

- TWIN Whitepaper

Reference Architecture

Trade Worldwide Information Network
Seamless Trading for All

Table of Contents

Table of Contents.....	2
Revisions.....	4
Credits.....	5
Rapporteurs.....	5
Technical vision.....	5
Other contributors.....	5
License.....	5
Executive Summary.....	6
Purpose and Scope.....	9
Introduction.....	10
The Challenges.....	10
TWIN's Approach.....	12
TWIN Architecture Overview.....	13
Preliminary Concepts.....	13
High-Level View.....	14
Application Plane.....	16
Data & Services Plane.....	17
Infrastructure Plane.....	18
Example Scenario.....	19
TWIN Reference Architecture.....	21
Design Principles.....	21
Service Anatomy.....	23
TWIN Trust Framework.....	26
Preliminary Concepts.....	26
Description.....	27
Enabling and Federation Services.....	29
Participant Onboarding.....	30
Visibility Services.....	31
Functional Overview.....	31
Core Visibility Services: Auditable Item Services.....	32
Extended Visibility Services: Event Management.....	36
Document Management Services.....	40
Functional Overview.....	40
Description and Functionalities.....	42
Data Exchange Services.....	48
Functional Overview.....	48

Publication and Discovery.....	50
Rights Management.....	52
Authentication and Authorization.....	57
TWIN Data Exchange Protocol.....	58
TWIN Data Space Connector Apps.....	60
Data Plane Functions.....	60
Data Plane interfaces.....	61
Push Data Transfers.....	62
Provider-initiated data transfers.....	62
TWIN Adaptor.....	63
Infrastructure Services.....	64
DLT and TWIN.....	66
Edge Devices and Connectors.....	67
Common Platform Services.....	69
Local User Management and Policies.....	69
Glossary.....	71
Technology-related.....	71
Identity and Credentials.....	71
Data Spaces.....	73
DLT.....	75
Domain-specific.....	76
Value Chains.....	76
International Trade.....	78
References.....	81

Revisions

Version	Date	Comments
Draft 1	December 2024	First version
Draft 2	January 2025	Second version. Reviewed by the IOTA team.
Draft 3	February 2025	Version for partner comments.
Draft 4	July 2025	<p>AIG component description refined to match reality. Event Management repurposed as an extended service and referencing other supply chain event standards, such as DCSA. Refinement of TWIN Data Space Connector, especially the subscriptions.</p> <p>Further development of the concept of TWIN Data Apps. Generalisation of the TWIN Adaptor concept. Multiple comments received from partners addressed.</p> <p>Refinement of Policy Management towards Rights Management and alignment with the IDSA policy enforcement architecture. Introduction of new services: archive storage, decentralized storage, and the TWIN Engine.</p> <p>TWIN ID separated from Ecosystem Admin tools.</p>
Release 1.0	January 2026	<p>Alignment with Eclipse DSP and W3C DCATv3 for catalogue resource descriptions.</p> <p>Removal of the Gaia-X Loire Service Offering and Data Resource. Addition of Data Space Contract Negotiation aspects.</p> <p>References to the Eclipse DCP.</p> <p>DS Connector Data Plane based on the W3C Linked Data Notifications and W3C ActivityPub, removing custom notifications mechanism.</p> <p>TWIN Adaptors section refined with latest practices.</p>

Credits

Rapporteurs

José Manuel Cantera Fonseca, David Philips (IOTA)

Technical vision

José Manuel Cantera Fonseca, Christoph Strnadl, Martyn Janes, Michele Nati, Isaac Odhiambo (IOTA)

Other contributors

Jens Lund-Nielsen, Ian Clark, Frank Dunsmuir, Andrew Brough, Antony Magayu, Adrian Grassl, Åsa Dahlborn (IOTA)

License

This work is licensed under a **Creative Commons Attribution Non-Commercial 4.0 International License**.



<https://creativecommons.org/licenses/by-nc/4.0/>

Executive Summary

TWIN (Trade Worldwide Information Network) is building the connecting backbone of digital value chain ecosystems, particularly in international trade and global supply chains. TWIN is based on open-source software and a suite of open protocols for data integrity, self-sovereign data management, and secure data exchange.

By leveraging *Distributed Ledger Technology*¹ (public, permissionless blockchain), *Data Spaces*², *DID*³, and *Verifiable Credentials*⁴, TWIN connects all parties involved in value chain ecosystems, ensuring:

- Scalable and cost-effective data exchange.
- Verifiable data integrity and transparency, as its decentralized ledger records and verifies all data exchanges.
- Confidentiality, privacy, and security.

TWIN technology provides a foundation for different solutions, such as:

- **Document and data exchange in international trade**, e.g., providing advance information throughout a consignment journey, including verifiable digital trade certificates to pre-clear consignments, minimising the frequency of manual document and physical inspections, and facilitating trade finance instruments.

¹ A public, permissionless DLT is a public ledger that is shared, replicated, and synchronized across a distributed and decentralized network of nodes, and where permissions are not required to maintain and operate a node.

² The term 'Data Space' refers to a type of data relationship between trusted partners who adhere to common standards and guidelines for data storage and sharing within one or many vertical ecosystems. In Data Spaces data is not stored centrally but rather at the source. Thus, data is only transferred through semantic interoperability as necessary.

³ A Decentralized Identifier (DID) is a unique, self-sovereign digital ID used to securely verify and manage identities without relying on central authorities.

⁴ Verifiable Credentials are secure, tamper-proof digital certificates that prove the authenticity of an issued credential concerning the attributes or qualifications of an individual, organisation or thing (e.g. a shipment).

- **Environmental and sustainability compliance⁵, declarations, and assessments** supported with traceable evidence – particularly in cases such as critical raw materials⁶ or capturing deforestation surveillance monitoring data – enabling the issuance and exchange of Digital Product Passports.
- **Supply chain visibility and optimization**, by increasing transparency and enabling data exchange and data verifiability through the use of Data Spaces and immutable ledger entries.

TWIN constitutes a digital pipeline that plays a crucial role in value chain ecosystems by fostering **collaboration** and **efficiency**. It allows stakeholders to access accurate and consistent data, enabling them to **make informed decisions** and **coordinate efforts seamlessly**. Additionally, it provides a comprehensive view of the ecosystem's performance, helping users to **identify issues early**, optimize operations, and innovate faster. In essence, TWIN empowers an **ecosystem** of users by promoting a unified, transparent, and data-driven approach.

TWIN's approach has been tested in use cases in East Africa via the Trade Logistics Information Pipeline Project⁷ (TLIP), and the UK's Border Trade Demonstrators⁸. Beyond the direct benefit of increased speed and reduced errors in cross-border goods movement, a TWIN-based solution brings the indirect benefit of touch-free information and document sharing. While the use cases named above focus on interactions with government and border agencies, the platform can integrate seamlessly with private sector platforms through interoperable, open interfaces and open source connectors.

To **minimise digitalisation costs** and unnecessary disruption, TWIN complements existing digital systems rather than replacing them, enabling them to evolve and integrate into a **broader ecosystem** while preserving data integrity, sovereignty, and privacy. It also democratises access to trade digitalisation for micro, small, and medium-sized enterprises (MSMEs), particularly in low and middle-income countries. This approach can help overcome export barriers, for example, by making it easier to provide verifiable evidence of compliance with ESG regulations⁹ aligned with [\[UN/CEFACT-Rec-49\]](#) recommendations.

⁵ Examples: EU's *ESPR Regulation* (Digital Product Passport) [\[EC-ESPR\]](#), the FDA's *Food Safety Modernization Act* [\[US-FSMA\]](#), or the *EU Deforestation Regulation* [\[EC-Deforestation\]](#) or several ESG standards and regulations.

⁶ <https://uncefact.github.io/project-crm/docs/about/purpose>

⁷ <https://www.tlip.io>

⁸ <https://mag.wcoomd.org/magazine/wco-news-103/developing-an-ecosystem-of-trust-at-the-uk-border/>

⁹ https://en.wikipedia.org/wiki/Regulation_of_ESG_rating_in_the_European_Union

From a technological perspective TWIN enhances the efficiency, transparency, traceability, interoperability, and trust of value chain ecosystems by digitalising processes. It offers open APIs and formats based on open software standards and recommendations from global trade and economic intergovernmental organizations. TWIN also pursues **technical compatibility** with the Gaia-X architecture [\[Gaia-X\]](#) and the International Data Spaces Association Reference Architecture Model [\[IDSA-RAM-4\]](#). Last but not least, the TWIN infrastructure incorporates an off-the-shelf DLT connector to the public, permissionless **IOTA** ledger¹⁰, a market-leading Distributed Ledger Technology (DLT) project and ecosystem.

An ongoing initiative, TWIN is the result of more than five years of research and development by the **IOTA Foundation** and its ecosystem. It has received sponsorship from non-governmental organizations – namely [TradeMark Africa](#), [Chartered Institute of Export & International Trade](#) (CIE&IT) – as well as the public funded research program [EU Horizon](#).

Aligned with the Principles for Digital Development¹¹, TWIN is committed to fostering innovation, preventing monopolistic practices, and enabling large-scale, global implementation. The TWIN-specific software is **open source** and implements several open standards, while key parts of the underlying software, such as IOTA DLT nodes, are also open source. At the same time, TWIN can complement proprietary software, allowing both models to coexist within its architecture.

To reinforce this strategy, TradeMark Africa, CIE&IT, and the IOTA Foundation were joined in May 2025 by the [World Economic Forum](#), the [Tony Blair Institute for Global Change](#), and the [Global Alliance for Trade Facilitation](#) to launch¹² the **TWIN Foundation**, tasked with shepherding TWIN's technology and business development and driving its adoption across the world.

¹⁰ <https://iota.org>

¹¹ <https://digitalprinciples.org/>

¹² <https://blog.iota.org/twin-foundation-launched/>

Purpose and Scope

This whitepaper presents the **Reference Architecture** for implementing TWIN. TWIN aims to enable Digital Public Infrastructures¹³ (DPIs) that connect participants across different value chain ecosystems, particularly in international trade and global supply chains. TWIN technology is based on open-source software and a set of open protocols that ensure data integrity, self-sovereign data management, and secure data exchange.

In this context, “architecture” refers to the fundamental organization of a system embodied in its components, as well as the relationships between the components to each other and to the environment. The architecture described here is a reference for any future DPIs and solutions built on TWIN. Therefore, this document is intended primarily for a technical audience – software architects, solution architects, and software engineers – who want to understand the whole vision, problem landscape, and proposed solutions. The [introductory chapter](#), however, provides a summary¹⁴ suitable for broader audiences.

The whitepaper outlines TWIN’s vision and long-term aspirations. However, at the time of writing, the TWIN software implementation is at varying levels of maturity. While some components are production-ready, notably those being used by the TLIP project, others are still in development or testing. During the coming months, the TWIN open-source codebase¹⁵ will continue to grow, and the open-source community governance models will launch. Future technical whitepapers may address the design landscape of specific components or deployment aspects of TWIN.

For business audiences, further background papers are planned, explaining how the TWIN technology is applied in different industries, such as international trade.

¹³ <https://www.undp.org/digital/digital-public-infrastructure>

¹⁴ An overview can be also found at <https://blog.iota.org/introducing-twin-technology/>

¹⁵ <https://github.com/twinfoundation>

Introduction

The Challenges

Modern **value chain ecosystems** are highly complex, involving not only many actors in both the public and private sectors but also different local and global regulations (*trading, sustainability, environmental, etc.*). For example, international trade ecosystems involve actors playing different business roles (*buyer, seller, exporter, importer, border control agent, health inspector, certification body, insurance companies, banks, etc.*), multiple procedures and rules (regulatory, transport, etc.), and the exchange of multiple documents (*Invoice, Import/Export Declaration, product certificate, Bill of Lading, Digital Product Passport,¹⁶etc.*) and business process data.

In addition to this complexity, value chains also face a fundamental technical challenge: the absence of digital ecosystems that can support a collaborative, sustainable, circular economy, allowing multiple participants to interact (i.e., exchange data/documents/credentials) without intermediaries and on a need-to-know basis. However, to achieve this, there are several technical challenges to be tackled [\[as described by Gaia-X\]](#), namely:

- **Interoperability:** The lack of standardized systems, interfaces, protocols, and semantics creates barriers to interoperability, making it difficult to deliver seamless, end-to-end solutions.
- **Integration** (particularly legacy systems): Integrating disparate systems and technologies across different countries and industries without a **common framework** can be costly, time-consuming, and prone to errors.

¹⁶ Digital Product Passports are part of new regulations and recommendations (EU Ecodesign for Sustainable Products [EU-ESPR], UN/CEFACT Recommendation #49. Transparency at scale [UN/CEFACT-Rec49]) where businesses can accredit sustainability compliance to border control agencies, market surveillance authorities, environmental control agencies or other businesses and consumers.

- **Data Sovereignty:** Ensuring that sharing data does not lead to potential risks like business leaks or legal and compliance issues.
- **Limited Resources and Complexity:** Companies, especially SMEs, often lack the resources to manage complex IT infrastructures or navigate intricate compliance requirements.
- **Vendor Lock-in:** The presence of ad-hoc integrations and bespoke applications makes it difficult to switch services or scale solutions.
- **Security and Compliance:** Mitigating the risks associated with data breaches, non-compliance, and third-party service providers.

To address the challenges listed above, several underlying factors must come into play:

- **Decentralization:** Designing a system with no single point of failure, no central information silos, and no single target of cybersecurity attacks, while harnessing network effects for simpler management.
- **Identity Management**, i.e., how different actors can be onboarded, issued an immutable digital identity, and start interacting in a trustworthy environment without the moderation of a central authority.
- **Information Management**, i.e., how to model and represent ecosystem information such as data about items, key business events, documents, etc., in a way that there is semantic interoperability and flexibility to accommodate different business needs and industry-specific particularities.
- **Data Exchange**, i.e., how actors can keep sovereignty of their data while establishing policies permitting **trusted** data exchange and collaboration without intermediaries or the need to make complex changes to existing systems.
- **Data Verifiability and Authenticity**, i.e., how actors can verify the provenance and authenticity of the exchanged data and trade documents.
- **Traceability, Immutability, and Transparency** in interactions among actors and concerning the assets they own. These are key to guaranteeing trust for frictionless collaboration that can solve disputes.

TWIN's Approach

Trade solutions built on centralized, proprietary platforms¹⁷ – controlled by a limited set of privileged actors – face not only scalability challenges but also resistance from the many participants in complex ecosystems, especially when sharing sensitive data such as IoT-generated information, trade volumes, commercially sensitive details, or personal data. Additionally, this concentration of control can grant too much power to a single infrastructure provider while creating an attractive target for malicious actors. Lastly, integrating different digital systems into a single proprietary infrastructure without widely adopted open interfaces exposes organizations to the **risk of vendor lock-in**.

Additionally, most value chains are not linear, but rather dynamic compositions that cross industry boundaries (e.g., the automotive sector is composed of metals, plastics, textiles, etc.) and jurisdictional boundaries (e.g., materials or parts can be sourced from multiple countries). Value chains **intersect at multiple points**, which poses interoperability challenges. Thus, actors have to gather data from different parties, and, as the number of stakeholders and value chains increases, the problem of data interoperability, verification, authorization, and authentication intensifies. This is critical for both governments (border control agencies, market surveillance authorities, etc.) and businesses (customs brokers, freight forwarders, carriers, ports, manufacturers, recyclers, importers, etc.).

Our proposed architecture revolves around a **TWIN Node**, a **modular** and **extensible** agent with open interfaces that facilitates participation in a **TWIN Ecosystem** without requiring ad-hoc integration between actors' IT systems. A TWIN Node encompasses all required infrastructure, including hardware and software infrastructure for processing, data, and object storage, including DLT nodes, platform software services, and core services essential for managing information in value chains. Additionally, a TWIN Node offers a public, open interface compliant with the Eclipse DataSpace Protocol **[Eclipse-DSP]** (Eclipse DSP) – implemented by a **TWIN Data Space Connector**¹⁸. On top of that, to tackle the specific needs of sectors, markets, or local regulations, specific software services referred to as **TWIN Data Apps** can also be deployed on a TWIN Node, particularly *TWIN Data Space Connector Apps*.

TWIN Nodes can run within a **Participant's** data centers, or a Participant can be onboarded and authorized to make use of the services offered by a TWIN Node provided by a third party (Node as a Service model).

However, it is noteworthy that there is no vendor or technology lock-in in our approach. It is not even necessary to deploy a TWIN Node to participate in a TWIN Ecosystem. It suffices to

¹⁷ <https://wisechainconsult.substack.com/p/8-why-did-tradelens-failed-and-and>

¹⁸ <https://internationaldataspaces.org/idsa-data-space-connector-report/#heading-main-5>

comply with the rules established by the [governance](#) of each TWIN Ecosystem and implement the **TWIN Data Space Connector protocol**. This unique approach enables existing supply chains and third-party platforms to seamlessly evolve from data silos to ecosystem Participants, leading to real, digital value chains for the benefit of governments, businesses, and consumers.

TWIN Architecture Overview

To understand TWIN and its vision, it is necessary to put different concepts into context.

Preliminary Concepts

In alignment with Gaia-X principles¹⁹ the following ecosystem-related definitions are in scope:

- **TWIN Ecosystem:** A value/supply chain ecosystem composed of: the *governance*, which defines the set of rules agreed upon by the ecosystem's parties and its implementation; and *infrastructure* - i.e., hardware and software for computing, storage, and network services, which adopts the rules defined by the governance.
- **Governance:** This defines the rights and duties of formal data management, ensuring quality and trust throughout a TWIN Ecosystem. This is mission-critical to TWIN, given that a central supervisory authority is missing by design.
- **Participant:** An actor who participates in a TWIN Ecosystem. Through a TWIN Ecosystem, Participants collaborate towards a common goal: the efficiency and effectiveness of global trade and supply/value chain processes. Participants adopt the *governance*, using the ecosystem's *infrastructures* “*to access and use data in a fair, transparent, proportionate and/non-discriminatory manner with clear and trustworthy data governance mechanisms.*”²⁰ The Participant can have one of two main roles:
 - **Consumer:** A Participant (or a service acting on their behalf) who searches for, requests, accesses, and consumes (via Data Services) Datasets offered by Providers.
 - **Provider:** A Participant who offers, negotiates for, and exposes Datasets via Data Services.

¹⁹ https://docs.gaia-x.eu/technical-committee/architecture-document/24.04/gx_conceptual_model/

²⁰ <https://interoperable-europe.ec.europa.eu/collection/semic-support-centre/data-spaces>

For example, a Provider can offer a Dataset corresponding to “phytosanitary certificate documents” issued in the UK, and expose it through a Data Service for use by a Consumer.

In addition, the following related technical concepts are paramount:

- **TWIN Node:** Infrastructure and platform software (agent) that enables Participants to interact within a TWIN ecosystem. A TWIN Node is extensible and provides the runtime environment for several services, including, but not limited to, data exchange services. The latter must be published to the [TWIN Catalogue](#) as Data Services. A TWIN Node is a logical construct and does not necessarily correspond to a single runtime process.
- **TWIN Adaptor:** The technical component that acts as a bi-directional **software bridge**, adapting the proprietary formats and API calls of an original IT system (that needs to take part in a TWIN Ecosystem) to the protocol and formats of a TWIN Data Space Connector. A TWIN Adaptor defines a **common framework** for integration with existing systems.
- **TWIN Data App:** Self-contained, self-descriptive software package that can be distributed via a TWIN-compatible package manager and deployed to **extend the functionality of a TWIN Node**.
- **TWIN Data Space Connector** (TWIN DS Connector): A set of technical core components that implement the TWIN Data Space Protocol, a key part of the [Data Exchange Services](#). It can be extended with one or more TWIN Data Apps, referred to as *TWIN DS Connector Apps*.
- **TWIN Catalogue:** A Federation Service that realizes a catalogue of compliant resources, namely Participants, and Dataset offerings (exposed through Data Services), enabling their registration and discovery.
- **TWIN Native Solution:** A tailored set of software services and applications designed to address specific business needs or problems by consuming and, where applicable, extending (via TWIN Data Apps) the core services offered by a TWIN Node.

High-Level View

[Figure 1](#) describes a basic overview of the TWIN Architecture gravitating around TWIN Nodes. From top to bottom, three main planes can be distinguished:

- The **Application Plane** consists of applications and services that deliver solutions by leveraging the capabilities provided by the *Data & Services Plane*. Existing systems or

third-party platforms that participate in a TWIN Ecosystem through a TWIN Adaptor are also part of the Application Plane.

- The **Data & Services Plane** is embodied by software services provided by TWIN Nodes, particularly Data Exchange Services, with the TWIN DS Connector as its core component.
- The **Infrastructure Plane** consists of the software infrastructure needed by TWIN Nodes to operate, which facilitates decentralization, data sovereignty and availability, discovery, and trust. A central element of this plane is the TWIN Federated Catalogue.

