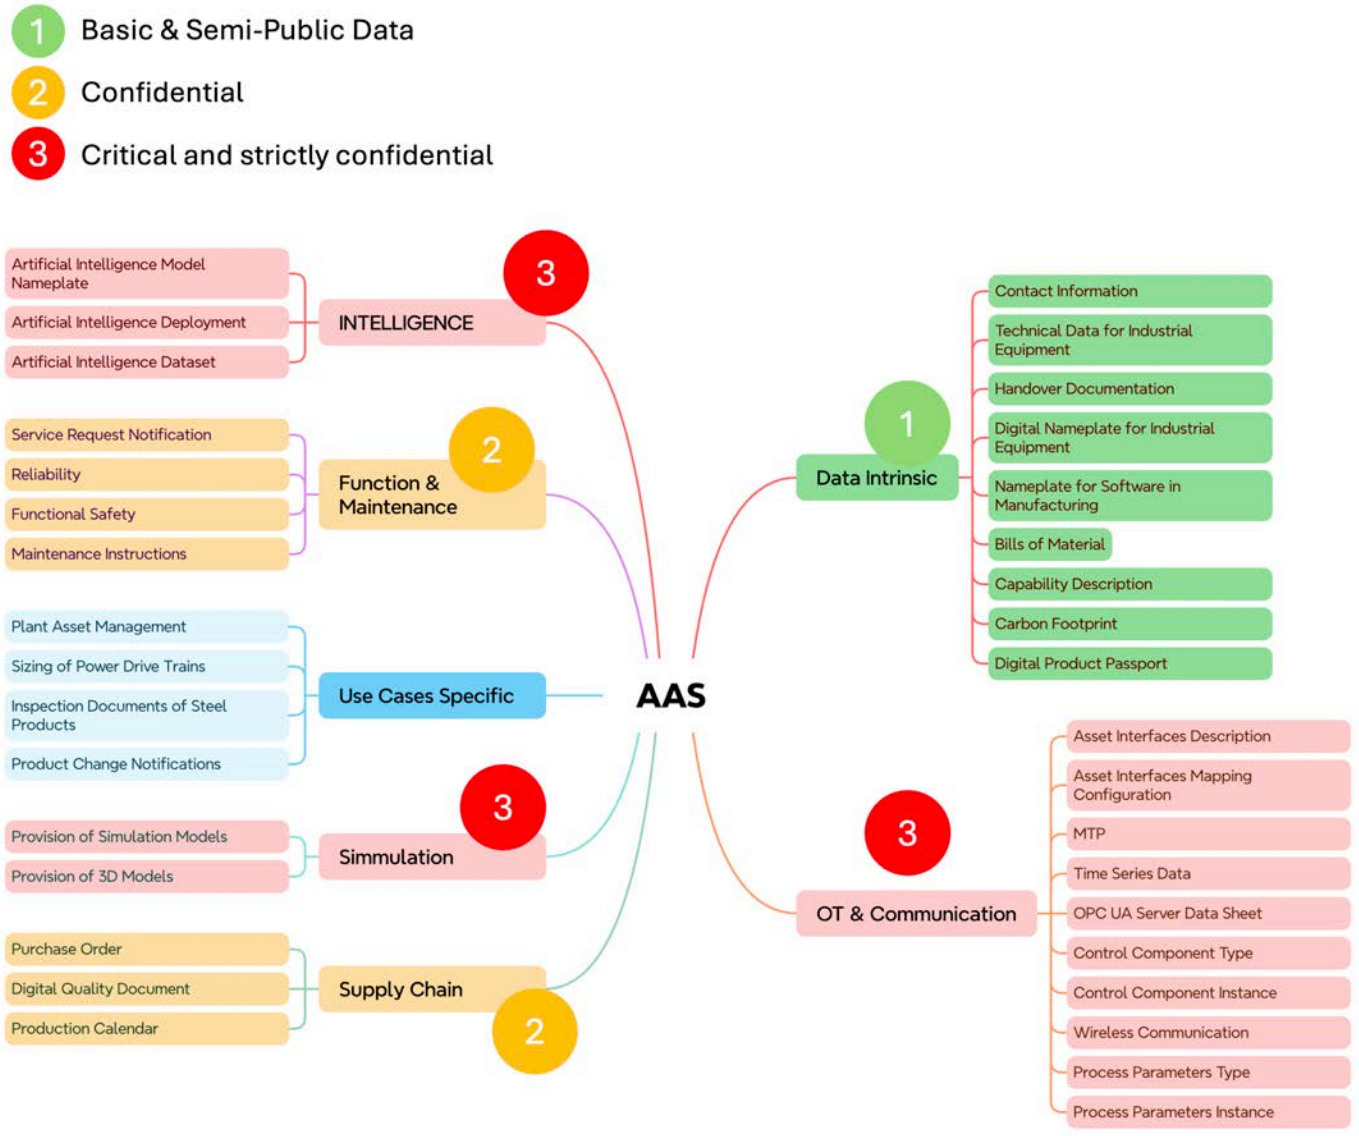


Figure 13: Example how AAS submodels make criticality assignment per use case easy to implement

The AAS Advantage: Submodels Are Purpose-Built for Use Cases

The Asset Administration Shell (AAS) fundamentally changes this equation. The AAS is a machine-readable digital twin of physical or digital assets used in the Industry 4.0 context. While the AAS can be used to describe any asset, process, or task (see Semiconductor-X #AAS-for-everything), its data are grouped in submodels. By definition, AAS submodels are

developed to support specific use cases. In the Eclipse Tractus-X / Manufacturing-X context, these use cases are realized through dedicated Tractus-X Kits, each of which maps directly to one or more standardized AAS submodels.

Because each submodel is designed with a concrete business purpose in mind, it already carries an implicit understanding of the data's sensitivity, sharing patterns, regulatory exposure, and protection needs. The mapping from use case → data criticality is therefore not an additional mapping exercise — it is native to the AAS structure by design.

This creates a powerful and elegant opportunity: instead of assigning security levels to raw data attributes, we can assign them at the AAS submodel level, where the semantics are already clear and stable.

Use Case → Eclipse Tractus-X KIT → AAS Submodel → Data Criticality Level → Security Level

This built-in correlation makes the assignment of criticality levels straightforward and systematic. The dataspace software (Industrial Data Hub and EDC) can use AI support to recognize the AAS submodels, verify the required data criticality and corresponding security level, and finally verify that the receiving side supports the same level — all without manual configuration of individual attributes.

Implementing early and systematic mapping of use cases, AAS submodels, and criticality levels in the Tranctus-X KITs simplifies configuration, facilitates the selection of appropriate security controls, and significantly reduces the main risk of misconfiguration that exists in today's complex attribute-level access control approaches. Figure 6 shows such an example mapping of AAS submodels used in Tractus-X to the three data criticality levels.

9 Risk-based Cybersecurity Management

9.1 Core Design Principles of a secure dataspace

The architectural design is guided by the following core principles, listed alphabetically:

- **Compliance** with regulatory demands – security mechanisms and governance models must support compliance with applicable legal, regulatory, and industry-specific requirements across jurisdictions, including auditability, transparency, and evidence generation.
- **End2End security** – security mechanisms are reaching from OT over dataspace technology into target systems. Since OT is purely the domain of the participant and very heterogeneous, however, hard enforcement of security policies is limited to higher security levels.
- **Federation model** – decentralized connectors combined with central federation services (identity, discovery, monitoring) without a central data lake.
- **Scalability of security mechanisms** – security controls must remain effective and manageable as the number of participants, connectors, use cases, and cross-dataspace interactions grows, without creating disproportionate operational overhead or complexity.
- **Security by Design** – security must be embedded into a product from the start of development throughout its entire lifecycle (planning, design, implementation, deployment, operation).
- **Zero Trust principles** – no entity (user, device, connector, or dataspace) is trusted by default; every interaction requires continuous verification, authentication, and authorization.

