Gaia-X Architecture Document - 3.1 Release

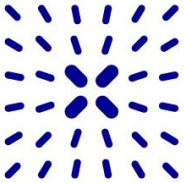

Description	Gaia-X specification to build trusted decentralised digital ecosystems.
Repository	https://gitlab.com/gaia-x/technical-committee/architecture-working-group/architecture-document
Author(s)	Gaia-X European Association for Data and Cloud AISBL
Copyright(s)	

Table of Contents

I About

1 Editorial Information
- 1.1 Publisher
- 1.2 Authors
- 1.3 Contact
- 1.4 Other format
- 1.5 Copyright notice

2 Introduction

3 Gaia-X Context
- 3.1 Basic Principles
- 3.2 Understanding Ecosystems and Data Spaces
  - 3.2.1 Federation of Ecosystems
- 3.3 Gaia-X high-level Positioning
  - 3.3.1 Trust Plane
  - 3.3.2 Management Plane
  - 3.3.3 Usage Plane
  - 3.3.4 Complementarity of Technical Compatibility and Compliance
- 3.4 Complementarity with additional
  - 3.4.1 Integrating with external
  - 3.4.2 Gaia-X 3.0 “Danube” software release
- 3.5 Gaia-X Alignment
  - 3.5.1 Aligning with Other Associations and Foundations
  - 3.5.2 Aligning with Other Initiatives
  - 3.5.3 Aligning with External Projects
  - 3.5.4 Aligning with Standards and Regulations

4 Gaia-X Trust Framework Architecture
- 4.1 Gaia-X Trust Framework
- 4.2 Using the Trust Framework
  - 4.2.1 Performing automated onboarding and offboarding
  - 4.2.2 Performing
  - 4.2.3 Identifying Ecosystem
- 4.3 GXDCH
- 4.4 Gaia-X Conceptual Model
  - 4.4.1 Terminology sources
  - 4.4.2 Definitions
  - 4.4.3 Model Core
- 4.5 Gaia-X Trust Protocol
  - 4.5.1 Problem Statement
  - 4.5.2 Gaia-X Trust Protocol Technical Specification
- 4.6 Cross-Ecosystem Interoperability
- 4.7 Inter-Ecosystem Interoperability
- 4.8 Services and Service Composition
  - 4.8.1 Resources and Service Offerings
- 4.9 Policies
  - 4.9.1 Policy Definition
  - 4.9.2 Policy Description
  - 4.9.3 Gaia-X Policy Reasoning Engine
  - 4.9.4 Policy Decision Point (PDP)
- 4.10 Ecosystem Trust Functions
  - 4.10.1 Rights Delegation
  - 4.10.2 Trust Indexes
- 4.11 Data Space Architecture using the Gaia-X Trust Protocol
- 4.12 Trusted Artificial Intelligence Environments
  - 4.12.1 Compliance of the
  - 4.12.2 Compliance of the underlying infrastructure and operators
  - 4.12.3 Compliance of AI services (e.g. AI-Agents, Gen AI, ML or Simulation services)

5 Gaia-X Implementation of Trusted Data Transactions
- 5.1 Data Transactions Conceptual Model

6 Gaia-X Technical Compatibility specifications
- 6.1 Defining Technical Compatibility
- 6.2 Understanding Identity and Identifier
  - 6.2.1 Using Identifiers in Gaia-X Credentials
- 6.3 Using Linked Data
- 6.4 Using Verifiable Credentials
- 6.5 Verifying Gaia-X Credentials
- 6.6 Gaia-X Credential Format
- 6.7 OpenID Connect for Verifiable Credentials
  - 6.7.1 OpenID Connect for Verifiable Credential Issuance

- 6.7.2 OpenID Connect for Verifiable Presentations
- 6.7.3 Usage
- 6.7.4 Cloud/Enterprise Wallet
- 6.8 Gaia-X Ontology
  - 6.8.1 Versioning
  - 6.8.2 DCAT
  - 6.8.3 CAP
  - 6.8.4 Gaia-X Schema
- 6.9 Managing Trust Services
  - 6.9.1 Trusted Service Operators
  - 6.9.2 Using Compliance Engine
  - 6.9.3 Using the Registry
  - 6.9.4 Using the Gaia-X Legal Registration Number (LRN) Notary

II Appendices

7 Supported Credential Formats and -Exchange Protocols, Wallets and
- 7.1 Credential Formats
- 7.2 Credential Exchange Protocols

8 Trust Indexes
- 8.1 Sub-Indexes
  - 8.1.1 Veracity
  - 8.1.2 Transparency
  - 8.1.3 Composability
  - 8.1.4 Semantic Match

9 Changelog
- 9.1 Release 3.1 (2026 June Release 26.06)
- 9.2 2025 November Release (25.11)
- 9.3 2025 May Release (25.05)
- 9.4 2024 April release (24.04)
- 9.5 2023 October release (23.10)
- 9.6 2022 October release (22.10)
- 9.7 2022 September release (22.09)
- 9.8 2022 April release (22.04)
- 9.9 2021 December release (21.12)
- 9.10 2021 September release (21.09)
- 9.11 2021 June release (21.06)
- 9.12 2021 March release (21.03)
- 9.13 2020 June release (20.06)

I. About

1 Editorial Information

1.1 Publisher

Gaia-X European Association for Data and Cloud AISBL
Avenue des Arts 6-9
1210 Brussels
www.gaia-x.eu

1.2 Authors

Gaia-X European Association for Data and Cloud

1.3 Contact

https://gaia-x.eu/contact/

1.4 Other format

For the reader's convenience a PDF version of this document is generated here.
Note, though, that this PDF version is neither optimized for layout nor for any off-line visualization user experience.

1.5 Copyright notice

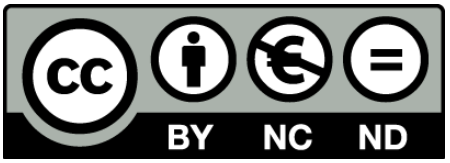

2 Introduction

This document is designed to guide a technical or technically-minded audience into understanding, applying, and using the software system-related parts of the Gaia-X Trust Framework. This includes roles such as enterprise IT architects, software architects, technical leads, lead developers, software developers, (software) product owners, as well as project managers and technical leadership functions (up to and potentially including also CTOs, CIOs, or IT managers). Using the concepts contained in this architecture document, organisations and enterprises will be able to define and implement the technical measures necessary to create or participate in federated trusted digital ecosystems whose trust layers are interoperable with other digital ecosystems.

The document structures Gaia-X's architectural vision across interconnected layers, beginning with the conceptual underpinnings of federated digital ecosystems and progressing to technical specifications enabling interoperable trust. Chapter 3 contextualizes the larger environment Gaia-X is operating in from a technical vantage point, including a positioning towards related initiatives and standardisation efforts. Chapter 4 outlines and explains the Gaia-X Trust Protocol, a generic specification for interoperable trust relationships across ecosystems, and the Gaia-X Trust Framework, which represents an instantiation of the Gaia-X Trust Protocol based on Gaia-X Compliance rules and Gaia-X-accredited trust service providers. These specifications enable automated onboarding, credential verification, and ecosystem services to operationalize sovereignty and compliance with business rules across distributed infrastructures. Chapter 5 details an implementation mechanism supported by the particular semantic model for data products. The complete specification of Gaia-X technical compatibility (the "heart") is then explained and expanded in Chapter 6.

Overall, the architecture's layered design converges towards our single vision at Gaia-X: empowering digital ecosystems with the technical and governance foundations supporting the implementation of an interoperable universal trust layer securing, amongst others, sovereignty and security for all participants and the service interactions between them.

3 Gaia-X Context

This section provides basic notions related to ecosystems and data spaces to set the context where Gaia-X is positioned. It also describes the Gaia-X mission and how Gaia-X is aligned with other relevant initiatives, stakeholders and ongoing standardisation efforts.

3.1 Basic Principles

A digital ecosystem is defined by a group of participants, who define the ecosystem purpose and a common set of (technical and commercial) rules governing the exchange of data and services, together with the underlying legal and commercial agreements. To ensure compliance with these rules, a compliance mechanism is built by relying on a set of designated trust service providers. The Ecosystem provides value to its participants by:

- defining and publishing the purpose and the underlying rule set
- offering services which enable the value creation inside the ecosystems, which may include catalogues, directories, registries and shared services which support specific use cases e.g.,
  - vocabularies defining common data models, digital twin registries, (pseudo-) anonymization services
  - mechanisms to negotiate and audit the use of data (incl. the management of data usage agreements)
- providing participant directories to find potential business partners
- enabling discovery and exposure via service and data directories, allowing participant to find and respectively publish own offerings

Generally, these services are available through a portal.

Gaia-X provides the mechanisms to translate the rulesets and roles to a digital framework allowing:

- automated participant onboarding
- assurance of compliance for ecosystem and participant services
- validation of claims by participants in individual digital negotiations

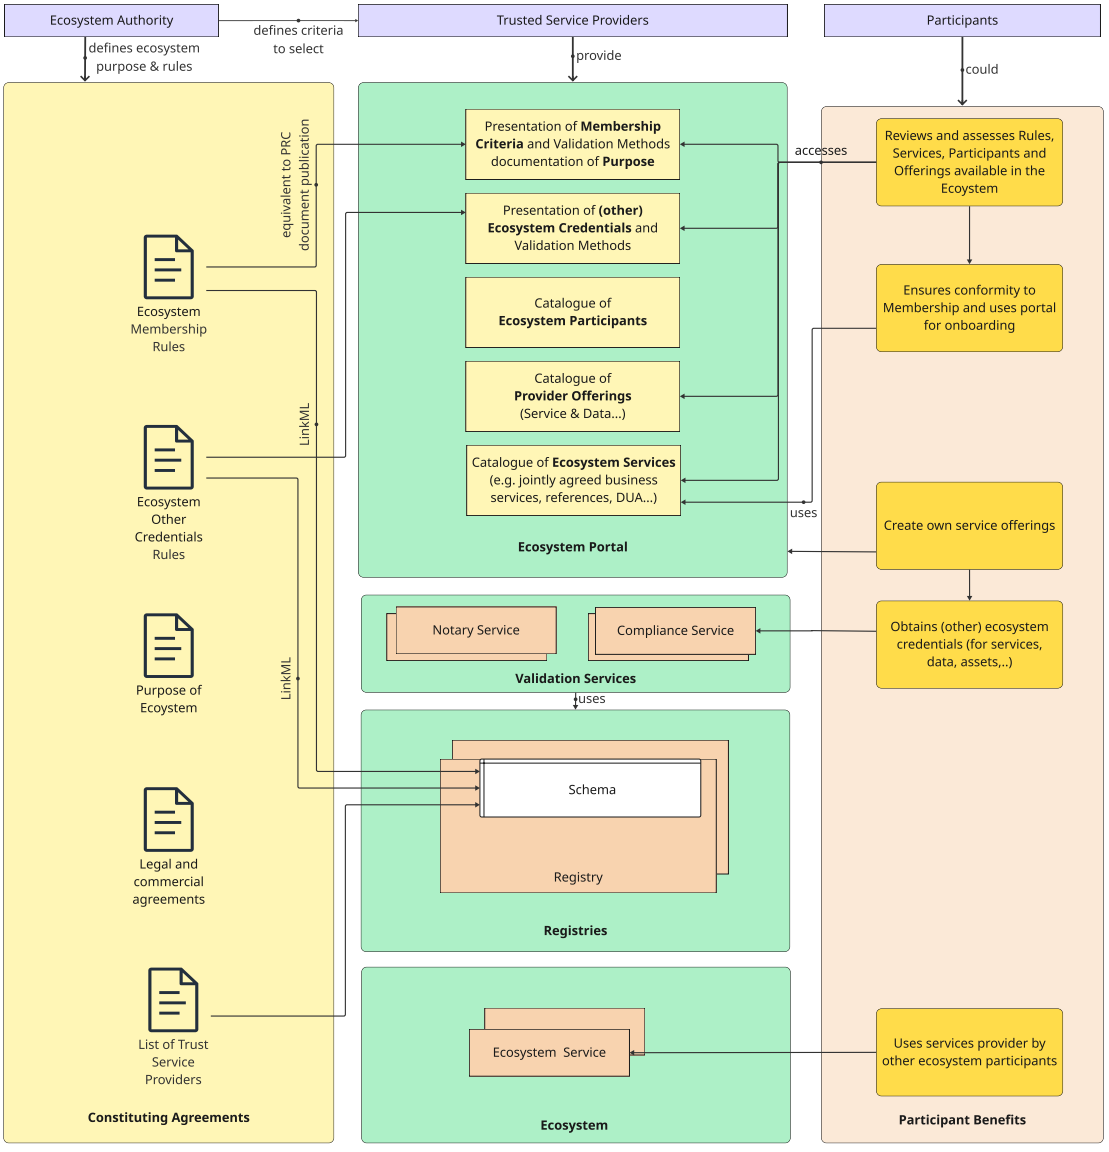

Figure 3.1 - Ecosystem overview

3.2 Understanding Ecosystems and Data Spaces

A digital ecosystem is a non-hierarchical organisational structure of a multilateral set of partners (or, equivalently, participants) that interact digitally in order for one or more focal value propositions to materialise.

The key characteristics are:

- Non-hierarchical means there is no single partner who is controlling or governing the whole ecosystem. While many, if not all, ecosystems feature something like a *center of gravity* for organising and structuring the rules and activities within the ecosystem. Any such "ecosystem governance authority" must not be controlled by a single partner. If that were the case, we cannot speak of a true ecosystem as it would simply be a hierarchy.
- Multilateral indicates that ecosystem participants have to be interrelated and organised towards achieving a common goal; just bringing together buyers and sellers on a (single) digital platform does not qualify as an ecosystem in our sense.
- We deliberately do not further restrict the notion of partner or participant: Typically, businesses, enterprises, and other legal persons form ecosystems, but natural persons are equally permitted to be prime participants in ecosystems. For instance, in health data spaces or the construction sector. Conversely, whole ecosystems or data spaces may also cooperate in order to achieve overarching, cross-ecosystem or cross-data space value propositions. In that sense, the *Mobility Data Space*, the *Smart Connected Suppliers Network* (SCSN), the *transport & logistics data space* and a *Smart City* ecosystem may form a "Super-Mobility-Transport-Supply"-ecosystem.
- The focal value proposition or overarching goal is holding and binding together the individual partners in any (digital) ecosystem. We do not restrict the range (broadness) or specificity of such a shared value proposition, although existing initiatives strongly suggest: the more concrete the value proposition, the stronger the ecosystem. For instance, the simple idea of "just sharing data in the [insert your favourite industry here]" is rather weak; forming a digital ecosystem to shorten the time for building new (nuclear) power plants by 50% is quite strong.

The four leading questions discriminating digital ecosystems from other forms of organisation and collaboration are:

1. Who are the partners or participants?
2. What is the common value proposition participants share?
3. Who is organising the set of partners?
4. What are the participants doing digitally?

Note that – in this *Architecture Document* – we intentionally do not ask about the commercial viability and sustainability of a digital ecosystem. This (fifth) question would run like

1. How do participants earn back the money they need to sustain the digital ecosystem?

Digital ecosystems can be established around infrastructures (and respective services, Infrastructure as a Service (IaaS)), applications and related services (Platform as a Service (PaaS), Software as a Service (SaaS)), data services, or any combination thereof. Participants within an ecosystem agree, through a formal governance body (often called *Ecosystem Governance Authority* as above), to a set of "Policy Rules" to which all participants must conform. Typically, these Policy Rules encompass a list of attributes, criteria for ecosystem conformity, methods for conformity attestation, and procedures for verification by mutually agreed-upon trusted parties.

Digital ecosystems typically operate within specific contexts:

- Regional: Participants must ensure compliance with regional regulations and legal requirements.
- Domain: Industry-specific standards and regulations are applicable.
- Sub-Domain: In addition to domain-specific standards, certain industries or sub-domains may impose further specific standards.
- Commercial Context: Interactions may be governed by particular commercial agreements.

A data space is defined as an "interoperable framework, based on common governance principles, standards, practices and enabling services, that enables trusted data transactions between participants." (source: DSSC Glossary) and is a particular instantiation of a digital ecosystem. The formal governance body of a data space is often called *Data Space Governance Authority* (DSGA).

While data spaces typically arise around domains or domain-specific value propositions, digital ecosystems are increasingly being recognised (also) on a higher organisational level as a broader infrastructure and governance layer uniting multiple data spaces. While this may sound confusing, our question 1 from above ("Who are the partners or participants of a particular ecosystem?") will immediately reveal the organisational level(s) at which a digital ecosystem operates.

> **Note**
>
> Enabling Services facilitate the operation of ecosystems. There are multiple technologies, products and implementations of each of the enabling services available.

In this context, a Gaia-X Ecosystem consists of entities which are Gaia-X technically compatible, as described later in the document. See also, Gaia-X Technical Compatibility Specifications

3.2.1 Federation of Ecosystems

Ecosystems can agree on common Trust Service Providers, ontologies, and shared services based on individual agreements. This allows a high level of interoperability between ecosystems and for participants in multiple ecosystems. Providers which are cooperating to provide interoperable services can agree on a common set of criteria, which they can use to create offering-specific "labels" (e.g. an infrastructure or middleware platform which has been validated for interoperability across different providers).

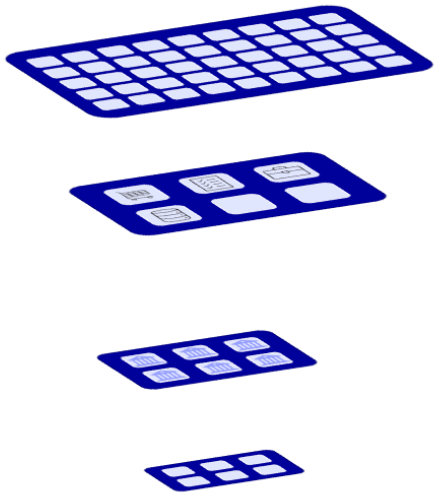

Participants offer and/or consume services based on individual negotiation and proof of conformity to the ecosystem rules they connect to

Ecosystem Operators can provide services to participants across different ecosystems if they adhere to the compliance rules of (all) the respective ecosystem (s)

Trust Service Providers (TSP) can provide services for multiple Ecosystem Governance Authorities (EGA) and to Participants from multiple eco-systems, based on the agreed definitions of the EGAs. TSPs can define trust relationships between each other

Ecosystem Governance Authorities can agree on common Policy Rules, Schemas and Trust Anchors and Compliance rules

- Rules applicable to all eco-systems of a particular ecosystem
- Individual agreements between ecosystems

These definitions are modular and extensible

Figure 3.2.1 - Ecosystem Federation

3.3 Gaia-X high-level Positioning

Gaia-X has the mission to create the *de facto* standard to enable federated and/or decentralized and trusted data and infrastructure ecosystems, by developing a set of specifications, rules, policies, and a verification framework.

This mission centers around establishing a *federated and/or decentralized* notion of *trust* in digital ecosystems as opposed to creating trust by relying on a single central party or contrary to zooming in on actual mechanisms or protocols related to data sharing or providing XaaS services (therefore our focus on *trust*). A *verification framework* is needed to automate trust by design: The sheer size of digital ecosystems – a single large airplane manufacturer has a supply network of around 10,000 enterprises; the whole automotive sector supply chain, comprising in excess of 250,000 companies, mandates this. Gaia-X not only defines the standards but also provides an (initial) open-source implementation of this *verification framework*. All components are collectively known as the Gaia-X Trust Framework (see Chapter 4 for a more rigorous definition and exposition).

The Gaia-X Trust Framework shall form the very basis for creating and ensuring interoperable, trusted relationships both in and between data spaces or digital ecosystems. It probably constitutes the only truly decentralized, cohesive, consistent, and future-proof set of standards needed to automate trusted digital transactions in arbitrary ecosystems. Of course, there exist several other partial solutions, many of which:

- feature a much more narrowly designed trust architecture
- do not discriminate between (technical) compatibility and (rule) compliance
- do not and will not provide a form of TCK (Technical Compatibility Kit)
- work on standardising the ontological underpinnings for defining an arbitrary ecosystem

We believe that using our Gaia-X Trust Protocol (see Chapter 4 for a detailed explanation) is the only way to allow cross-ecosystem interoperability at the level of trusted identities (of ecosystem participants) and trusted digital transactions (aka “service credentials”, or “data credentials” or any other “credential” an ecosystem may define) available today.

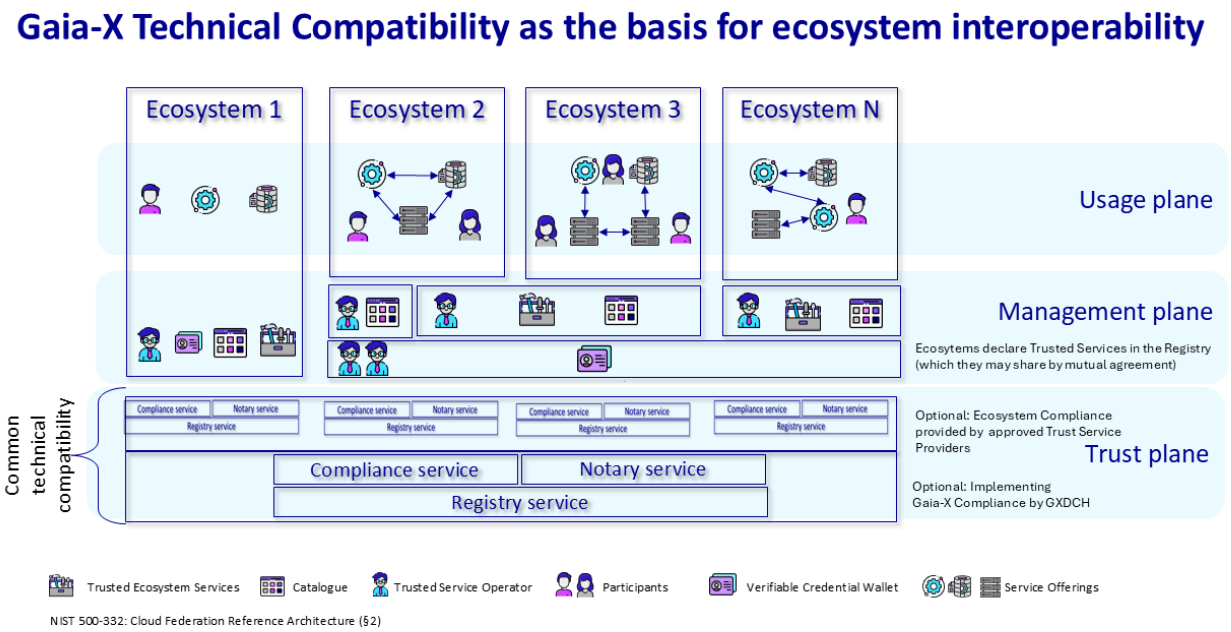

Figure 3.3 - The Gaia-X Technical Compatibility as an enabler for ecosystem interoperability

3.3.1 Trust Plane

Figure 3.3.4 features a *Trust Plane* (at the bottom) common to all ecosystems comprising two distinct features: 1. Technical compatibility: A mechanism by which ecosystems specify their particular notions of *trust*, for instance, which identity providers or catalogues to trust. Technically, this amounts to defining,

- (i) a common but extensible vocabulary (actually, it is an *ontology*) and
- (ii) a consistent set of standards on how to define and (later) verify and enforce compliance.
- Ecosystem-specific compliance rules: Using the "common language" (*lingua franca*) of point 1, ecosystems then are free to declare their individual rules for compliance, such as rules for participant onboarding, service provision or data assets management.

The Gaia-X Trust Framework is contributing to both features of the Trust Plane by providing the following deliverables:

No	Gaia-X contribution	Technical compatibility	Ecosystem specific compliance rules
1	An extended vocabulary to describe Ecosystem Participants (Consumers, Trusted Service Operators, and Providers), Service Offerings, ICT infrastructures, and resources used in service composition: the Gaia-X Ontology	X	
2	A vocabulary to describe data spaces and ecosystems: the Data Space Ontology	X	
3	Software libraries, software components, test beds, and a Gaia-X Technical Compatibility Kit (GX-TCK)	X	X
4	Compliance criteria and their deployment in production through the network of Gaia-X Digital Clearing Houses (GXDCH)	X	X

Ecosystems then become interoperable because they (a) use the same language for specifying their trust model and (b) by agreeing on a common set of trusted ecosystem services for those (few) services both deem relevant for becoming interoperable, for instance, agreeing on a set of mutually accepted identity providers, notaries, or conformity assessment bodies.