The TWIN architecture has an underlying **Decentralized Trust Framework**, which facilitates self-sovereign Participants and published resources, namely Dataset offerings and their access services (Data Services), onboarding, governed by rules specific to each ecosystem. This framework is designed to ensure **technical compatibility** with the IDSA, Gaia-X and EU Common European Data Space recommendations [\[EU-Data-Spaces\]](#). Thereby, it is built on **W3C Decentralized Identities** [\[W3C-DID-Core\]](#), **W3C Verifiable Credentials** [\[W3C-VC-DATA-MODEL\]](#), and **Trust Anchors**²¹. It enables the enforcement of data-sharing policies, ensuring secure and regulated data exchange.²²

²¹ https://docs.gaia-x.eu/technical-committee/architecture-document/24.04/component_details/#gaia-x-trust-anchors

²²

https://docs.gaia-x.eu/technical-committee/identity-credential-access-management/24.07/trustframework_implementation/

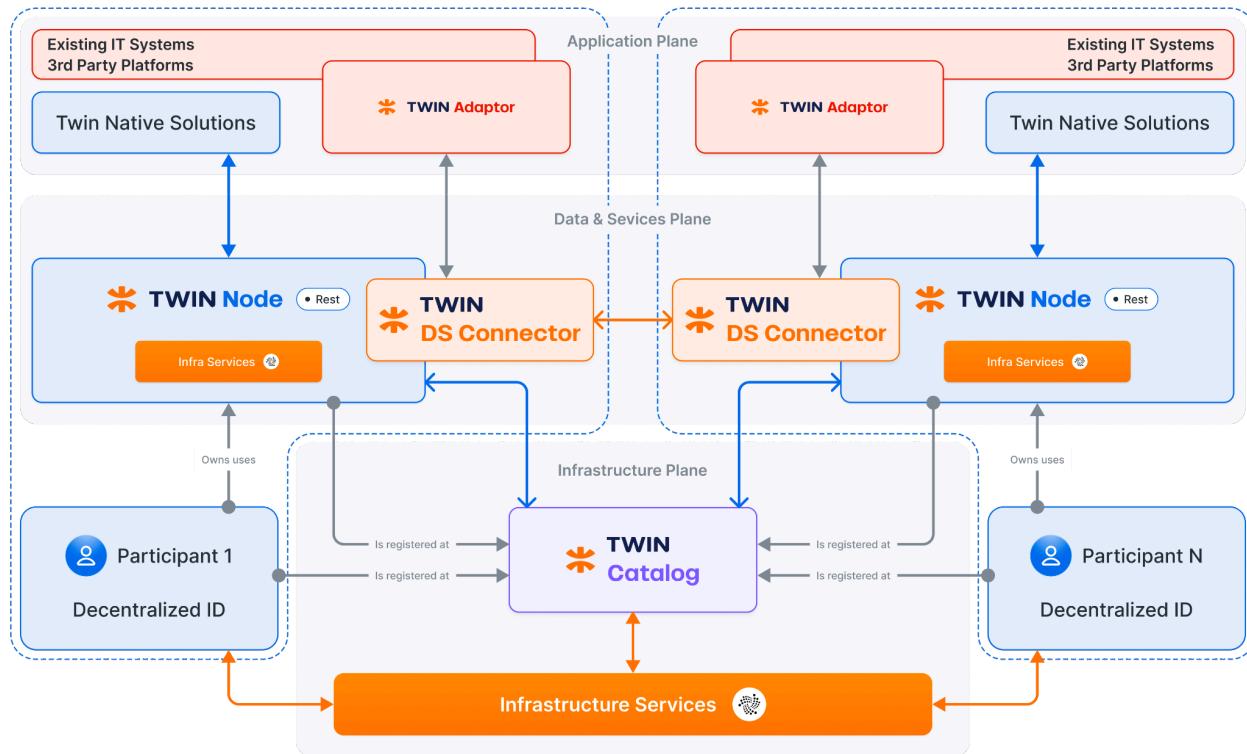


Figure 1 TWIN Node(s) and their interactions

Application Plane

The Application Plane includes:

- **Existing or legacy IT Systems** or third-party platforms that can participate in TWIN ecosystems through a TWIN Adaptor, i.e., a common integration framework. TWIN Adaptors' offered data may also be discoverable as Dataset offerings entries in the TWIN Catalogue. They could be government-owned Single Window Systems, trade operations systems, Enterprise Resource Planning systems, sustainability management systems (which are the source of sustainability data like carbon intensity), customs broker Solutions, automated compliance checkers, etc.
- **TWIN Native Solutions**, such as track and trace solutions, custom dashboards that allow users to gain visibility over items of interest, AI-based automated checkers for consignment clearance at borders, analytics tools, trade finance workflow tools, etc.

TWIN Native Solutions can discover services through the TWIN Catalogue or interact directly with a well-known TWIN Node, performing read and write operations. These interactions happen through open Web APIs that involve payloads represented by

industry-standard JSON-LD Vocabularies [\[W3C-JSON-LD\]](#) (GS1 Web Vocabulary²³, schema.org²⁴, UN/CEFACT BSP²⁵, etc.) The use of these standards enables semantic interoperability, so all exchanged data is machine-readable and unambiguous to Participants.

- **Third-party, cloud-native, digital platforms** can also complement a TWIN Ecosystem. These are IoT platforms that supply real-time data points along the trade journey and Track and Trace (traceability) platforms that supply real-time data points, giving visibility to the flow of goods and tracing products, materials, or components across a value chain. The service endpoints of these third-party platforms may also be discovered through the TWIN Catalogue.

Below this plane, the “Data & Services Plane” and the “Infrastructure Plane” are crucial to implementing a TWIN Node, the main software agent needed for ecosystem participation.

Data & Services Plane

This level is modular and composed of several software components that implement one or more services, which typically expose Web APIs (REST and/or WebSocket). These software components can be common platform services, information management services, or extensions (see below). The collection of these software components, when instantiated, constitutes, together with the necessary [infrastructure software services](#), a TWIN Node.

It is envisaged that general-purpose **extensions** can also be deployed as additional services to complement the core of a TWIN Node ([TWIN Data Apps](#)), for instance, to implement a custom protocol or format, such as GS1 EPCIS 2.0, for a certain supply chain subdomain. Another example is a TWIN Data App that is deployed to interface with an external OCR software as a service provider so that scanned documents’ data points can be extracted and persisted, as structured data, to the Node.

Therefore, the services exposed by **TWIN Nodes** (core and extended) are the backbone of the Data & Services Plane.

Apart from offering value chain visibility, document management services, and other extended functionalities implemented by TWIN Data Apps, TWIN Nodes can share data and documents with other TWIN Nodes and existing systems or third-party platforms. Actually, through TWIN Nodes, Participants exercise **sovereign control** over their data, i.e., they can choose which

²³ <https://ref.gs1.org/voc/>

²⁴ <https://schema.org>

²⁵ <https://vocabulary.uncefact.org>

documents or data they wish to share and with whom. Each business or government can own one or more TWIN Nodes. For discovery purposes, a TWIN Node's endpoints also need to be **registered** by Participants in the TWIN Catalogue. Each application may interface with one or more TWIN Nodes according to internal policies, data availability, etc.

At any time, an authorised Participant, or a client application on their behalf, can query a TWIN Node for data, for instance, about a specific item in the supply chain (e.g., a consignment or shipment). When needed, the services executing within a TWIN Node (for example, a custom TWIN Data App) can delegate a request to other TWIN Nodes or third-party applications (discovered through the TWIN Catalogue), effectively implementing broker services.

The interaction among different TWIN Nodes happens through a **TWIN DS Connector**, whereas existing IT systems can interact bi-directionally with TWIN Nodes by implementing a **TWIN Adaptor** that in turn speaks to a TWIN DS Connector. The data exchange transactions happening through a TWIN DS Connector (realized through Web APIs, i.e., REST or Websocket) are either directly related to fetching specific entities (documents', item's, digital twins, etc.) or receiving notifications of an "activity of interest". While the former does not mutate the state of the receiver Node, the latter could. For instance, notification of the availability of a new document will result in the mutation of an item and, optionally, the storage of such a document on the receiving Node.

The Datasets offered are registered in the TWIN Catalogue so that Participants or other services (possibly offered by TWIN Nodes) can discover them when needed. Registration will usually be accompanied by metadata, for example, the jurisdiction for which data/documents are provided, the associated offering [policies](#) for rights management purposes, the type of document/data items provided, etc. All these descriptions are represented using Linked Data Vocabularies – namely, the W3C DCATv3 Vocabulary [\[W3C-DCATv3\]](#), the GS1 Web Vocabulary, the UN/CEFACT BSP Vocabulary, or Schema.org.

Infrastructure Plane

On this plane, several infrastructure software services appear:

- The **TWIN Catalogue**, a decentralized registry of Participants and Dataset offerings. The latter is exposed by TWIN Nodes through a Data Space Connector that enables discovery, i.e., it can be queried to know who can provide which data, together with contract negotiation endpoints and offering policies.
- A **Verifiable Registry** that contains traceable and immutable objects, namely a Participant's Decentralized Identities. A default IOTA-based implementation is provided off-the-shelf, but the usage of other DLTs is possible if the corresponding [TWIN DLT Connector](#) is implemented.

- Datastores, object stores, key management systems, etc. that facilitate (secure) data storage and data availability, which are key for a TWIN Node.

Besides, **physical edge devices** (such as RFID Readers, scanners, printers, or mobile sensors) can also be part of the TWIN infrastructure, as they connect the physical world of trade items to the digital world, improving automated identification and data capture tasks. Devices manufactured by Zebra Technologies (namely fixed RFID Readers exposing the Zebra IoT Connector and Android scanners) have already been successfully tested²⁶ to interoperate with TWIN through their corresponding Edge Connectors (see [Edge Devices and Connectors](#)).

Example Scenario

The data exchange scenario can be better conceptualized through the example in [Figure 2](#), involving cross-border data/document exchange in international trade (one of the fundamental use cases of TWIN).

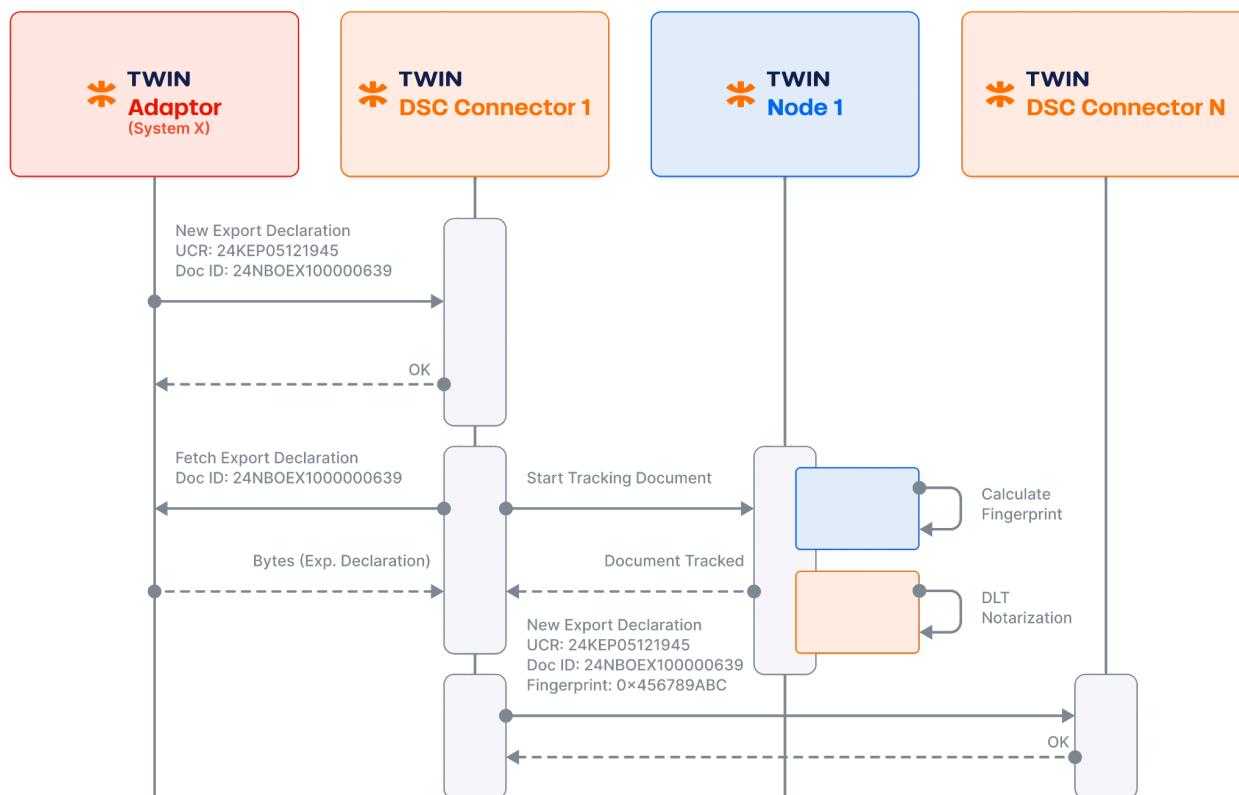


Figure 2 Scenario of interaction in a typical cross-border document exchange

²⁶ <https://developer.zebra.com/blog/introducing-zebra-iota-edge-sdk>

When a document is issued (such as an *export declaration* as in the example above), an external, existing IT System that has implemented a TWIN Adaptor (for instance, a Single Window System) can publish this activity information through a target Node's DS Connector. A TWIN Node can then start tracking the new document, preventing potential tampering (by first checking for the document's authenticity and then calculating and storing a fingerprint²⁷). It is important to note that documents can remain at the source and only be revealed (through the mediation of a TWIN Node) to participants authorized by the document owner. The offering policies, which specify who may be granted access to the data, are declared and published to the TWIN Catalogue by the document's owner using a standard, interoperable language – W3C ODRL [\[W3C-ODRL-22\]](#) (Open Digital Rights Language) – with the option to be referenced through a self-issued W3C Verifiable Credential.

On the other hand, we can imagine the destination country, namely border control agencies, being notified of the existence of an export declaration before the consignment's arrival, for the sake of efficiency and preparedness. The figure above shows another TWIN DS Connector that is also notified about the existence of the export declaration. This is feasible, as a TWIN DS Connector also offers push data transfer interfaces that enable authorized Participants to "follow" an activity stream of their interest, that will be received by their respective TWIN Nodes.

²⁷ There can be cases where the document is fully notarized in a DLT or even tokenized for further transfer operations.

TWIN Reference Architecture

Design Principles

The main technical design principles that guide the TWIN Reference Architecture are:

- **Interoperability:** To guarantee mainstream adoption, systems, and solutions “Powered by TWIN” must be interoperable. This enables TWIN value chain ecosystems to grow quickly and smoothly without incurring high integration costs, stimulating participation and removing entry barriers. Interoperability requires addressing the **architecture, protocols, payload, trust framework, and policy** aspects of TWIN as a software product, in alignment with relevant standards.
- **Data at the source:** Instead of creating new systems, TWIN aims at interworking with existing ones and facilitating their expansion by integrating off-the-shelf software components and libraries that expose standardised open APIs. As a result, the exchange or sharing of existing data, rather than the creation of new copies, can be achieved. The use of distributed ledger technology guarantees that exchanged data remains immutable and their source verified, thus increasing accountability and minimizing mistakes and fraud.
- **Data owner controls access:** Data Providers must be guaranteed direct control (**data sovereignty**) over who can access their data (and for which purpose), thus enforcing privacy and confidentiality. Unlike the widespread and problematic model of centralized data hubs, TWIN enables the design of neutral data-sharing solutions without central actors (such as platform owners) holding privileged positions. Data Providers can start new relationships intended to exchange data at their own will and thus create an attractive market for solution providers to compete in providing new services.
- **Confidentiality and privacy:** Data security must be ensured, allowing access only to rightful parties on a need-to-know basis. Additionally, safeguards should prevent publicly

shared value chain data from being exploited by other parties, such as competitors, for their advantage.

- **Decentralized data/document sharing and verification:** By the principles of decentralized technology, data and document sharing among stakeholders, as well as the verification of their authenticity and integrity, shall be achieved without any intermediaries. Participants will be able to discover one another and securely share exchange, and verify data/documents without the intervention of a central, privileged organization or intermediary.
- **Data minimization and selective disclosure:** For data and document sharing, it must be possible to expose only the minimal amount of information needed by other trade and supply chain Participants, for instance, by enabling selective disclosure at fine-grained levels.
- **Open-Closed alignment with international standards:** To ensure interoperability, TWIN does not intend to create new standards but rather to adopt existing, relevant ones while allowing for extensions. For digital twin representation, TWIN endorses GS1 Web Vocabulary and schema.org, among others, while being ready to support other vocabularies based on JSON-LD and Linked Data principles²⁸. For supply chain visibility, it is technically compatible with GS1 EPCIS 2.0, [\[GS1-EPCIS\]](#), which is inherently extensible. Concerning Digital Identity, TWIN adheres to W3C standards, and, for Data Space and ecosystem aspects, it is technically aligned with the Gaia-X and International Data Spaces Reference Architecture (IDSA) models and the Eclipse Dataspace Protocol.

²⁸ <https://www.w3.org/wiki/LinkedData>

Service Anatomy

[Figure 3](#) shows a functional overview of the architecture of TWIN aligned with the design principles outlined above and gravitating around the TWIN Node. The TWIN Node acts as a container of several software services, including core and extended ones (implemented through TWIN Data Apps), which expose several REST and WebSocket endpoints, using JSON-LD as the data representation format.

Through a TWIN Engine, which acts as a component manager, components are instantiated and their REST endpoints are set up to enable request processing. These exposed software services make use of the infrastructure (software) services, also included in the realm of the TWIN Node, to deal with data storage, DLT capabilities, etc.

The services exposed by a TWIN Node can be deployed (through Docker containers) as a monolith, i.e., one TWIN Engine hosts multiple components, or they can be split up into multiple microservices, via multiple TWIN Engines, each one executing in its own process space.

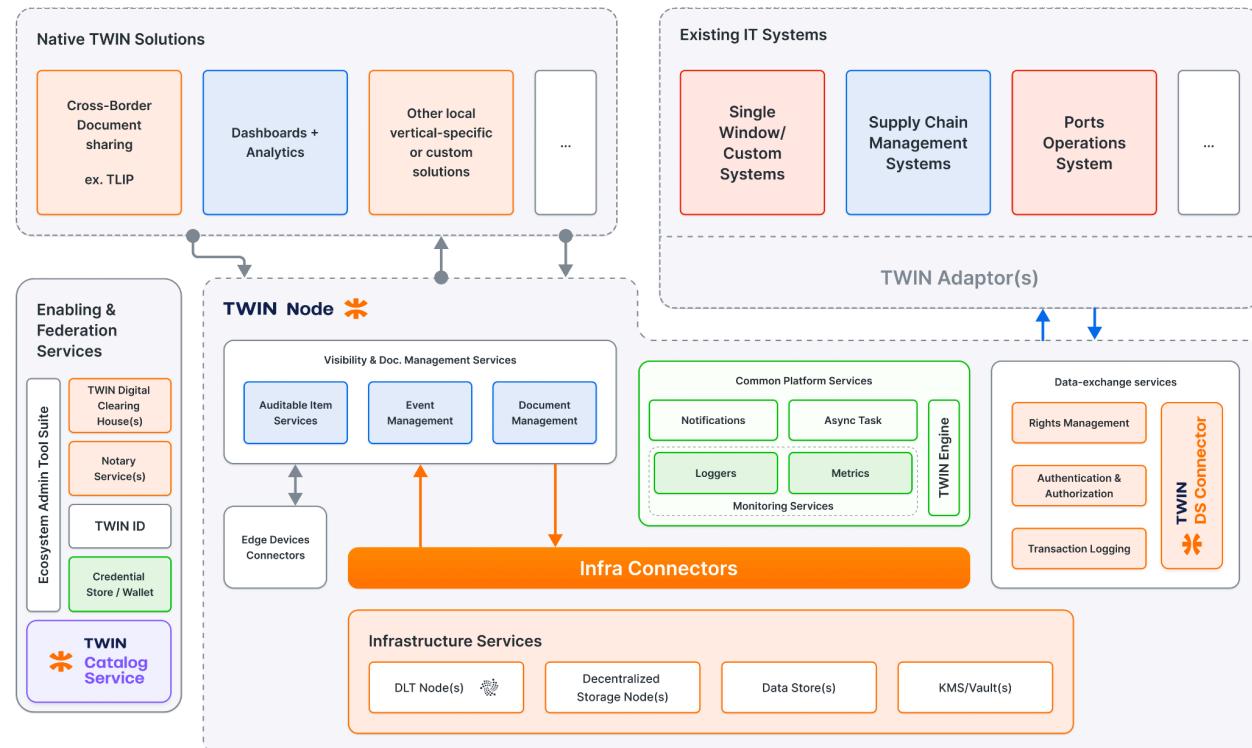


Figure 3 Anatomy of a TWIN Node and service categories

Different software services are classified under the following categories:

- **Enabling & Federation Services**, which realize the clearance, publication, and discovery of participants and the services they publish, so that decentralized interactions can take place. Authorization policies, depending on each ecosystem's governance, might also apply to these services for security and privacy reasons.
The *TWIN Catalogue*, depicted in [Figure 3](#) above, is a key component of TWIN's architecture as it contains all the information about Participants, Datasets, and their associated offering policies (represented using W3C ODRL).
- **Visibility Services** manage auditable digital representations of objects – descriptive digital twins – through their properties, relationships, business events (for instance, GS1 EPCIS 2.0) traceability data, and related resources (for instance, associated documents or external resources). Visibility Services are designed to be application-agnostic, enabling the usage of open standards based on Linked Data Vocabularies (schema.org, GS1, UN/CEFACT) and REST APIs.
- **Document Management Services** facilitate multi-version document discovery, location, retrieval, storage (if data cannot be kept at the source), document traceability, authenticity, and on-chain attestation NFT tokenization. Additionally, through TWIN Data Apps, data extraction and document transformation are also available. Besides, it facilitates multiple representations as per different industry standards (W3C VC, eInvoice, eBill of Lading, etc.²⁹). Document transfer, as per the Model Law on Electronic Transferable Records [\[UNCITRAL-MLETR\]](#), via IOTA DLT tokenization, is also under the scope of these services.
- **Data Exchange Services** facilitate the exchange of data and documents between the different ecosystem participants in accordance with negotiated contracts described via ODRL policies. In addition to rights management and policy enforcement functions, control plane services implement the Eclipse DSP Transfer Process Protocol³⁰, and data plane services offer query and publish/subscribe interfaces via REST endpoints, i.e. HTTP is used as a wire data protocol.
- **Infrastructure (Software) Services** realize the software infrastructure needed for a TWIN Node to operate, namely, public, permissionless distributed ledger technology, object storage, data stores, and cryptographic key stores (vault).