While the core principles guide the architectural design, other principles also influence the design:

- **Achievable Security** – security measures must be realistic, practical, and economically viable in real industrial environments. The architecture achieves the right balance between strong security and broad adoption by offering graduated protection levels.
- **Control Plane / Data Plane separation** – policy negotiation and enforcement (Control Plane) is strictly separated from actual data transfer (Data Plane).
- **Industrial Data Hub Concept** – Inside the enterprise, a central Industrial **Data Hub** acts as the data-centric orchestrator. It collects data from IT and OT systems, performs criticality assessment, validates intended use, checks for corruption or anomalies using AI, enforces contracts and authentication, and dispatches data to the appropriate EDC instance and security layer. The Industrial Data Hub ensures full interoperability between internal systems and the dataspace.
- **Defense-in-Depth** – multiple independent layers of protection (technical, organizational, and procedural).
- **Identity and trust** – decentralized identifiers (DIDs), verifiable credentials (VCs), and attestations provided by the Data Trust & Security KIT. Trust is explicit, machine-verifiable, and dynamic.
- **Interoperability with Security** – all layers must remain interoperable across European and international dataspaces while maintaining consistent security controls.
- **Least Privilege and Default-Deny** – access is granted only on a need-to-know basis and only for the agreed purpose and duration, while checking the actuality of privilege must be enforced.
- **Policy-driven data sovereignty** – the data provider retains full control over how, when, and by whom their data is used, enforced through machine-readable and enforceable policies. ODRL-based policies define not only access but also purpose, duration, and obligations after access.

9.2 Trust Infrastructure and Identity Management

Trust is established through decentralized, cryptographically verifiable mechanisms:

- Self-sovereign identity using **Decentralized Identifiers (DIDs)** and **Verifiable Credentials (VCs)**.
- Integration with the **Data Trust & Security KIT** (Eclipse Tractus-X) for exchange of trustable and verifiable data using verifiable credentials. Setting a second level of security over the existing Dataspace Protocol (DSP) + Decentralized Claims Protocol (DCP) trustable data exchange process.

- Support for **eIDAS 2.0** and other European trust anchors.
- Hardware-backed key device entities.
- Central trust services, such as monitoring services and trusted identity services.

Identity Management addresses four core questions: which identities exist, how they are provisioned, how they are authenticated, and how they are revoked.

The underlying protocol is the Decentralized Claims Protocol (DCP), which is based on two open standards: Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs). Every participant requires a wallet that supports DCP for credential exchange and presentation. Regarding decentralization, DCP takes a middle ground. Although credentials are issued by a central authority, querying this authority is not required during authentication itself. Verification is handled through trusted cryptographic signatures. The central authority thus acts as the trust anchor and governance body of the dataspace, establishing the foundation for decentralized interactions. A shared root of trust and governance is essential for participants to reliably trust each other and interoperate. In this context, developments under eIDAS 2.0 are particularly relevant, as they introduce the European Digital Identity Wallet framework and may provide a useful foundation for stronger and more interoperable trust infrastructures. Emerging wallet concepts for organizational identity (like EU business wallets, EUBW) should therefore be monitored closely. As these approaches mature and gain adoption among companies, they could strengthen identity assurance during participant onboarding by providing reusable and verifiable organizational identity information. Over time, this could support more interoperable trust frameworks across dataspaces, with wallet-based business identity as a common foundation and dataspace-specific credentials covering membership and access.

The primary identifier within a dataspace is the Business Partner Number (BPN), which is issued by the Dataspace Registry Authority (DRA). Based on this BPN, each participant generates DIDs to sign their Verifiable Credentials. These credentials allow a holder to prove to a verifier that they comply with the agreed contractual conditions for participation in the dataspace. Currently, only company-level identities are considered. It is recommended to extend this to technical and human users to enable more granular access control.

When two connectors communicate, they use the Dataspace Protocol (DSP). Every DSP communication is preceded by a Decentralized Claims Protocol (DCP) communication. While DSP describes the usage of data, it does not establish trust on its own. Trust is built through DCP, leveraging DIDs and the provided Verifiable Credentials. Upon successful completion of the DSP process, the consumer receives an access token granting access to specific data of the provider.

Verifiable Credentials can be revoked using the Verifiable Credential Status List. Implementing a BPN revocation mechanism through a central list could potentially contradict the decentralization principles of the protocol, but is a central means of participant governance. This is one of the major challenges at the moment: A unified identity management concept applicable to the various dataspace solutions that exist out there. Especially for trust building between two dataspaces, this is important.

Since eIDAS 2.0 is not fully rolled out (expected by 2027), especially for organizations, bridge technologies are used in many implementations. The current state of identity management in MX relies on DIDs derived from the Business Partner Number (BPN) and did:web as the DID method. Since these technologies are not fully mature yet, some aspects of ecosystem governance are not yet fully defined. For traditional X.509-based authentication, several guidelines and best practices exist (such as the CA Browser Forum), but this is still not fully the case for DIDs and did-methods (even if the underlying credentials are often X.509-certificate-based).

To enable proper authentication, higher-level trust establishment mechanisms, multiple security levels of components, and traceability of data flows, connector-specific identities need to be supported. If a valid assignment of connector identities to the membership credentials is not performed, authentication is purely based on this credential. A leakage of the associated key material would break the security of the whole participant gateway deployment. Due to the wallet-based, decentralized nature of identities, the governance for using DIDs and did-methods needs to be defined, and a number of best practices need to be adopted. Aspects to be considered when using did:web could be:

- did:web relies on HTTPS and DNS. So security of did:web depends on HTTPS and DNS, so its trust assumptions are those of the Web PKI and the domain name system.
- This implies the DNS-servers of the did:web controllers should be using DNSSEC for added security.
- There should exist a concise policy on how DNSSEC keys and servers are managed to make sure to avoid fraudulent configuration. This includes supported TLS versions and algorithms.
- For DNSSEC, it needs to be made sure that applications or services requesting a did-document do not fall back to OS-provided DNS resolvers.
- The DID documents are not necessarily integrity-protected or versioned.

There are methods in development that try to fix common problems (such as did:webvh to overcome some shortcomings of did:web). However, these are also still growing in maturity.

To enhance security for the ecosystem, several aspects need to be considered:

- Each gateway/EDC instance needs a separate identity (i.e., private/public key pair).
- Technical identities (such as gateway identities) must be governed by the participant company. This should adhere to best practices (e.g., generating and storing private keys on HSMs or TPMs).
- Governance rules need to be defined to decide who is entitled to request a membership credential for this technical identity (e.g., a company e-mail generally is not sufficient).
- Depending on the security level, access control should be regulated in a more fine-grained manner than just the organization (e.g., department level).

9.3 Security building blocks of the Eclipse Dataspace Connector

The proposed Manufacturing-X E2E security architecture provides a regulator-friendly view of a secure dataspace and its components. It describes the system in a structured way to make it clear, auditable, and implementable. The Eclipse Dataspace Components Connector (EDC Connector) offered in Eclipse Tractus-X EDC or in other MX-Port variants is the central technical building block of the Manufacturing-X E2E security architecture. Its core security building blocks are:

Implemented	Not yet implemented / Gaps
Cryptographic services: Signing, encryption, and key management are supported; post-quantum algorithm support is not yet included.	**Secure communication:** TLS is used at infrastructure level, but TLS 1.3 is not enforced as a minimum by the EDC itself; post-quantum cryptography and end-to-end payload encryption layers are not supported.
Identity & Wallet integration (DCP): Verifiable Credential exchange and wallet management via DCP are implemented and operational.	**Multi-Tenancy Isolation:** Logical and cryptographic separation between participant contexts in shared connector deployments is not yet fully implemented
Policy Engine & Usage Control enforcement: ODRL policy evaluation and contract-level enforcement are partially implemented; enforcement is limited to the point of transfer initiation and is not fully guaranteed even within the connector itself. Full technical enforcement throughout a blueprint lifecycle requires an additional layer beyond the EDC. This has been implemented in specific projects, such as the GLUE[2] project.	**Data Sovereignty Enforcement:** Technical enforcement of access and usage constraints is only partially realized within the EDC and does not extend beyond the transfer initiation point; full enforcement throughout the data lifecycle requires an additional dedicated enforcement layer based on remote integrity verification and support for Trusted Execution Environments and other hardware-based security mechanisms.
Attestation and integrity verification hooks: Software-based attestation and integrity verification hooks are present; hardware-backed attestation has been implemented in specific projects, such as the GLUE project.	**Full Integration:** While many of these aspects have reference implementations or are implemented to certain extents, there exists no singular platform that implements all requirements.
Monitoring and logging interfaces: Standardized logging and telemetry interfaces are available.	**Proof of Technical Enforcement:** The proof of technical enforcement based on remote attestation requires trusted reference values against which the gateway in question is compared. The verification of industry-specific data processing apps requires an authority attesting that their operations will comply with given usage control policies or behave in a
Access Control & Authorization: API key management and token-based authentication (X-API-Key, OAuth2/JWT Bearer) are implemented for the Management API.	