Gaia-X itself is an example for the application of this two-fold approach: Gaia-X provides a rulebook for Gaia-X compliant services as defined in the Gaia-X Compliance Document. Validation of Gaia-X service compliance is accomplished by the so-called Gaia-X Digital Clearing Houses (GXDCH), which are the approved Trust Service Providers in this case. This is indicated by marking these Gaia-X contributions as "optional" in the figure above.

3.3.2 Management Plane

In addition to the basic notions of trust shared (or not shared) by different ecosystems, the purpose of the Management Plane is to specify additional rules and embodiments or enforcement of these rules. Examples for this kind of *Trusted Ecosystem Services* are catalogues, wallet providers, or marketplaces. Depending on an ecosystem's needs, these services may also be *shared* with other ecosystems (cf. the catalogue service shared by `Ecosystem 2` and `Ecosystem 3` or the wallets shared by all ecosystems from `Ecosystem 2` until `Ecosystem N`).

3.3.3 Usage Plane

The notion of the *Usage Plane* refers to the level where individual participants of an ecosystem exchange data or engage in mutual service interactions with other participants of the same ecosystem or of another ecosystem.

3.3.4 Complementarity of Technical Compatibility and Compliance

As indicated above, trust frameworks discriminate between *technical compatibility* and *compliance* with certain rules. The following diagram depicts how the Gaia-X Trust Framework approaches this separation of concerns by showing the two main pillars, namely, Gaia-X Technical Compatibility on the left-hand side and Gaia-X Compliance in the middle. Elements of the Gaia-X Endorsement Programme (right-hand side) essentially target the adoption of our Gaia-X Trust Framework. A thorough and precise definition is given in Chapter 4.

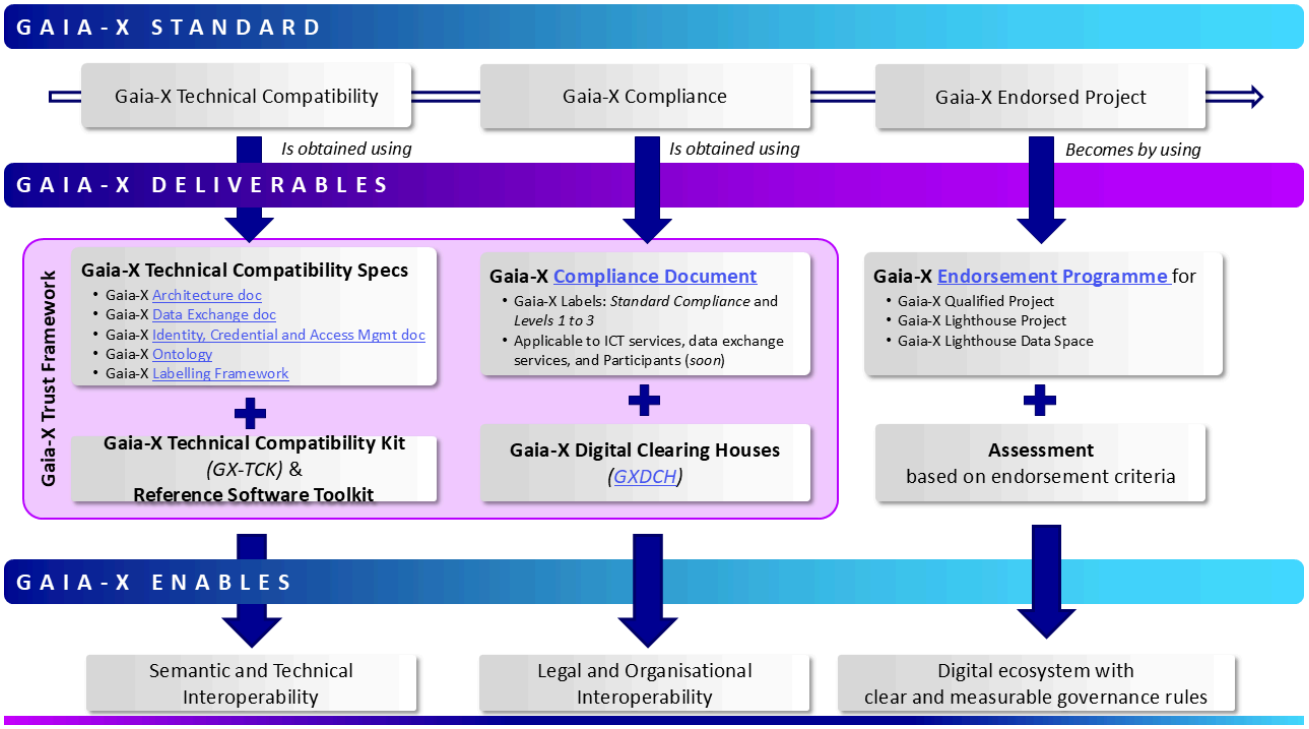

Figure 3.3.4 - The Gaia-X Framework

3.4 Complementarity with additional TSP and support in the Gaia-X OSS releases

3.4.1 Integrating with external trust frameworks

Ecosystems whose Governance Authorities are extending the set of compliance rules, choose to integrate external Trust Frameworks (e.g. eIDAS) or decide to trust TSPs from other ecosystems in general, require interoperability between the credential formats and the protocols for credential exchange.

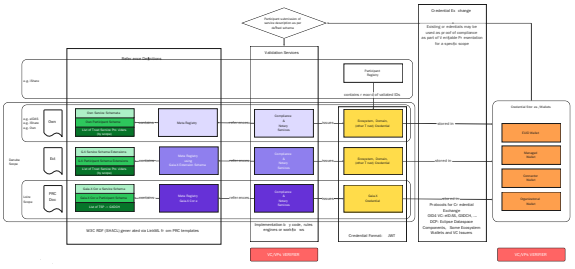

Figure 3.4 - Federated Trust scenario

Within such an environment, a compliance service can use credentials issued by another TSP as a basis for rule evaluations and issue credentials which in turn can be used across different ecosystems.

3.4.2 Gaia-X 3.0 “Danube” software release

With the Gaia-X 3.0 “Danube” release of the Gaia-X OSS, the code base introduces support to allow integration of arbitrary compliance rules or rule engines allowing any ecosystem to use specialised TSPs not limited to (i) the criteria specified in the Gaia-X Compliance Document or (ii) Gaia-X Digital Clearing Houses (GXDCHs) as sole TSPs for establishing compliance.

Based on an exhaustive re-engineering of the Gaia-X Loire components, the “Danube” release completely separates Gaia-X technical compatibility from any form of compliance, be that Gaia-X compliance or any other ecosystem compliance. This is based on the high-level requirements defined in the Geography and Domain Extension White Paper (available only to Gaia-X members), issued by the Gaia-X PRC in July 2025.

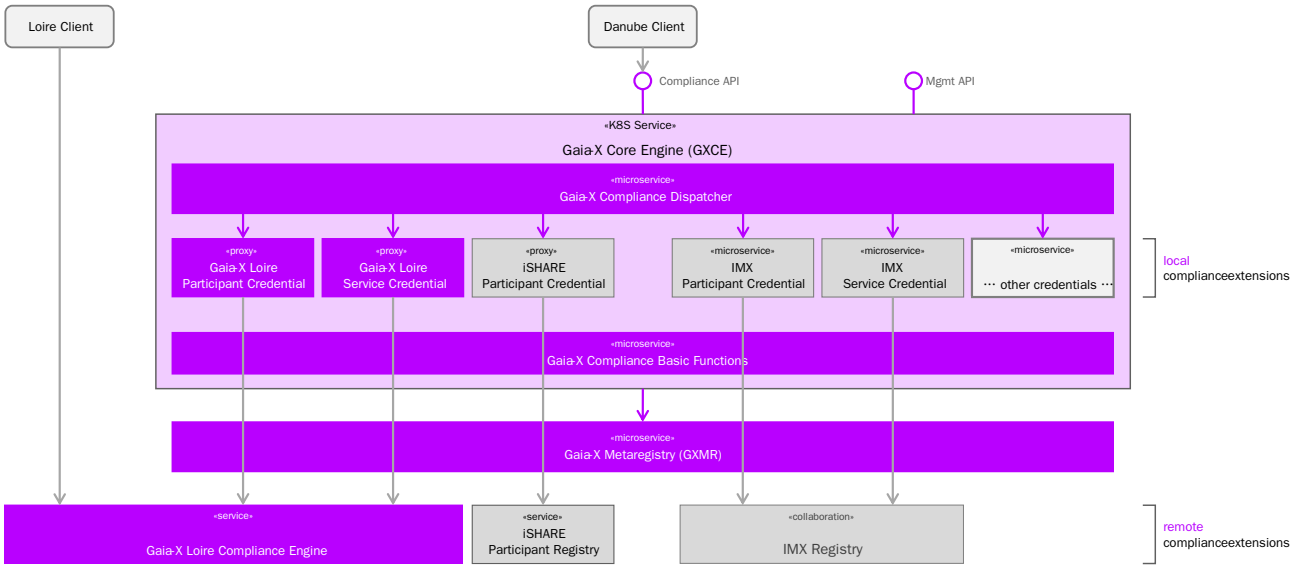

Figure 3.4.2 - Gaia-X 3.0 "Danube" architecture

The Danube architecture recognizes two major building blocks: - `Gaia-X Core Engine` - ensuring Gaia-X technical compatibility, hosting *local* compliance engines (called "local (compliance) extensions"), and providing means (`proxy`s) to access *remote* compliance extensions. - `Gaia-X Metaregistry` - exposing metadata characterizing different ecosystems, their respective trust services, and the verifiable credentials associated with these services.

From a compliance automation point of view, the `Gaia-X Core Engine` recognizes two forms to *run* or *access* compliance engines or rules engines called *extension types*:

Extension Type	Description	Example (in diagram)
local (compliance) extension	The component that checks compliance with a certain set of criteria runs in the same *local context* as the `Gaia-X Core Engine` (e.g., as a microservice within the same K8S cluster).	IMX Participant Credential IMX Service Credential
remote (compliance) extensions	The component that checks compliance runs outside the run-time control of the `Gaia-X Core Engine` and will be accessed through the network. Network access to the remote extension will be mediated within the `Gaia-X Core Engine` through a so-called `proxy`.	Gaia-X Loire Participant Credential Gaia-X Loire Service Credential iSHARE Participant Credential

> **Note**
>
> We acknowledge that the difference between local and remote is somewhat vague. The exact mechanism for adding local extensions and for accessing *remote* extensions may vary technically depending on the actual "Danube" release and the (future) requirements.

The `Gaia-X Core Engine` itself consists of the following four (types of) components:

Component	Description
`Gaia-X Compliance Dispatcher`	accepts incoming verifiable presentations and dispatches the verifiable credentials contained therein to the correct local or remote compliance extension. Local and remote extensions will have to be registered with the `Gaia-X Compliance Dispatcher` with their ecosystem identifier and the VC types they support.
`Gaia-X Basic Functions`	implements and exposes a set of commonly used functions such as - SHACL verification - resolving a `did:web` - validating and verifying a `did:web` and verification method - following the trust chain of X.509 certificates - signature verification
local compliance extension(s)	runs in the same *local context* as the `Gaia-X Core Engine` (e.g., in the same K8S cluster as a microservice)
`proxy` (s)	accesses remote compliance extensions

3.5 Gaia-X Alignment

This section briefly explains how Gaia-X aligns with several other associations, initiatives, external projects, and standards and regulations.

Note that the vignettes contained in this section neither attempt to give a thorough technical evaluation nor propose a specific course of (technical) convergence. The set of initiatives, standards, or regulations covered in this release of the Architecture Document also does not indicate any specific preference but rather reflects the current state of discussion in the various Working Groups or Committees of Gaia-X.

3.5.1 Aligning with Other Associations and Foundations

3.5.1.1 IDSA

Gaia-X and the International Data Spaces Association (IDSA) deliver complementary technical frameworks to enable secure, interoperable data spaces. Their collaboration spans decentralized trust, protocol design, and governance, anchored in shared standards such as ISO/IEC 20151 (Dataspace concepts and characteristics).

The two complementary scopes of work are:

- Gaia-X: Decentralized digital trust framework
- IDSA: Dataspace Protocol and the specific design concepts for data spaces

Both associations support the creation of the global standard ISO/IEC 20151 Dataspace concepts and characteristics, which is the foundation for a common understanding of data spaces.

Semantic and technical interoperability are supported by the Dataspace Protocol and by the Gaia-X Technical Compatibility specifications. The Dataspace Protocol was initiated by the IDSA and is currently maintained and developed within the Eclipse Dataspace Protocol project associated with the Eclipse Dataspace Working Group, under the governance of the Eclipse Foundation (EF). It defines the steps for sharing data between parties, including policy language, the use of catalogs for datasets, contract negotiation, and the transfer process.

The Gaia-X Technical Compatibility specifications define the use and combination of standards to automate compliance verification, perform Policy reasoning, and manage Identity, Credentials, and Access. They also address the discoverability of Catalogues and Registries.

The organisational interoperability is enabled by the concepts defined by IDSA within the IDS RAM (Reference Architecture Model for governance and technical requirements) and IDSA Certification Scheme (compliance requirements for infrastructure components and services) and at the same time by the compliance criteria and related Labels defined by Gaia-X and addressing ICT services, data exchange services and ecosystem participants.

At a higher level, both Gaia-X and IDSA support the designing and building of data spaces, respectively via the establishment of an Endorsement Program selecting and supporting initiatives in line with the Gaia-X principles and architecture, and through the definition of the IDSA Rulebook, including legal, business, and operational guidelines for data spaces.

3.5.1.2 FIWARE

FIWARE Foundation is a non-profit organisation supporting the adoption of open standards (implemented using Open-Source technologies) that ease the development of smart solutions across domains such as Smart Cities, Smart Energy, Smart AgriFood and Smart Industry, based on FIWARE technology. The goal of the foundation is also to support the evolution of data platforms into data spaces, aligning with the strategy and the specifications developed by both IDSA and Gaia-X with regard to the creation of a trust layer.

FIWARE provides OSS components for:

- value creation, such as sectoral platform components and marketplaces (using Linked Data through NSGI-LD, in the development of which the foundation is involved, DCAT, and the TMForum Product Model).
- data exchange, with the FIWARE Dataspace connector, an integrated set of components in line with IDSA and Gaia-X technical specifications, meant to be deployed by organisations participating in a data space to connect to the data space.

3.5.1.3 Ocean Enterprise

Ocean Enterprise Collective e.V. is a German non-profit organisation with international members that develops and maintains a free open-source enterprise-ready dataspace ecosystem software solution that enables companies and public institutions to securely manage and monetize software & AI & data products and services in a trusted and compliant environment.

Domain agnostic and collectively governed by an independent non-profit association, Ocean Enterprise Collective e.V. is shaping a new transparent era of the data economy and is already being used by leading data-driven businesses in aerospace, agriculture, manufacturing, industry 4.0, mobility, smart cities and more.

Ocean Enterprise provides FOSS dataspace components, also part of the DSSC Toolbox, for:

- Value Creation: Marketplaces and smart contract-based catalogues and participant agent tooling.
- Data and Digital Service Exchange: Connectors and automated smart contract-based contracting services.
- Privacy and IP-Protection: Compute-to-Data orchestration enables technical data sovereignty and the exploitation of sensitive data without unnecessary replication.
- Monetization: Integration of e-money for real-time payment and settlement, besides usual postpaid mechanisms.

Ocean Enterprise is the basis for the Pontus-X dataspace ecosystem and several interoperable Gaia-X lighthouse dataspaces and endorsed projects. It has continuously aligned with the Gaia-X Trust Framework since 2021 and supported every version of the Trust Framework by integration of Gaia-X Participant and Service Credentials and Gaia-X Digital Clearing Houses (GXDCH).

The alignment process with Gaia-X now results in a native integration of the Gaia-X Loire standards, VC-JWT Credential format, and standards for credential exchange, and more.

The Ocean Enterprise Catalogue and FOSS framework enables distributed, tamper-proof, self-sovereign storage of data, services, and other offerings descriptions. Metadata records are stored as signed Verifiable Credentials utilizing Ocean Enterprise smart contracts. Metadata is openly extensible to support domain-specific descriptions and standards, such as DCAT, Gaia-X, and others.

Ocean Enterprise and the Pontus-X reference dataspace ecosystem are fully committed to supporting Gaia-X de facto standards as the basis for future interoperability between different dataspaces, technology stacks and solutions.

3.5.1.4 BDVA

The Big Data Value Association (BDVA) is an industry-driven, international not-for-profit organisation committed to developing an innovation ecosystem that facilitates the data-driven and AI-enabled digital transformation of Europe's economy and society.

The Association actively advances and promotes key areas such as big data technologies and services, data platforms and data spaces, Industrial AI, and data-driven value creation. Currently, there is a particular emphasis on data value creation in conjunction with AI, including the application of generative AI for data knowledge management and enhancing semantic interoperability.

BDVA collaborates closely with Gaia-X, focusing on the decentralisation of ecosystems and data spaces, as well as the integration of computing and data ecosystems.

3.5.1.5 iShare

The iSHARE Foundation was established to address the data handling and exchange needs of the Dutch logistics industry. iSHARE introduced a Trust Framework comprising legal, operational, and technical agreements designed to create favourable conditions for data sharing.
Within a data space or iSHARE network, participants must comply with these agreements, which are role-specific and vary according to each participant's function. To ensure consistent adherence, all organisations in the iSHARE network, including Data Owners, Data Providers, and Data Consumers, accept the same Agreements and Terms of Use. Through the federated Authorisation Registry, Data Owners can grant consent for specific data attributes to selected Data Consumers using licences.

The iSHARE Trust Framework provides a standardized approach to identification, authentication, and fine-grained authorisation, ensuring that only verified and authorised parties can access business data.

While the Gaia-X and iSHARE Trust Frameworks share common principles and objectives and can be used in parallel, additional efforts in architectural and technical alignment, particularly regarding the adoption of common standards, are required to achieve full technical compatibility.

3.5.1.6 X-Road

X-Road is an open-source software and ecosystem solution designed to facilitate secure and standardised data exchange between Public Administrations and between Public Administrations and the private sector. It provides a unified and secure data exchange layer between information systems within a collaborative ecosystem, ensuring confidentiality, integrity, and interoperability among data exchange parties.
X-Road was initially launched in Estonia in 2001, with Finland adopting it in 2014. In 2018, the two countries federated their X-Road ecosystems, establishing cross-border interoperability. As of 2024, X-Road is implemented in over 20 countries worldwide.

The Nordic Institute for Interoperability Solutions (NIIS), a non-profit organisation established in 2017 by Estonia and Finland, manages the development and maintenance of the X-Road core technology. NIIS oversees source code management, documentation, and facilitates the X-Road open-source community and adoption of X-Road in various jurisdictions. In 2022, Gaia-X and NIIS formalized their collaboration through a partnership.

An X-Road ecosystem comprises an X-Road Operator, Member organisations, and Trust Service Providers. The Operator manages the ecosystem and defines its regulations, policies and practices. Ecosystems can be national or limited to specific groups or domains. X-Road includes an authorisation framework for managing access rights based on organisation and service-level identifiers. Each service provider retains ownership of its data and is responsible for managing access rights to its services.

To connect with the emerging data space scenario and increase interoperability with other data exchange ecosystems and data spaces, X-Road is moving from X-Road-specific protocols to the data space protocol stack. In X-Road 8 "Spaceship", scheduled to release in 2026, the X-Road custom protocol stack will be replaced by the dataspace protocol stack, and the X-Road Trust Framework will be interoperable with the Gaia-X Trust Framework. All in all, the aim is to make X-Road technically compatible with the Gaia-X specifications and make X-Road interoperable with other Gaia-X data spaces.

3.5.1.7 NeoNephos Foundation

Launched by Linux Foundation Europe and originating under the EU's IPCEI-CIS initiative, the NeoNephos Foundation is an initiative focused on advancing open-source cloud infrastructure to address the demand for secure, scalable, and transparent cloud solutions aligned with European digital sovereignty goals. It is funded by the EU's NextGenerationEU program and supported by the German Federal Ministry for Economic Affairs and Climate Action.

While NeoNephos prioritizes open-source cloud infrastructure development, it shares Gaia-X's vision of advancing digital sovereignty and interoperable ecosystems.

Both initiatives align in fostering transparent, secure, and interoperable digital environments, with NeoNephos' open-source components (e.g., Kubernetes-based tools) potentially complementing Gaia-X's framework to support a multi-provider cloud-edge continuum.

3.5.2 Aligning with Other Initiatives

3.5.2.1 EuroStack (original)

> **Note**
>
> Currently, there are two initiatives carrying the name EuroStack. The original one, initiated in 2023, is described in this subsection. The "new" initiative (of the same name) split off the original one around the change of years 2024/25 and is described in the next subsection

The (original) EuroStack initiative (https://euro-stack.info//), ideated in late 2023, prepared in September 2024, and launched in the beginning of 2025 by releasing a 128 pages report, is an industry-driven effort to establish the continent as a leader in digital sovereignty by designing and building a sovereign, open, and collaborative digital ecosystem in Europe comprising the following 7 layers and concomitant key components:

Layer	Key Component
data and AI	DataCommons and SovereignAI
software	EuroOS
cloud	SovereignCloud
internet of things (IoT) and devices	SmartEurope IoT
networks	EuroConnect
chips	EuroChips
critical resources: raw materials, energy, and water	-

By integrating digital infrastructure into a cohesive framework, the EuroStack initiative ensures that Europe's Single Market remains robust and adaptive to 21st-century challenges. Complete self-sufficiency (aka "sovereignty") is neither feasible nor desirable in an interconnected world, but by building the capabilities and control necessary to protect its interests and those of its member states, Europe can create a resilient and at the same time competitive digital ecosystem that still benefits from global exchanges. The EuroStack initiative is more than a strategy for reducing dependencies: it is a forward-looking plan to build a thriving digital future for Europe.

EuroStack focuses on five core strategic actions:

1. Develop a European common digital stack
2. Deploy high-impact digital services (in the form of MVPs first)
3. Foster sovereign AI and federated data spaces
4. Lead in next-generation technologies
5. Scale innovation through "Europe first" procurement and strategic investments; establish a European Sovereign Technology Fund

The EuroStack Initiative represents Europe's ambition to achieve digital strategic autonomy through a total investment of € 300 billion over ten years. This effort proposes the creation of a European Sovereign Tech Fund, which includes an initial € 10 billion earmarked for the development of digital EuroStack demonstrators. These demonstrators – selected through an open competition known as the EuroStack Challenge – will serve as minimum viable products to showcase Europe's capacity to innovate and scale foundational digital technologies.

The Gaia-X AISBL has not joined the initiative yet.

Gaia-X and EuroStack share common values, such as data sovereignty, transparency, and interoperability. Specifically, Gaia-X is focused on enabling trust, while EuroStack is focused on the much broader goal reducing the EU's digital dependence on external providers while at the same time and with the same measures increasing Europe's competitiveness.

The two initiatives could benefit from collaboration, particularly by integrating the Gaia-X Trust Framework into the ecosystem/s formed by EuroStack services and operators, and by enriching the Gaia-X Trust Framework, for example, with regard to criteria for Gaia-X Labels, according to specific requirements from the EuroStack initiative.

3.5.2.2 EuroStack (spin off)

> **Note**
>
> Currently, there are two initiatives carrying the name EuroStack. The original one, initiated in 2023, is described in the previous subsection. The "new" initiative (of the same name) split off the original one around the change of years 2024/25 and is described in this subsection.

This (newer) EuroStack initiative (https://euro-stack.eu/) spun off the original EuroStack movement around the 2024/25 year change and released a Pitch Paper in January 2025 condensing the broad range and deep stack of its parent (cf. the 128 pages of the original report, Bria/Timmers/Gernone (2025): EuroStack – A European Alternative for Digital Sovereignty; Bertelsmann Stiftung. Gütersloh; see the pervious sub-section) into 19 pages focusing much more (but not exclusively) in IT infrastructure and services (XaaS).