²⁹ A list of comprehensive key documents used in trade can be found at
https://www.dsi.iccwbo.org/_files/ugd/8e49a6_9f8444133fc64fc9b59fc2eaaca2888e.pdf

³⁰ <https://eclipse-dataspace-protocol-base.github.io/DataspaceProtocol/2025-1/#transfer-protocol>

- **Infrastructure Connectors** are generic technical core components of the architecture that abstract away the specific interfaces of infrastructure services from the rest of the services present in a TWIN Node. As a result, there is loose coupling and improved flexibility concerning the underlying infrastructure. One noteworthy Infrastructure Connector is the *IOTA DLT Connector*, which is key for trust and data verification within a TWIN ecosystem.
- **Edge Connectors** are connectors that bridge the physical world of trade items with their corresponding digital twin. *Edge devices* (such as RFID readers, mobile scanners, and mobile sensors) are part of the physical infrastructure, and Edge Connectors enable data to be captured seamlessly and object presence to be recognized, recorded, and attested on TWIN. Additionally, accessibility to trade item information held by TWIN can be improved (for instance, by reading or printing barcodes).
- **Common Platform Services** are general-purpose, reusable services that provide horizontal functions. Component management, Web API enabling, background tasks, notification delivery (messaging), metrics, and monitoring (telemetry) are some of the most noteworthy ones.

In addition to the above-referred software infrastructure, a TWIN Node must execute within a hardware infrastructure that encompasses computing, storage, and networking resources. Such a hardware infrastructure might be virtualized (IaaS) and be offered by cloud providers in conjunction with PaaS (containerization, clusterization, i.e., Kubernetes, etc.) capabilities. Nonetheless, TWIN Nodes, including the TWIN software infrastructure services (datastores, DLT, etc.), are also ready to be executed on-premise when required by organizations. In a nutshell, TWIN is not bound to any particular cloud provider or hardware platform. The deployment view of the architecture is out of the scope of this white paper and will be discussed in future documents.

TWIN Trust Framework

Additional and complementary definitions concerning this section can be found in the [glossary](#).

Preliminary Concepts

According to Gaia-X³¹, **Trust Frameworks** establish the rules that ensure minimum requirements are met for security, privacy, identification management, and interoperability through *accreditation and governance*. These operating **rules** provide a common framework for ecosystem participants, thereby increasing trust between them. The TWIN Trust Framework revolves around the following concepts:

- **Participant Identity:** A Decentralized Identifier (DID) plus other attributes related to a Participant.
- **Participant Attribute:** The identities of Participants in a TWIN Ecosystem rely on signed attributes, which can be requested and exchanged to gain individual trust from other participants. Participant Attributes are typically exchanged through Verifiable Credentials, whose subject is the Participant itself.
- **Participant Credential:** A type of Verifiable Credential that attests to Participant Attributes. For example, a *Legal Entity Credential* attests to the attributes of a legal entity, such as legal identity identifier, legal name, registered domain, residence country, etc. In other words, a Legal Entity Credential also implements a **cryptographic binding** between a DID and a legal entity identifier (EORI, DUNS, LEI, GLN, etc.).
- **Trust Anchor:** A Participant (such as a Conformity Assessment Body) accredited to issue attestations about specific claims. How the accreditation of Trust Anchors works depends on each TWIN Ecosystem's rules. There can be even regular Participants (such as prominent organizations in supply chains, for example freight forwarders) that also play the role of Trust Anchors.
- **Trust Anchor Credential:** A type of Verifiable Credential, its purpose is to provide a machine-readable representation of the accreditation of a Trust Anchor for a specific scope, vocabulary, or schema. There are several specializations:
 - *Organization Trust Anchor Credential* is held by an organization that can onboard other Participants within an ecosystem

³¹ <https://docs.gaia-x.eu/technical-committee/architecture-document/24.04/context/#gaia-x-trust-framework>

- *Ecosystem Trust Anchor Credential* is held by an organization that is entitled to define the rules of a particular ecosystem, for instance, actors promoting the creation of an ecosystem. Later, other actors may join if they have an interest in doing business within such a new ecosystem and abide by the rules.
- **TWIN Clearing House:** A service that assesses compliance of ecosystem entities, namely Participants and Dataset offerings, against predefined requirements. Upon successful validation, it issues a **Compliance Credential**. There can be multiple instances of a Clearing House operated by accredited ecosystem actors that abide by the same rules, available through a verifiable registry (DLT). It could be optional for lenient ecosystems that do not strictly require upfront verification of compliance.
- **Participant Compliance Credential:** A Verifiable Credential that attests that a Participant is compliant with the rules defined by a TWIN Ecosystem. Participants may need a Compliance Credential to be registered in a TWIN Catalogue.

Description

TWIN adheres to a **Trust Framework**, which enables the attestation of Participants' attributes and their seamless onboarding and interaction **without intermediaries**. The final aim is to ensure that all Participants in a TWIN Ecosystem are adhering to the policy rules agreed between the Participants of the Ecosystem itself.

The Trust Framework revolves around the following rules:

- Participants are identified by a *W3C DID* held in a Credential Wallet, either directly controlled or kept in custody by a third party. Even though TWIN **is not bound to any particular DID method**, it provides off-the-shelf support for IOTA Identity³². Thus, DLT infrastructure plays the role of a verifiable data registry.
- Participants' Attributes are attested by other Participants (**Trust Anchors**) through W3C Verifiable Credentials. Trust Anchors can be accredited by other Trust Anchors.
- Trust Anchors are defined by each TWIN Ecosystem. While there can be ecosystems with pre-defined, prominent Trust Anchors (for instance, government agencies, private institutions, banks, etc.), other ecosystems might be more lenient when it comes to attestation. For instance, an existing, compliant Participant might attest to the attributes of Participants it wants to interact with.

³² <https://docs.iota.org/iota-identity/>

- Participants must be compliant with each ecosystem's rules. Once a Participant has been attested to possess certain attributes, they may need to acquire a **Compliance Credential** through a TWIN Clearing House that checks for compliance. Stricter ecosystems require the presentation of Compliance Credentials in order to appear in the TWIN Catalogue.

Smart contracts are a good fit for the implementation of TWIN Clearing Houses. They guarantee clearance rules traceability, transparency onboarding processes, and immutability, enabling decentralization and seamless ecosystem governance.

- Ecosystem rules, which include participation rules, schemas, vocabularies, etc., could be registered on a **TWIN Registry** for traceability and immutability purposes, and be accessible to any TWIN Node and Clearing House.

External Trust Service Providers (also playing the role of Trust Anchor), through their controlled trusted data sources, are allowed to be part of the TWIN Trust Framework. Building on the concept of Gaia-X Notary³³, when these providers are not capable of issuing cryptographic material or signing claims directly, then a TWIN Ecosystem can accredit one or more **TWIN Notaries** to do so.

Example: The European Commission provides several APIs, including one to check the validity of EORI numbers. Unfortunately, those APIs are not returning Verifiable Credentials. A TWIN Notary service can be accredited by a TWIN ecosystem as the Trusted Data source for EORI validation and issue Verifiable Credentials to attest to EORI numbers.

- One advantage of Compliance credentials is that they, together with their evidence (also represented as Verifiable Credentials), are subject to **revocation**. As revocation lists are DLT registry entries, they can reach TWIN Nodes and be effective immediately. As a result, bad actors can be dismissed and offboarded from a TWIN ecosystem in a very efficient and effective manner.

³³

https://docs.gaia-x.eu/technical-committee/architecture-document/24.04/component_details/#gaia-x-trusted-data-sources-and-gaia-x-notaries

Enabling and Federation Services

The Enabling and Federation services encompass:

- The Credential Manager (**Wallet**) receives, stores, presents, and manages Verifiable Credentials and cryptographic key material, possibly using a Key Management System. TWIN plans to support external business wallets (for instance, the one under development in the EU³⁴) and cloud-based, custody wallets.
- The **TWIN Catalogue**, a federated or decentralized service that keeps a record of the different compliant Participants and the resources they publish, namely Dataset offerings.
- **Identity Service Providers** that can verify and attest the identity of Participants, enabling Know Your Customer (KYC).
- **TWIN Notary**, a service that, in collaboration with an external trust service provider, attests to a Participant's specific attributes.
- **TWIN Clearing House**, which checks the compliance of Participants and Dataset offerings per the ecosystem's rules (these rules are represented as records in the TWIN Registry).
- **TWIN ID**: An application that deals with and facilitates Identity Management for organizations, including trust anchors, credential schemas, legal entity attribute verification, etc. It can also play the role of a cloud-based organizational Wallet, storing keys and credentials in custody. Future documents will describe in detail the services offered by TWIN ID.
- **Ecosystem Administration** tool suite. UIs and service orchestrators are intended for seamless onboarding of Participants and the resources they publish into an ecosystem. Orchestration flows may involve different services such as the above-referred TWIN ID, TWIN Notaries, the TWIN Clearing House, Rights Management, etc.

The interaction among these services is depicted in [Figure 4](#):

³⁴ <https://www.webuildconsortium.eu/>

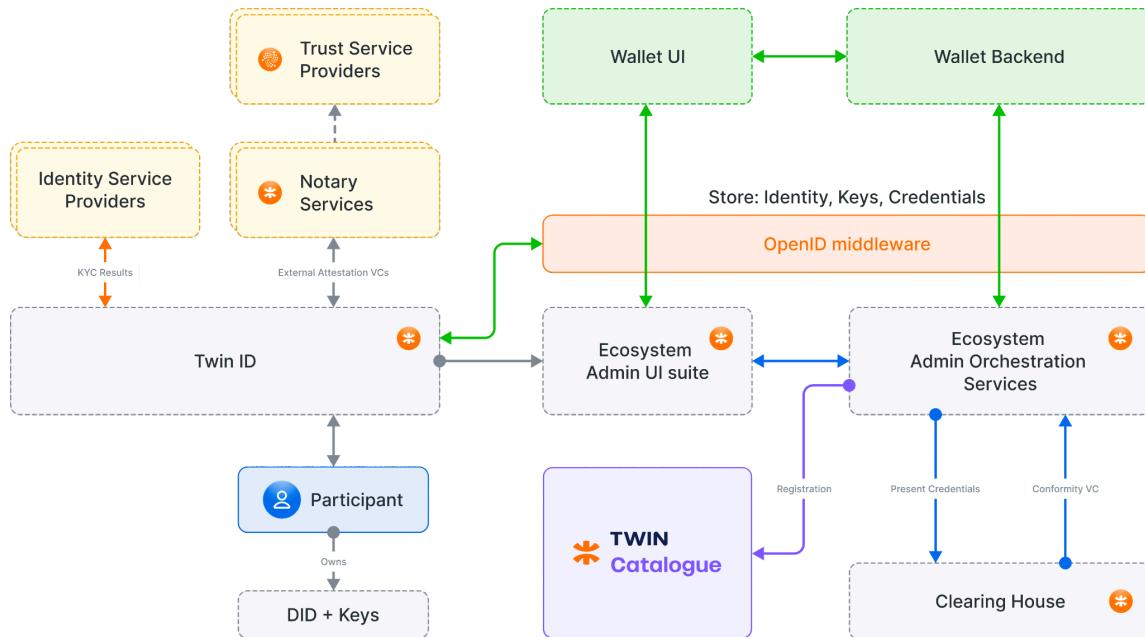


Figure 4 TWIN ID and onboarding overview

Participant Onboarding

To facilitate onboarding, TWIN offers a TWIN ID Application that can orchestrate the issuance or presentation of essential credentials for Participants.

For organizations that don't have a decentralized identity, onboarding starts with the minting and registering of a new DID (for example, by using IOTA APIs³⁵) along with the corresponding key material. Participants can then declare their attributes and, where necessary, solicit verification from external identity providers. Throughout the process, Participants can present evidence that will be compiled into a set of Verifiable Credentials signed by Trust Anchors (including [TWIN Notaries](#)), to attest to the Participants' attributes.

Optionally, the TWIN ID application can collaborate with **Identity Service Providers** to enable KYC Services. However, TWIN remains agnostic to any specific KYC methodology. Different onboarding flows may be required depending on the type of Identity Provider involved – ranging from formal procedures (e.g. bank or government-issued verification) to lighter alternatives such as domain or email ownership proofs. Additionally, TWIN Notary services can also act as Credential Issuers. The interactions needed for credential issuance could be based on the

³⁵ <https://docs.iota.org/developer/iota-identity/getting-started/wasm>

Eclipse Decentralized Claims Protocol [\[Eclipse-DCP\]](#) (Credential Issuance³⁶) or **OIDC4VCI** (Open ID for Verifiable Credential Issuance) flows [\[OID4VCI\]](#). This aspect is still under evaluation and prototyping.

When an organization's Verifiable Credentials are available, they can be presented to a Clearing House through the *Ecosystem Administration Tool*. A Clearing House can verify their compliance and issue a Compliance Credential once approved. For those more restrictive TWIN Ecosystems, this Credential is the ultimate proof that a Participant is entitled to be part of.

The presentation process adheres to a standard **OIDC4VP** (Open ID for Verifiable Presentations) flow [\[OIDC4VP\]](#), for the sake of interoperability. In this case, the Ecosystem Administration Tool backend will act as a mediator between the organization's Wallet or credential storage service and the Clearing House. Upon successful compliance, a new Credential will be issued, following a standard OIDC4VCI flow.

Visibility Services

Terminology concerning this section can be found in the [glossary](#).

Functional Overview

Visibility refers to the ability to fully understand and track the steps taken by an item along the value chain, thereby achieving transparency. The final aim is to enable ecosystem Participants to capture and query, on a need-to-know basis, information and business events concerning their circulating items of interest.

Just as Participants' identities should be *discoverable, resolvable, and verifiable*, so too should item identifiers. Items can refer to shipments, products, documents, locations, etc., and their related business events could be manufacturing, reuse, recycling, shipping, delivery, export, and import. etc.

Visibility is a key enabler for traceability, optimization, and efficiency. It enables Participants to always know where trade items are and where they come from, who has custody over them, when they will arrive, and whether they are being transported properly. Visibility benefits businesses, consumers, governments, and regulators in a wide variety of use cases such as compliance, provenance, authenticity/anti-counterfeiting, or N-Tier traceability.

³⁶

<https://eclipse-dataspace-dcp.github.io/decentralized-claims-protocol/v1.0.1/#credential-issuance-protocol>

Core Visibility Services: Auditable Item Services

The Auditable Item Services are at the core of visibility functionalities and are composed of:

- The **Auditable Item Graph** (AIG) is a service that allows the management of auditable entities. The underlying data model of the AIG is the **property graph** data model³⁷. Graph vertices represent entities and are labelled with JSON-LD annotation objects to capture entities' descriptive digital twin. Graph edges, which can also be annotated with JSON-LD, are used to represent entity relationships. Each entity, i.e., Auditable Item, has extra, well-known, annotated graph relationships that represent aliases (alternative identifiers), attached resources (for instance, trade item's digital images), or external data sources.

Auditability refers to the ability to store different versions (change sets) of an entity – its history – as it undergoes change, such as additions or removals of properties, edges, etc. This allows for tracing the entity's state over time.

- **Auditable Item Streams** (AIS) is a complementary service to the Auditable Item Graph, whose main functionality is to capture append-only, auditable **data streams** typically associated with Auditable Items. A data stream can just contain a few value chain business events or even Streaming Data. Streaming data could be emitted at high volume in a continuous, incremental manner with the goal of low-latency processing. Typically, organizations have thousands of data sources that simultaneously emit messages, records, or data (including location, event, and sensor data) that companies use for real-time analytics and visibility into many aspects of their business.

[Figure 5](#) depicts the high-level architecture and context within a TWIN Node of the Auditable Item Services. These services are wholly generic, application-agnostic, and are technically compatible with Linked Data Vocabularies from schema.org, GS1, and UN/CEFACT. These services primarily expose a REST API that can be consumed by TWIN Native Solutions, or a custom TWIN Data App, including TWIN DS Connector Apps. Such a REST API exposes methods for creating and updating an Auditable Item, i.e., a uniquely identified graph vertex, and adding/removing alias IDs, resources, graph edges, data streams, and so on. Auditable Item's annotation objects, edges, and resource identifiers, along with their metadata, can be stored in a regular datastore (MySQL, DynamoDB, etc.) using the proper data store connector. Blob files are stored in a blob store. Data streaming is also facilitated by a WebSocket interface.

The Auditable Item Services offer a subscription interface so that changes on Auditable Items of interest can be notified. For example, when a new Auditable Item (for instance, a document) is

³⁷ https://en.wikipedia.org/wiki/Property_graph

newly referenced from an existing Auditable Item, other TWIN Nodes (via a TWIN DS Connector) can get notified.

Finally, the audit capabilities of the Auditable Item Graph and Auditable Item Services enable Data Consumers to perform data verification. This allows them to verify whether a trade entity's data has been tampered with or altered in any way, helping to detect if a trade item has suffered any diversion in the supply chain. This is where the **DLT Connector** plays a critical role. Through this component, the AIG or AIS data is made verifiable through an object in the IOTA Ledger that represents a Data Integrity Proof [\[W3C-Data-Integrity\]](#). As a result, third parties can independently validate the authenticity of the data without relying on intermediaries.

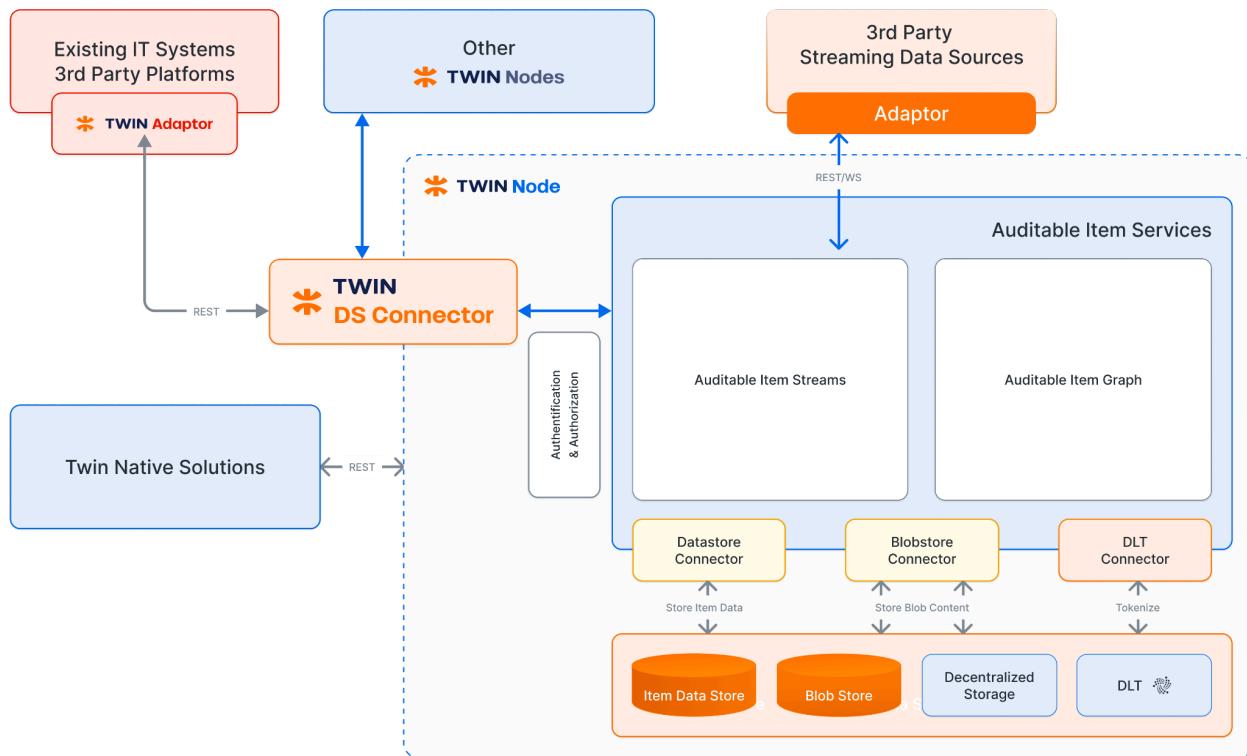


Figure 5 Auditable Item Services context diagram and functional overview

Trade Entity Information Model

[Figure 6](#) illustrates the information model that can be implemented through the AIG component. Following the W3C **Linked Data** recommendations and conventions, trade entities in TWIN (usually [documents](#) or [trade items](#)) are represented by:

- The entity identifier, a URI, ideally a resolvable Digital Link [\[GS1-Digital-Link\]](#).