[2] https://www.iosb.fraunhofer.de/en/projects-and-products/glue-trustworthy-distributed-usage-control.html

Secret & Credential Management: Secure storage and retrieval via external secret stores such as HashiCorp Vault is supported, preventing plaintext credential exposure.	trustworthy manner. To date, there exists no unified standard or trustworthy entity that fulfills such a role.
Audit Trail & Non-Repudiation: Basic tamper-evident logging of contract negotiations and data transfer events is present, though completeness and standardization across all policy decision points vary.	

Note that the EDC is not a turnkey secure product in its current implementation. Security depends on correct deployment, hardening, monitoring, and operator responsibilities. The open-source Eclipse Tractus-X EDC will always require ongoing hardening and lifecycle management, which applies to every open-source project. For a list of attack vectors, security controls, and requirements for the ongoing hardening, we refer the reader to Appendix C. The use at your own risk policies apply to every open-source software. In the end, companies that are willing to operate in production need to put in place the necessary security measures in order to make sure the software and operations are secured at all times. Some Security Measures implemented in the Eclipse Tractus-X Repository can be found in Appendix D.

To enable higher security levels, trust establishment is not only enabled by identity management. Here, remote attestation is used to provide information on the hardware and software stack of the remote entity before any data is exchanged. The signature of such an attestation report can be traced back to the Root of Trust for Measurement (RTM) in firm- or hardware, which in turn can be traced back to the manufacturer of the trustworthy hardware component. Trusted Platform Modules (TPMs) are a well-known example, and Confidential Computing Technologies also provide such mechanisms that can be traced back to the hardware vendors, e.g., Intel, AMD, or Nvidia. This way, each component of the hard- and software chain measures and attests to its following components, creating a chain of trust that can be verified by a remote party.

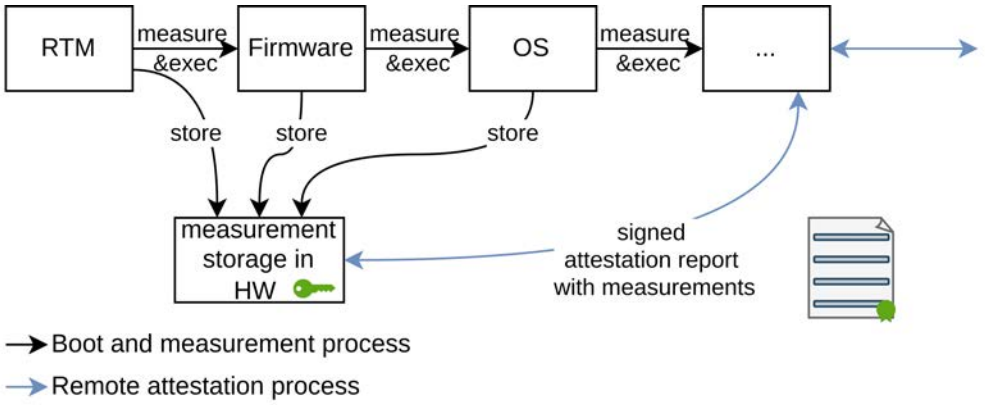


Figure 14: Exemplary Measured Boot and Remote Attestation Process of a Hard- and Software Stack

A verifier can then use the provided reference values to verify that the running software is approved by them or a trustworthy third party. The creation of these reference values is ideally automated and managed in the software development process. Alternatively, a central trust authority may review and determine certain software components to be trustworthy, e.g., that the software enforces usage control policies. Remote attestation relies on these reference values; without these, the process adds little additional security guarantees.

The remote attestation process can be used to link the hard- and software stack to the previously introduced identities, thus proving that the endpoint hosted by the identity enforces the provided usage control policies and processes the data in accordance with these. This way, usage control enforcement moves from a contractual obligation towards (full) technical enforcement.

9.4 Technical implementation recommendations

There already exist efforts by multiple institutions that aim to provide up-to-date recommendations on implementing specific security aspects. As such, this document does not aim to create another such list that must be regularly updated. In the following, we instead provide references to several of these documents as well as their publishing institutions.

Cryptography – Current cryptographic recommendations can be found in BSI –TR 02102[3] of the Bundesamt für Sicherheit in der Informationstechnik (BSI – Federal Office for Information Security, Germany). They include information on current ciphers and key lengths, TLS-specific recommendations, cryptography for IPSec and Internet Key Exchange, as well as for Secure Shell (ssh). The National Institute of Standards and Technology (NIST) offers multiple documents on cryptographic recommendations, for example, for HMACs in 800-224, elliptic curves in 800-186, cryptographic key generation with 800-133, as well as broader guides for designing cryptographic key management systems in 800-130.[4]

Manufacturers and Products – The BSI published technical requirements for cyber resilience of manufacturers and products as defined in the Cyber Resilience Act, in TR-03183[5]. It encompasses several parts, including general requirements as well as requirements for software bill of materials (SBOMs) and vulnerability reports and notifications. Part H is dedicated to expanding upon existing cybersecurity processes and capabilities to product development and vulnerability handling. Similarly, the NIST published a catalogue of security and privacy controls for information systems and organizations in 800-53, encompassing a wide range of threats and risks and addressing functionality and assurance. The ISO/IEC 27001[6], published by the International Electrotechnical Commission and the International Organization for Standardization, defines requirements for information security management systems.

Architectural Design and Security – The NIST offers multiple guides on how to design architectures and specific aspects, whilst taking security into account. For example, there exist security guidelines for storage devices (800-209), a zero-trust architectural model for access control in cloud-native applications (800-270A) based on their Zero Trust Architecture (800-207), an application container security guide (800-190), and building secure microservice-based applications using a service-mesh architecture (800-204A). The BSI publishes similar documents, for example as minimum standards and technical guidelines, such as the BSI minimum standards for the use and shared use of external cloud services.

Depending on the use case, different guides may be of use.

9.5 Secure Development Lifecycle for OSS and related artefacts

A secure development lifecycle provides a supporting framework to systematically develop components with suitable security features (cf. ISA/IEC 62443-4-1, ISO/IEC 27034). To achieve this, security-related activities are integrated into all lifecycle phases, i.e., planning, requirements engineering, design, implementation, testing, deployment, operation, continuous development, and maintenance. A core activity is Threat and Risk Management, which includes identification, assessment, prioritization, and treatment of threats according to the associated risk. The output of a risk analysis is a list of security requirements that must be implemented to achieve a level of security for the developed component that is appropriate, given the identified risks.

International standards, as mentioned above, target software product development in a company context. The key difference with respect to open-source software development (OSS) is its distributed nature. Hence, “trusting the process” is even more important than “trusting the person” when managing security in OSS projects. Security-related practices, i.e., security management, specification of security requirements, secure design, secure implementation, security testing, management of security-related issues, and update management, need to be adjusted for community governance, transparency, and distributed contributors.

Secure Development Lifecycle

Regarding security management, security policies and responsibilities should be maintained in the project's repo. This should include how vulnerabilities are reported to the project (coordinated vulnerability disclosure (CVD)) and addressed, within what time frame, and who is responsible for security-sensitive changes. Responsible maintainers should be defined for vulnerability management and for releases (signing releases). Security policies should also cover supported versions, and it should be clearly communicated that security properties claimed by the project only apply to these versions.

Threat and risk analysis is at the core of security requirements specification. This activity, therefore, is the foundation of "product" security. Identified requirements should be maintained as part of the security design documentation in the

[3] https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Standards-und-Zertifizierung/Technische-Richtlinien/TR-nach-Thema-sortiert/tr02102/tr02102_node.html

[4] https://csrc.nist.gov/publications/sp800

[5] https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Standards-und-Zertifizierung/Technische-Richtlinien/TR-nach-Thema-sortiert/tr03183/TR-03183_node.html

[6] https://www.bsi.bund.de/EN/Themen/Oeffentliche-Verwaltung/Mindeststandards/Externe_Cloud-Dienste/Externe_Cloud-Dienste_node.html

repository. Processes should be established to ensure they are reviewed and updated as necessary. Security requirements specification assumes an intended use of the components/system, defined system boundaries, operational assumptions, and a security context, i.e., which security capabilities are expected from the environment where the component/system is operated. The security requirements specification should further include the threat model (e.g., using STRIDE), which is typically derived from assets/data and according to security goals, and the assessment of the threats based on likelihood and impact (risk assessment), as well as the requirements that have been identified as mitigations (e.g., authentication requirements, cryptographic requirements, input validation expectations). Standards like ISO/IEC 27005 (Guidance on managing information security risks) or IEC 62443-3-2 (Security risk assessment for system design) cover how to conduct a threat and risk analysis.