It features only three layers (compared to the original seven):

Component	Description
Hard/Physical Infrastructure	Hard/Physical infrastructure requires investment of patient capital, targeted regional deployment to meet local needs, a strong research pipeline with networked universities, and public/private partnerships on licensing.
Soft/Logical Infrastructure	Soft/Logical infrastructure requires that we drive adoption by focusing on integrating components with attractive products that drive demand. We can sidestep the race to the bottom in data and energy with nimbler solutions designed to address users' needs rather than catering to tech fashion. The intention is to accelerate cloud and development ecosystems by creating governance and funding structures that share benefits with the Open-Source community.
Intermediation	Intermediation requires that we federate existing protocols and Open-Source implementations that require scaling and industrialisation. Building on prior experience elsewhere, we can develop and deploy Open Transaction Networks (OTNs) built atop a common core, across all domains of digital commerce and set up stakeholder governance for all intermediated services.

Within each layer, several components are identified:

- Hard/Physical Infrastructure
  - chips
  - data centers
  - connectivity
  - high-performance computing (HPC)
  - quantum technologies
  - supporting energy infrastructure
- Soft/Logical Infrastructure
  - identity
  - cloud
  - AI engines
  - browsers
  - operating systems
  - data spaces
- Intermediation
  - commerce
  - advertising
  - search & social

- app stores
- communications & productivity
- energy/green
- mobility

The constraints of capital and time render it impossible to build any alternative comparable to existing incumbents within a feasible timeframe. Given the urgency, EuroStack proposes to identify existing assets that can be integrated into federated networks, which are commercially and technically interoperable.

Nearly 500 European companies and organisations (as of May 2026) have signed an Open Letter to European Commission President Ursula von der Leyen and Executive Vice-President Hanna Virkkunen, supporting the EuroStack initiative and its proposed approach.

The Gaia-X AISBL did not join the initiative yet, while several of its members did (see the Open Letter).

Technically, the (new) EuroStack initiative and Gaia-X can, in principle, collaborate and complement each other's purpose on several fields including,

- enhancing trust and trust services for lowering compliance barriers
- interoperability from hyper-centralized to hyper-distributed
- identity where EuroStack capitalizes strongly on eIDAS 2.0
- data spaces

3.5.2.3 Data Space Adoption Forum (DSAF)

The Data Space Adoption Forum (DSAF) was established in the fourth quarter of 2025 with the mission to define, deploy, and deliver software and services for accelerating the adoption of data spaces at scale. It specifically targets the onboarding of SME members which – currently – is complex, expensive, and hard to scale but vital to any digital ecosystem.

DSAF is organized as a working group within IDSA and backed by Gaia-X, and CISPE, plus leading manufacturers and cloud providers such as Aruba, opiquad, nexyo, and other ecosystem providers and stakeholders (e.g., Catena-X, Sopra Steria, TNO, OPC Foundation, iSHARE, Microsoft, and several others).

The Forum is preparing initial onboarding deployments in the Catena-X automotive data space, with the ambition to onboard 10,000 companies in the next 18 months.

Technically, DSAF is currently working on providing a virtualized version of the Eclipse Dataspace Components project called EDC-V. Somewhat simplified, EDC-V features a tenant-based separation of the *Control Plane* from the *Data Plane* supporting one Control Plane governing several Data Planes. This provides the foundation for a scalable multi-tenant deployment of the EDC framework. Before EDC-V, managed service providers (MSPs) had to deploy a full set of all EDC components (which is a lot) for each and every single customer they want to serve.

Gaia-X is currently working on adding an OID4VC and Gaia-X compliance extension to the EDC framework with the strategy to move that to the core software base of the EDC project. This will enable all data spaces to simply select these two options when configuring their individual distribution of the EDC framework. Participating in DSAF is integral part of this strategy and secures strong and continuous alignment with the larger data space adoption community.

3.5.3 Aligning with External Projects

3.5.3.1 DSSC

The Data Spaces Support Centre (DSSC) is an initiative funded by the European Commission under the Digital Europe Programme. Its primary objective is to facilitate the development of common European data spaces that collectively establish a sovereign, interoperable, and trustworthy data-sharing environment.

The Data Spaces Support Centre (DSSC) is an initiative funded by the European Commission under the Digital Europe Programme to support the development and adoption of European data spaces. The first phase of the initiative (DSSC1, October 2022 - March 2026) produced key assets such as the DSSC Blueprint, providing a reference framework for the design and operation of interoperable data spaces based on shared concepts and reusable building blocks.

Since May 2026, the second phase of the initiative (DSSC2) has been advancing the adoption, operationalisation, and interoperability of data spaces across Europe, while supporting the broader objectives of the European Data Union and the integration of data spaces with emerging AI ecosystems.

Gaia-X participated in DSSC1 and continues its involvement as a consortium partner in DSSC2, contributing expertise in trust infrastructures, identity and attestation management, compliance mechanisms, and interoperability. The Gaia-X Trust Framework is closely aligned with the Data Sovereignty and Trust pillar of the DSSC Blueprint, and several Gaia-X services are included in the DSSC Toolbox. More concretely, the Gaia-X Association will directly engage in the following DSSC2 activities and work packages:

- direct technical expert support for selected data spaces from the set of Common European Data Spaces
- maintenance and further development of the technical Building Blocks of the DSSC Blueprint
- contributions to the DSSC Toolbox
- interfacing and aligning with data sharing communities, liaisons, standards developing organizations (SDOs), and collaborations
- communication and dissemination
- definition and hosting of online and offline events including the yearly Data Spaces Symposium

3.5.3.2 Simpl

Simpl is a programme commissioned by the European Commission to a consortium of private organisations to develop an open-source middleware platform to support data-sharing and service interoperability. The programme consists of three products:

- Simpl-Open (the open-source software development);
- Simpl-Lab (a playground environment for Simpl-Open to test interoperability);
- Simpl-Live (the adoption of Simpl-Open in specific data space implementations, related to the Common European Data Spaces).

Since the first public releases and architectural demonstrations presented in 2025, Simpl has continued to evolve its architecture and technical capabilities, supporting the implementation and interoperability needs of European data spaces.

Gaia-X AISBL is regularly aligning and exploring synergies with Simpl through multiple channels, including the Gaia-X Business and Technical Committees and through collaboration activities within the Data Spaces Support Centre (DSSC). The Simpl-Open project is slated to officially end in December 2026.

3.5.3.3 DOME

The Distributed Open Marketplace for Europe (DOME) project establishes a catalogue of trusted cloud and edge services spanning multiple domains.

DOME employs Verifiable Credentials to assert that all service providers within its marketplace are trustworthy and adhere to EU legislation. Notably, DOME has introduced the "LEARCredential"—used to delegate rights to employees—and the "Verifiable Certification," which attests to compliance with specified certifications.

Recognizing the shared objectives of DOME and Gaia-X in promoting transparency and compliance within service offerings to foster a trusted ecosystem, efforts are underway to facilitate the integration of Gaia-X-compliant services into the DOME Marketplace. This integration targets Participants who satisfy DOME's specific requirements.

Concurrently, initiatives are in progress to harmonize compliance criteria for providers adhering to EU legislation, particularly concerning the Label levels defined by Gaia-X. Furthermore, DOME is considering the adoption of Notaries as specified by Gaia-X.

Lastly, ongoing discussions are focusing on the evolution of business models for the forthcoming DOME Operator, responsible for managing the primary instance of the DOME Marketplace, and the Gaia-X Digital Clearing Houses (GXDCHs).

3.5.3.4 IPCEI-CIS and the 8ra Initiative

The IPCEI-CIS (Important Project of Common European Interest – Next Generation Cloud Infrastructure and Services) is an EU initiative aimed at developing a European interoperable and openly accessible multi-provider cloud-to-edge continuum.

IPCEIs are a framework guided by the European Commission, designed to support large-scale, collaborative projects that significantly contribute to the EU's strategic objectives. Their innovative aspect lies in permitting State Aid and fostering cooperation across Member States.

Specifically, the IPCEI-CIS encompasses 19 companies (each leading one project) from seven Member States—France, Germany, Hungary, Italy, the Netherlands, Poland, and Spain—as well as 90 ecosystem partners. These 19 direct projects are united under the umbrella of the 8ra Initiative, which focuses on:

- Cloud-edge infrastructure: edge nodes, clouds, and components to interconnect,
- Cloud-edge capabilities, e.g., software to operate cloud-edge resources,
- Advanced data processing tools and services, e.g., toolkits for software developers,
- Advanced applications and customer use-cases, e.g., autonomous driving cars.

The 8ra Initiative can leverage the Gaia-X Trust Framework to establish an organisational and governance foundation for the large-scale digital transition, enhancing interoperability and upholding key values of the initiative, such as environmental sustainability and digital sovereignty.

3.5.4 Aligning with Standards and Regulations

3.5.4.1 European Digital Identity Framework Regulation (eIDAS, eIDAS 2.0)

The European Digital Identity Framework Regulation ("eIDAS 2.0"; Regulation (EU) 901/2014) updates and strengthens the rules for electronic identification and trust services ("eIDAS", Directive 1999/93/EC) across the EU internal market, mandating mutual recognition of qualified trust services and the publication of Member State Trusted Lists alongside the Commission's LOTL (List of Trusted Lists).

Under eIDAS 2.0, each Member State must offer a notified European Digital Identity Wallet (EUDI Wallet) that lets citizens and businesses authenticate at a high-security level and store user-controlled attestations of attributes. Qualified Electronic Attestations of Attributes (QEAA) are issued by certified Trust Service Providers, who expose interfaces for mutual authentication with EUDI Wallets and for verifying attestations against Authentic Sources. Non-qualified Electronic Attestations of Attributes follow the same conceptual model but rely on an appointed body to define the attribute schemas and Trust Architecture—collectively the "Attestation Rulebook."

Gaia-X's mission as a trust framework for data spaces maps directly onto these concepts, offering a governance layer for issuing and validating Verifiable Credentials (VC) and managing Trust Anchors and trusted issuers.

Both Gaia-X and eIDAS 2.0 share a reliance on:

- Trust anchors and published issuer registries
- User-centric wallets for secure authentication and attribute presentation
- A rulebook or specification that governs credential issuance and verification

By aligning Gaia-X Credentials (as an Electronic Attestation of Attributes service) with the EUDI Wallet architecture and the eIDAS attestation rulebooks, Gaia-X can ensure seamless cross-ecosystem interoperability.

Gaia-X is exploring how to integrate with eIDAS 2.0, for example by performing mutual authentication between EUDI Wallets and the Gaia-X Trust Framework, defining a Gaia-X Attestation Rulebook for non-qualified attributes, and publishing a Gaia-X registry of trusted issuers conformant with eIDAS specifications.

On the technical level, the following features are worth noting:

- The current implementation of an eIDAS 2.0 conformant EUDI Wallet (see Architecture and reference framework, ARF) and also Gaia-X use the OID4VC specification for credential exchange
- Gaia-X's notion of Trust Service Provider (TSP) is compatible with the same notion of eIDAS 2.0 in the sense that an eIDAS 2.0 TSP may also be appointed as a Gaia-X TSP should an ecosystem Governance Authority choose to do so.

3.5.4.2 CEN/CENELEC Joint Technical Committee 25 (JTC 25)

The CEN/CENELEC Joint Technical Committee 25 (JTC 25) on "Data Management, Dataspaces, Cloud & Edge" develops European standards supporting interoperable and trustworthy data ecosystems. Within JTC 25, Working Group 2 (WG2 – Dataspaces) is developing the Trusted Data Transactions (TDT) (pr)EN 18235 multi-part European standard (correctly called a *European Norm*, EN), covering terminology, concepts and mechanisms, trustworthiness requirements, and interoperability requirements in support of the European Data Act. As of 15 June 2026, the status of the three parts is as follows.

Number	Name	Status
EN 18235-1	Trusted data transactions — Part 1: Terminology, concepts and mechanisms	officially published
FprEN 18235-2	Trusted data transactions – Part 2: Trustworthiness requirements	at harmonized standards assessment (HAS); final vote (FV) in 07/26, publication ca. 10-11/26
prEN 18235-3	Trusted data transactions - Part 3: Interoperability requirements	in ENQUIRY phase at the National Standards Bodies; incorporation of comments ca. 09-11/26, publication 1H'27

Gaia-X became an official *Category A liaison* to CEN/CENELEC JTC 25 in June 2025 which, after the former CEN/CENELEC Workshop Agreement on Trusted Data Transactions (CWA TDT) came to an end, is prerequisite to be able to work on standardization projects if one is not a delegate of a national standards body (NSB). Any liaison organisation, though, is not allowed to vote on new standards. This is the prerogative of the CEN/CENELEC members, the NSBs like AFNOR, DIN, ASI, NEN, UNE. Building on its previous participation in the CWA TDT, Gaia-X in particular contributes expertise on trust frameworks, policies, claims, and evidence, and aligns its architecture with emerging European standards for trustworthy and interoperable data ecosystems.

3.5.4.3 Eclipse Foundation - CAP

Within the Eclipse Dataspace Working Group (EDWG), Gaia-X actively contributes to the Eclipse Conformity Assessment Policy and Credential Profile (CAP), an open-source initiative providing a common model for conformity assessment schemes, credentials, attestations, and trust-related policies in data spaces and digital ecosystems.

CAP complements ongoing European standardisation efforts, including the Trusted Data Transactions (TDT) work within CEN/CENELEC JTC 25 WG2, by providing practical implementation building blocks for trust and conformity assessment mechanisms.

4 Gaia-X Trust Framework Architecture

Trust
As per ISO/IEC DIS 20151-1 as well as CEN/CENELEC EN 18235-1, trust refers to a decision by an entity "to assume that a product, service or entity will behave as expected for a given circumstance".

The Gaia-X Trust Framework architecture consists of two elements: - The Gaia-X Trust Protocol provides the generic means to allow ecosystems to establish mutual trust in a standardized interoperable way. - The Gaia-X Trust Framework provides the means to validate the assumptions underlying trust based on sets of criteria and rules defined by governance authorities for an individual ecosystem or data space.

In the following, we first describe the Gaia-X Trust Framework and subsequently introduce the Gaia-X Trust Protocol, which generalises the trust concepts and mechanisms used across ecosystems.

4.1 Gaia-X Trust Framework

The Gaia-X Trust Framework comprises the technical and organisational standards and open-source software for its operationalisation, allowing the establishment of trust in digital ecosystems.

It comprises two complementary elements:

- Specification of Gaia-X technical compatibility at the level of technical standards with an interoperable trust architecture, an open-source reference implementation (Reference Software Toolkit), and a corresponding compatibility test (Technical Compatibility Kit (TCK)).
  This specification includes blueprints for a digital ecosystem or data space authority to define its own governance framework and rulebook. See DSSC Blueprint v3.0: Data space governance framework/Data space rulebook.
- Specification of Gaia-X Compliance: A definition of specific criteria that ICT services and participants must fulfill and are evaluated by a Gaia-X Digital Clearing House (GXDCH), to be characterised as Gaia-X compliant with "Gaia-X Standard Compliance" and/or a "Gaia-X Label Level" for currently three label levels.
  This specification includes an extension mechanism ("Gaia-X Geography and Domain Extensions") allowing certain organisations, so-called *custodians*, to define additional Gaia-X "labels" and compliance regimes.

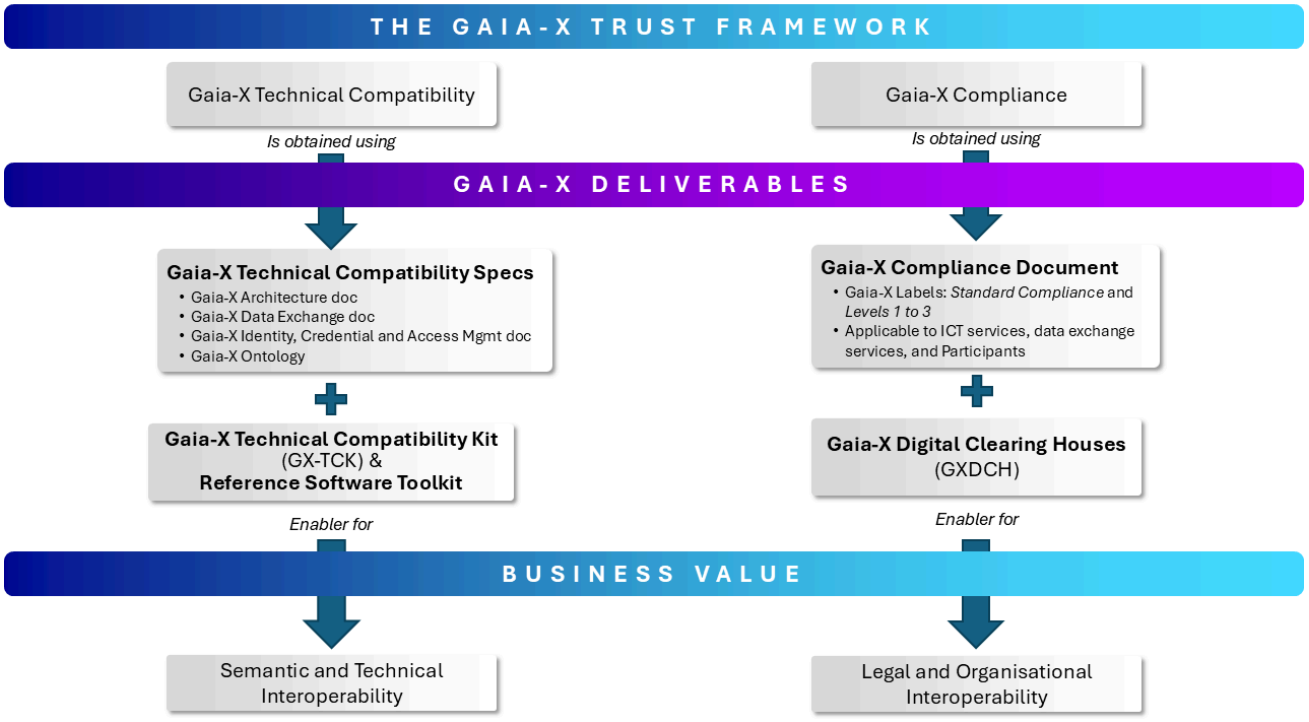

Figure 4.1 - The Gaia-X Trust Framework

Trust service providers (TSPs) such as the Gaia-X Digital Clearing Houses (GXDCH) are approved by ecosystem governance authorities and operationalize ecosystem compliance rules. Requiring these TSPs to provide their services conforming to Gaia-X technical compatibility ensures that attestations issued by the TSPs are easily acceptable by all ecosystem participants without any need for intermediation or translation.

Gaia-X Technical Compatibility is also a key enabler for technical interoperability at the trust level across different ecosystems with different rule sets (e.g. European data spaces versus Asian trade corridors) and constitutes an integral part of the Gaia-X Trust Protocol (see later in this chapter).

The Gaia-X Framework offers key advantages in the following areas:

Digital Ecosystems:

- Validates ecosystem compliance against established criteria
- Verifies that enabling and federation services meet ecosystem governance criteria, ensuring trust in these services

Individual Participants:

- Supports self-determination in data and service usage by collecting and verifying credentials demonstrating compliance with ecosystem rules
- Enables credential verification during negotiations between participants

Interoperability:

- Promotes agreement on common ontologies and trust service providers across ecosystems

Extension and Delegation:

- Provides mechanisms for extending and delegating trust to enhance ecosystem flexibility

4.2 Using the Trust Framework

The sections in this chapter describe some relevant applications of the Trust Framework.

> Digital Clearing House (DCH)" vs "Gaia-X Digital Clearing House (GXDCH)
>
> Note that unless a component name is prefixed with "Gaia-X", it refers to the rule-agnostic version of that component rather than the Gaia-X-specific implementation used to operationalize Gaia-X Compliance.
>
> Example: The Gaia-X Compliance Engine implements the Gaia-X Compliance criteria, and *a* Compliance Engine can implement non-Gaia-X-related criteria.

4.2.1 Performing automated onboarding and offboarding

Trust Frameworks can be used to automatically provide validation that all criteria of the onboarding process for participants and services, as defined by the ecosystem governance authority, are met.

The ecosystem governance authority defines criteria for onboarding, examples include:

- Gaia-X participant credential
- Identity credential issued by a regulatory authority (e.g. eIDAS in Europe)
- Ecosystem-specific criteria (e.g.):
  - Signature of a legal framework contract
  - Payment of the membership fee
- Attestations of conformity to specific standards or regulations
- Verifiable credentials issued by Trust Service Providers for which mutual trust has been declared

Proof of validation of criteria can be achieved through various methods:

- using the Gaia-X Trust Framework to validate Gaia-X Compliance with the help of the Gaia-X Technical Compatibility
- using the "Gaia-X Technical Compatibility Specifications" to validate ecosystem criteria
- Using an alternative Trust Framework (e.g. eIDAS, iShare,⋯) which provides an interoperable credential format.

> Acceptance of validations across trust service providers

Note that alignment on credential formats and credential exchange protocols allows mutual acceptance of validations across trust service providers.

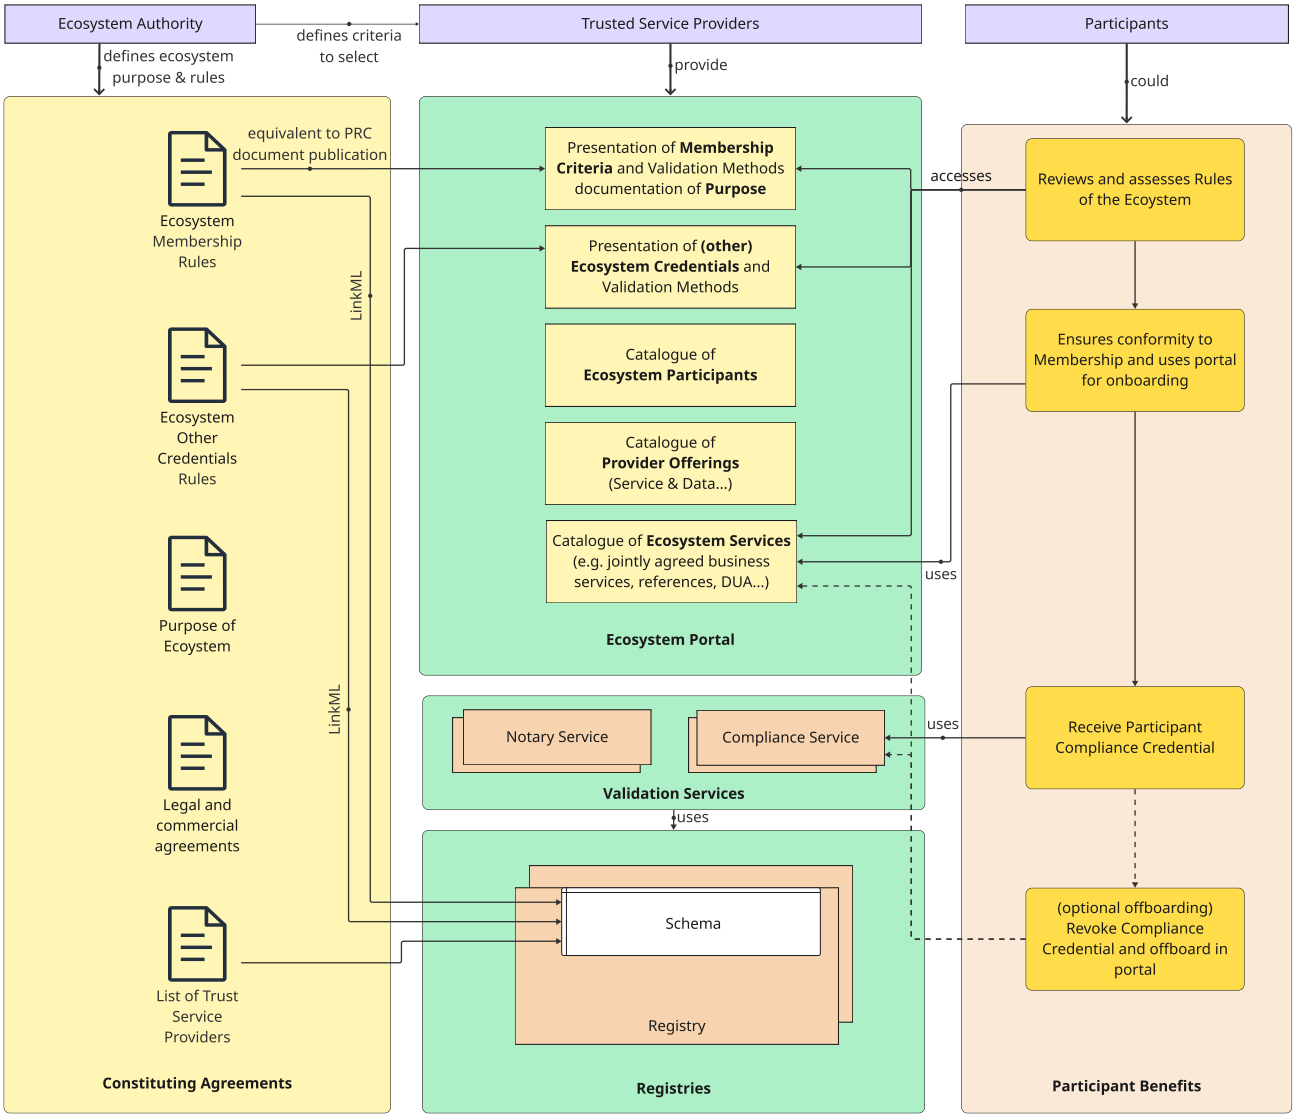

Figure 4.2.1 - Automated Onboarding

Use of the Trust Framework for automated onboarding requires:

- A specific schema that contains the criteria and the accepted trust service providers
- An API that makes the individual credentials linked to the participant or service available to the compliance engine evaluating the onboarding criteria
- In general, a Wallet (e.g. EUDI Wallet, managed wallets by ecosystem providers, individual organizational or personal wallets) providing the credential store that interacts via defined Credential Exchange Protocols (see the Appendix)
- A compliance engine evaluating the defined onboarding schema

Successful verification of the onboarding criteria results in the issuance of an ecosystem membership credential, which can be used to authorize participants to operate within the ecosystem.