- The entity type(s) (Product, Consignment, Shipment, Document, etc.), which are also identified by a URI, specify a semantic and unambiguous type(s).
- Properties (for instance, *weight*, *location*, *start date*, etc.). A property's fully qualified name is a URI, usually imported from the UN/CEFACT, GS1 Web, or schema.org vocabularies.
- Relationships (*is a child of*, etc.), also uniquely identified by a URI.
- Associated Resources, also uniquely identified by a URI, can correspond to different assets related to a trade entity:
 - Blobs (for instance, an image depicting a product's photography)
 - Documents managed by the [Document Management](#) services (see below) are also internally represented as Auditable Items.
 - Data Streams, managed by the [Auditable Item Streams](#) service, with streaming data associated with the Item, for instance, real-time location data with a fine granularity level.
 - Extra Datasets registered in the TWIN Catalogue that can be queried through a Data Service to obtain additional information, for instance, time series data for product tracking offered by third-party applications.

The main element (represented in [Figure 6](#) by a rectangle) is trade entities. These entities or items should ideally be **serialized** elements (**instances**) in a value chain – for example, a part of an electronic device. Apart from its main Item ID (a URI), a trade entity might be associated with other IDs named **Alias IDs**. This reflects the fact that, in a supply chain, different actors usually refer to the same trade entity through different identifiers.

Every trade entity has one or more **types** (represented in [Figure 6](#) by a rounded rectangle) (*isA* relationship). Entities can have **properties** (represented by an oval) – for example, the weight of the part, or the location of the part in a warehouse – and associated documents stored as blobs (represented in [Figure 6](#) as a quadrilateral), generalized as the trade item's **resources**.

Trade entities can be related to other entities, and these **relationships** are depicted by a diamond in [Figure 6](#). For example, in international trade, a *Consignment*³⁸ – related to a business agreement between a seller and a buyer – may be split into multiple *Consignment Items*³⁹ for delivery or customs purposes. From an information management perspective, this is represented as a single Auditable Item (vertex), labelled with an annotation object of type “*Consignment*”. This vertex has an edge that represents the “parent of” relationship, pointing to

³⁸ <https://vocabulary.uncefact.org/Consignment>

³⁹ <https://vocabulary.uncefact.org/ConsignmentItem>

multiple Auditable Items (vertices) labelled as a “Consignment Item” type. Conversely, each *Consignment Item* has an edge that represents the “child of” relationship with its parent *Consignment*.

Relationships can also have properties. A typical example is the date from which the “child Of” relationship applies, corresponding to the date the consignment was prepared. They are specified as properties of an annotation object bound to an edge in the AIG.

Trade entities can also be documents that may have relationships among them and among the trade items they reference. For example, a *Customs Declaration* can point to a *Commercial Invoice*, and a *Bill of Lading* might reference items formerly referenced by one or more *Custom Declarations*.

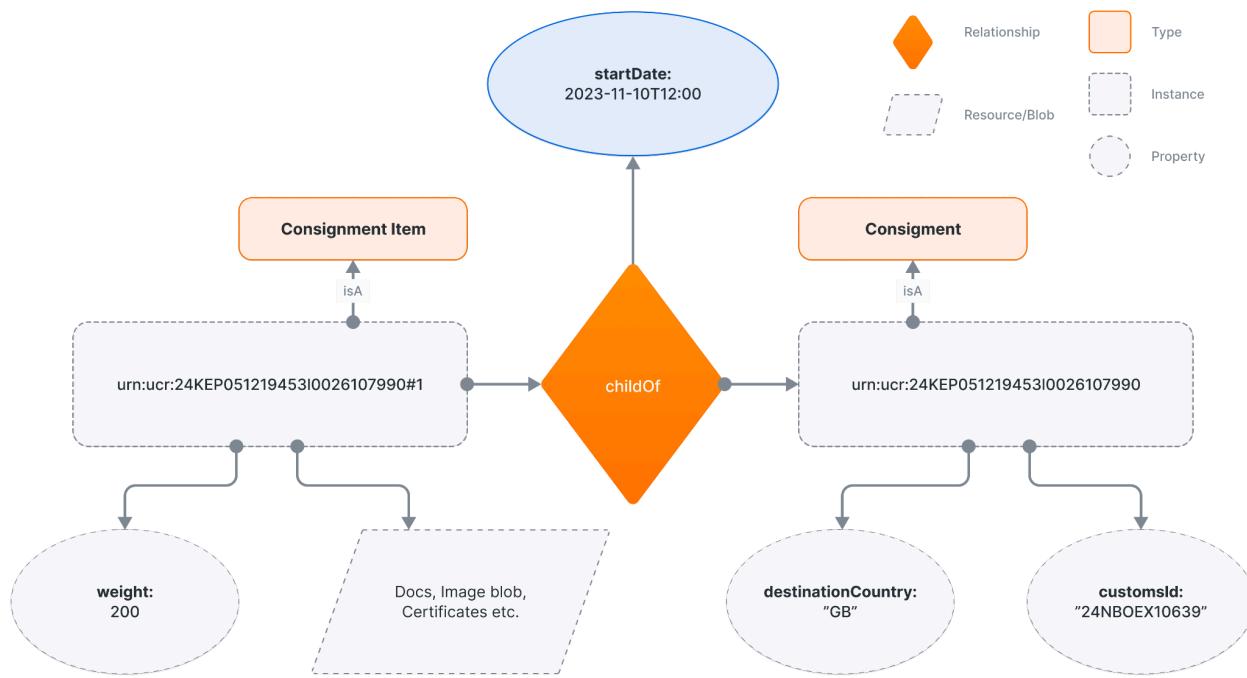


Figure 6 TWIN Information model of Trade Entities, implementable through the Auditable Item Graph

Extended Visibility Services: Event Management

In value chain ecosystems, particularly trade and supply chains, stakeholders often need advance information about trade items, including their nature, location, transport route, and condition. These data points constitute Events that can be transmitted as “signals” embodied in the data exchange between different TWIN Nodes. *Key Events* may include health checks, departure or arrival times, entry and exit records for ports, document availability, and other critical logistics details. See, for instance, the DCSA⁴⁰ track and trace specification.

Event Management refers to the capture, storage, retrieval (through queries or subscriptions), and timestamping of **key value chain Events** related to Auditable Items. Event management enables multiple business steps to be traced and can be applied to a wide variety of ecosystems.

Several characteristics of value chain Events can be exploited together with core TWIN capabilities to increase trust within an ecosystem:

- Events are **immutable** by definition. Once an Event has been disclosed to other Participants, it cannot be changed. Events are always **appended** and never deleted. If an Event proves to be erroneous, a further Event with an error declaration must be appended.
- Critical business Events can be linked to **proofs** that timestamp or attest to the legitimacy of an Event, i.e., that it has not been faked and its real occurrence time. TWIN advocates two types of proofs for Events:
 - DLT Commitments attest to the existence of one or more Events at a particular point in time. That feature is provided by the [Auditable Item Streams Service](#).
 - Physical Device Events can attest to the physical presence of an item at a particular location (see [Edge Connectors](#))

The following are minimal data points that should be captured by an Event (all have a standard representation in different Event data models):

- Event identifier, a URI that is constructed, for immutability, via a hash of the canonicalized representation of the event's data points.
- Event type(s), represented by a URI, which provides a semantic, unambiguous categorization of the event.
- Targeted item or items (**what**), identified by a URI.

⁴⁰ <https://dcsa.org/standards/track-and-trace>

- Timestamps (**when**):
 - When the event physically occurred.
 - When the event was recorded (for example, recorded digitally after its physical occurrence)
- Physical location (represented by geo coordinates or by an identifier – for instance, an UN/LOCODE⁴¹ location) where the event has taken place (**where**).
- A DID Identity of the Participant registering the event (**who**).

Figure 7 shows an example of an Event corresponding to an inspection performed on a target trade item at a particular business location (such as a warehouse or seaport) and originated by a Participant identified by a DID (for example, the customs agent performing the inspection).

	Event ID	ni:///sha-256;f12640477da404f845f5e4a2d4071fdecda7cb1af6fc0fd89462654cf2b94f43?ver=CBV2.0
What	Event Type	Visibility
	Target	https://id.gs1.org/01/09521987654327/21/202301
When	Observed at	2024-09-12T15:00:03.321Z
	Recorded at	2024-09-12T 15:05:03.321Z
Where	Location	ES VLC
Who	Actor	did:iota:0xb62afcd0150d048ea0679af61d28d0eb1ad1b969f411b03997194df232b27383
Why	Business Step	Inspecting

Figure 7 Conceptual overview of an Event

Another case is supply chain traceability represented as multiple business steps from manufacturing to distribution and finally retail. These steps are usually captured and timestamped using the GS1 *EPCIS 2.0 representation*, which may include additional data points such as:

⁴¹ <https://unece.org/trade/uncefact/unlocode>

- Read point (usually represented by a location identifier) that gives precise details of the location where the event occurred (for example, a specific door number in a warehouse).
- Environmental conditions of the capturing place, such as temperature or humidity.
- The device that served to capture the event (**how**), including its URI identifier and other details.
- Business process or business transaction involved (**why**).
- Details of the reported business process, such as completion status, start date, and end date.
- Participants concerned: generator (rapporteur) of the event, actor or actors involved.
- Other items involved, such as child items or attached sensors.
- Additional objects related to the event, such as documents involved.
- Target item-specific properties at the time of capture, for instance, temperature, weight, etc.

As there are multiple standards for representing Events, tailored to different domains and use cases, the TWIN Core Visibility Services remain agnostic, while at the same time offering capabilities to implement specific supply chain Event management standards packaged as TWIN Data Apps, i.e., as TWIN Node extensions.

[Figure 8](#) illustrates the context in which an Event Management Service can be implemented using TWIN Core Visibility Services.

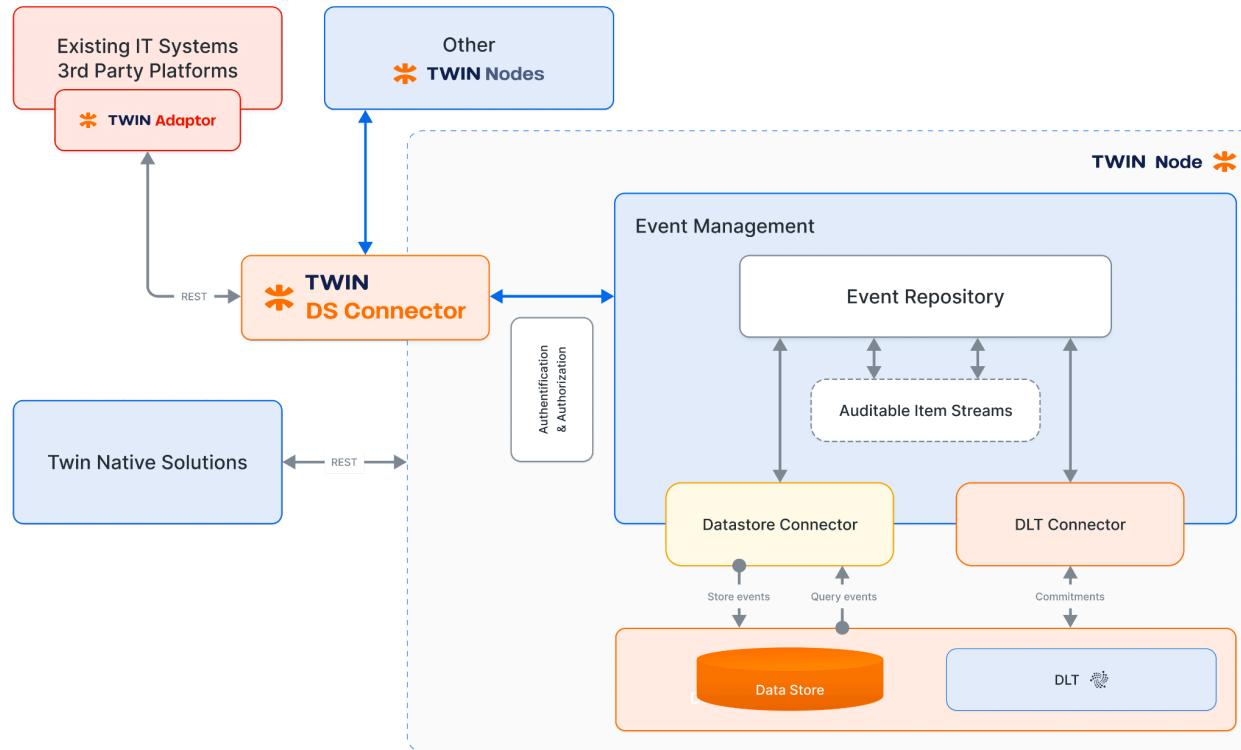


Figure 8 Event Management Service on TWIN: context diagram and functional overview

Events may arrive through REST requests generated by authorized TWIN Native Solutions or through the TWIN DS Connector. As Events are an essential data exchange item, many will come from existing IT systems or other TWIN Nodes. These Events have to be managed properly by a TWIN Node, typically involving two main software components:

- An **Event Repository** is in charge of exposing the fundamental REST API endpoints that allow Events to be captured (including syntax validation and indexation) for later retrieval. This component may support standard API definitions, such as the GS1 EPCIS 2.0 REST APIs⁴² or the DCSA Track and Trace APIs⁴³.
- The **Auditable Item Streams Service**, which offers the backend capability of grouping events (represented using JSON-LD) concerning the same trade entity under an object stream. This approach presents multiple advantages:
 - Different flavours of Event representation can be supported, in particular, the **GS1 EPCIS Linked Data Model**⁴⁴.

⁴² <https://ref.gs1.org/standards/epcis/openapi.json>

⁴³ https://app.swaggerhub.com/apis/dcsaorg/DCSA_TNT/2.2.0#/Events/get_v2_events

⁴⁴ <https://ref.gs1.org/epcis>

- Event store is data store agnostic as the AIS service utilizes the [TWIN Datastore Connectors](#).
- The Auditable Item Stream (AIS) concerning a trade entity can be linked to its Auditable Item, facilitating business event discovery.
- The entries of an AIS can be automatically timestamped through a DLT object for auditability and verifiability purposes, as formerly mentioned.
- Advanced features of the AIS service can be exploited to implement publish/subscribe API endpoints, such as the “query subscription endpoints” defined by the GS1 EPCIS 2.0 Open API⁴⁵.

Document Management Services

Definitions referenced in this section can be found in the [glossary](#).

Functional Overview

Documents are an integral data exchange asset in trade and value chain ecosystems. Different interactions critical among businesses, government, and consumers involve **key documents**:

- *Business to government (B2G) and government to government (G2G)*: Clearance of goods at borders requires the availability and authenticity of essential documents (such as a certificate of origin, health certificate, and export or import customs declaration).
- *Business-to-business (B2B)*: key trade documents⁴⁶ include bills of lading or commercial invoices that can be the subject of upper-level ecosystems such as *trade finance*.
- *B2G, B2B, and Business to Consumer (B2C)*: Value chain ecosystems and Digital Product Passports rely on documents like product certifications [\[UN/CEFACT-Prod-Certificate\]](#) issued by auditors (for example, eco-labels), test reports issued by manufacturers, product specifications and manuals, bills of materials to facilitate recycling operations, verifiable evidence of compliance with ESG Regulations, and so on.

Note the distinction between the two different classes of documents:

- *Pure, native eDocuments* (ePhyto, eInvoice, etc.) are represented using a semi-structured, machine-readable format (XML, JSON(-LD), CSV, etc.). If they include a digital signature, their authenticity and integrity can be checked through cryptographic

⁴⁵ <https://ref.gs1.org/standards/epcis/openapi.json>

⁴⁶ <https://www.legislation.gov.uk/ukpga/2023/38/section/1/enacted>

methods. Recent legislation, for instance, the *UK Electronic Trade Documents Act*⁴⁷, in alignment with MLETR, provides for certain electronic trade documents (including electronic bills of lading) to be accorded the same legal status as their paper equivalent if they meet certain relevant criteria.

- *Scanned, generated, plain documents*, a digital version of a paper document resulting from a scanning or rendering process, usually saved in PDF format. Even though they can be read by AI-based recognition tools, they are not ready for processing using simple techniques and are prone to error and fraud. Normally, they do not include any digital signature; thus, their verification requires human intervention (checking visually an official stamp, for instance).

From the point of view of document management, there are three key roles:

- **Issuer** (*issuerParty*⁴⁸ as per UN/CEFACT), an accredited entity (government agencies, conformity assessment bodies, regulators, auditors, accredited organizations, etc.) that issues a document to an interested party. For instance, a Chamber of Commerce issues Certificates of Origin (CoO). Traditional issuance methods rely on hand-made signatures and official stamps for authenticity.
- **Holder** (*senderTradeParty*⁴⁹ as per UN/CEFACT), a Participant who sends and presents a document issued by an issuer. For instance, the exporter in the case of a CoO. There are documents where the issuer and holder roles can be played by the same entity, for instance, an exporter that presents a commercial invoice to obtain export permits.
- **Verifier** (*recipientTradeParty*⁵⁰ as per UN/CEFACT), an entity that receives and verifies a document. The verification process implies checking document authenticity, consistency, and validity. For instance, a customs system that receives a CoO needs to verify it to apply reduced tariffs.

⁴⁷ <https://www.legislation.gov.uk/ukpga/2023/38/contents>

⁴⁸ <https://vocabulary.uncenfact.org/issuerParty>

⁴⁹ <https://vocabulary.uncenfact.org/senderTradeParty>

⁵⁰ <https://vocabulary.uncenfact.org/recipientTradeParty>

Description and Functionalities

TWIN Document Management services allow the storage and retrieval, on a need basis, of **multi-versioned, multi-format** *electronic* documents, ensuring *data availability, data sovereignty, data auditability, and tamper-proofness*. The Document Management services are ready to deal with both PDF-like, scanned, plain documents, and native eDocuments.

Even though the back-office processes that concern the issuance of documents are out of the scope of TWIN, a TWIN Node helps to transition from scanned, generated documents to native eDocuments. In fact, through a TWIN Node, a Participant, playing the role of document issuer, can announce (through an Event emitted by the corresponding IT system) the issuance of a key document (materialized, for instance, as a PDF) concerning a Trade Item. Once known by a TWIN Node, this new document can be bound to an Auditable Item; later, to facilitate verification by Participants fetching the document, a **data integrity proof** can be created on behalf of the document issuer. As a result, plain documents are turned into eDocuments, enabling the automation of document verification through digital means.

[Figure 9](#) shows a context diagram of the Document Management services. On top, there are different functional blocks related to the main functionalities offered (described above), all of them exposed as REST endpoints. On the middle layer is the Auditable Item Graph service: TWIN manages documents internally as Auditable Items, enabling seamless traceability and auditability. At the lower level are the different infrastructure services (including DLT) that constitute the fundamental substrate.