In the design phase of SDL, the goal is to ensure the development of a secure system architecture following a defense-in-depth concept, i.e., have multiple independent security measures to mitigate threats, and follow secure design principles, e.g., least privilege, separation of duties, and fail securely. Regarding software, risks are typically associated with interfaces. It should be specified which interfaces exist, how they are supposed to be used, whether access to the interface crosses a trust boundary, how users are authenticated, which roles/privileges are required for their use, which measures are in place to protect them, and how untrusted input is handled. A secure design review process needs to ensure that chosen mitigations handle risks appropriately. This process should also be triggered upon feature requests or security-related issues.

Secure implementation includes secure coding standards and security implementation reviews. Secure coding standards define how coding rules like input validation, memory safety, secret handling, etc., should be implemented. Secure coding standards should also be supported by tools in the build pipeline, like static analysis and dependency scanning. The OSS community further needs to define how security-critical pull requests are reviewed.

For security testing, the OSS release process should ensure that security-related functionality is covered by software testing procedures and additionally include methods like fuzzing, vulnerability scanning, or penetration testing.

This includes securing the software supply chain within the open-source components, starting with enforcing access rights and identity of contributors. In the case of Eclipse Tractus-X, which is a project governed by the Eclipse Foundation, the Identity of every committer can be verified using a governmental ID. This is a very new feature that provides trust in the Person behind, who holds the keys for maintaining open-source software. The Identity Verification feature is described in the project handbook from the Eclipse Foundation, a governance document for all Eclipse projects, including Eclipse Tractus-X.

For example, on the right side of Figure 14 (in the box of User Status), you can see if the committer has been fully verified against a valid governmental ID, in a publicly accessible view.

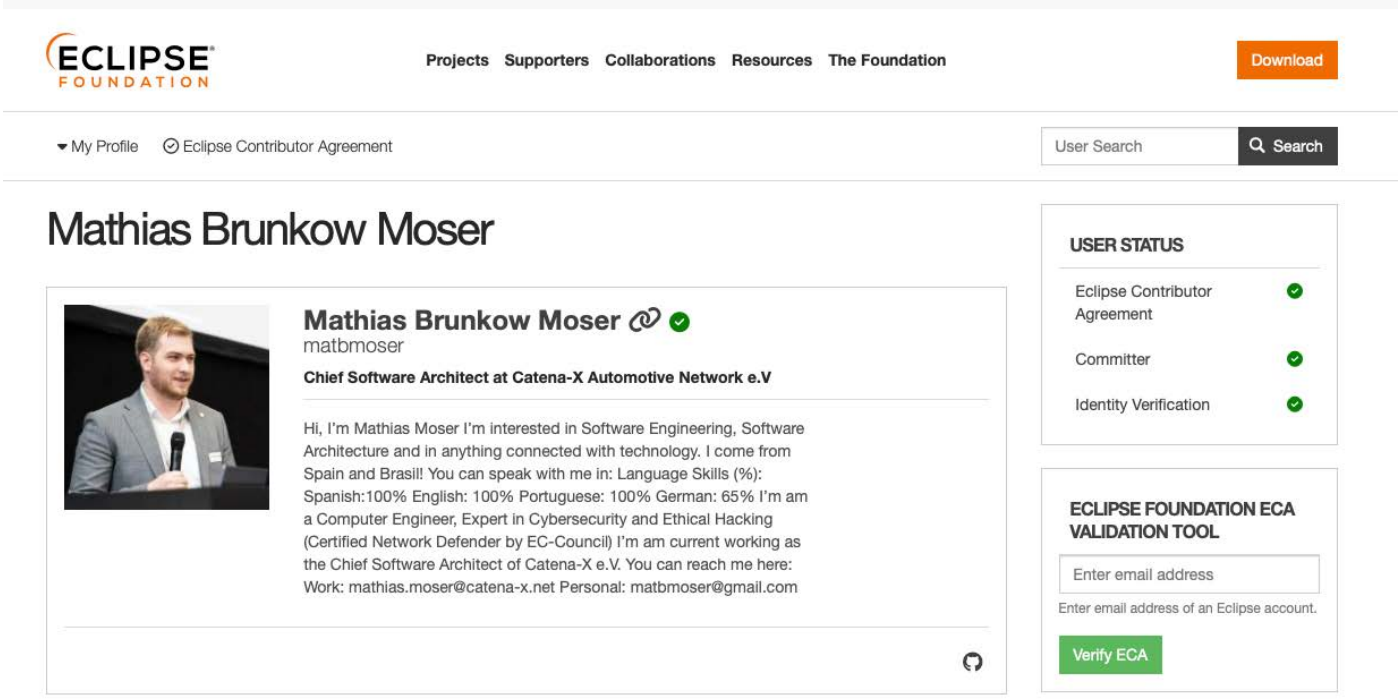


Figure 15 Example of Eclipse Committer with Identity Verification Checked, from the Eclipse Tractus-X Project.

9.5.1 Continuous Development & Maintenance

Regarding management of security-related issues, the OSS community needs to define how potential vulnerabilities are to be reported to the project and how vulnerabilities are identified by the project, e.g., by monitoring known vulnerabilities for dependencies based on an SBOM (Software Bill of Materials). Processes should be established to ensure that potential vulnerabilities are reviewed, assessed, and handled promptly.

Security update management is about preparing, testing, and publishing patches for security-related issues. Patches should be versioned, and patch releases should be signed so that users can verify the authenticity and integrity of a patch. Security fixes should be reflected in change logs, and it is also necessary to backport patches to older releases that are still supported.

If OSS is intended for commercial use, security-related activities as described above are also required by the European legal framework. The Cyber Resilience Act (CRA) aims to ensure an appropriate level of cybersecurity based on risk assessment and fulfill a set of essential cybersecurity requirements, e.g., products placed on the EU market must not have known exploitable vulnerabilities. Manufacturers also have to establish suitable processes for vulnerability handling, for which the CRA also introduces a set of requirements, e.g., manufacturers must apply effective and regular tests and reviews of the security of the product, and they have to share and publicly disclose information about fixed vulnerabilities.

Given these requirements, a community developing a dataspace connector would greatly benefit from a distribution provider, similar to what is seen with major Linux distributions, for example.

9.5.2 Threat Modeling Methodology

The primary methodology recommended for the Manufacturing-X dataspace is **STRIDE** (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege), supplemented by elements of **PASTA** (Process for Attack Simulation and Threat Analysis) where deeper risk quantification is required.

Key principles applied:

- Threat modeling shall be performed from the perspective of a full dataspace participant (including both **Company-Level** and **Centralized-Level** components).
- The scope includes the Eclipse Tractus-X EDC (Control Plane and Data Plane), DID/VC handling, contract negotiation, data transfer mechanisms, integration with Digital Twin Registries (AAS), and connections to OT/shopfloor systems.
- Modeling shall be conducted using Data Flow Diagrams (DFD) pointing to trust boundaries between participants, the dataspace infrastructure, and OT environments.
- Results shall be documented and maintained as a living artifact, with updates required after any significant architectural change, new use case onboarding, or major release of Eclipse Tractus-X components.

Threat modeling workshops shall be conducted at least annually and after every major system change.

10 Conclusions and Recommendations

The Manufacturing-X initiative aims to establish federated, sovereign industrial dataspaces for discrete manufacturing by extending the proven open-source foundation of Catena-X across broader ecosystems (Factory-X, Semiconductor-X, Aerospace-X, and others). The core technical building blocks are provided by the Eclipse Tractus-X KITs.

The core technical building blocks are provided by **Eclipse Tractus-X KITs**, documented at https://eclipse-tractusx.github.io/Kits/.