The ecosystem governance authority may set specific constraints. For example:

- Credentials may have a specific validity period and the membership credential's validity may be dependent on a rule set evaluating the different validity end dates
- A credential may include or be associated with one or more roles, enabling the holder to perform authorized operations within the ecosystem
- The compliance engine may periodically re-evaluate (verify) the individual credentials (e.g. check for revocations) and revoke an issued membership credential before the end of the original validity period (automated offboarding)

Prerequisites for the automated onboarding to work are:

- Interoperability of the Credential formats used by the different Trust Service Providers
- Interoperability of the Credential Exchange format between Credential Issuer, Wallet and Compliance Engine
- Interoperable Credential Verification methods

Please also refer to Chapter 3.4 for considerations on how compliance with ecosystem rules can be established by combining verifiable credentials issued by different Trust Service Providers.

4.2.2 Performing credential verification

Credential verification is executed by the Compliance Engine in conjunction with the Registry.

The ecosystem or data space defines credential schemas that specify the vocabulary for expressing claims about credential subjects. These schemas are represented as SHACL shapes, in accordance with the W3C Shapes Constraint Language (SHACL).

Credential schemas are published and maintained in the Registry to ensure consistent access and reuse.

At any point where Verifiable Credentials are issued or processed, the applicable SHACL shapes are known and constitute the shapes graph. Each Verifiable Credential constitutes a data graph. Verification is performed by validating the data graph against the shapes graph, as defined by the SHACL specification.

4.2.3 Identifying Ecosystem trust services

The discoverability of Ecosystem Trust Services can be enhanced using a catalogue of trustworthy registries following the Gaia-X Technical Compatibility specifications.

By facilitating the identification of ecosystem trust services, also enabled by the Gaia-X Trust Protocol, federation and cross-ecosystem integration of trust are supported.

Examples of ecosystem trust services are credential stores, Auditing and Observability services, Marketplaces, Authorization and Access services, together with standard services provided by Identity and Catalogue Providers.

4.3 GXDCH

The Gaia-X Digital Clearing House (GXDGH) is the mechanism adopted by the Gaia-X Association to provide running software to assert compliance without becoming a host or a point of centralization for the ecosystem.

Each GXDCH instance must be operated by a service provider according to rules defined and approved by the Gaia-X Association.
The instances are non-exclusive, interchangeable, and operated by multiple market operators from different geographical locations and different industries.

The GXDCHs are connected in a network to ensure both free access and selection by Gaia-X adopters, and consistency of the compliance data managed by these nodes.

Each GXDCH must provide public services to implement the compulsory elements necessary to achieve Gaia-X compliance under the sole governance of the Gaia-X Association. Furthermore, each GXDCH can offer (or resell) services to support the extended adoption of Gaia-X (out of the governance of the Gaia-X Association).

The mandatory components to be included in the GXDCH can vary from one release to another, as more services become available over time. The updated list of components is available here.

All the components that go into the GXDCH are open-source, either reused or developed by the Gaia-X Association.

The following page displays, for each Gaia-X Digital Clearing House, the version of its components along with their health status (UP or DOWN).

4.4 Gaia-X Conceptual Model

4.4.1 Terminology sources

The terminology employed in this section is informed by various sources, including the following recognised standards:

ISO/IEC 17000:2020 – Conformity Assessment — Vocabulary and General Principles This international standard, developed by the ISO Committee on Conformity Assessment (CASCO), provides standardized definitions and general principles related to conformity assessment activities, including the accreditation of conformity assessment bodies.

Verifiable Credentials Data Model v2.0 – World Wide Web Consortium (W3C) This specification defines a data model for expressing verifiable credentials on the Web in a cryptographically secure, privacy-respecting, and machine-verifiable manner. It outlines key concepts such as credentials, presentations, and associated roles like issuers, holders, and verifiers, facilitating interoperability and trust in digital credential ecosystems.

4.4.2 Definitions

Term	Definitions	Additional Information	Source
Accreditation	third-party attestation related to a conformity assessment body, conveying formal demonstration of its competence, impartiality and consistent operation in performing specific conformity assessment activities	N/A	ISO/IEC 17000:2020
Attestation	issue of a statement, based on a decision that fulfilment of specified requirements has been demonstrated	N/A	ISO/IEC 17000:2020
Certification	third-party attestation related to an object of conformity assessment, with the exception of accreditation	N/A	ISO/IEC 17000:2020
Conformity assessment	demonstration that specified requirements are fulfilled	N/A	ISO/IEC 17000:2020
Conformity Assessment Body (CAB)	a Conformity Assessment Body (CAB) is a body that performs conformity assessment activities, excluding accreditation	N/A	ISO/IEC 17000:2020
Conformity Assessment Scheme	is a set of rules and procedures that describes the objects of conformity assessment, identifies the specified requirements and provides the methodology for performing conformity assessment	The Conformity Assessment Scheme generates Label/Criteria based on the following conformity rules: List of permissible standards, Technically verifiable attributes (e.g., GLEIF) and Rulesets	ISO/IEC 17000:2020
Consumer	A Consumer is a Participant who searches Service Offerings and consumes Service Instances in the Gaia-X Ecosystem to enable digital offerings for End-Users.	Providers and Consumers within the ecosystem are identified and well described through their valid Credentials, which are initially created before or during the onboarding process. Providers define their Service Offerings and publish them in a Catalogue. In turn, Consumers search for Service Offerings in Gaia-X Catalogues that are coordinated by Operators and the Gaia-X Registry. Once the Consumer finds a matching Service Offering in a Gaia-X Catalogue, the Contract negotiation between Provider and Consumer determines further conditions under which the Service Instance will be provided. The Gaia-X Association does not play an intermediary role during the Contract negotiations but ensures the trustworthiness of all relevant Participants and Service Offerings.	N/A
Declaration	first-party attestation	N/A	ISO/IEC 17000:2020
DUA Notary	A Data Usage Agreement (DUA) Notary is a specialization of a Notary validating the existence of a legally binding (e.g., signed by both parties, not revoked) Data Usage Agreement (DUA) between a Data Producer and a Data Consumer.	Data Producers have to submit a DUA to a DUA Notary; Data Consumers then may (or, if an ecosystem Governance Authority mandates it, have to) check with the respective DUA Notary before a data access whether the DUA is still valid. See Chapter 5 for more details.	N/A
Entity	item relevant to the operation of a domain that has recognizably distinct existence	N/A	ISO/IEC 24760-1:2019
Evidence	can be included by an issuer to provide the verifier with additional supporting information in a verifiable credential.	This could be used by the verifier to establish the confidence with which it relies on the claims in the verifiable credential. It is expected that the credentials issued by the notaries contain the evidence of the validation process.	W3C Verifiable Credentials Data Model v2.0
Governance Authority	an ecosystem or data space Governance Authority (GA) is the body of a particular data space or ecosystem, consisting of participants, that is committed to the governance framework for the data space or ecosystem, and is responsible for developing, maintaining, operating and enforcing the governance framework	N/A	DSSC Glossary
Issuer	a role an entity can perform by asserting claims about one or more subjects, creating a verifiable credential from these claims, and transmitting the verifiable credential to a holder.	N/A	W3C, Verifiable Credentials Data Model v2.0

Term	Definitions	Additional Information	Source
Label	is a machine-readable, structured and signed document issued by the ecosystem-accredited Compliance services in case of a valid verification and validation of the criteria for a specific assessment scheme	The criteria for Gaia-X Labels are defined in the Gaia-X Compliance Document. Label / Criterion is validated by the following validation methods: Self-declaration, Validation by code, Conformity Assessment Body, and Existing credential/certificate	Gaia-X Compliance Document
Notary	Notaries are entities (not necessarily a Participant) approved by the Governance Authority, which perform validation based on objective evidence from Trusted Data Sources, digitalizing an assessment previously made.	The Notaries are converting “non-machine-readable” proofs into “machine-readable” proofs, i.e., verifiable credentials.	N/A
Object of conformity assessment	entity to which specified requirements apply	N/A	ISO/IEC 17000:2020
Orchestrator	An orchestrator is an entity within a digital ecosystem which is providing a certain set of services the ecosystem deems necessary or relevant.	These services may be of a strategic/business or technical nature and may be distributed to different orchestrators. These services facilitate the interplay of participants in a data space, enabling them to engage in (commercial) data-sharing relationships of all sorts and shapes. They can perform an intermediary role in the data space. Examples: (1) GXDCH instances are non-exclusive, interchangeable, and operated by multiple market operators from different geographical locations and different industries. (2) Trusted Service Operator (TSO) are Gaia-X Providers that have been approved by the ecosystem governance authority to authoritatively provide one or more services to the ecosystem. (3) Data space intermediary - A data space participant that provides one or more enabling services while not directly participating in the data transactions itself.	N/A
Participant	entity that is onboarded and has a Gaia-X Participant Credential. A Participant can take on one or more of the following roles: Provider, Consumer, and Operator.	N/A	N/A
Provider	a Provider operates Resources in the Gaia-X Ecosystem and offers them as services through Gaia-X Service Offering credentials.	For any such service, the Gaia-X Provider defines the Service Offering, including terms and conditions as well as technical policies. Furthermore, it provides the Service Instance, which includes a Credential and associated policies. A Gaia-X Provider is responsible for conformity with the claims made. If third-party data, services or infrastructure are part of the service offerings, the Gaia-X Provider can make those resources available (e.g., through Service Composition) but remains responsible and shall ensure coverage through appropriate back-to-back coverage.	N/A
Trust	Trust refers to a decision by an entity to assume that a product, service or entity will behave as expected for a given circumstance.	N/A	ISO 20151 and CEN/CENELEC prEN 18235-1
Trust Service Provider (TSP)	a Trust Service Provider is an entity (not necessarily a Participant) approved by the Governance Authority to authoritatively issue verifiable credentials for a given scope and purpose	- is itself identified by a Trusted Digital Identity, - issues (verifiable) credentials/certificates, - some TSPs act as Trusted Identity Providers and issue digital identities in this capacity, in case a Conformity Assessment Body (CAB) is capable of issuing verifiable credentials for a given scope and purpose, it also qualifies as a TSP. TSPs may use validation-by-code	N/A
Trusted Data Source	is a source of the information used by the issuer to validate attestations.	Notaries perform validations and issue attestations based on objective evidences from Trusted Data sources. The accepted Trusted Data Source and Notaries are determined within the Gaia-X Compliance document, while the detailed list of valid Trusted Data Sources and Notaries resides in the Gaia-X Registry.	N/A
Trusted Service Operator (TSO)	Trusted Service Operators are Providers that have been approved by the ecosystem governance authority to authoritatively provide one or more services to the ecosystem.	There can be one or more Trusted Service Operators for a given type of service, e.g., a catalogue service. The ecosystem participants therefore recognize the Operator as a trusted provider of that specific service.	Do not mix Trusted Service Operators with Trust Service Provider (TSP). TSO provides services you can trust, whereas TSP provides core elements of trust (in the form of verifiable credentials) for the ecosystem, e.g., identity or compliance credentials. See the Model Core later in this section.
Validation	confirmation of plausibility for a specific intended use or application through the provision of objective evidence that specified requirements have been fulfilled	N/A	ISO/IEC 17000:2020

Term	Definitions	Additional Information	Source
Verification	confirmation of truthfulness through the provision of objective evidence that specified requirements have been fulfilled	N/A	ISO/IEC 17000:2020

Note

the full list of the most relevant terms used in the Gaia-X deliverables is available in the Gaia-X Glossary. For terms defined within it, the Gaia-X Architecture Document remains the authoritative source. The Gaia-X Glossary is updated periodically to reflect the latest versions of Gaia-X deliverables.

4.4.3 Model Core

The diagrams describe the organisations and concepts foundational to the translation of the rules in the conformity assessment schema defined by the governance authority into digital credentials and their subsequent verification.

The (Ecosystem/Data Space) Governance Authority:

- Defines one (or multiple) conformity assessment schemes, which are translated into a machine-readable format and stored in the Registry
- Approves CABs as a validation method for specific scopes and purposes
- Approves Trust Service Providers (TSPs) for given scopes and purposes
- Selects and approves Trusted Digital Identities where a Trusted Digital Identity is bound to a cryptographic key
- Approve Notaries for given scopes and purposes

The Conformity Assessment Body:

- Acts as a validation method
- Has a scope
- May use a Notary
- If it can issue verifiable credentials, it will also act as a Trust Service Provider

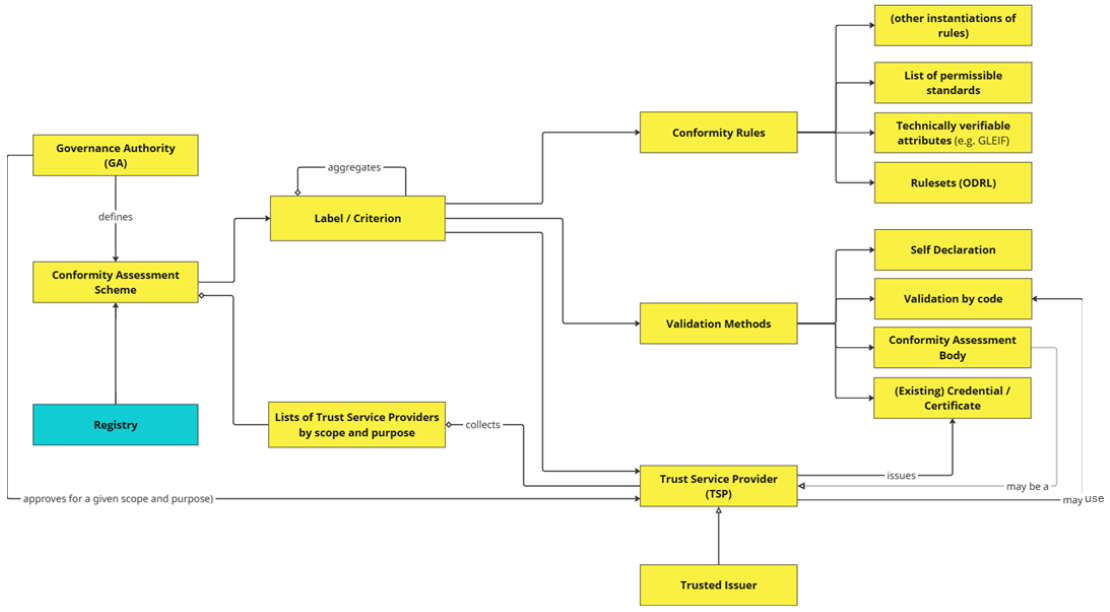

Figure 4.4.3 (a) - Conceptual Model-1

The definition of compliance criteria (e.g. for participants and services offered in the ecosystem/data space) encompasses both the definition of conformity rules and the definition of required validation methods. The definition of accepted validation methods may involve the definition of existing standards that can be used to assess conformity, or the identification of attributes that can be technically verified, and the definition of the type of assessment required (e.g. declaration or certification performed by a Conformity Assessment Body). While establishing a conformity assessment schema, the Governance Authority also defines the list of Trust Service Providers that are approved as trusted issuers for specific scopes and purposes.

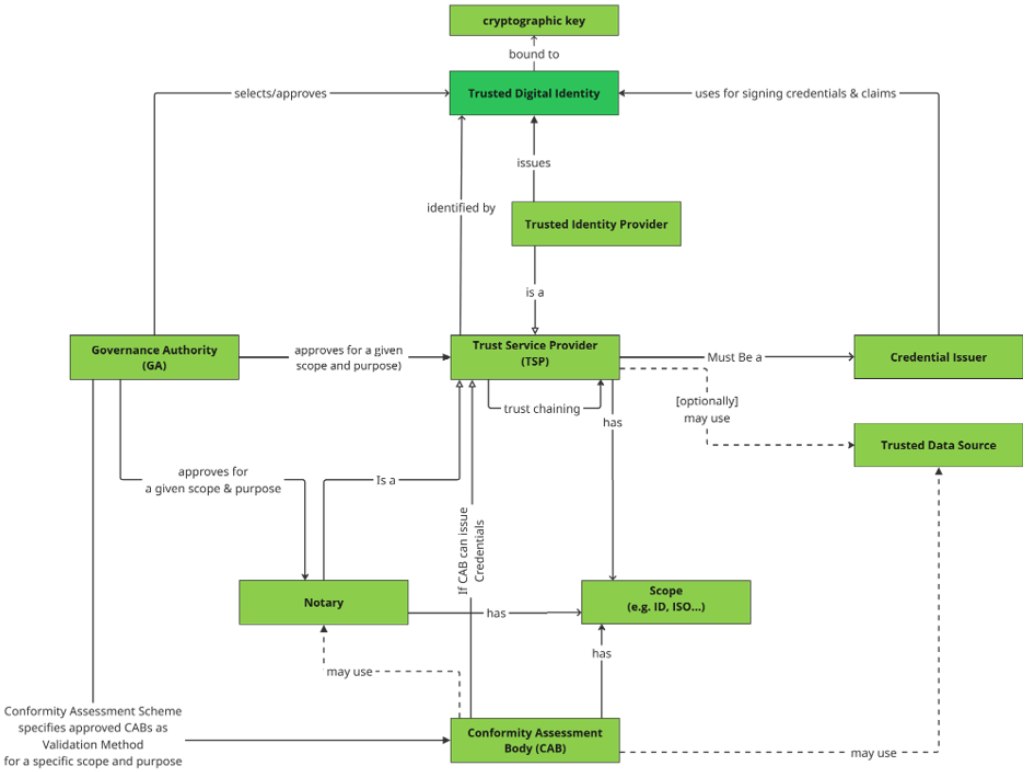

Figure 4.4.3 (b) - Conceptual Model-2

From an implementation perspective, the Participants issue claims on themselves and the services they provide according to the format defined in the Registry. These claims are signed by approved Trust Service Providers, and the validation of compliance criteria is performed by the Compliance Engine based on the information (credential schemas/shape graphs) that are maintained and made available through the Registry. The Notary service supports the issuance of credentials.

4.5 Gaia-X Trust Protocol

4.5.1 Problem Statement

Trust frameworks – including the Gaia-X Trust Framework described above – are typically defined by the ecosystem or data space governance authority for a specific ecosystem and its participants. Due to the interlinked nature of today's supply networks, however, current experiences and expectations indicate that a single organization participating in one data space or ecosystem will have to exchange data (and access services) simultaneously with many different ecosystems and their participants.

As ecosystems are completely autonomous when specifying their trust framework, the question then arises how these sovereign ecosystems can establish mutual trust with other ecosystems on behalf of their participants without any explicit bilateral or overarching governance. (This is often called the cross ecosystem trust dilemma.)

The Gaia-X Trust Protocol provides a generic way for an ecosystem governance authority

1. to define common rule sets for the ecosystem,
2. to define the methods to validate those rules,
3. to describe the methods to issue and verify digital attestations by which trust service providers attest conformance to these rules, and
4. to make these accessible to other ecosystems and their participants in a standardized and interoperable way.

It specifies a reference model to render trust-related communication interoperable between different ecosystems [3] without forcing ecosystems to accept anything from others. By complying with the Gaia-X Trust Protocol, different ecosystems may chose to accept digital attestations issued by one or more trust service providers of another ecosystem based on a common specification[4].

The Gaia-X Trust Framework is a specific instantiation of the Gaia-X Trust Protocol. In this case, the Gaia-X Policy Rules Committee as governance authority defines

- the rule sets for Gaia-X compliance (i.e., as stated in the Gaia-X Compliance Document with its associated schemas, validation rules and permissible standards) and
- the trust service providers (i.e., the Gaia-X Digital Clearinghouses, GXDCH).

4.5.2 Gaia-X Trust Protocol Technical Specification

4.5.2.1 Ecosystem Trust Profiles

The Gaia-X Trust Protocol is based on the premise that ecosystems need to be able to decide for themselves whether they trust (entirely or even just partially) the trust framework of another ecosystem or not. In order to allow the ecosystem to actually reach this decision, the "trusting" ecosystem (called *trustor*) needs to be able to obtain and understand the relevant elements of the other "trusted" ecosystem (called *trustee*). This is realizable in a scalable way when each *trustee* makes available the core elements of its individual trust framework in a common "language". Potential *trustors* then are able to access and understand a foreign trust framework and base their decision to trust (parts of this) trust framework on this information.

Based on our work with several ecosystems, the smallest set of technical concepts required to allow this decision process to take place comprises a set of triples with the following elements:

No	Element	Description
1	trust scope	range or characteristics of the objects of the conformity assessements covered by all credential types collected within this trust scope. The trust scope is a slight generalization of the term scope of attestation of ISO/IEC 17000:2020.
2	credential types	List of types of verifiable credentials (using W3C Verifiable Credentials Data Model 2.0) for a given trust scope which are *accepted* by ecosystem participants as being *equivalent*. The notions *acceptance* and *equivalence* of a conformity assessment result are taken from ISO/IEC 17000:2020.
3	trust service providers (TSPs)	List of TSPs the ecosystem accepts for validation and issuing of verfiable credentials for a given trust scope.

We shall call one such triplet a trust proposition and the list of all the trust propositions (= triplets) for a specific ecosystem an ecosystem trust profile.

4.5.2.2 Example

For the Gaia-X Trust Framework itself, the Gaia-X ecosystem trust profile recognizes the following elements.

Ecosystem Trust Profile Element	Gaia-X Trust Framework Values
trust scope	The Gaia-X Trust Framework recognizes participant and service credentials: - `gx:ecoParticipantScope` - `gx:ServiceAttestationScope`
credential types	Currently, the Gaia-X Compliance Document only recognizes a single (verifiable) credential type for both trust scopes, namely `gx:LabelCredential`
TSPs	The accredited Gaia-X Digital Clearing Houses (GXDCH). All GXDCHs are equal in the sense that Gaia-X participants will trust *any* individual GXDCH. In ISO/IEC 17000:2020 terms this means that the conformitity assessment results of the GXDCHs are *equivalent* and *accepted* by all Gaia-X participants.