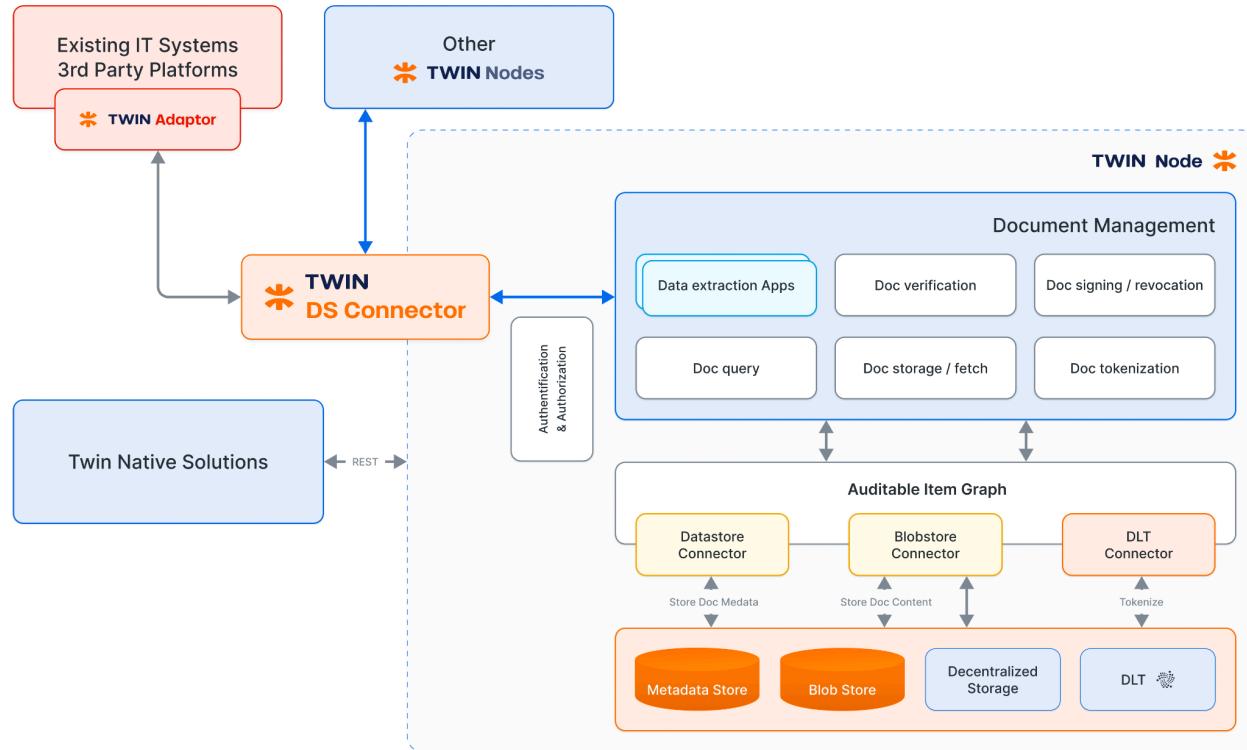


Figure 9 Document Management services context diagram and functional overview

There are two options for **storing a document's content**:

- The document's content is stored *at the source* (usually the document owner's external IT system). The Document Management services at a TWIN Node only keep metadata properties (see below) and a digital fingerprint that prevents tampering.

Typically, an external IT system storing document content will be proxied through a TWIN Adaptor. Such an Adaptor will be incarnated by a TWIN Node's **Data Service** published to the *TWIN Catalogue* (together with its describing Dataset), enabling other Participants to fetch (and verify) documents on a need-to-know basis (provided offering policies are matched and contract negotiation succeeds).

- The document's content and metadata are stored at the TWIN Node, typically when the Participant lacks storage capabilities. In this case, the document can be stored in a blob store managed by a TWIN Node or in a decentralized storage system (typically [IPFS](#)). If stored in a decentralized system, it may be encrypted to ensure that only authorized Participants can access it, following predefined sharing policies that govern the distribution of encryption keys.

A TWIN Node can store various metadata properties related to eDocuments (**document metadata properties**) through the annotations of an Auditable Item bound to the document. This flexible approach follows Linked Data models, enabling the use of vocabularies such as UN/CEFACT or schema.org. [Figure 10](#) provides an illustrative example.

Name	Tag	Description
Details	Document ID	urn:docId:337:b0ff65af9768c9a24b9579c953c8a856b5f3e197b265af3e662f0c572e23f923
	Document Type	unece:DocumentCodeList#851
	Revision Number	2
	Item Referred	https://id.gs1.org/01/09521987654327/21/202301
Content	Content size	5671
	File Format	application/xml
	Content Locator	urn:fs-blob:7df6bdb4cb31ad118c9dfb3053f8cd671330e9538bfe03027c1cd5ec7c268d9d
Timestamps	Issued at	2024-09-12T 15:05:03.321Z
	Created at	2024-09-12T 15:05:03.321Z
Party	Issuer party	did:iota:0xb62afcd0150d048ea0679af61d28d0eb1ad1b969f411b03997194df232b27383
	Sender party	did:iota:0xd7258832d2c578426f2d33ce320cf485b3707ae99b90fdcf25ca5b60b30381c3
Proof	digestSRI	sha384-c6b6f54aecdec6f37de08ae4a2375eceb1d69115c28c2738c915587362714d8c

Figure 10 Document metadata example

Although TWIN is not limited to a specific set of metadata properties, the following are the most essential document metadata properties:

- Document identifier (URI).
- Document type code (as per the UN/CEFACT document code list⁵¹).
- Content MIME type (XML, PDF, JSON, etc.).
- Content length (number of bytes).
- Content locator, a URI that points to a blob that contains the document's content.
- Issuer Identity (DID).
- Revision number (multiple document revisions can be stored and fetched).
- Digital Fingerprint (a digest to prevent tampering when the document remains at the source).
- Timestamps:
 - Creation (date and time when the document was registered on a TWIN Node).
 - Revision (date and time when a document revision was registered on a TWIN Node).
 - Issuance (date and time when the document was issued).

Other metadata properties:

- The Trade Item(s) referred **to** by this document (URI). However, there can be documents that, initially, might not refer to a specific Trade Item – for example, the result of a veterinary inspection performed on poultry. Later, it might be necessary to associate the original veterinary inspection with a food shipment that contains the poultry's meat.
- Document Description: a textual description of the document.
- Revised Document: the **URI** of the previous document version that this document updates.
- Holder Identity (DID).
- Data integrity Proof (can be pre-created or created on demand when a document is fetched).

⁵¹ <https://vocabulary.uncefact.org/DocumentCodeList>

- Validity Period.
- Document Status, including revocation status (a link to a location where a verifier can check to see if a document has been revoked).
- Other properties (see UN/CEFACT *Document*⁵² class and schema.org *DigitalDocument*⁵³ class).

A TWIN Node provides a Data Extraction service off-the-shelf but is also ready to host custom **Data Extraction applications** (actually [TWIN Data App](#) instances) that cater to the needs of two different scenarios:

- **eDocuments:** TWIN Document Management services are agnostic to specific document standards (ePhyto, elInvoice, etc.). As a result, a data extraction TWIN Data App may be required to process the information contained within an eDocument automatically.
- **Plain documents:** In this case, data extraction can be performed on PDF-like documents using OCR or other AI-driven recognition systems. The goal is to obtain a canonical, machine-readable version of a document, which can be transformed into an eDocument if required, using the [TWIN Data App](#).

In both scenarios, the final purpose of data extraction is:

- To capture relevant document metadata (such as validity period, issuance timestamp, or signatures) so that they can be properly processed by business processes.
- To capture additional data about associated trade item(s) and their journey, while also checking for consistency to detect fraud or suspicious activities.

Document fetching is a simple query operation, subject to data exchange policies and exposed by a TWIN Node via a REST API. To retrieve a document, it is necessary to know either its identifier or type code⁵⁴. Alternatively, the document's identifier can be discovered through its referred Auditable Item(s) annotation objects. The result of a document fetching operation includes the document's metadata along with the latest revision of its content. The content can either be transmitted directly to the client or provided as a reference to its location in decentralized storage. There can be multiple copies of a document if the data exchange policies allow different nodes to retain documents.

⁵² <https://vocabulary.uncerfact.org/Document>

⁵³ <https://schema.org/DigitalDocument>

⁵⁴ <https://vocabulary.uncerfact.org/documentTypeCode>

Document/Revision addition is a simple operation exposed through a REST API. The client needs to provide the document's

- Metadata.
- Fingerprint or signature.
- Content in case the document needs to be stored in a TWIN Node.
A TWIN Node is responsible for ensuring consistency within metadata properties across revisions. Other consistency checks can be performed by upstream, application-specific services.

Document signing can be subject to the following scenarios:

- **Pre-Signed Document (Verified Signature)**
If a document (whether a plain PDF-like file or an eDocument) is already signed by the issuer, whose signature can be verified by a TWIN Node or a TWIN Data App, no further action is required. This is the simplest scenario, as the TWIN Node only verifies the existing signature.
- **Unsigned or Unverified Signature**
If a document lacks a signature or has one that cannot be verified by a TWIN Node, the Node can sign it on behalf of the Participant, who can either be the issuer (higher trust level) or the holder (lower trust level). If the TWIN Node does not directly manage signing keys, it must interact with a Wallet that holds the relevant keys to complete the signing process.

A document's signature can be stored as part of its metadata properties, following the W3C Data Integrity Proof Recommendation.⁵⁵ Alternatively, the signature can be calculated on demand when another Participant requests the document.

Document verification occurs when TWIN Nodes or external systems, via TWIN Adaptors, exchange documents. This process ensures that the disclosed fingerprint matches the original or that the original signature can be verified.

Document timestamping and auditing follow the same mechanisms used for the [Auditable Item Graph](#), as documents within a TWIN Node are internally represented by an Auditable Item.

Document tokenization involves binding a trade document to a Non-Fungible Token (NFT) on-chain via a DLT Connector. While the document's content and metadata remain on a TWIN Node, its immutable entry on a distributed ledger enables [electronic record transfers](#), following the recommendations of the [MLETR](#) legislation. This capability is envisaged for eBills of Lading,

⁵⁵ <https://www.w3.org/TR/vc-data-integrity/>

eInvoices, digital promissory notes (DPNs), and other key trade documents, bringing life to [Trade Finance](#) ecosystems. This feature is being developed, tested, and experimented with at the time of writing. The results and vision are planned to be reported in future whitepapers, including a comparison and interoperability study with similar initiatives⁵⁶.

Readers may wonder about the role and relevance of W3C Verifiable Credentials (VCs) standards in this context [\[UN/CEFACT-VC-Trade\]](#). Simply put, VCs are also eDocuments that follow a particular W3C-defined data model with clear semantics. In addition, like other eDocuments or TWIN-signed plain documents, they can be cryptographically verified for authenticity and integrity using data integrity proofs. The bottom line, when a Participant adds a new document represented as a VC, a TWIN Node can act as a VC verifier, extracting data from such a VC. Alternatively, a TWIN Node can also be a VC issuer, on behalf of a Participant, when transforming a regular document into a VC as required by other services. In this regard, it is noteworthy that the [UNVTD⁵⁷](#) (*UN Verifiable Trade Documents*) initiative, a new project promoted by UN/CEFACT devoted to providing guidance, samples, and test services for the implementation of any trade document as verifiable credentials. TWIN plans to implement the resulting specifications once they are deemed as stable.

Data Exchange Services

Definitions useful to understand this section can be found in the [glossary section for Data Spaces](#).

Functional Overview

Data Exchange services are enablers for trusted data sharing among Participants and imply:

- **Publication to the TWIN Catalogue** of Dataset offerings compliant with the Catalog Protocol⁵⁸ so that they can be subject to Contract Negotiations (see below).
- **Vocabularies** used by TWIN to describe Participants, such as the Gaia-X Ontology⁵⁹, and W3C DCATv3 for Dataset offerings, conforming to the profile⁶⁰ defined by the

⁵⁶ <https://www.tradetrust.io/developer/technical-guides/>

⁵⁷ <https://un.opensource.unicc.org/unece/uncefact/spec-unvtd/>

⁵⁸ <https://eclipse-dataspace-protocol-base.github.io/DataspaceProtocol/2025-1/#catalog-protocol>

⁵⁹ <https://w3id.org/gaia-x/development>

⁶⁰ <https://eclipse-dataspace-protocol-base.github.io/DataspaceProtocol/2025-1/#catalog-protocol>

Eclipse DSP. To describe access rights, TWIN uses the W3C ODRL policy Vocabulary⁶¹, expressed as Linked Data.

- **Rights management**, so that Participants can unambiguously declare rules, i.e. policies that establish which data is shared with whom (*permissions*) and under which terms and conditions (*data usage policies*). Rights Management also concerns **Contract Negotiations**, which are a set of interactions between a Data Provider and a Data Consumer that establish a data sharing *Agreement* policy.
- **Authentication and authorization services** so that Participants can authenticate with a TWIN Node during data exchange interactions.
- **Data Exchange protocol**, which in compliance with the Eclipse DSP, defines a control plane and a data plane. At the control plane, a **Transfer Process** encompasses the set of interactions between a Provider and Consumer intended to give access to a Dataset under the terms of an Agreement. At the data plane, data is exchanged through HTTP via JSON-LD.
- **Transaction logging** (optional) to provide an auditable framework for data exchange transaction observability, so that there is a digital ledger that has the last word in case of dispute. As a TWIN Node always interacts with a DLT Node (by default, an IOTA Node), this is a differential functionality in our roadmap.

To clarify the concepts defined above, [Figure 11](#) shows a functional overview of the architecture that describes in detail how Data Exchange Services play together. For a data exchange to happen the following preconditions must be met:

- Participants are registered on the TWIN Catalogue.
- An ODRL policy expresses an “Agreement⁶²” to share data belonging to an offered Dataset. Such an agreement would typically be reached through a Contract Negotiation performed through the Eclipse DSP.

It is usually also necessary to setup a transfer process⁶³ via the Control Plane, so that the provider and consumer agree on the final data exchange endpoints and authentication details.

Once the above preconditions are met, the data exchange occurs through the Data Plane. As such, a Data Plane is simply the means to obtain data; it is usually necessary to delegate to the

⁶¹ <https://www.w3.org/TR/odrl-vocab/>

⁶² <https://www.w3.org/TR/odrl-model/#policy-agreement>

⁶³ <https://eclipse-dataspace-protocol-base.github.io/DataspaceProtocol/2025-1/#transfer-protocol>

Visibility and Document Management services when handling queries and persistence. On the other hand, TWIN DS Connector Apps implement custom functionality that might depend on the particularities of each ecosystem.

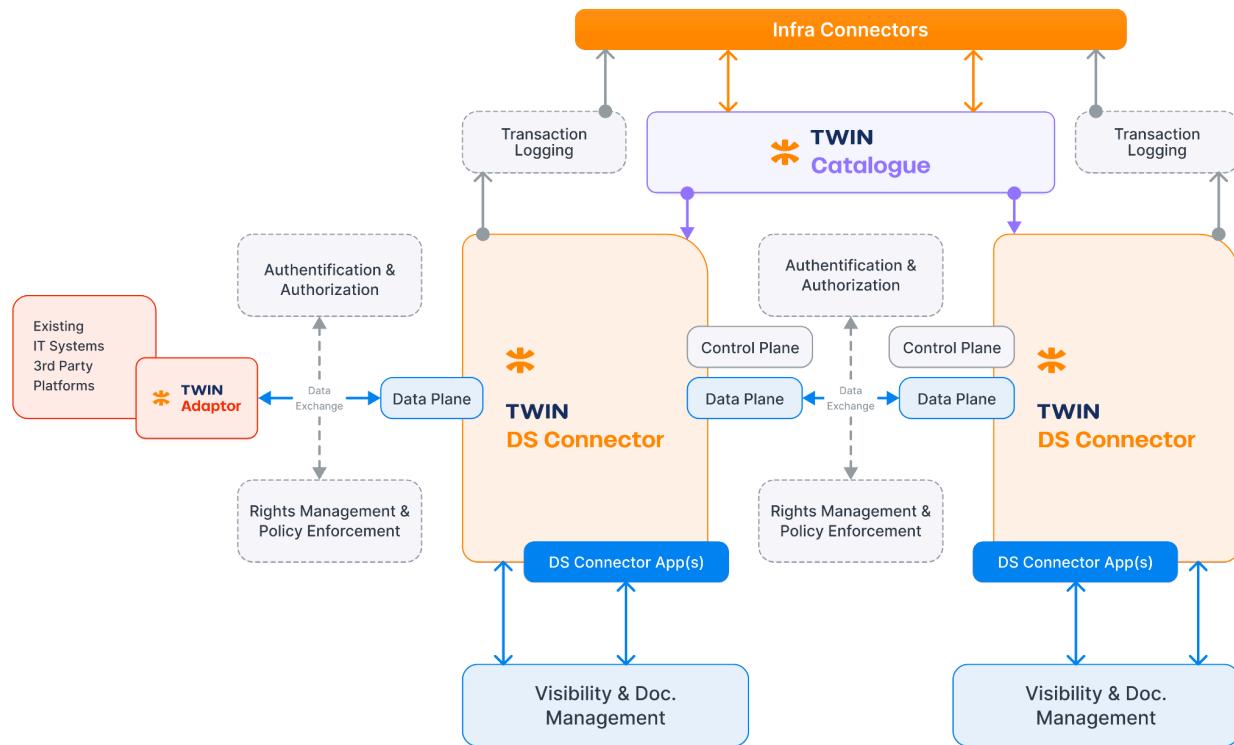


Figure 11 Data Exchange Services architecture overview

Publication and Discovery

The **TWIN Catalogue** is a decentralized service that stores and exposes descriptive metadata about compliant Participants and the Datasets they offer. It is a fundamental component that enables discovery during data exchange transactions.

For this purpose, Dataset offerings play a fundamental role as they announce in a machine-readable way the willingness of Participants to share data. The metadata of Dataset offerings must be declared by their publishers (a compliant Participant), possibly using self-issued Verifiable Credentials. As with Participants, each Dataset offering may undergo a **compliance** check by a [TWIN Clearing House](#) (depending on each ecosystem's rules). If a

Dataset offering meets the relevant ecosystem rules, a *Compliance Credential*⁶⁴ can be issued. This Credential can then be presented to the TWIN Catalogue to complete the onboarding process.

A *Dataset offering* contains metadata expressed using the W3C DCATv3Vocabulary that describes it.

Additionally, each TWIN Node will usually register at least one Data Service⁶⁵ that will implement the endpoints needed to conduct contract negotiations and data transfer processes.

An example of the metadata needed to describe a **Dataset offering** is provided below:

Name	Tag	Description / Example
Details	DatasetID	https://my-offerings.example.org/dataset1
	Type	https://vocabulary.uncefact.org/Consignment
	Title	Dataset offering example
	Publisher	did:iota:0xb62afc0150d048ea0679af61d28d0eb1ad1b969f411b03997194df232b27383
Distribution	Media Type	application/ld+json
	License	URL: https://tcs.example.org/1234 digestSRI: sha384-56A123
	Access Service	https://my-twin-node.example.org/dataspace-connector
	Format	HTTP Pull
Policy	Type: Offer Permission: Read	See policy description below

⁶⁴

https://docs.gaia-x.eu/technical-committee/identity-credential-access-management/24.07/credential_format/#gaia-x-compliance-inputoutput

⁶⁵ https://www.w3.org/TR/vocab-dcat-3/#Class:Data_Service

Name	Tag	Description / Example
	Constraint: <i>destination</i> <i>Country is UK</i>	

Figure 12 Dataset offering expressed using the W3C DCAT Vocabulary

Rights Management

Rights management encompasses:

- The management and run-time enforcement of **policies**, i.e. **rules**, that represent permitted and prohibited actions (uses and re-uses of data) over a certain asset, as well as the obligations required to be met by Providers and Consumers. In addition, policies may be limited by constraints assigned by the Data Provider (e.g., data, temporal or spatial constraints) and duties (e.g. payments) may be imposed on permissions. Policies are represented using **W3C ODRL**.
- **Contract Negotiation**⁶⁶ (CN), which are the interactions between a Provider and a Consumer that establish an *Agreement* policy from an initial *Offer* policy. An *Offer* policy represents rules that convey what is being offered from Data Providers. An *Agreement* policy represents rules that have been finally granted by a Data Provider (assigner) to a Data Consumer (assignee). It is important to note that the CN does not automatically lead to an immediate data or algorithm transfer. The result of a CN is a data contract agreement, which then can be executed at a later point in time.

ODRL Policies

The Open Digital Rights Language (ODRL) is a **policy expression language** that provides a flexible and interoperable information model, vocabulary, and encoding mechanisms for representing statements about the usage of content and services [[W3C-ODRL](#)]. In the case of TWIN, it concerns:

⁶⁶

https://docs.internationaldataspaces.org/ids-knowledgebase/idsa-rulebook/idsa-rulebook/3.-functional_requirements/3.8-data_sharing#contract-negotiation

- The informational resources, i.e., assets intended to be shared (documents, a trade item's digital twin, Events, etc.). These are usually expressed as constraints over a Dataset offering, such as “all items of type ‘Consignment’ whose destination country is the United Kingdom (UK)”.
- The Consumers (Participants, Services), which can get access to the subject resources. These can be enumerated by ID (DID, URI, etc.) or through constraints that the Consumer's Attributes must meet. For instance, “all Participants whose headquarters are in Kenya and whose role corresponds to a ‘government border agency’”.
- Certain environmental conditions (*optional*) concerning the operational, technical, or situational environment in which the information access occurs, for example, the time of day at which the data exchange can take place.
- The terms and conditions (*optional*) – for instance, if the data or documents shared can be retained or archived by a Consumer.