Key components:

- **Connector KIT** (mandatory): Uses the **Eclipse Tractus-X EDC** (reference implementation of Eclipse Dataspace Components). It implements the **Dataspace Protocol (DSP)** for contract negotiation, **Decentralized Claims Protocol (DCP)** for verifiable credential exchange, **ODRL** for usage policies, and separates **Control Plane** from **Data Plane**. This enables fully decentralized, peer-to-peer, and sovereign data exchange.
- **Data Trust & Security KIT**: Adds verifiable data using **W3C Verifiable Credentials (VCs)**, **Decentralized Identifiers (DIDs)**, and **Verifiable Presentations (VPs)**. Independent Attestation Providers issue signed attestations (e.g., for PCF, Digital Product Passport). Integrates with the Connector and Digital Twin Registry (AAS-based).
- **Support kits:** Data Governance KIT, Digital Twin KIT, Industry Core Foundation, Business Partner KIT, and manufacturing-specific kits.

The proposed E2E Security Architecture extends the existing framework with five foundational elements:

- A **three-level security model (**Basic, Trust, Critical) directly aligned with data criticality and use-case sensitivity, enabling participants to implement only the security level they require.
- The **Industrial Data Hub and Industry Core Hub** (having as one implementation, the Industry Core Hub) as the intelligent, AI-augmented interface between internal IT/OT systems and the external dataspace. It performs automated criticality assessment, validates intended use, enforces policies, and ensures interoperability via AAS submodels.
- **Native, use-case-driven security** through Asset Administration Shell (AAS) submodels. Because Tractus-X KITs and their corresponding AAS submodels are purpose-built for specific business scenarios, data criticality classification becomes systematic and significantly less error-prone than traditional attribute-level approaches.
- **Systematic risk-based cybersecurity management** is applied both in the **open-source** development of dataspace components (within the Eclipse Tractus-X community) and in **operational environments**. This includes threat modelling for each KIT, tiered control catalogues aligned with the three security levels, coordinated vulnerability management, and structured incident response mechanisms.
- **Operation and Governance Framework –** to provide guidelines for secure participant onboarding, secure operations of gateways and best practices for deployment in different infrastructures.

This architecture delivers the long-promised “one fits all” connector model while making security achievable — realistic, practical, and economically viable — for both SMEs and large enterprises. By combining Security by Design, Zero Trust principles, Defense-in-Depth, and a clear separation of Control Plane and Data Plane, the proposal addresses the most pressing challenges identified in this white paper: complexity, misconfiguration risk, OT integration, cross-dataspace interoperability, and emerging AI-driven threats.

10.1 Recommendations

10.1.1 When defining a use case / KIT provide a Security Subchapter in the Operation View

When creating a **new Use Case / KIT**, during the process of ideation of the architecture, try to think about how it will be operated in the future, and which possible configuration requirements will your use case / KIT have. Therefore, in the Operations View (one of the chapters of the KIT Framework Architecture), you should provide more content than just some operation guidelines. There is a place to provide Security Guidelines and analysis for your use case. Ask yourself questions on whether this use case KIT will become an App, and it will be operated in a productive scenario:

- **What are the risks and threats?** > Place these points in a simple diagram showing the possible interactions with external and internal systems, and enumerate the Risks in a table. Congratulations, you will be starting to take the first step of a threat modelling exercise! In Appendix B, you will find a template table you can use to structure your risks; this will help you to support your "business continuity" when your setup suffers an attack.
- **Which data in my use case is considered Critical and in which context?** > Here, the recommendation is to analyze which data models my company and other companies will be exchanging in these use cases and between which business partners. For example, when you will be exchanging Product Carbon Footprint data, and to do your calculation, you need to have access to very sensitive data from your supplier, you should provide a table in the security subchapter of your KIT (under the adoption view), which describes the criticality of the data exchanged between each participant of your use case.

Example:

Data Model / Semantic ID (with version)	Attribute (with path if needed)	Classification Level	Context
urn:samm:io.catenax.pcf:9.0.0#Pcf	pcfAssessmentAndMethodology. PcfMethodology. GwpCharacterizationFactorDetails. ipccCharacterizationFactors	Critical / Confidential	When exchange between Tier-1 suppliers and OEMs. This attribute / data endpoint may include information from other suppliers. Recommendation: EDC Exchange via VPN.
…	…	…	…

Currently, most of the KITs do not provide a proper operation view, for example, KIT. This results in Adopters ignoring the risks and consequences of not defining a proper security architecture for their industry dataspace data sharing setup. Therefore, if a proper security analysis is provided at the Operation View, KIT adopters can prepare and set up a secure dataspace setup, since they are now aware of the risks and criticality of the declared data endpoints.

10.2 Next Steps and Priority Actions

To translate the proposed architecture into practice, the following priority actions are recommended:

1. **Threat Modelling and Control Catalogues (High Priority)**

   Introduce systematic threat modelling for every Tractus-X KIT. Develop tiered, catalogue-based cybersecurity concepts and baseline controls aligned with VDE-AR-E 2802-20 (Horizontal cybersecurity concept for Industry 4.0). These catalogues should be structured according to the three Gateway Security Levels (Basic, Trust, Critical) defined in this document.

2. **OSS Risk Management and Secure Development Lifecycle**

   Establish a joint risk management process for Eclipse Tractus-X components (EDC, Identity Hub, etc.) together with the Tractus-X community. This includes coordinated vulnerability disclosure, SBOM management, and post-market surveillance processes in line with the Cyber Resilience Act (CRA).

3. **Certification and Assessment Framework**

   Develop certification and assessment templates for cybersecurity and governance, differentiated by security level. This should include self-assessment toolkits for participants and conformity assessment schemes for connectors and Data Hub implementations.

4. **Industrial Data Hub and AAS-Driven Implementation**

   Prioritize the development and reference implementation of the Industrial Data Hub (as an extension of the Eclipse Tractus-X Data Exchange Orchestrator) with built-in AI criticality assessment and AAS submodel awareness. This is considered the key enabler for both security and broad adoption, especially for brownfield scenarios.

5. **Governance and State Actor Involvement**

   As dataspaces grow in strategic importance for European industry and critical infrastructure, initiate a structured discussion on the appropriate role of state and governmental actors (e.g., BSI, ENISA) in threat intelligence sharing, coordinated incident response, and the protection of national and European interests.

6. **Cross-Dataspace Security Coordination**

Establish mechanisms for security intelligence sharing, joint incident response, and consistent policy enforcement across European dataspaces (Catena-X, Factory-X, Semiconductor-X, etc.) and with international initiatives.

10.3 Call to Action

The concepts presented in this white paper provide a coherent, regulator-friendly, and implementation-oriented path toward secure and trusted industrial dataspaces. Success now depends on coordinated execution across the Manufacturing-X projects, the Eclipse Tractus-X community, and industry stakeholders.

We recommend that the **Manufacturing-X Guidance Board**, in close collaboration with the **Platform Industrie 4.0/AI** and leading security experts, establish a dedicated implementation roadmap and working group to drive the priorities outlined above.

As industrial dataspaces scale and become part of Europe's critical infrastructure, the role of **state and governmental actors** must be actively discussed and progressively defined. Dataspaces are no longer purely commercial or technical constructs — they are increasingly strategic assets that can be targeted by state-sponsored actors and are essential for the resilience of European supply chains and industrial sovereignty.

Key aspects that require structured dialogue include:

- The integration of national and European threat intelligence (e.g., from BSI, ENISA, and national cybersecurity authorities) into dataspace monitoring and incident response processes.
- The establishment of coordinated, cross-dataspace incident response mechanisms that can operate effectively during large-scale or geopolitically motivated attacks.
- The potential need for secure information exchange channels between dataspace operators and state authorities, while preserving the decentralized and sovereign nature of the dataspace model.
- Long-term considerations regarding the protection of critical manufacturing infrastructures and the possible role of governmental bodies in certification, oversight, or crisis coordination when dataspaces reach a level of systemic importance.

This discussion must be conducted carefully and transparently. It should respect the core principles of data sovereignty, decentralization, and private-sector operational responsibility, while acknowledging that certain threats — particularly advanced persistent threats and AI-driven attacks at a national scale — may require capabilities and coordination that go beyond what individual companies or even dataspace operators can provide alone.

We therefore recommend initiating a dedicated working stream on **"State Actor Integration and Strategic Resilience"** under the umbrella of the Manufacturing-X Guidance Board, involving relevant governmental bodies, to develop a clear position and governance model for the coming years.