The full version of the Gaia-X ecosystem trust profile may be retrieved from the Gaia-X Meta-Registry here.
The ontology subset for specifying an ecosystem trust profile is part of the Gaia-X ontology and defined here.

As is evident from the example, the Gaia-X Trust Protocol defines the generic means for linking the (verifiable) credentials of arbitrary content to different validation methods (e.g. as defined in ISO/IEC) using various ecosystem-specific criteria and verification methods.

The implementation of these elements is supported by the Gaia-X ICAM Semantic Model (ref. ICAM Document), which provides the semantic constructs required to describe trust scopes, credential types, trusted issuers, and trust relations in a consistent and interoperable way.

4.5.2.3 Establishment of Mutual Trust

Mutual trust, where one ecosystem (the *trustor*) commits itself to trust another ecosystem (the *trustee*), is then simply established by having the *trustor* – in full autonomy – select and accept a suitable subset of the TSPs of the *trustee* using some particular credential types for one or more trust scopes. The *trustor* will then simply add these "triplets" to its own ecosystem trust profile. Participants of the *trustee* ecosystem will then be able to recognize that specific credentials issued by specifics TSPs for a trust scope are also approved by participants of the *trustor* ecosystem.

4.5.2.4 Gaia-X Meta-Registry

In order to allow different ecosystems and, especially, their participants, to easily identify potential areas of mutually shared trust, all trust propositions (the triplets) of the ecosystem trust profile adopted by a particular ecosystem are collected into and published through a dedicated Gaia-X Meta-Registry.

Retaining their autonomy, ecosystems may publish their own trust propositions to the Gaia-X Meta-Registry. Other ecosystems participants will then be able to identify which of their own TSPs for a specific trust scope are at the same time also accepted by another ecosystem they want to exchange data with.

The Gaia-X Meta-Registry specifications and implementation will be further extended and refined in future releases based on requirements gathered from ecosystems and adopters. This includes governance, publication and validation processes, provenance and authenticity of published artefacts, and trust establishment mechanisms.

The current version of the Gaia-X Meta-Registry can be accessed here. Instructions on how to integrate an ecosystem trust profile into the Gaia-X Meta Registry are available here.

4.5.2.5 Gaia-X Trust Protocol Architecture

The resulting cross-ecosystem trust framework architecture is depicted in the following figure. The critical observation here is that, since ecosystems and data spaces may vary in their individual compliance rules at the intra-ecosystem trust layer, successful establishment of cross-ecosystem trust requires consensus on a common Gaia-X technical compatibility layer (as specified in this Architecture Document) as well. Without it, transport and syntactical interoperability of the trust propositions shared between different data spaces cannot be automatically ascertained. Compliance with the Gaia-X technical compatibility specification is assured by conforming to a Gaia-X technical compatibility (test) kit (indicated as "Gaia-X TCK" in the figure below). This can be attested by a suitable "Gaia-X technical compatibility" credential (forthcoming).

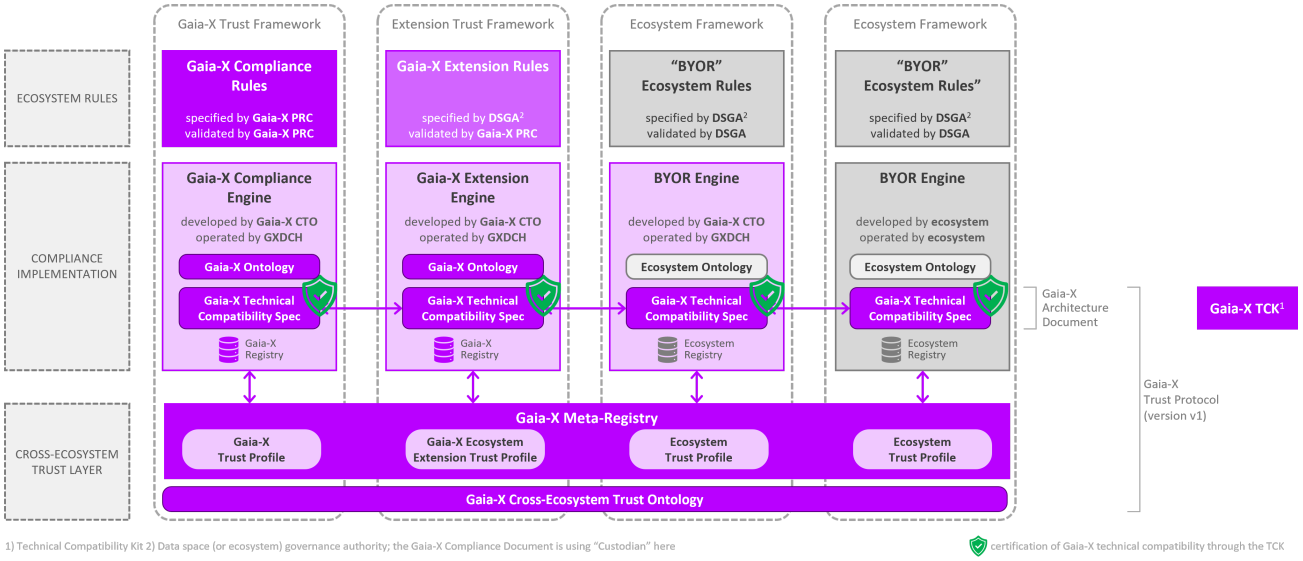

Figure 4.5.2.5 - Gaia-X Trust Protocol

4.5.2.6 Evolution of the Gaia-X Trust Protocol

Evidently, the Gaia-X Meta-Registry conflicts somewhat with the assumed autonomy of ecosystems. While the Gaia-X Meta-Registry does not constitute any form of supra-ecosystem governance mechanism, its (current) logical and organizational centralization within the Gaia-X Association is at odds with the independence of the individual ecosystems submitting their trust profiles to the Gaia-X Meta-Registry through the Gaia-X Association.

Recognizing this we suggest replacing the Gaia-X Meta-Registry with a suitable Gaia-X Trust Registry Protocol enabling different ecosystems to exchange information contained in their respective (ecosystem trust) registries (see the figure below).

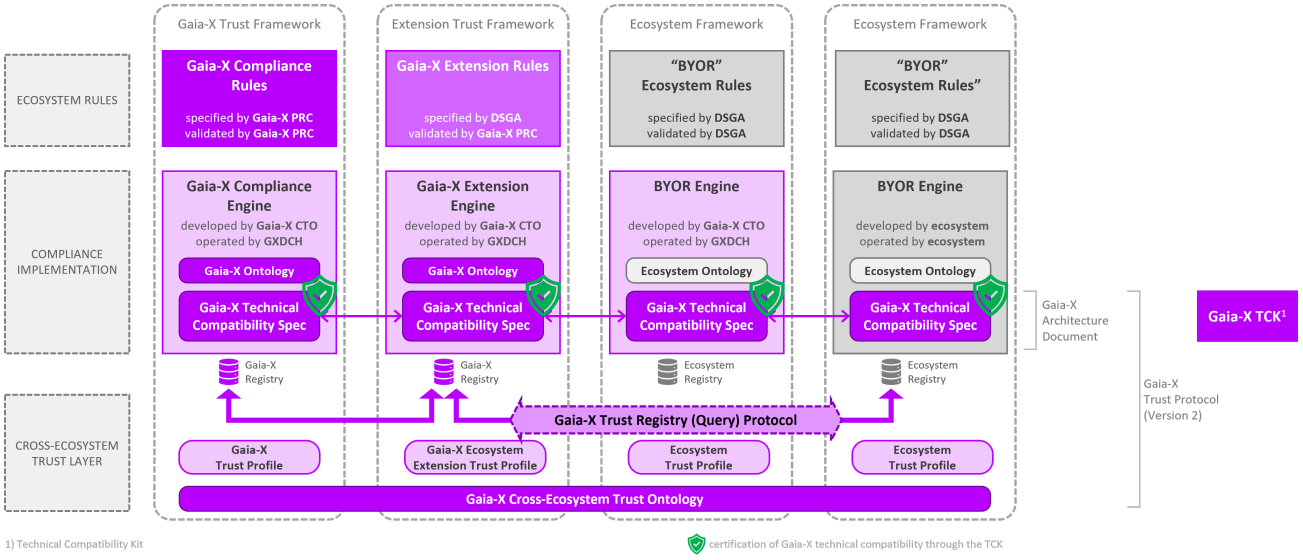

Figure 4.5.2.6 - Gaia-X Trust Protocol - evolution

However, from an adoption-oriented point of view we believe that the Gaia-X Meta-Registry is a much needed intermediate step towards deploying the Gaia-X Trust Protocol version 2. Our experiences during several events indicate that ecosystems very much appreciate the concrete embodiment of the abstract concept of ecosystem trust profiles in the form of a browsable catalog, i.e., the Gaia-X Meta-Registry. During further adoption of cross-ecosystem trust using the Gaia-X Trust Protocol version v1, ecosystems will gradually start to understand and realize the (even more abstract) solution to this centralized element present in version v1 of the Gaia-X Trust Protocol.

4.5.2.7 Model for Federated Ecosystems

The Gaia-X Trust Protocol establishes the foundation for interoperable trust across different ecosystems and domains. It enables ecosystems to federate their trust frameworks by adhering to this common model. Conformance with the Gaia-X technical compatibility specification (as captured in this Architecture Document) then allows any ecosystem to make use ("operationalize") the trust information contained in the Gaia-X Meta-Registry.

As a mechanism for federation between two ecosystems, the governance authority of Ecosystem *A* shall agree on the equivalence of the scope of specific properties between the individually defined schema, and ecosystems *A* and *B* shall mutually agree on the Trust Service Providers (and their validation methods) for the respective scope.

Mutually trusted services may be published in joint catalogues.

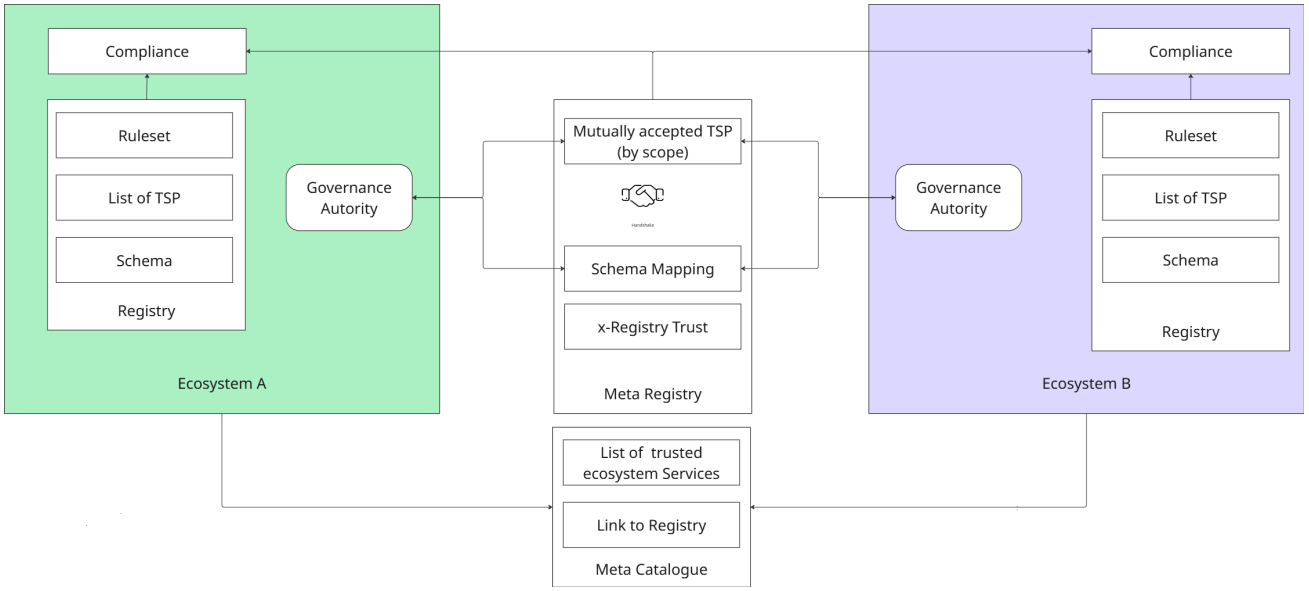

Figure 4.5.2.7 - Ecosystem Federation

4.6 Cross-Ecosystem Interoperability

Cross-ecosystems interoperability can be defined as “the ability of participants to seamlessly access and/or exchange data and services across two or more ecosystems.

This covers the cases where a use-case requires implementation across domains and geographies, across isolated environments (e.g. a factory shopfloor or a hospital), and an integration of data across domains to provide an integrated view of a product:

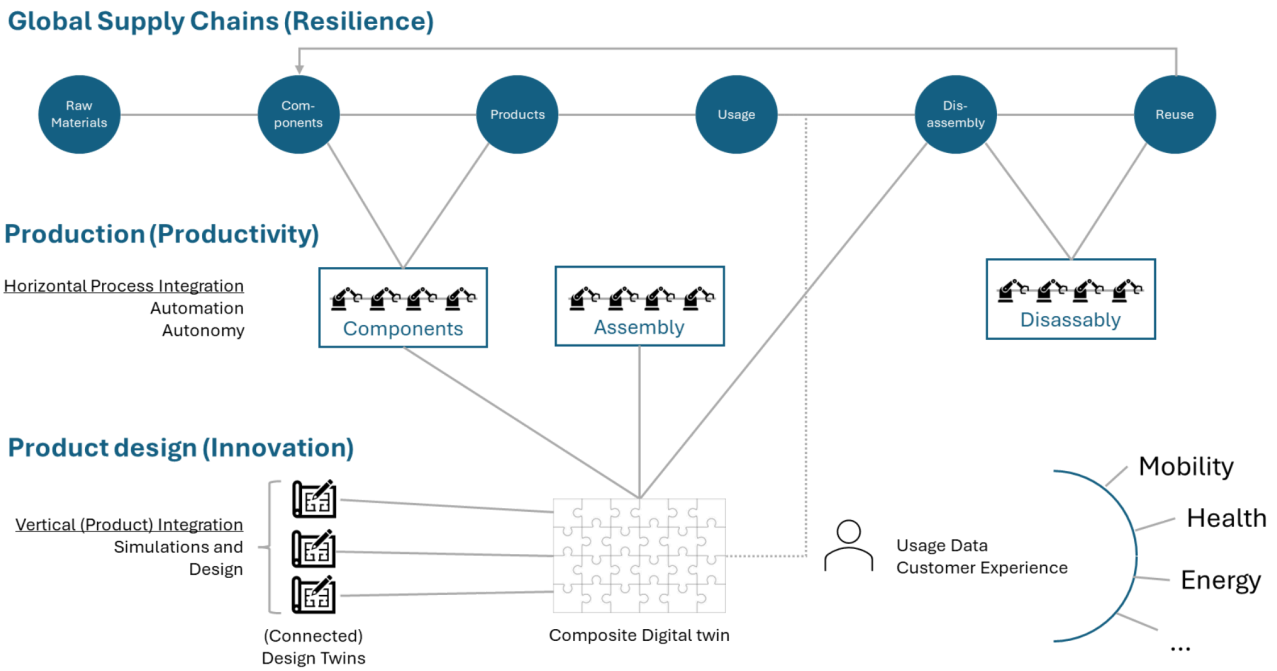

Figure 4.6.a - Cross-Ecosystem Interoperability - Use case scenarios

Cross-ecosystems interoperability addresses the governance, business and technical frameworks to interconnect multiple ecosystem instances seamlessly.” (adapted from the definition of “cross-data spaces interoperability” in the Glossary of the DSSC Blueprint v3.0).

This includes:

- trust in services that exchange data across ecosystems
- trust in services that are provided by a single ecosystem, but relied upon by other ecosystems
- mutual acceptance of identities

The Meta Registry defined in the Gaia-X Trust Protocol provides a common and interoperable way to publish and access ecosystem trust profiles and related information.

Mutually recognized Trust Service Providers and the capability to accept validations across ecosystems serve as enablers for cross-ecosystem interoperability.

Interoperability across ecosystems is further supported by adopting the Gaia-X Trust Framework, which provides specifications facilitating both technical and semantic interoperability. It also offers multiple sets of compliance criteria that can be selected, combined, and extended by ecosystem governance authorities to meet their specific needs.

Finally, the Eclipse Conformity Assessment Policy (CAP) Ontology, which provides a standardised framework for expressing and verifying conformity assessment policies, represents another important asset for interoperability in data-sharing ecosystems.

4.7 Inter-Ecosystem Interoperability

Trust Frameworks in general - not just limited to the Gaia-X Trust Framework - intend to increase interoperability at the following four layers (ref. European Interoperability Framework):

- legal interoperability: specifying criteria related to jurisdictions or domain- or ecosystem-relevant legislation (e.g., GDPR, Data Act, Data Governance Act)
- organizational interoperability is enhanced by better aligning data transactions with (i) business processes and their (user) requirements and (ii) with common domain or ecosystem governance rules
- semantic interoperability is supported by ensuring that the meaning of the exchanged information is preserved throughout exchanges between parties (for instance, by the definition of semantic models and common ontologies)
- technical interoperability derives from the use of international standards widely adopted and recognised

In the approach chosen by Gaia-X, these four layers are addressed in a two-pronged configuration:

- Gaia-X Compliance addresses legal and organisational interoperability (of trust)
- Gaia-X Technical Compliance addresses semantic and technical interoperability (of trust)

Depending on the particular configuration of an ecosystem's trust framework, an ecosystem may also implement certain elements of organisational and legal interoperability using Gaia-X technically compliant mechanisms. For instance, an ecosystem may extend the Gaia-X core ontology for ecosystem participants to include certain credentials (e.g., a *Business Partner Number*) allowing the automated recognition of an organisation as a valid ecosystem participant; such an extended ontology may also include credentials specifying certain roles a participant may hold within an ecosystems, e.g., service provider or service consumer.

4.8 Services and Service Composition

4.8.1 Resources and Service Offerings

Resources describe, in general, the goods and objects of the Gaia-X Ecosystem.

Resource Categories

A Resource can be a:

- Physical Resource: it has a weight, a position in space and represents a physical entity that hosts, manipulates, or interacts with other physical entities
- Virtual Resource: static data in any form and necessary information such as a dataset, configuration file, license, keypair, an AI model, neural network weights, and so on
- Instantiated Virtual Resource: an instance of a Virtual Resource. It is equivalent to a Service Instance and is characterised by endpoints and access rights.

A Service Offering is a set of Resources, which a Provider aggregates and publishes as a single entry in a Catalogue.

A `Service Offering` can be associated with other `Service Offerings`.

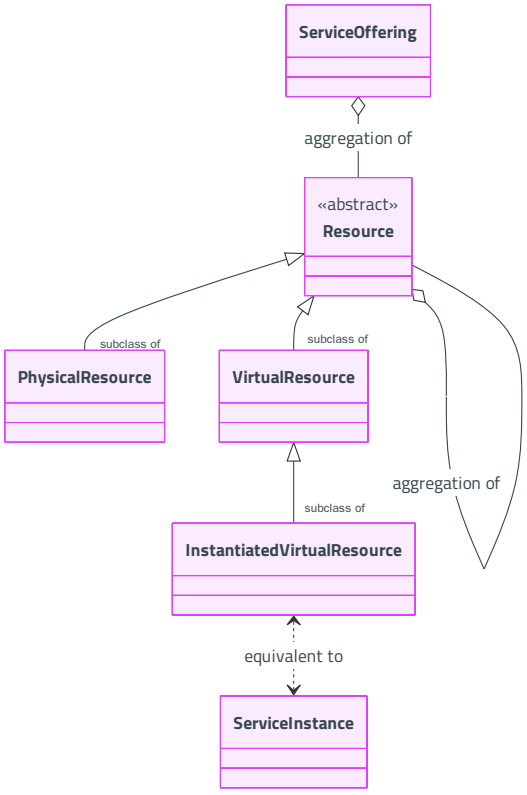

Figure 4.8.1 - Resources and Service Offerings

Services can be selected and composed across different ecosystems which rely on a common Trust Framework. This approach enables the creation of new service offerings that meet specific compliance requirements, as these are satisfied by the constituent services.

For example, a software provider might offer a service that utilizes multiple Cloud Service Providers, each of which complies with the requirement to use only data centers certified under ISO/IEC 27001.

4.9 Policies

4.9.1 Policy Definition

A Policy is a statement of objectives, rules, practices, or regulations governing the activities of participants within a given ecosystem or data space. From a technical perspective, Policies are statements, rules or assertions that specify the correct or expected behaviour of an entity[12].

The Gaia-X Compliance Document defines the general Gaia-X Policies applicable to Participants and Service Offerings. These policies cover, for example, privacy and cybersecurity requirements and are expressed indirectly as attributes of Resources, Service Offerings, and Service Instances.

4.9.2 Policy Description

The general Policies defined by Gaia-X form the basis for detailed Policies for a particular Service Offering, which can be defined additionally and contain particular restrictions and obligations defined by the respective Provider or Consumer. They occur either as a Provider Policy (alias Usage Policies) or as a Consumer Policy (alias Search Policy):

- A Provider Policy/Usage Policy constrains the Consumer's use of a Resource. *For example, a Usage Policy for data can constrain the use of the data by allowing to use it only for x times or for y days.*
- A Consumer Policy describes a Consumer's restrictions on a requested Resource. *For example, a Consumer gives the restriction that a Provider of a certain service has to fulfil demands such as being located in a particular jurisdiction or fulfilling a certain service level.*

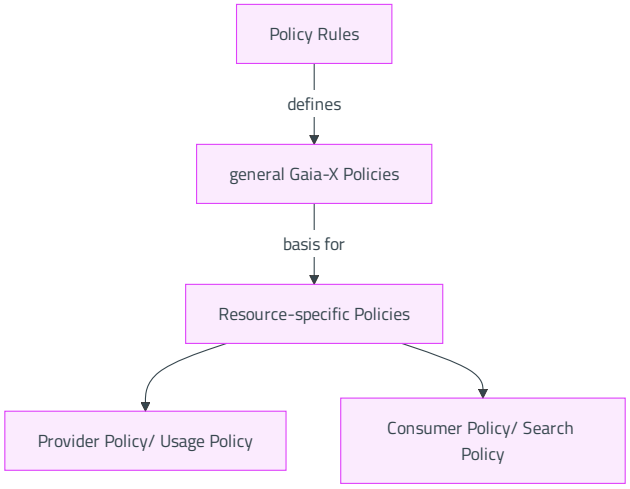

Figure 4.9.2 (a) - Policy Description

`Policies` are expressed by one or more `parties` about one or more `assets`. An `asset` can also be a `party`, making this simple model recursive and capable of handling even the most complicated usage scenarios, with multiple providers, consumers, federators, data intermediaries, data subjects, business legal representatives, employer/employee relationships, and much more.

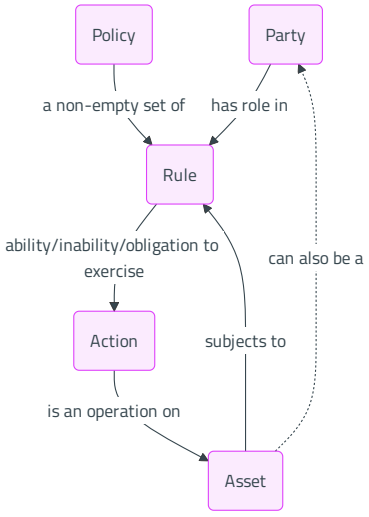

Figure 4.9.2 (b) - ODRL Workflow model

The workflow above is the generic representation of the Open Digital Rights Language (ODRL) model.

The main aspects to be considered from a technical point of view are:

- Policy representation: interoperable access and usage policies that are specified in a human and machine-readable format. Policies generally express three possible restrictions: prohibitions, obligations, and permissions. Constraints defining a rule can be combined into more complex rules, which then form the applicable policy.

- Decision-making/policy engine: during the execution of a data transaction, the policies need to be evaluated. This decision typically requires context information. With the decision context, the policy engine will decide whether the request or usage is permitted. The evaluation process is handled by the policy engine, which is instantiated by the data product provider and the data consumer or a trusted third party.
- Enforcement and execution of policies: the enforcement and execution of policies is a key capability which needs to be implemented for both access and usage policies.