These rules are associated with an action (*read*, *modify*, *annotate*, etc., as per ODRL) that is allowed or prohibited. In data exchange processes, the usual action will be “read”. An example of a Policy can be found below ([Figure 13](#)).

Name	Tag	Description
Details	Policy ID	https://my-policies.example.org/policy1
	Type	Agreement
	Profile	https://twindev.org/odrl:profile:01
	Description	Agree to share consignment details whose destination country is the UK with the UK Border Force Participant
Who	Assigner	did:iota:0xb62afcd0150d048ea0679af61d28d0eb1ad1b969f411b03997194df232b27383
	Assignee	did:iota:0x3edd589d510d58c211237c79755314d81c228b8bd95559f9ad35911736b3a4aa
What	Target Asset	https://my-offerings.example.org/dataset1

Name	Tag	Description
	Action	read

Figure 13 Policy structure example

Contract Negotiation Protocol

TWIN adheres to the Eclipse DSP Contract Negotiation protocol. Under this protocol, contract negotiations (CN) are typically initiated by a Data Consumer that discovers an *Offer* (expressed in ODRL) published by a Data Provider through the TWIN Catalogue. The negotiation process is guided by a state machine, with state transitions triggered by the receipt of well-defined messages (via REST endpoints). During the negotiation, the *Offer* Policies are evaluated against the attributes of the requesting participant. If successful, the negotiation concludes with an *Agreement*, captured in ODRL and acknowledged by both parties. This Agreement is subsequently referenced to initiate a data transfer process over the Control Plane of the Data Space Connector.

Policy Enforcement Implementation

Regarding the services in charge of managing and enforcing Policies, TWIN adheres to IDSA⁶⁷, Gaia-X, NIST, and OASIS recommendations (see [Figure 14](#) below for an illustration).

⁶⁷

https://docs.internationaldataspaces.org/ids-knowledgebase/ids-ram-4/layers-of-the-reference-architecture-model/3-layers-of-the-reference-architecture-model/3_4_process_layer/3_4_6_policy_enforcement

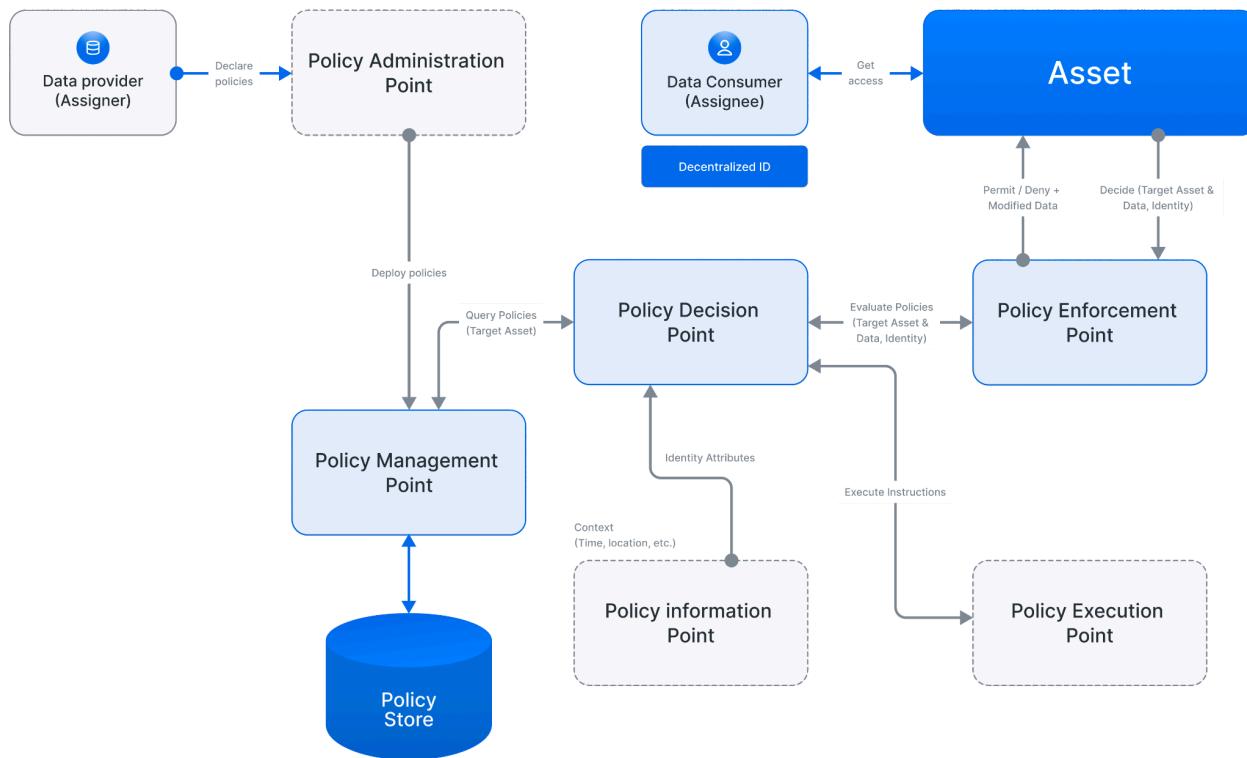


Figure 14 Functional view of the policy enforcement architecture in TWIN

Two main services are involved in policy enforcement in a decision process illustrated by [Figure 15](#):

- The **Policy Enforcement Point (PEP)** is the entry point for enforcement, where data or metadata is stopped and transferred to the Policy Decision Point (PDP). The PDP makes a decision and returns it to the PEP. Then the PEP will subsequently manipulate or lock the data according to the decision.

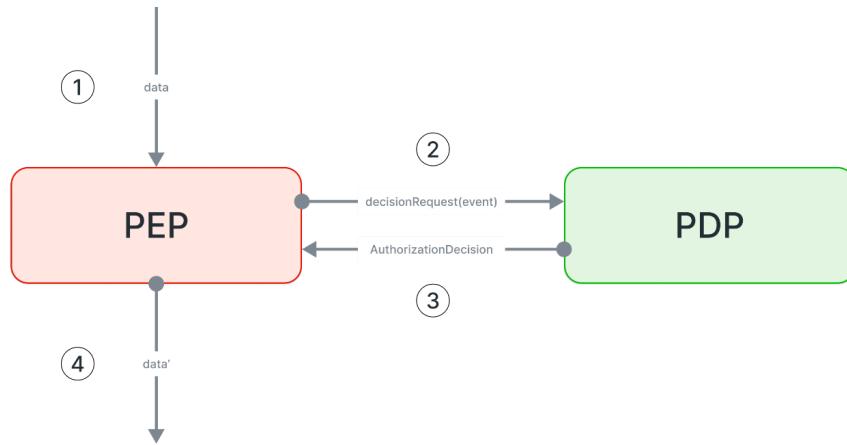


Figure 15 The decision process for policy enforcement. Source IDSA

- The **Policy Decision Point (PDP)** makes a decision based on:
 - The target asset, the data requested, and the identity of the Data Consumer.
 - The registered policies, obtained through the Policy Management Point (PMP), that specify permissions, prohibitions, and obligations, and their constraints.
 - Other optional information offered by a Policy Information Point (PIP) including:
 - The attributes of the requesting identity (Data Consumer).
 - Context conditions, for instance, time or location.

Other services involved are:

- The **Policy Execution Point (PXP)** is the component for implementing instructions or requirements. These can be before a decision, and their successful execution can be included as a condition, or they can be executed after a decision has been made, for instance, for logging purposes.
- The **Policy Information Point (PIP)** supplies Identity Attributes or context conditions necessary for the PDP to make decisions. This role can be fulfilled by various entities, including the **TWIN Catalogue**. Additional attributes may be disclosed through the **SD-JWT** [[IETF-SD-JWT-2025](#)][IETF-SD-JWT-2025](#) used for authentication.

There are additional components that are not directly needed for enforcement but are important for the specification and management of usage policies:

- **Policy Management Point (PMP)** The PMP, as the name implies, is responsible for the management or handling of the policies. It makes the policies available to the PDP, and activates, deactivates, and deletes them.

- **Policy Administration Point (PAP).** The PAP is used to support the creation and specification of usage policies, often via a user-friendly graphical interface.

The process flow of policy enforcement, involving all described components, is depicted in [Figure 16](#).

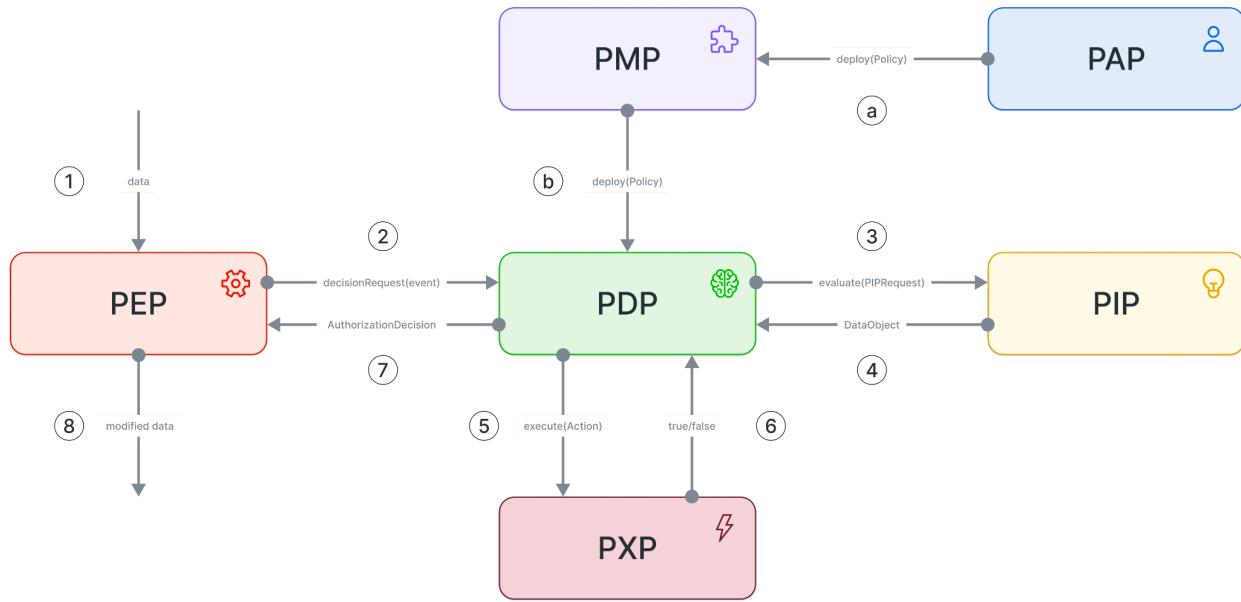


Figure 16 Process flow of policy enforcement. Source IDSA.

Authentication and Authorization

To perform data exchange processes, TWIN data exchange services must first determine who wants to perform a certain action (for instance, who wants to fetch a document), i.e., *authentication*. Authentication can be achieved using an (SD-)JWT token, which the requesting party must present to prove, at a minimum, control over a specific identity (DID). Once the Identity is verified through a verifiable registry (such as the IOTA DLT), the Policy Decision process can take place as depicted in the process flow described in [Figure 16](#).

At the time of writing, the TWIN community is evaluating the Eclipse DCP [[Eclipse-DCP](#)] specification, which specifies interoperable mechanisms for communicating participant identities and credentials to secure data access. The specification defines a set of protocols for asserting participant identities and issuing and presenting verifiable credentials using a decentralized architecture for verification and trust. However there is still no consensus whether DCP represents a preferable alternative to extending existing OpenID flows, such as SIOPv2, OpenID4VP and OpenID4VCI.

TWIN Data Exchange Protocol

The TWIN Data Exchange Protocol extends the Eclipse DSP through a TWIN Data Space Protocol profile, defining minimal extensions to support custom ODRL constraints and Data Transfer formats that comply with TWIN's wire data protocol. It is designed to handle both push and pull data transfer modes.

The protocol is implemented via the **TWIN Data Space Connector**, which provides two primary interfaces: the Control Plane and the Data Plane.

Control Plane

The Control Plane is based on the Eclipse DSP (Transfer Process Protocol) and extends it with two data transfer formats:

- **HTTP Pull Query:** allows Consumers to pull data via a REST endpoint that accepts a query and renders matching entities as JSON-LD.
- **HTTP Activity Push:** Allows Consumers to receive Dataset entities as an Activity Stream in push mode.

Data Plane

[Figure 17](#) depicts an overview of this key component of the TWIN architecture and its relationship with [TWIN Adaptors](#), as these two components must collaborate closely.

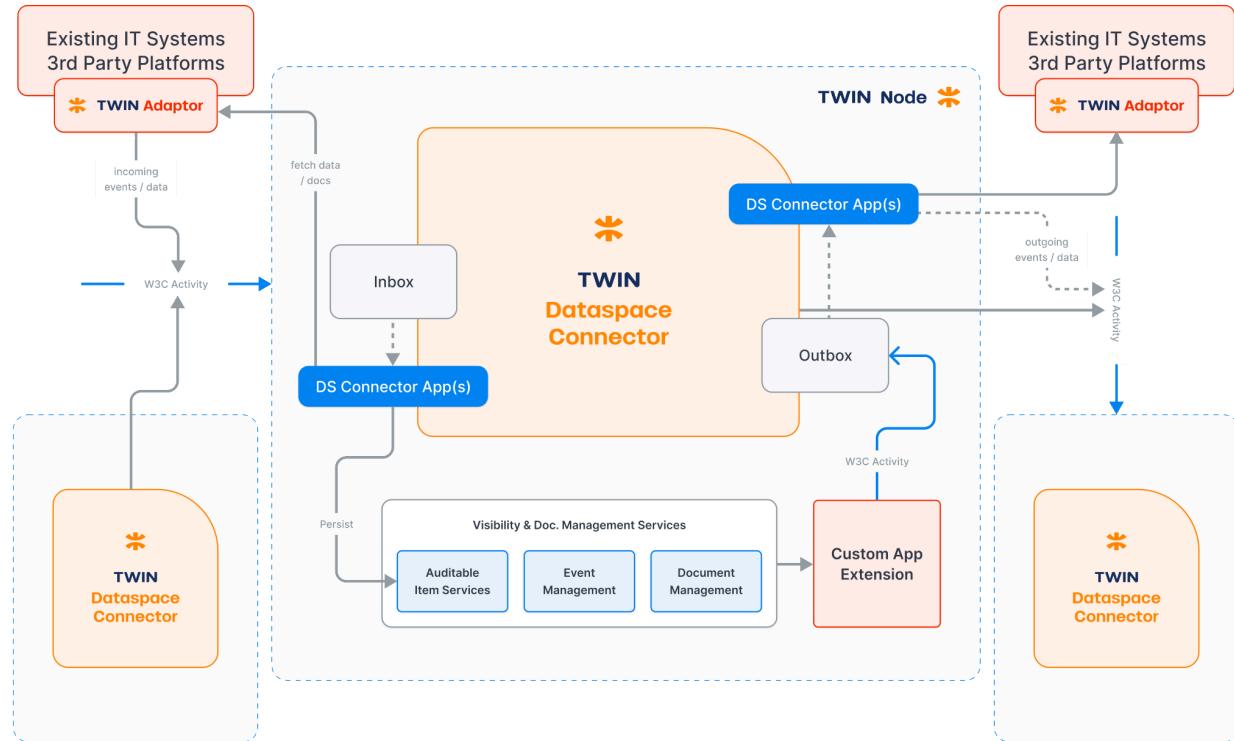


Figure 17 Overview of the Data Plane of the TWIN Data Space Connector.

The **Data Plane** facilitates collaboration between the TWIN DS Connector and TWIN Adaptors. It reuses three W3C Social Web standards originally designed for decentralized social networks, adapted to value chain data sharing:

- **W3C Activity Streams** [\[W3C-Activity-Streams\]](https://www.w3.org/TR/activitystreams-vocabulary/#dfn-activity). Defines the data model and JSON-LD representation for expressing “activities of interest” within value chain ecosystems. Each **Activity**⁶⁸ includes the generator (rapporteur), the actor (typically a Participant), the action type⁶⁹ (e.g., “Add”, “Request”, “Create”), and the associated **object** and **target** (e.g., a *Document*⁷⁰ linked to a *Consignment*). Objects can also be **Events**.
- **W3C ActivityPub** [\[W3C-ActivityPub\]](https://www.w3.org/TR/activitystreams-vocabulary/#activity-types): Provides extended property definitions and formalizes the concepts of “activity inbox” (all received Activities) and “activity outbox” (all generated Activities).

⁶⁸ <https://www.w3.org/TR/activitystreams-vocabulary/#dfn-activity>

⁶⁹ <https://www.w3.org/TR/activitystreams-vocabulary/#activity-types>

⁷⁰ <https://vocabulary.uncif.org/Document>

- **W3C Linked Data Notifications** [[W3C-LDN](#)]: Defines the Web endpoints to perform activity notification among Participants. Activities are delivered through an **inbox**⁷¹ endpoint that follows W3C-LDN.

TWIN Data Space Connector Apps

As described throughout this document, the fundamental TWIN services – **Auditable Item Services** and **Document Management** – expose APIs that facilitate the full lifecycle management of trade items, business events, data streams, and documents. In addition, a TWIN Node exposes a data plane via the TWIN DS Connector as a specialized service designed for data exchange within value chain ecosystems.

The **TWIN DS Connector's data plane** can be extended to include ecosystem-specific processing logic or data formats. This is because different ecosystems may require unique formats or custom reactions to activity streams propagated through a Data Space Connector. To support this, **TWIN DS Connector Apps** (aligned with the International Data Spaces architecture) allow integration and deployment of complementary features on top of the generic data exchange processes, such as data processing, data format alignment, and data exchange protocols.

A TWIN DS Connector App can implement custom processing of incoming or outgoing Activities (those received through the Node's inbox or Node's outbox) and custom handling of queries received through the query interface, (as described in [point B](#)). The latter is a viable option to implement a simple TWIN Adaptor that can expose to the TWIN Data Space documents kept at the source by an external system (see below).

Data Plane Functions

The TWIN DS Connector's data plane performs the following key functions:

- **Receiving and validating “activities of interest”** (playing the role of a data sink⁷²). The incoming Activity stream (**inbox**) is an ordered collection that is generated by other Data Space Connectors or TWIN Adaptors, which play the role of a data source. Examples of such activities include the creation of a consignment, the addition of a document, the finalization of an inspection, and so on.
- **Processing the incoming Activity Stream (**inbox**)**: may involve the collaboration of TWIN DS Connector Apps to:

⁷¹ <https://www.w3.org/TR/ldn/#receiver>

⁷² [https://en.wikipedia.org/wiki/Sink_\(computing\)](https://en.wikipedia.org/wiki/Sink_(computing))

- Create new objects (e.g., Auditable Items or Auditable Item Streams) in the target TWIN Node.
- Update existing Auditable Items in the target TWIN Node.
- Retrieving the referenced Activity's objects via the Data Space Connector of the Activity's actor, if permitted by policy (e.g., for verification purposes).
- Deriving additional "activities of interest" for the outgoing Activity Stream (outbox), typically for external delivery (see below).
- **Processing the outgoing Activity Stream (outbox)**, playing the role of "source Data Space Connector". It might involve the collaboration of specific TWIN DS Connector Apps. The processing of an outgoing Activity might result in:
 - Delivering outbox activities as-is or transformed to other TWIN Nodes (data sinks) via their corresponding target Data Space Connectors.
 - Transforming outbox activities into notifications to be delivered to external IT systems (via a TWIN Adaptor), TWIN Solutions, or other relevant entities.
- **Handling queries over Datasets** involved in pull and push data transfers, in collaboration with TWIN DS Connector Apps.

Data Plane interfaces

A TWIN DS Connector's data plane offers:

- A. **REST endpoint for receiving Activity objects.** This REST endpoint is compliant with the W3C Linked Data Notifications and consumes push data transfers.
- B. **REST endpoint that exposes an interface for querying Datasets managed by a TWIN Node**, such as those containing trade items, documents, or Events. It handles pull data transfers. The data retrieved is generally represented using JSON-LD. This allows maximum flexibility and opens the door to derivatives such as NGSI-LD [\[ETSI-NGSI-LD\]](#), which allow an interoperable representation of context information as property graphs.
- C. An interface enabling TWIN DS Connector Apps to register as handlers of certain Activities (incoming or outgoing) and queries (pull or push) over Dataset offerings .
- D. An interface enabling outgoing Activities to be put in the *outbox* for later processing and delivery. This interface is also used by TWIN DS Connector Apps to handle push data transfers.