Appendix A: Relevant standards of design, operation, and governance

The following table provides an overview of the most relevant standards and regulations for the secure design, operation, and governance of industrial dataspaces. The final column indicates where each standard is primarily referenced in this white paper.

Standard / Regulation	Type	Short Description	Relevance to Industrial Dataspaces	Priority for Manufacturing-X	Main Chapter Reference
NIS2 Directive (EU) 2022/2555	EU Regulation	Mandatory cybersecurity risk management, incident reporting (24h), supply chain security, and management accountability for essential and important entities.	Applies to dataspace operators, identity providers, onboarding platforms, and connector-as-a-service providers.	Very High	Chapter 3, Chapter 6
Cyber Resilience Act (CRA)	EU Regulation	Security-by-design, SBOM, vulnerability management, post-market surveillance, and lifecycle requirements for products with digital elements.	Directly affects commercially distributed connectors and dataspace software components.	Very High	Chapter 3, Chapter 5.6
GDPR	Regulation	Data protection requirements for personal data processed in dataspaces.	Relevant for ensuring lawful data processing, purpose limitation, and data minimization in policy definitions and connector configurations.	High	Chapter 4
EU Data Act	Regulation	Regulation enabling fair access to and sharing of data across sectors and establishing rules for data portability, access rights and data sharing.	Highly relevant for industrial dataspaces as it governs data sharing rights between participants.	Very High	Chapter 4
IEC 62443 series	International Standard	Comprehensive series for Industrial Automation and Control Systems (IACS) security, including zones & conduits secure development and component requirements	Core standard for OT integration and secure industrial data exchange	Very High	Chapter 4, Chapter 5
IEC 62443-4-1	International Standard	Secure product development lifecycle for industrial automation components.	Highly relevant for EDC and connector development.	High	Chapter 5.5
IEC 62443-4-2	International Standard	Technical security requirements for IACS components	Product standard for EDC and connector development.	High	Chapter 5.5
DIN SPEC 27070	Technical Specification	Reference architecture for secure industrial data exchange gateways.	Directly applicable to EDC/IDS-style connectors in industrial environments.	High	Chapter 4
TISAX (VDA ISA based)	Assessment Scheme	Trusted Information Security Assessment Exchange – automotive information security assessment.	Already required in Catena-X for many provider roles; highly relevant for automotive MX projects.	Very High (Automotive)	Chapter 3, Chapter 7
ISO/SAE 21434	International Standard	Cybersecurity engineering for road vehicles.	Relevant for automotive-related use cases and OT integration.	High	Chapter 3
IEC 81001-5-1	International Standard	Security activities in the health software product lifecycle.	Key standard for medical/health dataspaces (Health-X) – extends IEC 62304 with cybersecurity.	High (Medical)	Chapter 5.6
ISO 14971	International Standard	Application of risk management to medical devices.	Risk management foundation for medical use cases (including cybersecurity risks).	Medium-High	Chapter 5
ISO/IEC 27001 + 27002	International Standard	Information Security Management System (ISMS) and controls.	Baseline for organizational security governance of operators and participants.	High	Chapter 6
ISO/IEC 27005	International Standard	Information security risk management.	Methodology support for risk assessment in dataspaces.	High	Chapter 5

Standard / Regulation	Type	Short Description	Relevance to Industrial Dataspaces	Priority for Manufacturing-X	Main Chapter Reference
International Data Spaces (IDS) RAM	Reference Architecture	Reference Architecture Model for sovereign data sharing (usage control, identity, policy enforcement).	Foundational conceptual model for industrial dataspaces.	High	Chapter 4
Eclipse Tractus-X / EDC Security	Open Source + Best Practices	Security guidelines and implementation patterns for EDC connectors (identity, policy, attestation).	Practical baseline for secure connector deployment in MX.	High	Chapter 4.7, Chapter 5.5, Chapter 6, Chapter 8.5
ETSI EN 303 645	European Standard	Baseline cybersecurity requirements for consumer IoT.	Sometimes referenced for edge devices and gateways.	Low – Medium	Chapter 3

A.1 Dataspace-Specific & Industrial Data Exchange Standards

Standard / Framework	Type	Relevance	Priority for MX
DIN SPEC 27070	Technical Spec	Reference architecture for secure industrial data exchange gateways (directly relevant to EDC/IDS connectors).	High
International Data Spaces (IDS) RAM	Reference Architecture	Foundational model for sovereign data sharing. Security concepts (usage control, identity, policy enforcement).	High
GAIA-X	Framework	European data infrastructure with strong sovereignty and security requirements.	Medium-High
Eclipse Dataspace Components (EDC) / Eclipse Tractus-X Security	Open Source + Best Practices	Practical security implementation for connectors (identity, policy, and attestation). Eclipse Tractus-X sig-security group.	High (implementation)
ISO/IEC 20151 - Dataspace Concepts and Characteristics	Standard	Defines the foundational concepts and essential characteristics of dataspaces, including governance, trust, and interoperability requirements. Builds on the IDSA Rulebook.	Medium-High
ISO/IEC DIS 26450 - Eclipse Dataspace Protocol (DSP)	Protocol Standard	Standardized protocol for secure, policy-controlled data exchange between dataspace participants. [LINK]	Medium-High
ISO/IEC DIS 26451 - Eclipse Decentralized Claims Protocol	Protocol Standard	Standardized protocol for verifiable credentials and decentralized trust assertions in dataspace ecosystems. [LINK]	Medium-High
Eclipse Tractus-X Release Guidelines (TRGs)	Rulebook & Governance Documentation	DevOps, Free and Open-Source Development (FOSS), Security, Documentation (KITs) rules of the Eclipse Tractus-X project, mostly indicated by the Eclipse Foundation and other open-source experts in the community. These are release guidelines which every product included in the changelog of the Eclipse Tractus-X release MUST follow, thanks to these rules the delivered products can remain in a high quality, and the project can deliver mature solutions. [LINK].	High
Catena-X Standards (ex. CX-0018, CX-0149, CX-0152)	Standard Collection	The Catena-X standards, published by the company members of several Expert Groups and Committees at the Catena-X Association. They are an important reference and example of productive usage of dataspace in production. They are currently the most relevant, broad, and mature standards from all productive dataspaces available at the moment. The CX-0018 is a very good example for a connector standard, CX-0149 defines the identity and wallet requirements and finally CX-0152 defines mature policy constraints to be used to secure the usage control of data endpoints behind connectors. [LINK]	Very High
Eclipse Tractus-X KIT	Open-Source Specifications	The Eclipse Tractus-X KITs (Keep it together) is a collection of open-source tool-box specifications available at the Eclipse Tractus-X project, and as the name says, which is governed by the Eclipse Foundation. They offer a significant level of implementation paths for different use cases across multiple industries (Automotive, Chemical, Semiconductor, Manufacturing, etc.) [LINK]	Very High
Catena-X Operating Model	Standard / Normative Document	The Catena-X Operating Model is a normative document for the Operative Catena-X dataspace. It sets the visibility, roles, and	High

Standard / Framework	Type	Relevance	Priority for MX
		governance of the ecosystem. Also, it has many good references when it comes to security in dataspaces. Including requirements like TISSAX certification by dataspace roles. [LINK]	

A.2 Industrial / OT Security Standards (Highly Relevant)

Standard	Relevance to Dataspaces	Priority for MX
IEC 62443 series	Industrial Automation & Control Systems (IACS). Zones & conduits, secure development (4-1), component requirements (4-2), system security (3-3). Core for OT integration.	Very High
IEC 62443-4-1	Secure product development lifecycle. Especially important for EDC and connector development.	High
VDI/VDE 2193	Asset Administration Shell (AAS). Security aspects of digital twins / AAS in Industrie 4.0.	High

A.3 Cross-Cutting IT Security & Risk Management

Standard	Relevance	Priority for MX
ISO/IEC 27001 + 27002	Information Security Management System. Baseline for organizations and operators.	High
ISO/IEC 27005	Information security risk management . Risk assessment methodology.	High
ISO/IEC 15408 (Common Criteria)	IT product security evaluation. Can be used for connector/gateway certification.	Medium
ETSI EN 303 645	Consumer IoT (baseline). Sometimes referenced for edge devices/gateways.	Low-Medium