Realization of Policy Definition and Policy Enforcement follows the W3C ODRL specifications; the definition of the execution components follows NIST (see PDP and PEP)

4.9.3 Gaia-X Policy Reasoning Engine

The Gaia-X Policy Reasoning Engine allows for a comparison between policies set up by the provider of a service and usage intentions declared by a consumer, by leveraging the following standards:

- W3C ODRL (Open Digital Rights Language) to express policies
- W3C Verifiable Credentials
- JSON Path to evaluate the credentials
- RDF to represent the policies as triples (subject, predicate, object)
- SPARQL to query the RDF triples

An open-source library to perform policy reasoning is provided by the Gaia-X Lab.

4.9.4 Policy Decision Point (PDP)

4.9.4.1 Reference implementation

A specific ODRL profile is built by the Gaia-X Lab with the purpose to refer in a clear and precise way to verifiable credential claims in an ODRL Policy. This gives assignors of policies a way to enforce policies using trustworthy and verifiable claims from an assignee, and in doing so, having more trust and confidence in the enforcement of the policy.

4.9.4.2 Policy Enforcement

The Policy Enforcement Point (PEP) intercepts the request from the data product provider. To verify the request, the decision context is set up and processed through the Policy Decision Point (PDP). After retrieving the relevant context information, the data request is verified. If the request is invalid, the data product provider sends a rejection, which is processed by the data consumer. If the request is valid and certain actions are required, these actions are executed before granting data access or starting the transfer. The response is then processed by the data consumer. If the request is valid, the same interception and processing steps as on the Data Product Provider side are executed (involvement of PAP, PDP, PEP, PIP). In the last phase, the data consumer provides proof of the implemented policy enforcement, which the data product provider verifies. In case of an invalid proof, the data product provider can and must reject the interaction and may initiate further actions. If everything is valid, the transaction is complete.

The enforcement phase starts when the actual data transaction is executed, and it takes place throughout the data transaction. The goal of the enforcement phase is to evaluate the relevant rules of the policy and decide whether the data transaction may proceed unless an agreement or contract has already been negotiated, in which case, only the contract's validity must be verified.

For different types of rules, evaluation may occur at different stages of the data transaction:

- Access rules are primarily evaluated by the data provider before any data is exchanged. The data consumer may also check these rules when receiving data to ensure compliance with access conditions.
- Usage rules are evaluated by the data consumer when the data is used. Depending on the usage rule, evaluation might be required at initial data usage and constantly throughout the lifecycle of the data usage.
- Consent rules require the evaluation of consent of third parties, which might be revoked during the data transaction or data use. Therefore, evaluation should occur on both the data product provider and consumer sides.

4.10 Ecosystem Trust Functions

In addition to the core elements, trust frameworks typically contain supplementary trust-related functions that accomplish or facilitate more specific tasks. Here, we explain two functions included in this version of the Architecture Document, namely:

- means for rights or trust delegation and consent management
- computation of indexes providing interoperability metrics and information on the potential trustworthiness of an entity or other ecosystem element.

4.10.1 Rights Delegation

The delegation of rights takes place within a Trust Scope, which defines the scope in which credentials, trusted issuers, identity providers, and associated semantic definitions are considered valid. It also identifies the trusted issuers, vocabularies, and schemas that participants use to interpret and validate credentials within that scope.
The rights are delegated based on the type of credential and its usage.
The Party Credential is based on the Verifiable Credentials Data Model v2.0 and is the basis for all IAA Parties such as Natural Persons, Services, and Legal Persons. The general purpose Party Credential can be extended into specialised credentials, for example:

- Private Party Credential: Party Credentials containing PII are not considered published or reachable via their identifier by everyone. Instead, they are intended to be stored in secure storage such as a wallet, secure storage device, secure vault storage, etc. An example of this type of credential is the NaturalPersonCredential, issued by a Legal Participant to one of his users/employees for the purpose of authorising a Natural Person to interact with Relying Parties (RP) in a certain context.
- Public Party Credential: The Party Credential contains data that can be publicly accessed and queried.
- Trust Scope Credential: This credential enables the usage of Party Credentials by defining their accredited issuers. Furthermore, the Trusted Scope Credential supports the cooperation and interoperability between organizations/ecosystems/data spaces, by easing the use of external Trust Service Providers.

Credential Statuses

A Verifiable Credential can have one of the following statuses:

- expired if the `validUntil` attribute is older than the current datetime or the certificate containing the key used to sign the claim has expired.
- revoked
- if the keypair used to sign the credential is revoked.

- if the credentialStatus has the statusPurpose property set to “revocation” and the value of status at position credentialIndex is true
- suspended if the credentialStatus has the statusPurpose property set to “suspension” and the value of status at position credentialIndex is true
- deprecated if another verifiable credential with the same identifier and the same signature issuer has a more recent issuance datetime.
- active only if none of the above conditions apply.

For more details, refer to the ICAM document.

4.10.2 Trust Indexes

Traditional assessment schemes defined by ecosystem authorities often face limitations:

- They do not automatically adapt to market dynamics, which typically evolve faster than established rules
- They fail to capture the complexities of organizational and semantic interoperability, effectively reducing interoperability to the most basic commonly accepted denominators.

To address these challenges, four Trust Indexes are introduced:

- veracity
- transparency
- composability
- semantic match

These indexes are designed to be used by all stakeholders — licensors, licensees, producers, providers, and consumers — as metrics for assessing interoperability and trust relative to other offerings within ecosystem catalogues. The intention is for the ecosystem to assist the market by providing measurement tools, enabling market participants to converge towards optimal solutions.

A more detailed explanation of the four indexes and the way to compute them is provided in an Appendix to the document. (see Appendices - Trust Indexes)

4.11 Data Space Architecture using the Gaia-X Trust Protocol

Data Spaces can leverage the Gaia-X Trust Protocol to onboard participants by verifying compliance with their Data Space Rulebook and issuing corresponding proofs of membership and roles to facilitate trusted data exchange.

As illustrated in the figure below, the Data Space Governance Authority defines the Data Space Rulebook, aiming to operationalize regulatory, business, and/or technical requirements.

The Data Space Rulebook must specify, in a machine-readable format, the Criteria, Compliance Rules, Validation Methods (via first-, second-, or third-party assessments, potentially relying also on technical means), and the accepted Trust Service Providers authorized to sign the required claims.

The Registry stores the lists of accepted Trust Service Providers along with the schemas for data space credentials, which serve as the vocabulary for claims about credential subjects (e.g., data space offering credentials and data space membership credentials). The Rules Engine validates data graphs against the shape graphs stored in the Registry to assess compliance with the Data Space Rulebook.

The assessment outcome is issued as a credential under the authority of the Data Space Governance Authority. Participants can store this credential using a Wallet Service (also referred to as a Credential Store) and use it as proof of compliance or for further qualification under a data space Conformity Scheme.

The setup of the Data Space is completed and can be enhanced by utilizing additional Trust Services accessible through a catalogue adhering to the Gaia-X technical compatibility specifications, along with protocols designed to enable contract negotiations and trusted data transactions.

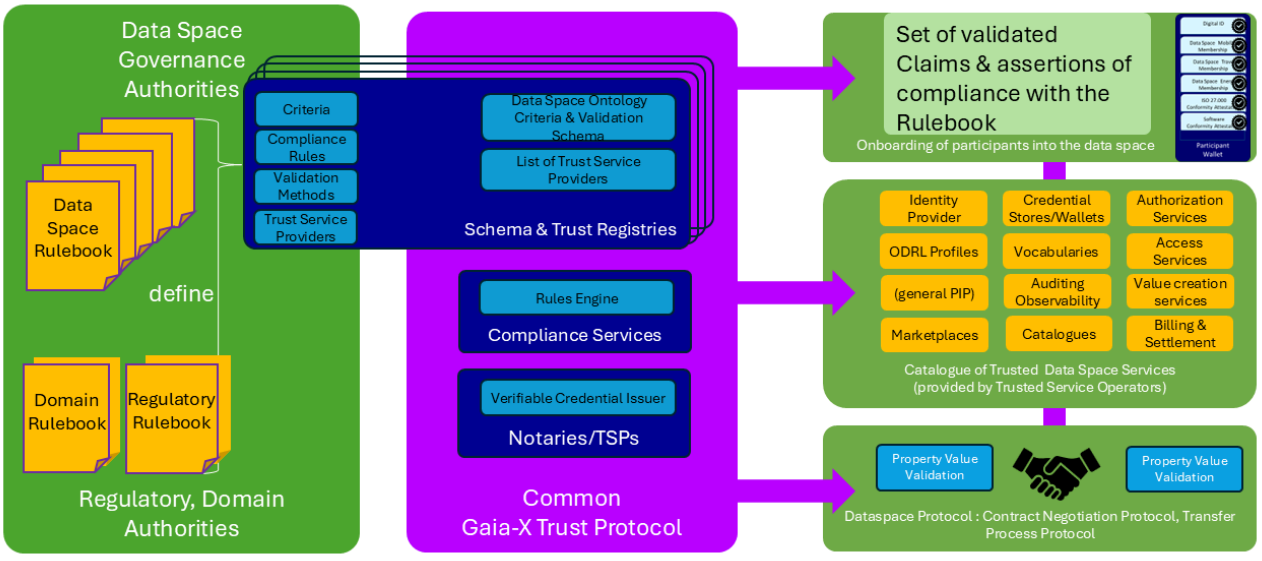

Figure 4.11 - Data Space architecture enabled by the Gaia-X Trust Protocol

4.12 Trusted Artificial Intelligence Environments

One scenario which combines the elements defined above is that of an environment where artificial intelligence (AI) services operate on controlled data sets and infrastructures. The following example illustrates how the concepts and mechanisms described in this chapter can be applied to support trust, compliance, and interoperability in AI-enabled environments.

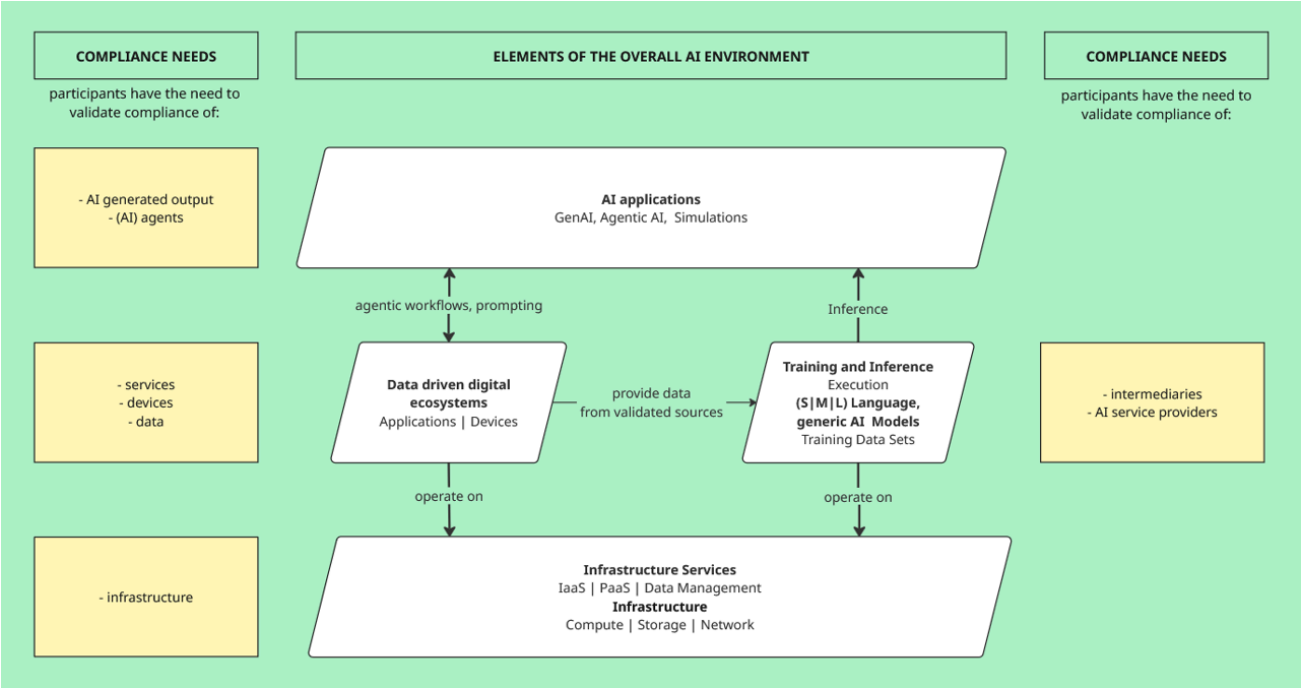

Figure 4.12 - Scenario of a controlled AI environment

4.12.1 Compliance of the data used in AI use cases

All machine learning (ML) and AI models rely on (ideally large) data sets for training purposes. This has become even more entrenched with the advent of generative AI (genAI) algorithms and the associated *foundation models* featuring literally billions of parameters which all have to be trained and calibrated. These data sets are usually provided by different sources (discounting synthetic data for training purposes). Critical applications often rely on data sets provided by a federation of data owners of a specific domain (e.g. Industrial AI, Healthcare…). In such a context, the Gaia-X Trust Protocol can be used to ensure:

- Provenance
- Authorship
- Digital Rights
- Compliance (e.g. Personal use of health data)
- Quality
- Validated correctness and precision
- Scientific value

of the data which is being provided.

To ensure improvements through fine-tuning and reasoning, as well as to avoid degradation and hallucinations, the prompting and inference should also be assessed with respect to the elements listed above.

4.12.2 Compliance of the underlying infrastructure and operators

Proof needs to be provided that the companies comply with the regulatory requirements of their region (e.g. AI Act) and domain (e.g. Healthcare), with respect to the underlying infrastructure. This can include:

- Requirements as defined in the Gaia-X Compliance
- Extensions (as defined in Sovereign Cloud Frameworks or specific domain requirements)

Companies training and operating trustworthy AI are required to have measures in place to allow

- Explainability of the output generated
- Validation of ethical principles and guardrails
- Compliance with data protection and privacy regulations

4.12.3 Compliance of AI services (e.g. AI-Agents, Gen AI, ML or Simulation services)

- Agents need to be clearly associated with the liable company (e.g. an Agent needs to have a unique ID linking it to the Agent owner, which can be achieved by chaining an Agent ID to an Organizational ID using the Trust Scope Credential Mechanism).
- Proof of permission from the connecting systems or persons to allow Agent operations (e.g. similar to DUA in Data Exchange: for example, an “Agent Usage Agreement” where a patient would authorize an agent to access and process their data and provide a diagnosis).

- Proof that an Agent has been validated by a CAB to operate according to the principles of the AI Act.

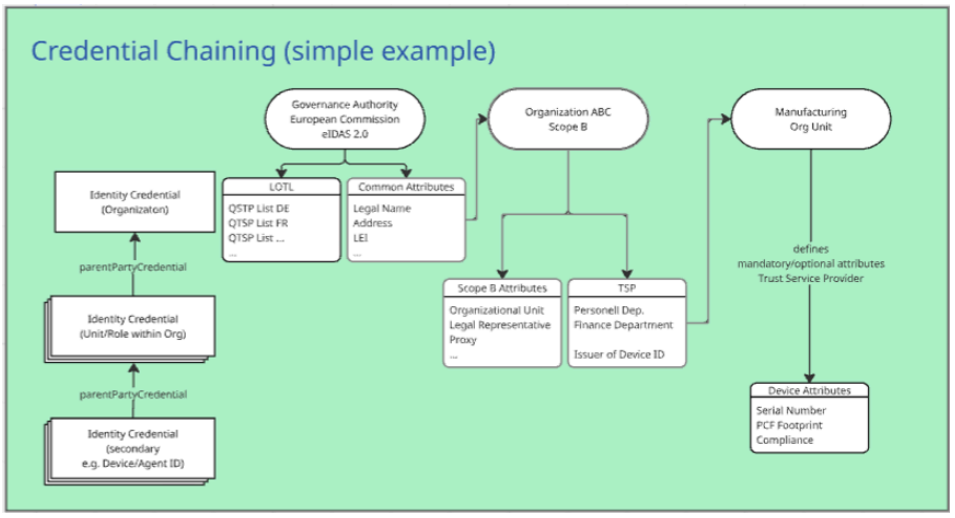

Figure 4.12.3 - Agent ID chained from Organization Credential

1. Singhal, A., Winograd, T., & Scarfone, K. A. (2007). Guide to secure web services: Guide to Secure Web Services - Recommendations of the National Institute of Standards and Technology. Gaithersburg, MD. NIST. https://csrc.nist.gov/publications/detail/sp/800-95/final https://doi.org/10.6028/NIST.SP.800-95 ↩
2. Oldehoeft, A. E. (1992). Foundations of a security policy for use of the National Research and Educational Network. Gaithersburg, MD. NIST. https://doi.org/10.6028/NIST.IR.4734 ↩
3. While the usage of the term “protocol” may sound unfamiliar in the context of cross-ecosystem communications, it fully falls under accepted definitions of the term such as IEEE Std. 610.12 (1990): “A protocol is a set of conventions that govern the interaction of processes, devices, and other components of a system” or ISO/IEC 2382-2015 “a set of rules that determines the behavior of functional units in achieving communication”. ↩
4. We have described the cross-ecosystem trust dilemma and its solution in the form of the Gaia-X Trust Protocol in more detail in the forthcoming book chapter: C F Strnadl & G Del Buono (2026), “Trust but Verify. Defining and Automating Intra- and Inter-Ecosystem Trust with the Gaia-X Trust Framework”, in: T Tuikka, B Otto, E Curry (eds.), *Developing and Deploying Data Spaces*, Springer (to appear). ↩

5 Gaia-X Implementation of Trusted Data Transactions

This chapter outlines how Gaia-X Data Exchange concepts and trust mechanisms support the implementation of trusted data transactions, with particular reference to the CEN-CENELEC prEN 18235 “Trusted Data Transactions” standardisation work. Note that organisations need to complement the trust-related elements described in this chapter with additional mechanisms for controlling and performing the actual data exchange, such as those being developed by the Eclipse Dataspace Working Group.

Enabling digital transformation and developing innovative services requires timely access to relevant data, potentially aggregated from multiple sources or suitably transformed, to generate valuable insights.

Note

Data means any digital representation of acts, facts or information and any compilation of such acts, facts or information.

However, data sharing across organisations is often hindered by stakeholder resistance, governance policies, lack of appropriate tools, and challenges in addressing regulatory constraints:

- For Data Producers, sharing personal or non-personal data may involve legal, industrial, and reputational risks, while the local benefits of sharing are often difficult to demonstrate.
- For Data Consumers, usage rights and restrictions are frequently unclear, insufficiently formalised, or expressed in legal terms that are difficult to verify and enforce automatically.
- For legal and compliance teams, data access and governance processes are often fragmented, making the verification of lawful data usage complex and resource-intensive.

Note

As per the Gaia-X Glossary, a Consumer is a participant who searches service offerings and consumes service instances in the Gaia-X Ecosystem to enable digital offerings for end users.

In this service-centric perspective, a Data Consumer simply is a Consumer who is consuming services pertaining to the sharing of data. CEN/CLC prEN 18235, on the other hand, places the data sharing transaction and associated roles, actors, and concepts in the center. Both viewpoints, though, are fully compatible with each other.

This complexity often results in overly risk-averse decisions, delaying data sharing and limiting opportunities for digital transformation and innovation.

Overcoming these barriers requires trustworthy mechanisms throughout the data transaction lifecycle. Data Rights Holders need to be able to define how their data may be used and have confidence that these constraints will be respected. Data Consumers require assurance that the data is authentic and that its use is authorised. Likewise, Data Providers need assurance that recipients are entitled to receive the data.

The Gaia-X Data Transaction concept and the associate models provide such mechanisms, enabling Data Rights Holders to retain control over how their data is used and by whom (data sovereignty), while supporting compliance with European data regulations, including the GDPR and the Data Act.

5.1 Data Transactions Conceptual Model

The Gaia-X Data Product conceptual and operational models provide the trust mechanisms required to support trusted data transactions and enable Data Rights Holders to control how their data are used and by whom (data sovereignty). They also support compliance with European data regulations, including the GDPR and the Data Act. The concepts described below are fully detailed in the Data Exchange Document and summarised here for architectural purposes.

A Data Product represents the unit of data sharing within the Gaia-X ecosystem, packaging data, metadata, and any associated licence terms. Data Products include, without being limited to, metadata describing the data product and data licence terms as defined by the Data Rights Holder, or by the Data Provider authorised to do so.

Data is furnished by Data Producers to Data Providers who compose them into a Data Product to be used by Data Consumers.

According to the concepts defined in CEN-CENELEC prEN 18235-1 “Trusted Data Transactions: Terminology, concepts and mechanisms ”, a data transaction relates to three main phases:

- Granting Rights and Publication, which is the provisioning phase leading to the publication of metadata and data policies associated with a Data Product;
- Discovery and Negotiation, which leads to an agreement between a Data Provider and a Data Consumer regarding a Data Product;
- Data Exchange and Data Usage, which operationalises the agreement through a data transaction and includes access to and usage of the Data Product by the Data Consumer.

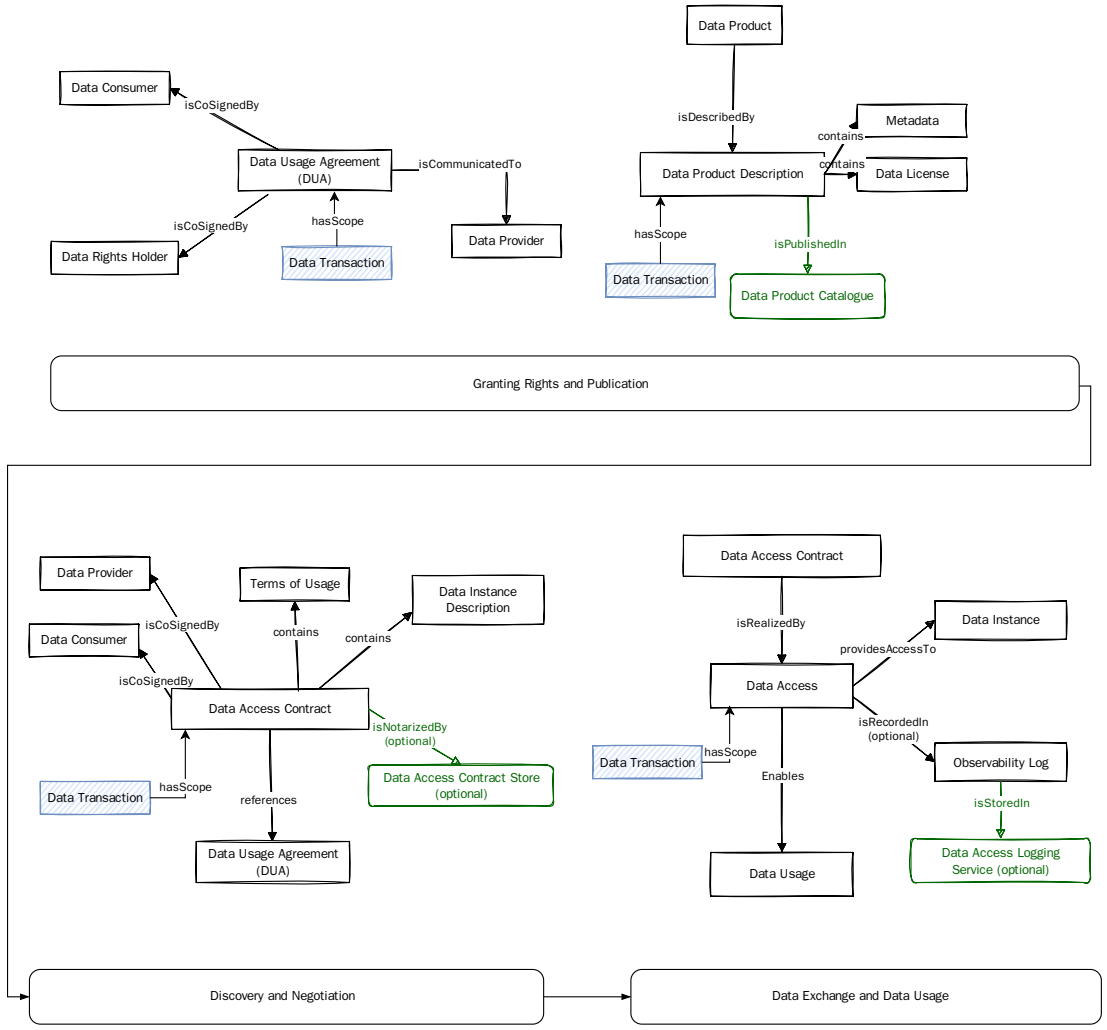

Figure 5.1 - Data Transaction phases

The Gaia-X model relies on two complementary contractual artefacts: the Data Access Contract (DAC) and, where licensed data are involved, the Data Usage Agreement (DUA). The DAC is established between a Data Provider and a Data Consumer and governs the provision and consumption of the Data Product. It is based on the Data Product Description published by the Data Provider and incorporates the terms agreed by the parties during the negotiation process.