Push Data Transfers

For **push transfers**, data is delivered through the Consumer's **inbox Activity Stream** following a "subscription" via an Activity of type "Follow"⁷³ sent to the Provider's inbox. The object of this Activity is a Data Transfer previously established at the **Control Plane**. Optional query conditions can also be specified. To stop receiving data, a Consumer can send an Activity of type "Undo"⁷⁴ to the Provider's Activity Stream inbox, eliminating the need for separate subscription payloads.

As already described, push notification payloads are represented as **JSON-LD Activities** and delivered via a webhook (HTTP POST) to the recipient's inbox in accordance with the W3C Linked Data Notifications specification. A TWIN DS Connector can therefore receive Activities generated by another TWIN DS Connector.

The delivery of Activity notifications requires authentication by the corresponding TWIN Node, referencing the ongoing data transfer process that uniquely identifies the intended Data Consumer. The JSON-LD Activity is typically received at the Consumer's TWIN Node inbox. However, the Data Transfer Process Protocol is flexible enough to allow the Data Consumer to specify an alternative endpoint for a specific data exchange, which may be external to a TWIN Node.

Provider-initiated data transfers

Following the **W3C ActivityPub** specification, all Participants in a TWIN Ecosystem can publish their "inbox" endpoints⁷⁵ in the TWIN Catalogue. This enables Provider-initiated data exchanges interactions, which are useful in value chain ecosystems where proactive communication is required.

In this setup a Data Provider can send an Activity to a Data Consumer, typically under a pre-negotiated, out of band generic data sharing agreement.

For example, in **Digital Product Passport (DPP)** ecosystems, actors such as repairers – who may not be known in advance – can communicate actions taken on a product, such as an EV battery, to the product owner or manufacturer, ensuring the product's DPP is updated in real time.

⁷³ <https://www.w3.org/TR/activitystreams-vocabulary/#dfn-follow>

⁷⁴ <https://www.w3.org/TR/activitystreams-vocabulary/#dfn-undo>

⁷⁵ <https://www.w3.org/TR/activitypub/#actor-objects>

TWIN Adaptor

A TWIN Adaptor is a technical component that acts as a bidirectional **bridge** between a TWIN Node and an external system.

This bridge propagates “activities of interest” to a Node’s TWIN DS Connector (i.e., the Node’s inbox) when specific business workflows occur in external systems, such as the issuance of documents or the creation of new items. Conversely, when activity occurs within a TWIN Node and is published to the Node’s outbox, a TWIN Adaptor may propagate it to the connected external system, provided the activity is relevant to that system. In addition, a TWIN Adaptor may delegate query interface operations of the data plane of a Data Space Connector – for example, to retrieve and supply trade documents by interfacing with the connected external system.

TWIN Adaptors are an important component of the TWIN Architecture. Within the international trade IT ecosystem, a range of software packages – such as Asycuda World⁷⁶, VUCE⁷⁷, or IPPC ePhyto⁷⁸ – are commonly used by multiple countries, with the caveat that they are mostly siloed data hubs. The development of a TWIN Adaptor for such widely adopted systems enables, by building on existing systems rather than replacing them, a real data pipeline for the sharing of documents and vital customs information across countries and economic zones, realizing the ultimate TWIN vision.

A related example is the EU eFTI⁷⁹ *Regulation*, which requires businesses to exchange freight transportation information with authorities through eFTI gate systems. A TWIN Adaptor can enable eFTI gates to query freight transportation data associated with trade items, allowing TWIN Nodes to also play the role of an eFTI Platform.

There are several architectural options for implementing a TWIN Adaptor:

- An “external TWIN Adaptor” implemented as a standalone software component and executed entirely outside of a TWIN Node. This approach offers maximum flexibility but may require additional development and maintenance. It may also be the only viable option when using a Node-as-a-Service offering.
- An “in-node TWIN Adaptor” implemented as a TWIN Data App or TWIN DS Connector App i.e. deployed as an extension within an existing TWIN Node. In this case, the

⁷⁶ <https://asycuda.org/en/>

⁷⁷ <https://aduananews.com/vuce-panorama-en-paises-mercosur/> (In Spanish)

⁷⁸ <https://www.ippc.int/en/ephyto/>

⁷⁹ https://transport.ec.europa.eu/transport-themes/logistics-and-multimodal-transport/efti-regulation_en

Adaptor fully operates within a TWIN Node and benefits from its off-the-shelf infrastructure services and Connectors. However, careful consideration must be given to resource consumption and scalability, which may require moving the TWIN Adaptor to a separate runtime service.

- A “hybrid TWIN Adaptor” where some functionalities are executed entirely outside a TWIN Node, while others are part of a TWIN DS Connector App, thus executing within a Node. For instance, periodically polling an external system or listening to middleware queues may be handled by an external service, while reacting to incoming or outgoing activities and handling queries can be performed in-node

Given their critical role in enabling interoperability and adoption, TWIN Adaptors are a priority within the TWIN technical roadmap. As such, the roadmap includes the release of a **TWIN Adaptor Blueprint**, providing an off-the-shelf starting point for adaptor development and promoting best practices, as well as sound architectural and design patterns.

Infrastructure Services

Definitions related to this section can be found in the [glossary of terms](#).

Infrastructure software services provide the essential substrate that enables the functionality of a TWIN Node. Essentially, they include:

- Datastores (SQL or NoSQL).
- Blobstores, i.e. large binary object stores (to store images, documents, etc.).
- Key Management Services (KMS) and Secret Management systems (like *Hashicorp Vault*⁸⁰) to store encrypted information, particularly keys, with a high degree of security.
- Distributed Ledger Technology (IOTA) as a verifiable registry with code execution (smart contracts) capability at Layer 1 (Move Virtual Machine⁸¹) and Layer 2 (Ethereum Virtual Machine⁸²).

TWIN is developing the following software components, which are key to keeping a level of agnosticism concerning infrastructure services:

⁸⁰ <https://www.vaultproject.io/>

⁸¹ <https://docs.iota.org/developer/iota-101/move-overview/>

⁸² <https://evm.iota.org/>

- **TWIN Datastore Connector:** A software component intended to store and query data/metadata concerning the entities managed by a TWIN Node (Auditable Items, Documents, Events, Participant Registrations, etc.).
- **TWIN Blobstore Connector:** A software component intended to store and fetch documents (trade documents, product certificates, etc.) managed by a TWIN Node.
- **Secret Management Connector or Vault Connector:** A software component intended to store secrets encrypted, manage their lifecycle (rotation, revocation, etc.), and control access to them. It is concerned with secrets needed by the different systems within TWIN, including those kept in custody, such as the keys of a Participant or service.
- **TWIN DLT Connector:** A software component intended to settle transactions (with or without smart contract intervention) or to read entries from a distributed ledger for data traceability, verifiability, or timestamping purposes.

TWIN offers a DLT Connector for IOTA, a Vault Connector for Hashicorp Vault, a blobstore connector for IPFS⁸³ and S3⁸⁴, and multiple flavours of datastore connectors, including DynamoDB⁸⁵ and ScyllaDB⁸⁶ for maximum data scalability.

On top of the services described above, the following advanced functionalities are also in scope:

- **Decentralized Entity Storage** enables the storage and sharing of a set of data entities across multiple TWIN Nodes. It is a key component for implementing the TWIN Catalogue and TWIN Registry or any other decentralized and verifiable registry used by ecosystems.

There are two infrastructure services involved in the implementation of decentralized entity storage in TWIN: a decentralized storage service (typically based on IPFS) and a DLT. The former stores snapshots and changes sets of entities so that they are available to any entity reader. The latter acts as a trust layer that facilitates:

- The coordination among different entity writers, ensuring there is no data lost.
- The initial discovery of the latest snapshot to bootstrap entity data so that a local copy can be created from scratch.
- The periodic discovery of the latest changes to be applied to a local copy.

⁸³ <https://ipfs.tech/>

⁸⁴ <https://aws.amazon.com/s3/>

⁸⁵ <https://aws.amazon.com/dynamodb/>

⁸⁶ <http://scylladb.com/>

- **Archive Storage** enables the storage of data entities and documents for archiving purposes after they have left an operational window. For example, data can be archived a few days after the physical goods have been delivered or when an item is decommissioned. As value chain ecosystems continuously generate new data, an archive storage service facilitates data management efficiency. The retention policy and access to archived content could depend on regulatory aspects for each ecosystem.

DLT and TWIN

TWIN fully exploits the four key characteristics of a public, permissionless DLT:

- **Transparency**, the append-only ledger is auditable by the whole network.
- **Immutability**, as data cannot be easily tampered with.
- **Traceability and nonrepudiation**, as each network participant cryptographically signs each transaction issued in the immutable ledger.
- **Decentralized execution** of immutable instructions, i.e., smart contracts.

In addition, the rise of decentralized applications (*dApps*) has created a need for standardized methods of representing information on DLTs. One of the most widely used approaches is *token* representation, where information recorded on a DLT represents a specific right, such as ownership of an asset, access to a service, receipt of payment, etc. Fungible tokens are commonly used for second-layer cryptocurrencies, including stablecoins, as they maintain uniform value and interchangeability. In contrast, Non-Fungible Tokens (NFTs) are utility tokens designed to represent and transact with unique tangible or intangible assets on DLTs. Unlike fungible tokens, each NFT is distinct and non-interchangeable.

The five key DLT features explained above are applied to TWIN through an IOTA DLT Connector⁸⁷ as follows:

- The Ledger plays the role of a Verifiable Registry, avoiding the need for trusted centralised actors, facilitating among others:
 - Identity (DID) registration, resolution, verification, and traceability.
 - Credential revocation lists can be stored on-chain as bitmap strings.
 - The coordination of decentralized entity storage related to the TWIN Catalogue and TWIN Registry functionalities.
 - Trust anchor registration on-chain, including traceability of their public Verifiable Credentials.

⁸⁷ Other DLT Connectors are also feasible and could be developed by the community.

- TWIN Clearing House functionalities (through smart contracts and NFTs).
- Objects managed by TWIN can be made auditable, improving Participant trust in traceability. Each critical event or relevant change could be immutably recorded and timestamped on the Ledger, ensuring transparency and verifiability.
- The Ledger facilitates dispute resolution, recording transactions within a TWIN ecosystem, such as data exchange transactions.
- Key trade documents that are transferable records (BoL, invoice, etc.) can be tokenized (attested) through an NFT on the Ledger. This ultimately enables the transfer of ownership, rights, and obligations and facilitates building Web3 dApps for Trade Finance facilitation.
- Luxury items tracked by TWIN can be tokenized as NFTs, facilitating product authentication, brand image, customer engagement, and so on.

In a nutshell, dApps built with TWIN leverage the verifiability of information stored on the distributed ledger and authentication mechanisms based purely on cryptographic primitives.

The operational models can vary:

- Each TWIN Participant may have their own ledger account owning the different assets and funding gas fees of each DLT transaction.
- TWIN Nodes can be delegated by Participants to have control and ownership over certain assets, including those that allow funding transactions.

TWIN also supports hybrid approaches where a TWIN Node can sponsor transactions using an IOTA Gas Station⁸⁸, while Participants remain the actual asset owners. Furthermore, some TWIN Nodes may also function as a DLT Validator, earning validation fees. The final setup may depend on business, subscription, or operational models, a topic outside the scope of this whitepaper.

Edge Devices and Connectors

Edge devices deal with automatic identification and data capture (AIDC). They play an important role: bridging the physical world of trade items with their corresponding digital twin represented through an Auditable Item. On one hand, as physical resources, Edge Devices are owned and operated by Participants and can form part of the underlying infrastructure. In cases where

⁸⁸ <https://www.iota.org/products/gas-station>

compliance tracking is required, they may even be onboarded through **Compliance Credentials** to verify their compliance status or device credentials.

On the other hand, through their **Edge Device Connector**, these devices can be characterized as services, offering events and data to a TWIN Node. For example, an RFID Reader can report that a set of trade items has been read at a particular location. The Edge Connector can listen to the Reader's low-level events (device events) and transform them into higher-level business Events, recording them on a TWIN Node. Afterwards, the corresponding Auditable Items will be updated, and new Events registered on an Event Repository, identifying the source as the Reader in question, etc. Additionally, DLT-based event notarization could be used to certify whether an item was physically in the possession of its claimed owner.

This architecture is not limited to RFID Readers but can be extended to **mobile sensors, scanners, printers**, and other **Edge Devices**, enabling seamless integration into the **TWIN ecosystem**.

To validate this architecture, initial experiments were conducted by the IOTA Foundation in collaboration with Zebra Technologies using the **FX9600 RFID Reader**⁸⁹. This device incorporates the Zebra IoT Connector⁹⁰ software that enables applications to be built that leverage the Reader's capabilities without having to develop specific software to be run on the device itself. The Zebra IoT Connector supports WebHooks (or MQTT) for data transmission. For instance, when a reading cycle begins and several RFID tags are detected, a JSON payload will be sent through a WebHook (HTTP POST) with the details of the tags read (EPC, TID, etc.). That is a Tag Data Event that could be transformed into another Event by the corresponding TWIN Edge Connector that reacts to the Webhook notifications. Finally, Events will be registered in a TWIN Node via an Edge Device Connector that will have to authenticate itself on behalf of the Edge Device or on behalf of the Device's owner.

A similar architecture can be put in place with other Edge Devices such as mobile sensors, scanners, printers, etc.

⁸⁹ <https://www.zebra.com/gb/en/products/rfid/rfid-readers/fx9600.html>

⁹⁰ <https://www.zebra.com/gb/en/software/rfid-software/iot-connector.html>

Common Platform Services

The following services offer cross-cutting functionalities to other services within a Node.

- **TWIN Engine**, which includes:
 - The ability to instantiate component packages and their runtime dependencies within a TWIN Node (including extended components that implement TWIN Data Apps).
 - The Web API enablement Engine, which facilitates exposing component functionalities as service instances that offer Web APIs (REST or WebSocket).
- **Telemetry**, which includes **logging** and **metrics** functions. Logging involves keeping a record of events that occur within TWIN software services, such as problems, errors, or information on current operations. Metrics relate to the quantitative measures used to evaluate the performance and efficiency of the various software components. A TWIN Node exposes Open Telemetry⁹¹ interfaces so that it can be easily monitored using Prometheus and Grafana.
- **Background Tasks**: A service that allows scheduling asynchronous tasks that have to be executed immediately or at a scheduled time.
- **Messaging**: A service that can deliver notifications through different messaging channels (SMS, email, etc).
- **Local User Management**: An internal service within a Node that allows the creation of different user or service accounts **internal** to a Node. Those accounts are not visible to other Nodes.

Local User Management and Policies

It is important to differentiate between local user management and Participant management through the TWIN Catalogue. A Participant registered on the Catalogue might interact with any TWIN Node by proving their Identity through a (SD-)JWT token generated through a Credential Wallet. Depending on the services offered through the Catalogue, their policies, and the Participant's attributes, the Node will either respond with data or documents or deny the request.

Additionally, to interact with a specific Node, a Participant may own additional user or service accounts added by a Node operator. For instance, a Participant may have deployed multiple

⁹¹ <https://opentelemetry.io/>

TWIN Native solutions, and each of these client services may have their own service account in their Node. Going further, there can be clients of a TWIN Node that are not even registered on the TWIN Catalogue because they are private Data Producers behind the façade of a Node. The bottom line, certain user or service accounts might not be visible to other Nodes, being purely local instead.

Concerning authentication, a TWIN Node could play the role of resource server within OpenID Connect and enable local authentication of clients together with an OpenID Provider (for instance, supplied by a third party such as cloud-based identity platforms or the Active Directory service). Concerning authorization, there might be local policies associated with each account that dictate local rules for the use of the Node services by the authorized Participants. Policies can be described and executed as already described for [Data Exchange Services](#).

Glossary

The following definitions have been adapted (and sometimes taken literally) from different sources cited at the end of this Whitepaper.

Technology-related

Identity and Credentials

- **Attribute Service Provider:** A type of [Trust Service Provider](#) with the role of collecting, creating, checking, or sharing pieces of information that describe something about a Participant. Attribute Service Providers can share their attributes with relying parties and Identity Service Providers, subject to the rules of obtaining the Participant's agreement being followed.
- **Certificate Authority:** The entity in a Public Key Infrastructure (PKI) that is responsible for issuing public-key certificates and enforcing compliance with a PKI policy. Also known as a Certification Authority.
- **Credential:** A set of one or more claims made by an issuer.
- **A Credential Dataset** defines the data (claims) about a subject that is to be included in a Credential.
- **Credential Format:** A Data Model used to create and represent Credential information. This format defines how various pieces of data within a Verifiable Credential are organized and encoded. TWIN supports both JWT and JSON-LD (using the W3C VC Data Model).

- **Credential Manager:** A synonym for “Credential Wallet”.
- **Credential Issuer** (or Issuer): An entity that issues Verifiable Credentials.
- **Credential Status List:** A mechanism used by a Verifiable Credential issuer where a verifier can check to see if a credential has been suspended or revoked.
- **Credential Wallet:** An entity used by the Holder to request, receive, store, present, and manage Verifiable Credentials and cryptographic key material.
- **Decentralized Identifier** (DID): A type of entity identifier that is globally unique, resolvable with high availability, and cryptographically verifiable. DIDs are used to identify Participants within a TWIN Ecosystem.
- **Holder:** An entity that receives Verifiable Credentials and has control over them to present them to the Verifiers as Presentations.
- **Identity:** A set of attributes related to an entity.
- **Identity Service Provider:** A type of [Trust Service Provider](#) with the role of proving and/or verifying Participant identities. They can do this by using online (for instance, through a Trusted Data source) or offline channels, or a combination of both. An Identity Service Provider can be a public or private sector organization (for instance, a bank).
- **Key Management System:** A system for the management of cryptographic keys and their metadata (e.g., generation, distribution, storage, backup, archive, recovery, use, revocation, and destruction). An automated key management system may be used to oversee, automate, and secure the key management process.
- **KYC:** Know your Customer. It is a process intended to verify the identity of new Participants. There can be multiple mechanisms and providers of KYC services.
- **Party Credential:** A Verifiable Credential that attests to the attributes of a Data Service, Dataset, or Participant.
- **Public Key Infrastructure:** A way to implement secure electronic transactions over insecure networks, such as the Internet. It is used to authenticate identities for data encryption and signing.
- **Relying Party (or Verifier):** A role an entity performs by receiving one or more Verifiable Credentials, optionally inside a Verifiable Presentation for processing. A TWIN Node might play this role when processing requests issued by other TWIN Nodes.
- **Trust Framework:** A legally enforceable set of specifications, rules, and agreements that govern a multi-party system established for a common purpose, designed for

conducting specific types of transactions among a community of participants, and bound by a common set of requirements.

- **Trust List:** Collection of trusted certificates used by Relying Parties to authenticate other certificates.
- **Trust Service Provider:** A role performed by an external entity to the TWIN ecosystem by offering identity verification or attribute attestation.
- **Verifiable data registry:** A role a system might perform by mediating the creation and verification of identifiers, keys, and other relevant data, such as verifiable credential schemas, revocation registries, issuer public keys, and other rules. For example, verifiable data registries include trusted databases, decentralized databases, government ID databases, and distributed ledgers. Often, there is more than one type of verifiable data registry utilized in an ecosystem.
- **Verifiable Credential:** A tamper-evident credential that has authorship that can be cryptographically verified. Verifiable credentials can be used to build **verifiable presentations**, which can also be cryptographically verified.
- **Verifiable Presentation:** A tamper-evident container of data derived from one or more verifiable credentials, issued by one or more issuers, that is shared with a specific verifier by a holder.
- **Verifier:** An entity that requests, receives, and validates Presentations.

Data Spaces

- **Authentication:** The process of verifying the identity of a requesting party, as a prerequisite to allowing access to resources.
- **Authorization:** The process of verifying whether a requesting party is allowed to access resources.
- **Creator:** Also known as *Data Producer*, the Creator generates data, e.g., by generating data such as from a sensor or accessing data in backend IT systems. In international trade, an exporter filling out an export declaration plays the role of “Data Producer”.
- **Data** here is synonymous with Data Asset, i.e., content exposed for exchange by a Participant acting as a Data Provider.