A.4 Sector-Specific Standards

Sector	Standard	Relevance	Priority for MX
Automotive	TISAX (based on VDA ISA)	Mandatory or strongly expected in Catena-X and automotive MX projects.	Very High
Automotive	ISO/SAE 21434	Cybersecurity engineering for road vehicles.	High
Medical / Health	IEC 81001-5-1	Security activities in the health software lifecycle (complements IEC 62304 + ISO 14971).	High (for Health-X)
Medical	ISO 14971	Risk management for medical devices (including cybersecurity risks).	High
Semiconductor	SEMI standards + supply chain security	Relevant for Semiconductor-X.	Medium

A.5 Emerging New Topics

Topic	Relevant Work / Standards	Notes	Priority for MX
Post-Quantum Cryptography	NIST PQC standards, BSI recommendations	Should be considered for long-term data protection.	Very High
AI Security	Emerging (EU AI Act + specific guidelines)	Important for Federated AI use cases.	Very High
Supply Chain Security	NIS2 Art. 21, CRA, VDMA recommendations	Increasingly critical.	Very High
Zero Trust	NIST SP 800-207, BSI guidance	Architectural principle more than a formal standard.	Very High

Appendix B: Risk Management Template Table

Here is an example of a risk management template:

Risk ID	Risk Description	Business Area	Likelihood	Impact	Overall Risk Rating	Existing Controls / Mitigations	Action Required	Owner	Status
R-001	Dataspace Connector Misconfiguration	DevSecOps	High	High	High	Cloud hosting, IAM integration	Test manual workaround and recovery process	Security Officer	Open
R-002	Sensitive Data Leakage	Dataspace Operations / SaaS Platform	Low	High	High	Platform monitoring, cloud infrastructure redundancy, incident alerts	Define process to monitor internal data access based on identity	Platform Operations	Open
R-003	Unauthorized Deletion of Data	Dataspace Operations / Trust & Access Management	Medium	Low	Medium	API with IDP Access Controls	Restrict usage of methods like “DELETE” in the Connector Assets	Security / Dataspace Operations	Open
R-XXX	...	...	...	...	...	...	...	...	...

Appendix C: Attack Vectors and Baseline Security Controls

C.1 Priority Attack Vectors for the Manufacturing-X Dataspace

The following attack vectors have been identified as the highest priority based on the decentralized architecture, the use of Eclipse Tractus-X EDC, and the specific characteristics of manufacturing environments (including OT integration and long-term data confidentiality requirements).

Priority	Attack Vector	Description	Relevance to Eclipse Tractus-X EDC
1	Policy Bypass & Contract Tampering	Manipulation or bypass of ODRL usage policies during contract negotiation	Core to data sovereignty
2	Verifiable Credential Forgery / Replay	Creation of fake VCs or replay of valid presentations to gain unauthorized access	High – affects all authentication flows
3	Lateral Movement from EDC to OT	Compromise of the EDC used as a stepping stone to attack manufacturing systems	Very High – specific to Manufacturing-X
4	Supply Chain Compromise of EDC	Malicious code introduced via dependencies or updates to Eclipse Tractus-X EDC	High – open-source component risk
5	Data Exfiltration via Data Plane	Abuse of push or pull transfer mechanisms to extract sensitive data	Medium-High – especially for IP and quality data
6	Denial of Service against Connector	Flooding of catalog requests, contract negotiations or management APIs	Medium – can disrupt business operations
7	VC Revocation Bypass	Exploitation of weak or missing revocation mechanisms for Verifiable Credentials	Medium – critical for long-term trust

These priority vectors shall be reviewed and updated as part of the annual threat modeling cycle and after any significant incident.

C.2 Baseline Security Controls and Requirements

As part of the ongoing hardening of the Manufacturing-X dataspace architecture, we propose a first list of cybersecurity controls as an initial approach. This list is not intended to be exhaustive or final; it must be further extended, refined, and tailored to the specific operational and technical requirements of the Manufacturing-X ecosystem over time. They should also be mapped to each of the proposed Security Levels of the architecture. When developing this framework, we reviewed several established control catalogues and normative references. However, many commercial or licensed control frameworks impose restrictions that limit their practical usability. For example, the Secure Controls Framework (SCF) offers a basic version that can be used free of charge, but its license explicitly prohibits AI-assisted generation or automated processing. Similar copyright and licensing restrictions apply to various other control libraries and normative documents. In light of these limitations, we started to develop our own, independent list of controls tailored to the Manufacturing-X context. In an open to future development approach, we support the use of AI in future Risk Assessment processes. The following list gives a first approach and will require further extension and adjustment work:

1. Governance, Risk Management & Compliance Framework

Objective: Establish accountability and risk-based decision-making.

Key Controls:

- Formal risk assessment (dataspace-specific threats: policy bypass, VC forgery, connector-to-OT lateral movement, supply-chain compromise).
- **Information Security Management System (ISMS)** scoped to dataspace operations (especially for Core/Enablement Service Providers), covering risk identification, security measures, monitoring, and verification of effectiveness (per NIS2 Implementing Act).
- Defined roles: Connector Operator, Data Provider/Consumer, Attestation Provider, Security Officer.

- Regulatory mapping matrix (NIS2, CRA, Data Act, AI Act, GDPR).
- Third-party assurance for managed EDC services or cloud providers.
- Clear distinction between Company-Level and Centralized-Level controls.
- **Implementation in Eclipse Tractus-X**: Adopt the Data Trust & Security KIT rulebook template. Map Manufacturing-X use cases to risk scenarios.

2. Secure Architecture & Design (Security-by-Design/Default)

Objective: Eliminate vulnerabilities through architecture.

Key Controls:

- Threat modeling (STRIDE/PASTA) covering DSP/DCP flows, Control/Data Plane separation, VC paths, and OT integration.
- Secure defaults (no open management APIs, strict policy validation).
- Defense-in-depth with **IEC 62443 zones & conduits** between EDC, enterprise IT, and shopfloor OT.

3. Identity, Authentication & Authorization (IAM)

Objective: Cryptographically strong, decentralized trust.

Key Controls:

- Mandatory **Self-Sovereign Identity (SSI)**: Use of DIDs. Prepare the future implementation of eIDAS 2.0.
- Hardware-backed key management (HSM or FIPS 140-3 compliant KMS) for DID private keys and VC signing.
- Verifiable credential validation and revocation proof before every contract or data access.
- Multi-factor / phishing-resistant authentication for human operators.
- Attribute-Based Access Control (ABAC), combining Manufacturing-X identifiers with VC claims.
- **Implementation in Eclipse Tractus-X**: Enforce via Eclipse Tractus-X EDC + Data Trust & Security KIT. Use Managed Identity Wallet and Business Partner KIT for onboarding.

4. Encryption, Cryptography & Privacy Enhancing Technologies (Updated – BSI TR-02102-1:2026)

Objective: Ensure all cryptographic mechanisms are BSI-compliant, quantum-safe, and suitable for long-term protection of intellectual property.

Key Controls (Company-Level & Centralized-Level):

a. All inter-organizational data transmission and storage shall use BSI-approved algorithms per TR-02102-1.

b. Symmetric Encryption: AES-256-GCM (or equivalent AEAD mode) — mandatory and fully recommended.

c. Asymmetric Cryptography / Key Exchange / Signatures:

- **Baseline**: Elliptic Curve Cryptography using approved curves with ECDH and ECDSA.
- **RSA-2048 is immediately deprecated**. All existing RSA-2048 usage must be identified and migrated to approved Elliptic Curve curves without delay.

d. Post-Quantum Cryptography (Hybrid Mandatory for Long-Term & IP Data):

- Given that IP-related use cases share the same cryptographic implementation as other data flows, hybrid post-quantum cryptography must be adopted as soon as technically feasible.
- For data requiring long-term confidentiality: Implement **hybrid post-quantum schemes**
- Pure classical asymmetric mechanisms acceptable timeline to be defined.

e. Cryptographic key management, rotation, and destruction shall follow BSI TR-02102 guidelines. Use HSMs for all critical keys.

f. All TLS connections (DSP, DCP, management APIs, data plane) shall use BSI-approved cipher suites, prioritizing hybrid post-quantum where supported.

g. Storage encryption on each participant's infrastructure using BSI-approved algorithms.

h. Privacy Enhancing Technologies (PETs):

- Anonymization/pseudonymization of personal data.
- Differential privacy, where individual records must not be traceable.
- Secure Multi-Party Computation (SMPC) for computations without data disclosure.
- Homomorphic encryption for computations on encrypted data.