The DUA is established between a Data Rights Holder and a Data Consumer and defines the conditions under which the underlying data may be used. In particular, it defines the usage constraints associated with the data and the obligations accepted by the Data Consumer. When licensed data are involved, the DAC and the DUA complement each other: the DAC governs access to the Data Product, while the DUA governs the authorised use of the underlying data.

Together, these contractual artefacts provide the legal and operational foundation for trusted data transactions.

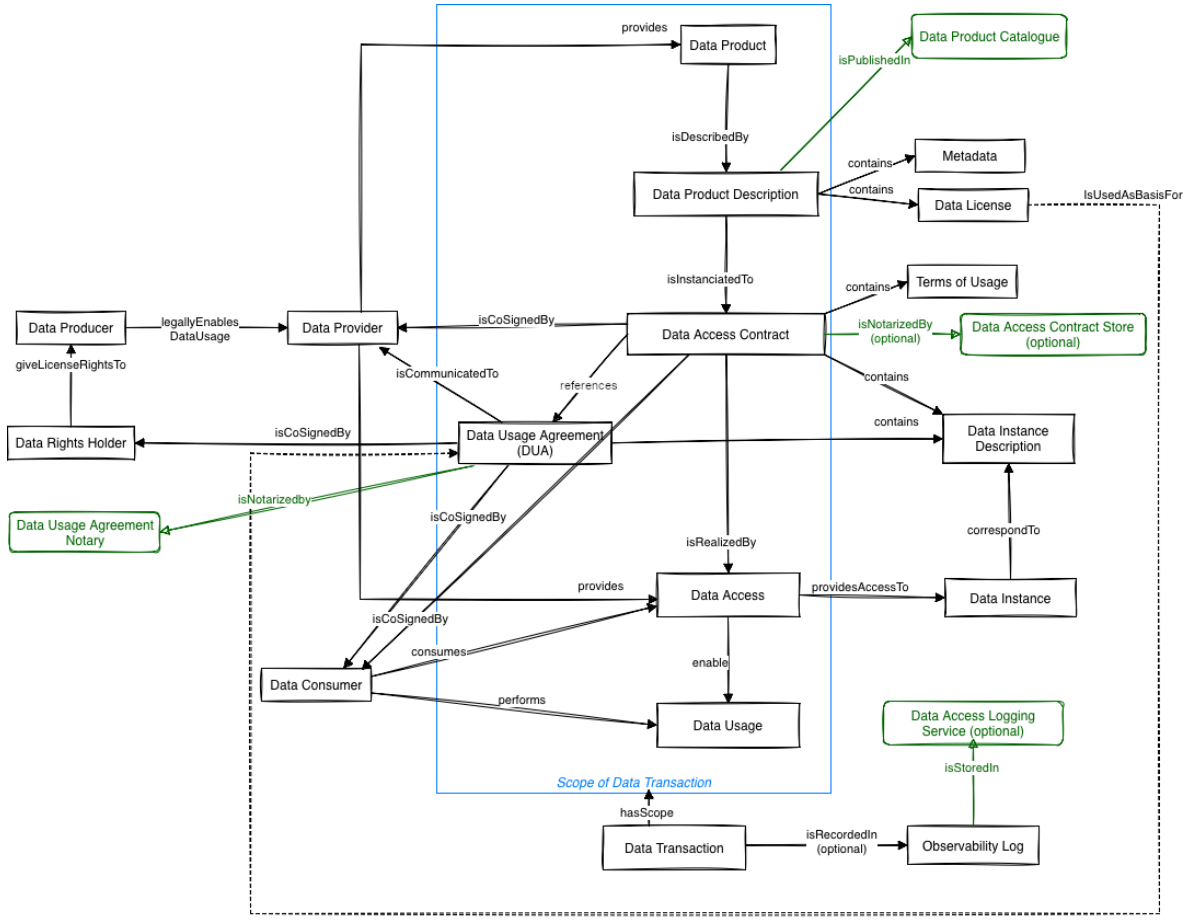

Figure 5.2 - Data Transactions conceptual model

Mapping of Gaia-X concepts with concepts used in EU data regulation

The following table maps the Gaia-X concepts with the concepts used within the different European regulations around data (GDPR and the EU acts on data – DxA):

European regulations concepts	Gaia-X concepts
data processor in GDPR	Data Provider
data subject in GDPR / user in DxA	Data Rights Holder
consent in GDPR / permission or authorization in DxA	Data Usage Agreement
recipient in GDPR / DxA	Data Consumer

For complete mapping of Gaia-X concepts used in EU data regulations, refer Data Exchange Document.

6 Gaia-X Technical Compatibility specifications

6.1 Defining Technical Compatibility

This document uses the term *technical compatibility* in the following way:

> **Note**
>
> Definition: A software component or system is Gaia-X technically compatible if and when it fulfils all requirements of a published release of this Architecture Document. In case a Gaia-X Technical Compatibility (test) kit (GX-TCK) is available, this must additionally be demonstrated by passing all (test) cases and requirements of the GX-TCK associated with the respective version of the Architecture Document.
>
> A GX-TCK is suitable software that supports the testing of software components or systems to ensure that they are compatible with the Architecture Document.

Our definition is (loosely) derived from the Eclipse Foundation's usage of the term *compatibility* in the context of the Eclipse Foundation Specification Process. Be aware, though, that the Eclipse Foundation's use of the term TCK is slightly different and means both, documented (i.e., written) requirements *or* testing software.

We stress that this is markedly different from other definitions encountered in the software engineering realm where compatibility often refers to the "ability of two or more systems or components to perform their required functions while sharing the same hardware or software environment "(IEEE) or to the "degree to which a product, system or component can exchange information with other products, systems or components, and/or perform its required functions while sharing the same common environment and resources" (ISO 25010).

We specifically talk about *technical compatibility* in order to differentiate this form of conformity to a technical specification from the notion of *Gaia-X compliance* which can be understood as conformity of participants and services to the criteria captured in the Gaia-X Compliance Document (e.g., the 3.0.0 version of the Gaia-X Compliance Document). In line with our separation of the two pillars of our Gaia-X Trust Framework, Gaia-X technical compatibility refers to rule agnostic elements, whereas Gaia-X compliance encompasses the technology agnostic requirements.

The requirements that software components or systems need to demonstrably fulfill are specified in this Chapter of the *Architecture Document*.

6.2 Understanding Identity and Identifier

An Identity is composed of a unique Identifier and an attribute or set of attributes that describe an entity within a given context.

> **Note**
>
> Gaia-X uses existing Identity standards and does not maintain them directly.

In the Gaia-X context, Identities describe participants (natural persons, companies) and resources (e.g., machines, interconnection or data endpoints) uniquely. An identifier is a unique attribute that is part of the Participant's identity.

Conditions to trust an identity at the time of the negotiation are:

- The identifiers are under the control of the credential issuer
- The VC signature can be verified
- The VC is cryptographically bound to the credential holder (See key-binding), preventing credential theft and replay attacks.
- The VC issuer strictly conforms to the KYB/KYC rules associated with the verification of the identity.
- The identifier issuers are trustworthy parties based on the Gaia-X Registry and optionally other Verifiable Data Registries that extend the Gaia-X rules.

> **Technical Specification**
>
> More info on https://docs.gaia-x.eu/technical-committee/identity-credential-access-management/24.07/TA_credential_and_party_credential/

6.2.1 Using Identifiers in Gaia-X Credentials

An identifier is a language-independent label, sign or token that uniquely identifies a Participant (cf. prEN 18235-3 Trusted Data Transactions - Part 3: Interoperability Requirements). Each Gaia-X Participant provides a self-chosen unique identifier (did:web) with cryptographic prove that the chosen identifier is really under their control (e.g., this is established through the EV SSL validation process by an external third party). The identifier is then used to strongly associate several attributes with the participant (e.g., country of origin of its head quarters). These attributes, in turn, are evaluated by (other) Participants to implement, enforce, and monitor permissions, prohibitions and duties when negotiating a contract for services or resources.

Each of the following MUST have a unique identifier:

- a Verifiable Presentation
- a Verifiable Credential
- the subject of a Verifiable Credential

> **Note**
>
> Gaia-X Credentials may reference other Gaia-X Credentials, if any. For example, a ServiceOffering that is provided by a Provider, is a composition of other ServiceOfferings, or is an aggregation of Resources.

6.3 Using Linked Data

Linked Data can be used to reference information contained in other Verifiable Credentials.

Linking the data only enables the use of IDs that are known and resolvable within the set of presented VC and always have the claims associated in its credential. This approach can result in large VP requests, which can be managed using compression algorithms.

Example of two linked Verifiable Credentials:

When a verifier requests VC1, the holder should embed in VP1 all VCs referred to in VC1's credentialSubject object, recursively, i.e., VC2.

Note: A Uniform Resource Identifier (URI) is not necessarily resolvable.

```
{
  "@id": "https://example.com/VP1",
  "verifiableCredential": [
    {
      "@id": "https://example.com/VC1",
      "credentialSubject": {
        "@id": "https://example.com/VC1/CS1",
        "hasProperty": {"@id": "VC2/CS1"}
      }
    },
    {
      "@id": "https://example.com/VC2",
      "credentialSubject": {
        "@id": "https://example.com/VC2/CS1"
      }
    }
  ]
}
```

The Linked Data principles and RDF data model's core concepts are:

- the structure of the information is a set of RDF triples called an RDF graph
- each RDF triple consists of three components: subject, predicate, and object
- each component is an IRI (which is an URI allowing arbitrary unicode charaters and not just a subset of US-ASCII as an URI). The object can also be a literal.

6.4 Using Verifiable Credentials

A credential is a set of one or more claims made by the same entity. Credentials might also include an identifier and metadata to describe properties of the credential, such as the issuer, the validity date and time period, a representative image, verification material, status information, and so on.

A verifiable credential is a set of tamper-evident claims and metadata that cryptographically prove who issued it. Examples of verifiable credentials include, but are not limited to, digital employee identification cards, digital driver's licenses, or digital educational certificates.

Gaia-X Credentials are Verifiable Credentials containing claims using the Gaia-X Ontology in specific context. Claims are expressed in Resource Description Framework (RDF).

Claims validation can be performed either by using publicly available open data and performing tests or using data from Trusted Data Sources. The underlying trust is then provided by Ecosystem Credentials.

VCs are used to share and exchange information between entities. VCs are managed via a wallet, a registry or an agent service run by the holder or by a party the holder has delegated rights to, called custodian wallet.

A Verifiable Credential includes:

- Metadata
- A subject
- A list of claims
- A list of evidence
- Verifiable proof

> **Note**
>
> - The holder is always in control with whom or what VCs are exchanged.
> - A holder can decide to delegate the management of its wallet to a 3rd party.
> - To ease the exchange of VC, it is expected for the VC identifier URL to be resolvable at a service endpoint which can implement access and usage control.
> - If the URL is not resolved, the holder is responsible to find means to exchange the VC(s).

> **Note**
>
> - Credentials that do not use Gaia-X ontology are not in Gaia-X scope.
> - The format of Gaia-X Credentials is defined in Identity, Credential and Access Management (ICAM) Document.
> - A holder can combine multiple Gaia-X credentials to build a Verifiable Presentation (VP).
> - Evidence can be included by an issuer to provide the verifier with additional supporting information in a verifiable credential. This could be used by the verifier to establish the confidence with which it relies on the claims in the verifiable credential. It is expected that credentials issued by the notaries contain the evidence of the validation process.

Example Verifiable Credential

If we use the following credential:


```
{
  "@context": [
    "https://www.w3.org/2018/credentials/v2",
    "https://w3id.org/gaia-x/development#"
  ],
  "@type": [
    "VerifiableCredential",
    "LegalParticipant"
  ],
  "@id": "https://example.org/legal-participant/68a5bbea9518e7e2ac1cc75bcc8819a7edd5c4711e073ffa4bb260034dc6423c/data.json",
  "issuer": "did:web:example.org",
  "validFrom": "2024-04-01T12:26:22.601516+00:00",
  "validUntil": "2024-01-01T12:26:22.601516+00:00",
  "credentialSubject": {
    "id": "https://example.org/legal-participant-json/68a5bbea9518e7e2ac1cc75bcc8819a7edd5c4711e073ffa4bb260034dc6423c/data.json",
    "type": "gx:LegalPerson",
    "gx:legalName": "Example Org",
    "gx:legalRegistrationNumber": {
      "id": "https://example.org/gaiax-legal-registration-number/68a5bbea9518e7e2ac1cc75bcc8819a7edd5c4711e073ffa4bb260034dc6423c/data.json"
    },
    "gx:headquarterAddress": {
      "gx:countrySubdivisionCode": "FR-75"
    },
    "gx:legalAddress": {
      "gx:countrySubdivisionCode": "FR-75"
    }
  }
}
```


6.5 Verifying Gaia-X Credentials

To ensure a Gaia-X Credential's integrity and authenticity, its claims MUST be cryptographically signed by the Issuer to prevent tampering and enable verification of the origin of the claims. The Gaia-X supported verification method is a did:web containing a JSON Web Key. See W3C JSON Web Key.

A Verifiable Credential is Gaia-X Compliant if:

- it is signed by a trusted issuer present in the Gaia-X Registry
- its issuer has a verifiable identity from one of the Trust Service Providers
- it complies with the Gaia-X Ontology SHACL Shapes
- it uses the cryptographic proof mechanism specified in the ICAM document
- it follows the rules specified in the Compliance Document

Note

The Gaia-X Compliance verification will fail if there is no link between the issuer's verification method and an approved Gaia-X Trust Service Provider.

To be able to assess the chain of trust,

- The publicKeyJwk property MUST include either the RFC7517 x5c (X.509 Certificate Chain) parameter or RFC7517 x5u (X.509 URL) parameter.
- The x5u parameter should be resolvable to a X509 .crt or .pem file which contains a complete chain to a valid Gaia-X Trust Service Provider eligible for the signed claims.
- To ensure the correct cryptographic tools are used with the public key, the alg property MUST be specified, and the value must comply with the JSON Web Algorithms RFC7518 alg.

6.6 Gaia-X Credential Format

This section extends the W3C Verifiable Credentials Data Model v2.0 to specify how it shall be applied in the scope of Gaia-X.

A holder can combine several Gaia-X credentials to build a Verifiable Presentation (VP).

A Verifiable Presentation contains one or more Verifiable Credentials with individually disclosed claims, packaged in such a way that the authorship of the data is verifiable. It SHOULD be extremely short-lived, and bound to a challenge provided by a verifier. Each Verifiable Credential that might have been issued by multiple issuers contains signed claims about one or more subjects.

Example of Gaia-X Credential Format

An example of a Gaia-X Credential document before becoming a Verifiable Credential is provided below:

Example Verifiable Credential

document.json


```
{
  "@context": [
    "https://www.w3.org/ns/credentials/v2",
    "https://w3id.org/gaia-x/development#"
  ],
  "@type": [
    "VerifiableCredential",
    "LegalPerson"
  ],
  "@id": "https://example.org/legal-participant/68a5bbea9518e7e2ac1cc75bcc8819a7edd5c4711e073ffa4bb260034dc6423c/data.json",
  "issuer": "did:web:example.org",
  "validFrom": "2024-01-01T12:26:22.601516+00:00",
  "validUntil": "2024-04-01T12:26:22.601516+00:00",
  "credentialSubject": {
    "id": "https://example.org/legal-participant-json/68a5bbea9518e7e2ac1cc75bcc8819a7edd5c4711e073ffa4bb260034dc6423c/data.json",
    "type": "gx:LegalPerson",
    "gx:legalName": "Example Org",
    "gx:legalRegistrationNumber": {
      "id": "https://example.org/gaiax-legal-registration-number/68a5bbea9518e7e2ac1cc75bcc8819a7edd5c4711e073ffa4bb260034dc6423c/data.json"
    },
    "gx:headquarterAddress": {
      "gx:countrySubdivisionCode": "FR-75"
    },
    "gx:legalAddress": {
      "gx:countrySubdivisionCode": "FR-75"
    }
  }
}
```


Additional information on Gaia-X Credentials can be found in the following documents:

- the Gaia-X Compliance Document, which defines the Gaia-X conformity assessment schemes and the requirements for the respective Trust Service Providers.
- the Gaia-X Ontology, which contains the models used to support the automation of Gaia-X Compliance.

6.7 OpenID Connect for Verifiable Credentials

OpenID Connect for Verifiable Credentials (OID4VC) is a collection of OpenID Connect specifications allowing the exchange of Verifiable Credentials and Verifiable Presentations.

The two main specifications that are relevant in the Gaia-X Ecosystem are:

- OpenID Connect for Verifiable Credential Issuance (OID4VCI)
- OpenID Connect for Verifiable Presentations (OID4VP)

6.7.1 OpenID Connect for Verifiable Credential Issuance

This specification is based on the OAuth 2.0 specification and allows an issuer to communicate with a holder and its wallet to issue Verifiable Credentials in a secure manner.

It supports both machine-to-end-user credential issuance and machine-to-machine issuance as it leverages OAuth 2.0's battle-tested and widely adopted standards.

Verifiable Credentials will be exchanged using the VC-JWT format.

> ✅ The OID4VCI specification has reached its first approved version. The Gaia-X Lab implements OID4VCI v1.0 final, using the pre-authorized code flow: the issuer generates a credential offer, the wallet exchanges the pre-authorized code for an access token, and then requests the credential.

6.7.2 OpenID Connect for Verifiable Presentations

Just like OID4VCI, OID4VP is based on the OAuth 2.0 specification to allow a holder and its wallet to present one or multiple Verifiable Credentials to a verifier through a Verifiable Presentation.

As per OAuth 2.0 standards, this protocol supports both machine-to-end-user and machine-to-machine interactions.

Verifiable Presentations will be exchanged using the VC-JWT format.

> ✅ The OID4VP specification has reached its first approved version. The Gaia-X Lab implements OID4VP v1.0 final, which uses a DCQL query to request credentials and the did client identifier scheme so the wallet can verify the signed Authorization Request against the verifier's DID document.

6.7.3 Usage

Both these protocols will be used each time a Verifiable Credential needs to be issued or consumed by the Gaia-X Clearing House, thereby securing the exchange through authentication and proof-of-possession mechanisms.

6.7.4 Cloud/Enterprise Wallet

As most of the exchanges will be in a machine-to-machine environment, a cloud or enterprise wallet will be used, although no specific solution has yet been selected.

6.8 Gaia-X Ontology

The Gaia-X Ontology is designed with extensibility and interoperability in mind.

It offers multiple ways to use and extend the ontology:

- Using the ontology through w3id.org - SHACL shapes, JSON-LD Context and OWL Ontology can be retrieved by using specific queries
- Extending the ontology through LinkML

6.8.1 Versioning

The Gaia-X Ontology versions are easy to identify and reference. To do so, the context name contains the referenced version in its URI.

The Gaia-X Ontology namespace follows this URI pattern: `https://w3id.org/gaia-x/{version}#`. For the 2.1.11 version, the specific URI is https://w3id.org/gaia-x/v2111

> **Note**
>
> The . in the version number is removed for convenience and technical reasons in the URI.
>
> For more details on how to use the ontology, click here

6.8.2 DCAT

The Data Product Catalogue (DCAT) Services are mandatory services that publish Data Product Descriptions (including metadata) and support search functionality. They use DCAT as the core protocol for Data Product description.

Data Catalog Vocabulary allows publication and discovery of catalogues with defined vocabularies.

> **Note**
>
> - A catalogue provides mechanisms to publish Data Product Descriptions (metadata) and support searching and querying the descriptions A catalogue can be implemented as a centralized or decentralized service, but the capability can also be implemented as a distributed functionality.
> - The Data Exchange Document mandates the use of DCAT-3.

6.8.3 ODRL

The Open Digital Rights language (ODRL) allows to express policies that can be evaluated for access and usage control or in contract negotiations between participants.

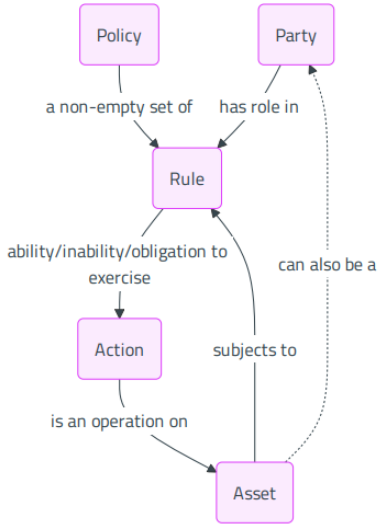

Figure 6.8.3 - Representation of ODRL Model

> **Note**
>
> Realization of Policy Definition and Policy Enforcement follows the W3C ODRL specifications.
>
> A specific ODRL profile is built by the Gaia-X Lab with the purpose of being able to refer in a clear and precise way to verifiable credential claims in an ODRL Policy.
>
> To enable automated processing of Data License constraints (cf. https://www.w3.org/TR/odrl-model/), Gaia-X mandates the use of Open Digital Rights Language (ODRL), in combination with an ontology typically defined by the ecosystem.

6.8.4 CAP

The Eclipse Conformity Assessment Profile (CAP) Ontology developed within the Eclipse Dataspace Working Group (EDWG), provides a standardized framework for expressing and verifying conformity assessment policies, leveraging verifiable credentials and aligning with ISO/IEC 17000:2020 standards. By defining key concepts such as certifications, attestations, evidence, and requirements, CAP enhances interoperability and trust in data sharing ecosystems.

6.8.5 Gaia-X Schema

Gaia-X members define the Schema for Gaia-X Credentials. It is used as the vocabulary of the claims about credential subjects and must be available in three formats: * SHACL shapes (cf. the W3C Shapes Constraint Language SHACL 3) * JSON-LD Context, and * OWL Ontology.

Whenever Credentials are created or received, they form a data graph. To ensure compliance with Gaia-X and/or specific ecosystem extensions, this data graph must be validated against the given SHACL shapes graph according to the SHACL specification. The current version of the Gaia-X Schema is maintained and is available through the Gaia-X Registry Service.

6.8.5.1 Generation via LinkML

Linked Modeling Language (LinkML) is a modeling language designed to define data schemas that can be used across a wide range of formats and technologies. It enables the creation of a single, platform-independent source of truth that can be automatically transformed into multiple serializations, including JSON-LD, OWL, SHACL, Python, and TypeScript.

By defining classes, properties, constraints, and relationships in a LinkML YAML schema, members can generate semantic web artifacts like OWL for ontological reasoning or SHACL for data validation. Additionally, JSON-LD outputs facilitate integration with linked data systems. This multi-target capability allows to maintain consistency across data modeling, validation, documentation, and code generation workflows.

As a result, LinkML streamlines the process of creating interoperable, semantically-consistent data models which are both machine-readable and easy to maintain.

6.8.5.2 Schema Extensions

The Gaia-X ontology extensions can be implemented using LinkML's import keyword.