- **Dataspace** governance framework and supporting services to build trustworthiness and enable the sharing of data through an agreed set of Policies, semantic models, protocols and processes. Source [\[ISO/IEC 20151\]](#).
- **Data Sovereignty:** The ability of a natural or legal person to exclusively and sovereignly decide concerning the usage of data as an economic asset.
- **Data Exchange:** Data exchange takes place in the vertical cooperation between organizations to support, enable, or optimize value chains and supply chains.
- **Dataset:** A collection of data of interest to ecosystem participants, published by a Data Provider, exposed through a TWIN DSConnector, and available for access or download in one or more formats. It must declare access rights in terms of ODRL Offers. See also Dataset Credential. It is equivalent to a “dataset” description in terms of the [\[W3C-DCAT-v3\]](#) standard and must comply to the Eclipse DSP Catalog Protocol
- **Dataset Credential:** Dataset description cryptographically attested by its Provider, which follows the corresponding TWIN Schema. Its claims are validated by the TWIN Clearing House issuing a Dataset Compliance Credential. This compliance credential can be presented to a TWIN Catalogue for registering a new Dataset. See also Dataset.
- **Operator:** Providers that have been approved by the ecosystem governance to operate Federation Services and the Federation, which are independent of each other. There can be one or more Operators per type of Federation Service.
- **Policy:** A set of rules, duties, and obligations that define the terms of use for a Dataset.
- **Policy Enforcement:** System functionality intended to execute policies so that Providers or Consumers are ensured to meet the rules associated with Datasets.
- **Rule:** An abstract concept that represents the common characteristics of Permissions, Prohibitions, and Duties.
- **TWIN Schema:** A JSON Schema or SHACL⁹² JSON-LD aimed at validating different data elements within TWIN (Participants, Dataset offerings, etc.). TWIN Schemas are compliant with the Eclipse DSP with extensions when needed.
- **TWIN Registry:** A verifiable data registry that captures Trust Anchors, TWIN Schemas, and Ecosystem (participation) Rules, which are key to complying within a TWIN Ecosystem. For decentralization purposes, TWIN aims at a DLT-based (IOTA) implementation.
- **TWIN Data App Provider:** The developer or distributor of a TWIN Data App.

⁹² <https://www.w3.org/TR/2017/REC-shacl-20170720/>

- **Vocabulary** ontologies, reference data models, or metadata elements that can be used to annotate and describe participants, datasets, usage policies, apps, services, data sources, etc.

DLT

- **Asset:** A representation of value.
- **Content-addressable storage:** A way to store information so it can be retrieved based on its content, not its name or location.
- **Decentralized Storage:** A mechanism for storing data, split into small pieces, across multiple computers or nodes connected to a P2P network like the InterPlanetary File System (IPFS) protocol.
- **Decentralized Storage Node:** A Node participating in a Decentralized Storage system.
- **Decentralized System:** A distributed system wherein control is distributed among the persons or organizations participating in the operation of the system.
- **Distributed Ledger:** A type of ledger that is shared, replicated, and synchronized in a distributed and decentralized manner.
- **DLT Commitment:** A record stored by an application on a Distributed Ledger that enables proving data stored externally has not been tampered with. A DLT commitment can also be used for timestamping purposes.
- **DLT Node:** device or process that participates in a distributed ledger network.
- **Fingerprint:** A cryptographic hash of the content of an object (file, document) that allows integrity checking.
- **NFT:** (Non-fungible token) a unique digital representation of an asset.
- **Timestamping:** A timestamp proves that a message existed before some point in time. Timestamping, also known as notarization, is the act of creating a timestamp.
- **Public distributed ledger system:** A distributed ledger system that is accessible to the public for use.
- **Permissionless distributed ledger system:** A distributed ledger system where permissions are not required to maintain and operate a node.

- **Smart Contract:** A program written on the distributed ledger system that encodes the rules for specific types of distributed ledger system transactions in a way that can be validated and triggered by specific conditions.

Domain-specific

Value Chains

- **Automatic Identification and Data Capture (AIDC):** The automated process of identifying and capturing data about trade items, facilitating tracking, processing, and inventory processes. AIDC technologies include different variants of barcodes, QR Codes, smart cards, biometrics, RFID, NFC, or even machine learning/deep learning techniques.
- **Auditable Item:** An informational resource, usually bound to a trade entity, whose history is explicitly captured by an information system. Each historical state's fingerprint of an Auditable Item can be recorded on a Distributed Ledger for timestamping purposes, enabling actors to perform external data verification.
- **Auditable Item Graph:** A graph in which vertices correspond to Auditable Items and edges to relationships among those Auditable Items (child of, parent of, etc.).
- **Global Trade Item Number (GTIN):** A Trade Item identifier scheme for which GS1 is the authority. The GTIN can be used to identify types of products at any packaging level (e.g., consumer unit, inner pack, case, pallet).
- **Circularity:** An economic model that aims to minimize environmental impact by reducing waste and maximizing reuse. A useful conceptual reference is the “9R framework⁹³” (*refuse, rethink, reduce, reuse, repair, refurbish, remanufacture, repurpose, recycle, and recover*).
- **Document** (used in trade): Written, printed, or electronic matter that is referenced and concerns a Trade Item. From an information management perspective, documents can have multiple representations, JSON-LD documents, PDF, XML, etc. Often, documents carry key information referring to their related Trade Item.

⁹³ EC "Categorization System for the Circular Economy", March 2020, available at https://circularconomy.europa.eu/platform/sites/default/files/categorisation_system_for_the_ce.pdf .

- **Digital Link:** A Trade Item Identifier encoded as a persistent URL that, when resolved, leads to machine-readable information about the concerned item. A GS1 Digital Link is a Digital Link that encodes a GS1 EPC.
- **Digital Product Passport:** A structured collection of product-related data with predefined scope and agreed data management and access rights conveyed through a unique identifier that is accessible via electronic means through a data carrier. The intended scope of the DPP is information related to sustainability, circularity, value retention for reuse, remanufacturing, and recycling.
- **ESPR Regulation:** *Ecodesign for Sustainable Products Regulation (EU)*. A framework of environmental sustainability requirements for European goods, making mandatory a Digital Product Passport for certain products on the EU market. Under the framework of the EU Green Deal, the European Commission adopted the Ecodesign for Sustainable Products Regulation (ESPR) in 2024. The overall aim of the regulation is to reduce the lifecycle environmental impacts of products through efficient digital solutions.
- **FSMA Regulation:** *Food Safety Modernization Act (US FDA)* A regulation for the way foods are grown, harvested, and processed. It includes several rules, such as the Preventive Controls Rules for Human and Animal Food, the Produce Safety Rule, and the Foreign Supplier Verification Programs (FSVP) rule.
- **Supply Chain:** All upstream activities and processes of the value chain of the product, up to the point where the product reaches the end-user (customer).
- **Supply Chain Event:** Data record concerning one or more Trade Items to enable visibility, within organizations as well as across an entire ecosystem of Participants. It helps answer the questions “what, when, where, why, and how” of Trade Items, enabling the capture and sharing of key information such as status, location, movement, and chain of custody.
- **Trade Item:** Items that represent products or services that are priced, ordered, or invoiced at any point in the supply chain. From an information management point of view, Trade Items are represented by a descriptive Digital Twin that captures their type, properties, and relationships with other items.
- **Trade Item Identifier:** A collection of alphanumeric characters that can be used to identify a Trade Item. A GS1 Electronic Product Code (EPC) is a Trade Item Identifier represented using the GS1 identification scheme.
- **Trade Document:** A [document](#) used in Trade that is [transferrable](#).
- **Traceability:** According to the United Nations Global Compact and Business (UNGBC), traceability is the ability to identify and trace the history, distribution, location, and

application of products, parts, and materials, to ensure the reliability of sustainability claims, in the areas of human rights, labor (including health and safety), the environment and anti-corruption.

- **Value Chain**⁹⁴: All activities and processes that are part of the life cycle of a product, as well as its possible remanufacturing.

International Trade

- **Airway Bill**: The Air Waybill (AWB) is a critical air cargo trade document that constitutes the contract of carriage between the shipper and the carrier (airline).
- **Bill of Lading**: A trade document with contractual value, issued to the shipper, which confirms the carrier's receipt of the cargo, acknowledging goods being shipped or received for shipment, and specifying the terms of delivery (as one of the evidences of the contract of carriage). The Bill of Lading is usually prepared based on shipping instructions, including cargo description, given by the shipper on forms issued by the Carrier, and is the title to the goods and can be a negotiable document. Historically, a nautical term, "Bill of Lading" is used in the context of other forms of transport (air, train, or truck shipments) for similar documents such as the airway bill and CMR (road transport).
- **Buyer**: The Participant to whom goods or services are sold as stipulated in a sales order contract.
- **Carrier**: The Participant who provides transport services.
- **Certificate of Origin**: A document that attests to the country of origin/originating status of a Trade Item. It is an attestation that can be used to claim the Trade Item satisfies the applicable origin criteria.
- **Commercial Invoice**: A trade document that contains a demand for payment (concerning a Shipment) made from the invoice issuer (the invoicer, usually the Seller) to the invoicee (usually the Buyer).
- **Consignment**: A separately identifiable collection of Trade Items to be transported or available to be transported from one **consignor** to one **consignee** in a supply chain via one or more modes of transport, where each consignment is the subject of one single transport contract.

⁹⁴ See also <https://www.cisl.cam.ac.uk/education/graduate-study/pgcerts/value-chain-defs>

- **Customs Declaration:** A document whereby a Participant indicates in the prescribed form and manner the request to place Trade Items under a customs procedure (outward Export or inward Import procedures). There must always be an Export and an Import customs declaration, and these are made in the countries of dispatch and receipt. Typically, the declarations are made by different parties.
- **Export Declaration:** A type of Customs Declaration. The export declaration is made by the seller/despatcher of the goods or their customs agent. It is submitted to gain authorisation for the goods to leave a country and is mainly focused on compliance (e.g., any standards that the goods might need to meet, any restrictions that apply to exporting particular types of goods, or the country that the goods are being sold to/moved to).
- **Exporter:** The Participant who makes the export declaration, or on whose behalf the export declaration is made, and who is the owner of the goods or has similar rights of disposal over them at the time when the declaration is accepted.
- **EUDR:** Regulation on Deforestation-free Products (EU). Rules to guarantee that the products EU citizens consume do not contribute to deforestation or forest degradation worldwide.
- **Freight Forwarder:** The Participant who undertakes on behalf of a shipper the forwarding of goods by provision of transport, logistics, formalities services, etc., by liaising with Carriers.
- **Import Declaration:** A type of Customs Declaration. The import declaration informs goods compliance procedures for the receiving country and will be made by the buyer/receiver of the goods. There is also a financial element for an import declaration, should customs import duty or sales tax apply.
- **Importer:** The Participant who makes, or on whose behalf a customs clearing agent or other authorized person makes, an import declaration.
- **MLETR Regulation Model Law on Electronic Transferable Records (United Nations):** It aims to enable the legal use of electronic transferable records both domestically and across borders. The MLETR applies to electronic transferable records that are functionally equivalent to transferable documents or instruments.
- **Packing List:** A document that provides the exporter, international freight forwarder, and ultimate consignee with information about a Consignment, including how it's packed, the dimensions, and the weight of each package.
- **Phytosanitary Certificate:** A document that attests the necessary sanitary and phytosanitary measures have been taken for the protection of human, animal, or plant life or health.

- **Seller:** The Participant selling goods or services as stipulated in a sales order contract.
- **Shipment:** A shipment is an identifiable collection of one or more trade items (available to be) transported together from the seller (original consignor/shipper) to the buyer (final/ultimate consignee). A Shipment may form part or all of a Consignment or may be transported in different Consignments. It is a synonym for Trade Delivery.
- **SPS regulations,** Sanitary and Phytosanitary regulations — government standards to protect human, animal, and plant life and health, to help ensure that food is safe for consumption.
- **Trade Finance:** The financial instruments and products that are used by companies to facilitate international trade and commerce. It makes it possible and easier for importers and exporters to transact business through trade.
- **Transferable documents or instruments:** Paper-based documents or instruments that entitle the holder to claim the performance of the obligation indicated therein and that allow the transfer of the claim to that performance by transferring possession of the document or instrument. Transferable documents or instruments typically include bills of lading, bills of exchange, promissory notes, and warehouse receipts.
- **UCR:** *Unique Consignment Reference*⁹⁵, a reference number, applied to all international goods movements for which Customs control is required, with the following characteristics: used only as an access key for audit, consignment tracking and information, and reconciliation purposes; unique at both national and international level; applied at consignment level; issued as early as possible in the international transaction.

⁹⁵ <https://www.wcoomd.org/en/topics/facilitation/instrument-and-tools/tools/ucr.aspx>

References

[CIRPASS] CIRPASS Project. DPP in a nutshell. URL: <https://cirpassproject.eu/dpp-in-a-nutshell>

[GS1-Standard] GS1 General Specifications Standard. Release 24. URL: <https://ref.gs1.org/standards/genspecs/>

[GS1-EPCIS] GS1. EPCIS standard. Version 2.0. URL: <https://ref.gs1.org/standards/epcis/2.0.0/>

[GS1-Digital-Link] GS1. Digital Link standard. URI Syntax. Version 1.5.0. URL: <https://ref.gs1.org/standards/digital-link/uri-syntax/>

[DCSA-Shipping-Glossary] Digital Container Shipping Association (DCSA). Shipping Glossary. URL: <https://dcsa.org/standards/shipping-glossary>

[EC-ESPR] European Commission. REGULATION (EU) 2024/1781 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 13 June 2024. Establishing a framework for the setting of ecodesign requirements for sustainable products. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202401781

[EC-Deforestation] European Commission. REGULATION (EU) 2023/1115 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 31 May 2023. Making commodities and products associated with deforestation and forest degradation available on the Union market and as exports from the Union. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32023R1115>

[Eclipse-DSP] Eclipse Dataspace Protocol (November 2025 edition). URL: <https://eclipse-dataspace-protocol-base.github.io/DataspaceProtocol/2025-1>

[Eclipse-DCP] Eclipse Decentralized Claims Protocol (November 2025 edition). URL: <https://eclipse-dataspace-dcp.github.io/decentralized-claims-protocol/v1.0.1>

[ETSI-NGSI-LD] ETSI GS CIM 047 V1.1.2 (2024-12). Context Information Management (CIM);NGSI-LD API. URL:

https://www.etsi.org/deliver/etsi_gs/CIM/001_099/009/01.08.01_60/gs_CIM009v010801p.pdf

[EU-Data-Spaces] European Commission. COMMISSION STAFF WORKING DOCUMENT of January 2024. Common European Data Spaces. URL:

<https://ec.europa.eu/newsroom/dae/redirection/document/101623>

[US-FSMA] 'FDA Food Safety Modernization Act'. PUBLIC LAW 111-353—JAN. 4, 2011. URL:

<https://www.govinfo.gov/content/pkg/PLAW-111publ353/pdf/PLAW-111publ353.pdf>

[Gaia-X] Gaia-X Consortium. What is Gaia-X?. URL: <https://gaia-x.eu/what-is-gaia-x/>

[Gaia-X-Architecture] Gaia-X Consortium. Gaia-X Architecture Document v24.04. URL:

<https://docs.gaia-x.eu/technical-committee/architecture-document/24.04/>

[Gaia-X-Credentials] Gaia-X Consortium. Gaia-X Identity, Credential and Access Management Document v24.07. URL:

<https://docs.gaia-x.eu/technical-committee/identity-credential-access-management/24.07>

[IDSA-RAM-4] International Data Spaces Association. Reference Architecture Model v4.

<https://docs.internationaldataspaces.org/ids-knowledgebase/ids-ram-4>

[ITU-T-DLT] ITU-T Focus Group on Application of Distributed Ledger Technology. Technical Specification FG DLT D1.1 Distributed ledger technology terms and definitions. URL:

<https://www.itu.int/en/ITU-T/focusgroups/dlt/Documents/d11.pdf>

[ISO-20151] ISO/IEC 20151 Information technology — Cloud computing and distributed platforms — Dataspace concepts and characteristics (under development). URL:

<https://www.iso.org/standard/86589.html>

[OIDC4VCI] OpenID Foundation. OpenID for Verifiable Credential Issuance. Draft 15. December 2024. URL: https://openid.net/specs/openid-4-verifiable-credential-issuance-1_0.html

[OIDC4VP] OpenID Foundation. OpenID for Verifiable Presentations. Draft 23. December 2024. URL: https://openid.net/specs/openid-4-verifiable-presentations-1_0.html

[IETF-SD-JWT-2024] Fett, D., Yasuda, K., and B. Campbell, "Selective Disclosure for JWTs (SD-JWT)", Work in Progress, Internet-draft, draft-ietf-oauth-selective-disclosure-jwt-14, 15 November 2024, URL:

<https://datatracker.ietf.org/doc/html/draft-ietf-oauth-selective-disclosure-jwt-14>

[UN/CEFACT-BSP] United Nations Economic Commission for Europe. Centre for Trade Facilitation and Electronic Business. (UN/CEFACT). Buy–Ship–Pay (BSP) Reference Data

Model. Summary presentation. URL:

https://unece.org/fileadmin/DAM/cefact/brs/BuyShipPay_BRS_v1.0.pdf

[UN/CEFACT-Rec49] UN/CEFACT. Draft Recommendation No. 49 - Transparency at Scale.

URL: https://unece.org/sites/default/files/2024-07/ECE-TRADE-C-CEFACT-2024-06E_0.pdf

[UN/CEFACT-VC-Trade] UN/CEFACT. White Paper eDATA Verifiable Credentials for Cross Border Trade. September 2022. URL:

https://unece.org/sites/default/files/2023-08/WhitePaper_VerifiableCredentials-CrossBorderTrade_September2022.pdf

[UN/CEFACT-Prod-Certificate] UN/CEFACT. White Paper on Digital Product Conformity Certificate Exchange. October 2023. URL:

https://unece.org/sites/default/files/2023-10/WhitePaper_DigitalProductConformityCertificateExchange.pdf

[UNCITRAL-MLETR] United Nations Commission on International Trade Law (UNCITRAL). Model Law on Electronic Transferable Records. (MLETR). URL:

https://uncitral.un.org/sites/uncitral.un.org/files/media-documents/uncitral/en/mletr_ebook_e.pdf

[UNTP] UN/CEFACT. United Nations Transparency Protocol (UNTP). Technical specification.

URL: <https://uncefact.github.io/spec-untp/docs/about>

[W3C-Activity-Pub] ActivityPub. 23 January 2018. W3C Recommendation. URL:

<https://www.w3.org/TR/2018/REC-activitypub-20180123/>

[W3C-Activity-Streams] Activity Streams 2.0. 23 May 2017. W3C Recommendation. URL:

<https://www.w3.org/TR/2017/REC-activitystreams-core-20170523/>

[W3C-DCAT-v3] Data Catalog Vocabulary version 3. 22 August 2024. W3C Recommendation.

URL: <https://www.w3.org/TR/2024/REC-vocab-dcat-3-20240822/>

[W3C-Data-Integrity] Verifiable Credentials Data Integrity 1.0. W3C Candidate

Recommendation. 26 January 2025. URL: <https://www.w3.org/TR/vc-data-integrity/>

[W3C-DID-Core] Decentralized Identifiers (DIDs) v1.0. Core architecture, data model, and representations. 19 July 2022. W3C Recommendation. URL:

<https://www.w3.org/TR/2022/REC-did-core-20220719/>

[W3C-JSON-LD] JSON-LD 1.1. A JSON-based Serialization for Linked Data. 16 July 2020.

W3C Recommendation. URL: <https://www.w3.org/TR/2020/REC-json-ld11-20200716/>

[W3C-LDN] Linked Data Notifications. 2 May 2017. W3C Recommendation. URL:

<https://www.w3.org/TR/2017/REC-ldn-20170502/>

[W3C-ODRL-22] ODRL Information Model v2.2. 15 Feb 2018. W3C Recommendation. URL: <https://www.w3.org/TR/2018/REC-odrl-model-20180215/>

[W3C-VC-DATA-MODEL] Verifiable Credentials Data Model v1.1. 3 March 2022. W3C Recommendation. URL: <https://www.w3.org/TR/vc-data-model-1.1/>

[WCO-Guidelines] World Customs Organization (WCO). Guidelines on certification of origin. <https://www.wcoomd.org/-/media/wco/public/global/pdf/topics/key-issues/revenue-package/guidelines-on-certification.pdf>

[WCO-Instruments] World Customs Organization (WCO). Instruments and tools. URL: <https://www.wcoomd.org/en/topics/facilitation/instrument-and-tools.aspx>

[WCO-Handbook] World Customs Organization. WCO Handbook on Inward and Outward Processing Procedures. URL:

https://www.wcoomd.org/-/media/wco/public/global/pdf/topics/facilitation/instruments-and-tools/tools/wco-handbook-on-inward-and-outward-processing-procedures/pc_handbook.pdf

[WTO-Sanitary] World Trade Organization (WTO). Agreement on the Application of Sanitary and Phytosanitary Measures. URL: https://www.wto.org/english/docs_e/legal_e/15sps_01_e.htm

[WTO-Glossary] World Trade Organization (WTO). Glossary of terms. URL: https://www.wto.org/english/thewto_e/minist_e/min99_e/english/about_e/23glos_e.htm#ag

[WTO] World Trade Organization (WTO). Technical Information on Rules of Origin. URL: https://www.wto.org/english/tratop_e/roi_e/roi_info_e.htm

- TWIN Whitepaper

Trade Worldwide Information Network
Seamless Trading for All