5. Secure Software Development & Component Lifecycle

Objective: Prevent vulnerabilities in code and dependencies.

Key Controls:

- Secure SDLC with threat modeling, secure coding, peer reviews, SAST/DAST.
- Mandatory SBOM for all Eclipse Tractus-X EDC deployments and custom components.
- Implement dependency vulnerability scanning.
- Implementation of automated AI-assisted vulnerability scanning using the latest Cybersecurity models.
- Coordinated vulnerability disclosure.
- Security requirements traceability for Manufacturing-X extensions.

6. Vulnerability Management, Patching & Configuration Hardening

Objective: Maintain a secure state continuously.

Key Controls:

- Continuous CVE monitoring for Eclipse Tractus-X EDC and dependencies.
- Risk-based patching SLA (critical: immediate).
- System hardening following BSI recommendations (disable unnecessary services, least privilege, secure configurations).
- Antimalware and exploitation protection tools on all company services (kept up to date).
- Regular attack surface reduction (unused APIs disabled, strict CORS, rate limiting).
- Hardware hardening (secure boot, runtime attestation).

7. Logging, Monitoring, Detection & Incident Response

Objective: Detect anomalies early and respond within regulatory timelines.

Key Controls:

- Structured, tamper-protected logging of all DSP negotiations, DCP exchanges, contracts, transfers, and policy evaluations.
- Integration with SIEM/EDR with correlation rules for suspicious patterns.
- 24/7 monitoring with defined escalation.
- Documented Incident Response.
- NIS2 compliant notification timelines and CRA vulnerability reporting. To be enhanced with immediate response patterns due to upcoming AI-Threats.
- Forensic readiness (preserve contract artifacts).
- Central incident reporting (ticketing) system.

8. Supply Chain, Third-Party & Ecosystem Risk Management

Objective: Secure the entire value chain.

Key Controls:

- Supplier due diligence and contractual security requirements.
- Continuous monitoring of Eclipse Tractus-X / Eclipse supply chain.
- Assurance via SBOM and audits.
- Secure onboarding/offboarding of participants (VC issuance/revocation).

9. Business Continuity, Resilience & Recovery

Objective: Ensure availability and rapid recovery.

Key Controls:

- High-availability EDC deployment with tested failover.
- Regular backup of critical configuration, policies, and contract state.
- Business continuity plan covering connector outage and alternative paths.
- Resilience testing (chaos engineering, tabletop), including OT failure modes.

10. Auditing, Penetration Testing, Continuous Improvement, Physical Security & Metrics

Objective: Demonstrate and continuously improve security posture.

Key Controls:

- Regular independent penetration testing and red-team exercises (focus: EDC endpoints, VC verification, policy enforcement, WAF/IDS/IPS effectiveness).
- Internal audits against the controls framework with external certification (ISO 27001, IEC 62443-2-1/3-3).
- Security metrics dashboard (policy strength, VC validation failures, patch compliance, MTTD/MTTR).
- Lessons-learned process feeding back into risk assessment and architecture.
- **Physical Security** (Company-Level):
    - Physical protection of server rooms and data storage locations.
    - Biometric scanners or access cards + video surveillance.
    - Consider standards: DIN EN 60839-11-1/11-2, ISO 27001 Annex A.9, BSI-TR 03126-5.
- **Web Application Firewall (WAF)** for all centralized web services, APIs.
- Intrusion Detection/Prevention Systems (IDS/IPS).

Appendix D: Current Security Measures Implemented in Eclipse Tractus-X

In Eclipse Tractus-X, an open-source project governed by the Eclipse Foundation, the maintainers list the following organizational measures and processes. Some of the processes are currently operated and imposed by the Eclipse Foundation, and other rules & security measures are imposed by the project members.

Process Security Measures:

Private Channel for Reporting Vulnerabilities available at:

- Eclipse Security: https://gitlab.eclipse.org/security/vulnerability-reports/
- Eclipse HelpDesk: https://gitlab.eclipse.org/eclipsefdn/helpdesk/-/issues
- Eclipse Tractus-X Sig-Security: https://github.com/eclipse-tractusx/sig-security/issues/new/choose
- Security Advisory at each Repository, for example: https://github.com/eclipse-tractusx/industry-core-hub/security/advisories

Processual Security measures:

Security Team: In the past, Eclipse Tractus-X had a dedicated security team. Currently, because the security team is understaffed, all the committers have responsibility regarding the Security of Repositories, especially the project leads. All contributors and security team members are mentioned on the Eclipse Tractus-X project website. It is mentioned that all committers are members of the security team.

- **Improvement Point: Security Team Volunteers are Needed!** If Eclipse Tractus-X would have a dedicated staffed team again, older procedures which were in place can be executed more often and with more level of detail, like performing Threat Modellings, Monitoring Vulnerabilities Across all repositories, putting in place security measures and executing risk assessments / pentesting in specific components which require a “higher” quality. For example, the Eclipse Tractus-X EDC has been pentested multiple times, and with that result the quality of the open-source solution has increased.

Eclipse Security Group: The Eclipse Security Group is always available if support is needed with implementing security practices in Eclipse Tractus-X, finding vulnerabilities, and supporting incident management (already triggered in the past). There is a dedicated staffed Eclipse Foundation manager responsible for it: https://accounts.eclipse.org/users/mbarbero

The support of this group was already used in several incidents that have occurred in Eclipse Tractus-X, for example, Trivy Vulnerability: https://mikael.barbero.tech/blog/post/2026-03-24-stop-trusting-mutable-references/

Project Leads also get notified by the Eclipse security team when a critical vulnerability is detected in one of the repositories of the Project under the Eclipse Foundation.

Identity Verification: https://www.eclipse.org/projects/handbook/#identity-verification Committers have the option to verify their identity: https://www.eclipse.org/projects/handbook/#identity-verification

Eclipse Tractus-X Security Release Guidelines (TRGs): In Eclipse Tractus-X, there are obligatory security guidelines for every single repository, this are the TRGs 8: https://eclipse-tractusx.github.io/docs/category/trg-8---security,

TRGs / Rules for deprecation of “unsecure” repositories: In the following TRG 7: https://eclipse-tractusx.github.io/docs/release/trg-7/trg-7-09 it is described that any product/repository that is unmaintained and do not have the latest security updated on their third-party libraries will get deprecated.

Eclipse Tractus-X Software Release Lifecycle: The Eclipse Tractus-X project has strict rules what can be include in the release changelog, therefore there is an open release process in place: https://eclipse-tractusx.github.io/docs/oss/release-process, this process includes an Open Planning (so there is no surprise features), an E2E Testing Phase which tests in a live environment the functionality of the systems (with documentation from the test cases being done in an Open Source JIRA [LINK/DOCS])

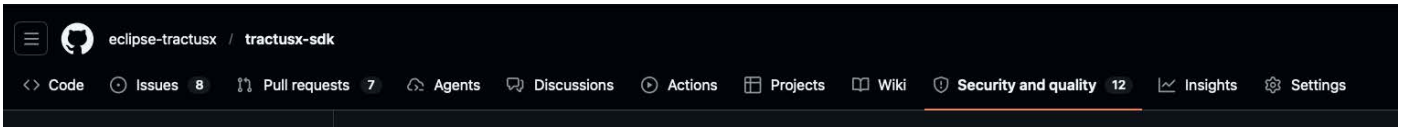
Automatic Security Scans:

In Eclipse Tractus-X the following security scans are in place:

Scan Type	Technology Used	TRG / Docs
Static Code Analysis	CodeQL	TRG 8.01

Infrastructure as Code (IaC) Static Code Analysis - GitHub Actions, Kubernetes, etc.	KICS	TRG 8.02
Secret Scanning	Trufflehog	TRG 8.03
Dependency / Container Vulnerability Scanning	Trivy	TRG 8.04
Dependency Automated Updates	Dependabot	TRG 8.05
Code Quality Scans	SonarQube	https://sonarcloud.io/organizations/eclipse-tractusx/projects
Mutable References Scanning (Pinable Actions)	Zizmor	Example EDC
API Pentesting	OWASP ZAP	Example: Tractus-X Identity Hub

All these report findings are collected at each repository security tab, which is only visible for committers (maintainers & security team) for security reasons:



Knowing the vulnerabilities of their products, committers can adapt and improve their solutions to support the latest secure third-party libraries.