To produce the same mygx:LegalPerson entity as defined in the SHACL Shapes section, the following LinkML snippet can be used:

```
id: https://my-gaia-x.eu#my-ontology
name: my-ontology
default_prefix: mygx
prefixes:
  gaia-x: https://w3id.org/gaia-x/development/linkml/
  mygx: https://my-gaia-x.eu#
imports:
  - gaia-x:types
classes:
  MyLegalPerson:
  title: "My Legal Person"
  is_a: LegalPerson
  description: A custom definition of the Gaia-X LegalPerson
```

6.9 Managing Trust Services

6.9.1 Trusted Service Operators

Services such as Credential storage, vocabulary, authorisation, and other value-added services, when provided by trusted operators within specific ecosystems, must comply with the general requirements for technical compatibility described in this section.

6.9.2 Using Compliance Engine

The Gaia-X Compliance engine implements the rules defined in the Gaia-X Compliance Document from high level objectives to low level requirements. It performs checks to determine:

- whether the received Verifiable Presentation is Gaia-X Compliant,
- and whether a Gaia-X Compliance credential can be issued and at which conformity level.

The Gaia-X Compliance Engine is the main Gaia-X Digital Clearing House entry point. It exposes REST endpoints that allow participants to obtain a Gaia-X Compliance credential by submitting credentials describing themselves and the service offerings they provide.

Several levels of Gaia-X conformity exist, from Standard Compliance to Gaia-X Label Level 3 depending on the criteria that have been validated.

Each conformity level has specific criteria, some of which may be shared across levels, as defined in the Compliance Document. A filter chain is defined within the Compliance Engine to validate these criteria. Each filter may be responsible for validating one or multiple criteria.

Before going through the filter chain, the input VC-JWT (Verifiable Credentail - JSON Web Tokens) must be validated and verified. This process is explained in the next chapter.

6.9.2.1 Checking JWT Headers

The input of the Compliance Engine is a verifiable presentation in `application/vp+jwt` format containing multiple verifiable credentials each in `application/vc+jwt` format, describing the participant's service offering. This input must be verified according to the Securing Verifiable Credentials using JOSE and COSE specification.

The JWT headers must be checked to ensure the content-type and type fields conform to the specification. Next, the issuer `iss` and key ID `kid` headers are used to retrieve the corresponding DID document and public key. The algorithm used to sign the JWT is indicated in the `alg` header.

```
{
  "alg": "ES256",
  "typ": "vp+jwt",
  "cty": "vp",
  "iss": "did:web:gaia-x.eu",
  "kid": "did:web:gaia-x.eu#key-0"
}
```

6.9.2.2 Verifying and Decoding JWT

Using the previously extracted headers, the JWT can be decoded and verified using any widely used JOSE library.

During the verification process, the JWT signature is checked to ensure that the content of the payload is not tampered. The signature is verified using the issuer's public key from the verification method of the DID document in conjunction with the signing algorithm to make sure the JWT was signed with the issuer's matching private key. This enforces non-repudiation of the Compliance Engine input.

The aforementioned verification method in the `kid` retrieved from the DID document must also contain a certificate delivered from a Trust Service Provider.

Furthermore, claims are checked to make sure that the JWT is valid regarding the validity period using the `exp` (expiration date) claim. This value must be later than the current date as stated in the JWT specification. Since the claim is optional, this check is skipped if the claim is not present in the JWT.

> **Note**
>
> The JWT's expiration date can be different from the verifiable credential's validFrom and validUntil attributes defined in the Verifiable Credential Data Model v2.0 specification. These attributes represent the validity of the information carried by the verifiable credential, whereas the JWT expiration date claim expresses the token's validity.
>
> For example, one can present a driver's license valid until next year in a JWT that expires a few minutes after it has been issued to avoid it being compromised and reused.

6.9.2.3 Converting to Verifiable Credential List

Once the input Verifiable Presentation is verified, its content can be converted from a list of Enveloped Verifiable Credentials to a list of Verifiable Credentials that can be passed through the filter chain.

6.9.3 Using the Registry

The Gaia-X Registry is one of the mandatory components to perform compliance checks. It acts as the source of truth for the ecosystem and stores the files used by the compliance engine.

The following elements are critical for the Gaia-X Registry Service:

Shapes

The registry serves the shapes that are used by the compliance engine to ensure the validity of the structure of the Verifiable Presentations that are sent to it.

Trust Service Providers

The Registry contains the root certificates of all parties allowed to issue credentials or sub-certificates based on the trust framework.

The Registry also contains the list of TSPs approved for a particular scope.

Currently, it contains the list of eIDAS national trust service providers as well as the list of accredited EV-SSL certificates issuers from Mozilla Certificate Authority Store.

Revocation lists

To exclude a malicious actor, the registry contains a list of revoked certificates.

Provided APIs

`/api/trustAnchor`

The endpoint allows to check whether a certificate belongs to a trust anchor and has not been revoked.

> **Note**
>
> References to the term Trust Anchor will be updated in the document once the corresponding changes are included in a software release.

`/api/trusted-issuers`

The endpoint exposes the list of accredited credentials issuers for Gaia-X credentials, e.g. the Gaia-X Digital Clearing House (GXDCH).

6.9.3.1 Gaia-X IPFS Pinning Service

The Gaia-X IPFS Pinning Service supports Gaia-X efforts to manage, update, and distribute essential artefacts via the IPFS network. These artefacts includes trust framework schemas and shapes, trusted issuers lists, and revocation lists. These elements are critical for the operation of the Gaia-X Registry Service.

This service uses DNS TXT records to broadcast the root content ID, which is an addressable hash in IPFS, to GXDCHs. It is used to publish, share and update artefacts, across the ecosystem and specifically with Gaia-X registries that are also encouraged to seed their files through IPFS to other GXDCHs and the broader community.

Using IPFS as a storage layer for the registry helps eliminate single points of failure (SPOFs) by distributing data across a decentralized network, ensuring high availability and resilience. It also provides content-addressable storage, which guarantees data integrity and verifiability, fostering a secure and reliable network.

6.9.4 Using the Gaia-X Legal Registration Number (LRN) Notary

The Gaia-X Legal Registration Number (LRN) notarisation service plays a crucial role in validating legal registration numbers. Its primary function is to verify the authenticity and existence of the registration numbers provided and subsequently issue signed Verifiable Credentials.

The VCs serve as tangible proof of the registration number's validity and existence. It is important to note that while the VC affirms the existence of the registered number, it does not establish a direct association between the credential holder and the respective company. This is because the verification process allows anyone to confirm the validity of a registration number, making it a valuable tool for various purposes.

The Gaia-X LRN (Legal Registration Number) service focuses on three specific types of legal registration numbers, namely VAT ID, EORI, LEI code, and potentially a TAX ID. These are key identifiers used in the business and legal domains. When entities present these numbers for verification, the Gaia-X LRN notarization service confirms their existence, providing a level of trust and reliability to these critical identifiers.

Example of a VC issued by a notary:

The Notary exposes an API endpoint(s) that require as input:

- The `vcid` and `subjectId`, which will be used respectively as a credential ID and as credential subject ID in the Verifiable Credential response
- The type of the requested legal registration number
- The value of the legal registration number property to verify

> **Note**
>
> In JSON-LD, the `id` attribute should be resolvable and ideally should lead to the future storage location of the Verifiable Credential. This means that the provided `id` should resolve to the designated storage location for the VC.

If the legal registration number is valid, you will receive a Verifiable Credential. This VC serves as tangible proof of the number's authenticity and existence, enhancing trust and reliability. If the provided number is not valid, the API will return an HTTP error with a message.

II. Appendices

7 Supported Credential Formats and -Exchange Protocols, Wallets and DID

7.1 DID Methods

	did:web	did:ebsi	did:elsi	did:key	did:keri	did:op
Tagus	X					
Loire	X					
Danube	X [1]	via extension [2]	via extension [2]	X [1]	via extension [2]	via extension [2]

Note

1. The did method is directly supported by the `Gaia-X Basic Functions` component of the Danube `Gaia-X Core Engine`.
2. The did method may be directly supported by a suitable *extension* of the `Gaia-X Core Engine`.

In the table, the following did methods are mentioned:

- did:ebsi (European Blockchain Services Infrastructure; now the Europeum EDIC)
- did:elsi (linked to eIDAS)
- did:key (equivalent to allowing completely anonymous participants)
- did:keri (GLEIF)
- did:op (Ocean Protocol) or other DLT-based methods

7.2 Credential Formats

	JWS 2020	VC-JWT
Tagus	X	
Loire		X
Danube		X

where: - JWS 2020: W3C JSON Web Signature 2020 - VC-JWT: Verifiable Credential as JSON Web Token as specified in Securing Verifiable Credentials using JOSE and COSE

Rationale: JWS uses a signature *embedded* in the JSON file of the verifiable credential payload whereas VC-JWT takes the whole VC payload and *wraps it into a JWT envelope*. This separates the signature process from the actual VC contents. VC-JWT is also compatible with the eIDAS 2.0 EUDI Wallet and EUBW (Business Wallet).

7.3 Credential Exchange Protocols

Currently, the Gaia-X 3.0 “Danube” platform and other Gaia-X technically compatible compliance engines need to support the following credential exchange formats.

Version	VC-API	OID4VC	Remarks
Tagus	X		will not be upgraded to OID4VC
Loire	X	X	The Loire Compliance Engine itself does not support OID4VC. OID4VC access to Loire compliance is accomplished through the Gaia-X 3.0 “Danube” platform
Danube	X	X	Gaia-X 3.0 “Danube” platform supports both protocols for Loire compliance and any extension

where:

- VC-API: simple HTTP/REST API typically exposed at `https://compliance-engine.company.com/v2/docs` (example)
- OID4VC: Open ID Foundation set of standards for verifiable credentials. Gaia-X currently uses OID4VCI (Issuance) and OID4VP (presentation)

Note on Gaia-X 3.0 “Danube” platform
The Danube platform supports OID4VC through its `Gaia-X Compliance Dispatcher` component.

Note on potential additional credential exchange protocols
We are well aware of various other credential exchange protocols such as - CHAPI: Credential Handler API - DCP: Decentralized Claims Protocol - DIDComm: Decentralized Communication - WACI: Wallet and Credential Interactions

In the current architecture of the Gaia-X 3.0 “Danube” platform, all in-bound API requests to the Danube platform are handled by the `Gaia-X Compliance Dispatcher` component (`GXCD`). That means that any additional credential protocol needes to be implemented directly in the `GXCD`.

8 Trust Indexes

The ecosystem authority can define different conformity assessment schemes, but those:

- Do not automatically adapt to the market, which always evolves faster than the rules.
- Do not capture the subtleties of organisational and semantic interoperability complexity, *de facto* lowering the interoperability to the basic commonly conceded denominators.

To address this challenge, additional scoring tools named Trust Indexes are developed as Veracity T_V, Transparency T_T, Composability T_C and Semantic-Match T_{SM} indexes.

The exposition here, admittedly, is somewhat mathematical-technical but still tries to be understandable. If you need the rigorous formal purely mathematical definitions and formulae, turn to the following paper here.

Example

Those indexes enable the ecosystem parties to:

- As a consumer, to compare objects or offerings with a more granular ranking that wouldn't necessarily fit into one of the predefined conformity schemes.

Example

Given two compliance schemes, A and B, one (B) built on top of the other (A). In the normal binary logic of (service) compliance, a service offering can either be compliant with the base compliance scheme A or with the extended scheme B (which implies that the service offering is compliant with the base scheme A as well). In this logic, there is no room to say that the service offering is 100% compliant with A but only 50% compliant with B.

Trust indices – here called service compliance overlap, "SO" – can solve this problem by assigning to a service offering a score r_{SO} in the range $1.0 < 2.0$. $r_{SO} = 1$ then means that a service is only compliant with scheme A and 0% with scheme B, and a score $r_{SO} = 1.6$ indicates that a service offering is 60% also compliant with scheme B.

- The providers to compare their offering with existing ones and help them to improve their interoperability.

??? example
As a provider _"is my service interoperable and composable with 10% or 90% of the other offerings in the ecosystem catalogues ?"_

The proposed Trust Indexes T_V, T_T, T_C, T_{SM} are meant to be used by all parties - licensor, licensee, producer, provider, consumer, ... - as a measure of distance for interoperability and trust with regards to the other offerings in the ecosystem catalogues. The intent being that the ecosystem helps the market by providing measurement tools and letting the market actors themselves converge to an optimum solution, similar to gradient descent algorithms where an error metric is given and a process is iteratively applied to converge to an optimum solution.

Note

A provider that declares its services or products by providing only the bare minimum information without machine readable claims nor policies may still meet the mandatory ecosystem compliance criteria but will have a low interoperability score with other offerings from the ecosystem catalogues and it would be in the provider's own interest to improve the quality and quantity of the provided information.

The Trust Index T is a function of several sub-indexes, all using VC (`VC`) as inputs. The output is a number in $\mathbb{R}$.

$$T : T_V, T_T, T_C, T_{SM} \rightarrow \mathbb{R}$$

with each sub-index T_i as $T_i : \{\texttt{VC}\} \rightarrow \mathbb{R}$. For a thorough (mathematical) introduction see the paper Gaia-X Trust Indices.

Sample code included in the source of this page

All the graphs in this page are dynamically computed in JavaScript.

If you are looking for an index implementation, look at the source code of this page.

8.1 Sub-Indexes

8.1.1 Veracity

The Veracity index T_V is a function of the n chains of the public keys from the issuer to the Trust Anchor.

$$\forall C_k \in \{Chains\}, T_V = \frac{1}{n} \sum_{k=1}^{n} \alpha^{length(C_k)-1}$$

Tip

$\alpha = 0.9$ is recommended to keep a meaningful T_V value even with long machine-to-machine chains.

Evolution of the index with α

Smaller is α, more aggressive is the index with regards to long chains.

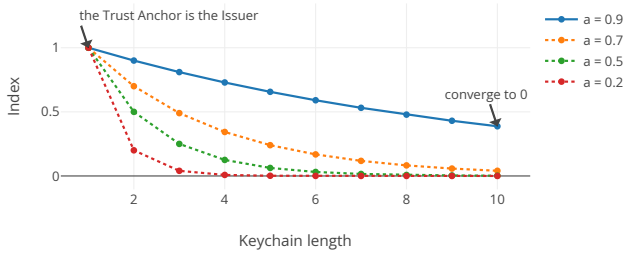

If there is one chain then the value is set to 1 else the more chains there are, higher ↗ is the veracity index.

Example

For a given claim, the number of signatures on that claim (root → ⋯ → leaf → VC → subject ← VC ← leaf ← ⋯ ← root)

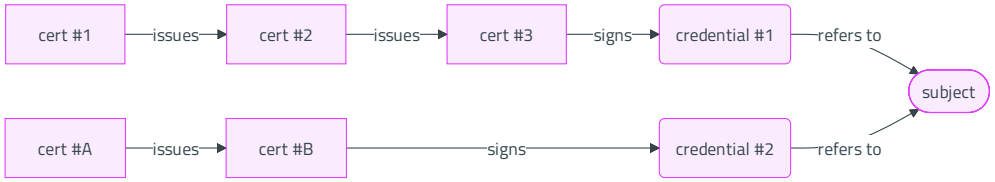

If the issuer is the Trust Anchor then the value is set to 1 else longer the chain is, lower ↘ is the veracity index.

Example

For a given claim, the length of the signing key chain (root cert → intermediate cert → ⋯ → leaf → VC → subject)

8.1.2 Transparency

The goal of the transparency index is to reward parties that expose information.

Transparency vs Compliance

The transparency index is not linked to the notion of compliance. It's possible to have a high transparency index and not be compliant, because mandatory properties are missing or in the wrong format.

Transparency vs VC type

The transparency index can only be compared across VC of the same type.

The transparency index of a credential is based on the number of properties (more correctly: *claims*) contained therein, the cardinality of the number of information items accepted for characterizing a property in more detail, and (most importantly) the number of information items for a given property actually supplied by the service provider in the credential. The transparency index addresses primarily optional information items and properties of a credential to allow a more detailed and deeper comparison of two fully compliant credentials for the same service type.

T_T is computed in a two-step approach:

1. T_{T_n} values for each VC n of the graph G.
2. an overall T_T value from the previous T_{T_n} values, for the VC n_0 representing the object of interest.

T_T and T_{T_n} are characterised by:

- $0 \leq T_T \leq 1$
- $T_T = 0$ when $|\{p\}| = 0$, ie the cardinality of the set $\{p\}$ of property p is 0, ie no property are filled in.
- $\lim_{|\{p\}| \to +\infty} T_T = 1$ when more properties p are filled in.

Example

Example of the same offering - an API endpoint returning a fortune from the BSD packet fortune - with an increasing ↗ transparency index, $T_{T_2} > T_{T_1} > T_{T_0}$

Minimum requirement T_{T_0}

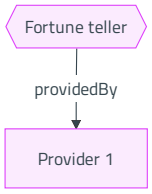

Service Offering

```
name: Fortune teller
description: API to randomly return a fortune
providedBy: url(provider1)
termsAndConditions:
- https://some.url.for.terms.and.condition.example.com
```

Provider 1

```
registrationNumber: FR5910.899103360
headquarterAddress:
country: FR
legalAddress:
country: FR
```

Improved requirement $T_{T_1} > T_{T_0}$

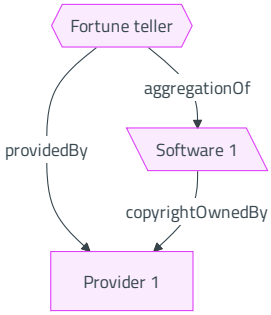

Service Offering

```
name: Fortune teller
description: API to randomly return a fortune
providedBy: url(provider1)
aggregationOf:
- url(software1)
termsAndConditions:
- https://some.url.for.terms.and.condition.example.com
```

Software 1

```
name: api software
copyrightOwnedBy:
- url(provider1)
license:
- EPL-2.0
```

Improved requirement $T_{T_2} > T_{T_1}$

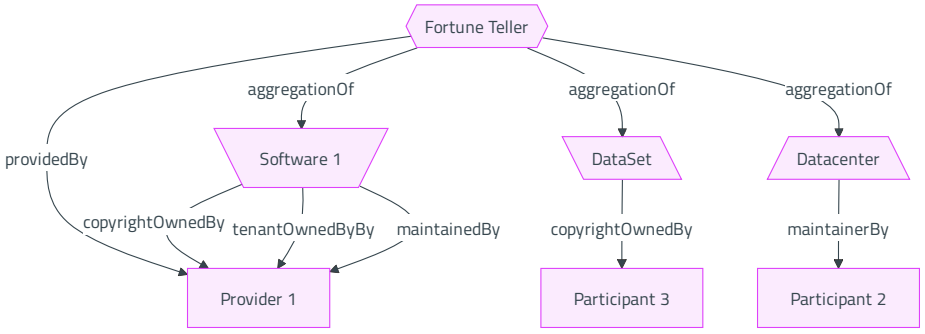

Service Offering

```
name: Fortune teller
description: API to randomly return a fortune
providedBy: url(provider1)
aggregationOf:
- url(software1)
- url(dataset1)
- url(datacenter1)
termsAndConditions:
- https://some.url.for.terms.and.condition.example.com
policies:
- type: opa
    content: |-
    package fortune
    allow = true {
        input.method = "GET"
    }
```

API 1

```
name: api software
maintainedBy:
- url(provider1)
tenantOwnedByBy:
- url(provider1)
copyrightOwnedBy:
- url(provider1)
license:
- EPL-2.0
```

Dataset 1

```
name: fortune dataset
copyrightOwnedBy:
- name: The Regents of the University of California
    registrationNumber: C0008116
    headquarterAddress:
    state: CA
    country: USA
    legalAddress:
    state: CA
    country: USA
license:
- BSD-3
- https://metadata.ftp-master.debian.org/changelogs//main/f/fortune-mod/fortune-mod_1.99.1-7.1_copyright
```

Participant 2

```
name: Cloud Service Provider
registrationNumber: FR5910.424761419
headquarterAddress:
country: FR
legalAddress:
country: FR
```

Datacenter 1

```
name: datacenter
maintainedBy: url(participant2)
location:
- country: FR
```

8.1.2.1 For each VC of the graph

To meet the objectives of the Transparency Index set in the previous section, each VC with its associated SHACL shapes must be analysed as follow:

- with $countTotal(n)$ the number of available properties for a VC n, every property p is given a weight w_p.

$$w_p = \frac{1}{countTotal}$$

- each property p is given a value v_p depending on its cardinality and the number of given values c

$$\forall p \in n, T_{T_n} = \sum_p w_p v_p$$

> **Example**
>
> for a VC with 5 properties, each property p has a weight $w_p = 0.2$

For a cardinality with known boundaries `n..m`

$$v_p = \begin{cases} 0, & \text{if } c \leq 0 \\ \frac{c}{m}, & \text{if } 0 < c < m \\ 1, & \text{if } c \geq m \end{cases}$$

Example with a known maximum

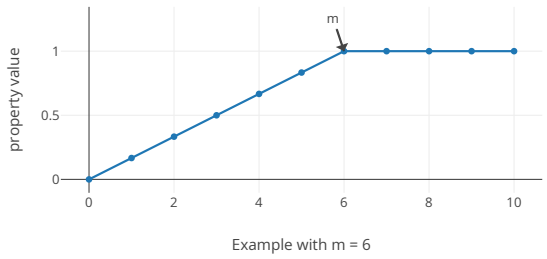

Example with m = 6

For a cardinality with an unknown upper boundary `n..*`

$$v_p = \begin{cases} 0, & \text{if } c < n \\ 1 - \beta^{c-n}, & \text{if } c \geq n \end{cases}$$

Tip

$\beta = 0.9$ is recommended to keep a meaningful v_p value even with a long array of values.

Example with a unknown maximum and evolution of the value with β

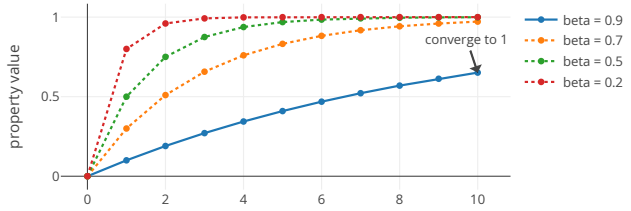

Example without upper boundary

Example of the same VC type for different level of information

In the example below, we have a VC schema with 6 properties with different cardinality and 7 different VC instances, each with more or less defined values.

property with max cardinalty	VC 1	VC 2	VC 3	VC 4	VC 5	VC 6	VC 7
prop1, maxCard = 1	0	0	1	1	1	1	1
prop2, maxCard = 1	0	0	0	0	0	1	1
prop3, maxCard = 1	0	1	1	1	1	1	1
prop4, maxCard = 2	1	1	1	1	1	1	1
prop5, maxCard = *	0	0	0	1	2	1	2
prop6, maxCard = *	0	1	1	1	1	1	6
T_{T_n}	0.083	0.267	0.433	0.450	0.465	0.617	0.693

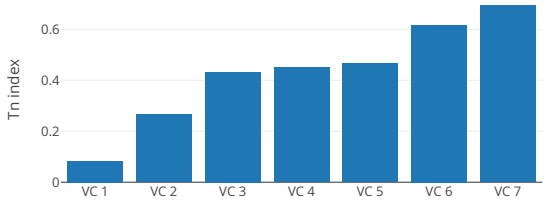

8.1.2.2 For the node of interest

Once all the T_{T_n} values for each VC n from the graph G are calculated, an overall T_T value is calculated taking into account the length of the path $d_n = dist(n_0, n)$ from the VC of interest n_0 to a VC n in the graph G, with the principle that further away from n_0 is a VC, less impact it has on the overall T_T index value.

If there are several VCs at the same distance from n_0, the average $\overline{T_{T_d}}$ of the T_{T_n} indexes at the distance d is computed.

$$\exists n_0 \in G, \forall n \in G, T_T = \sum_{d_n} \gamma(1-\gamma)^{d_n-1}\overline{T_{T_d}}$$

Info

The above formula was built to normalise the index T_T in the range $[0, 1]$.

$$\text{for } 0 \leq \gamma \leq 1, \int_1^\infty \gamma(1-\gamma)^{x-1}dx = 1$$

Tip

$\gamma = 0.5$ is recommended to have a balance between VC with and VC without linked VC.

Evolution of the value with γ

When γ is close to 1, linked VC have a low impact on the overall Transparency index T_T.

When γ is close to 0, linked VC have a high impact on the overall Transparency index T_T.

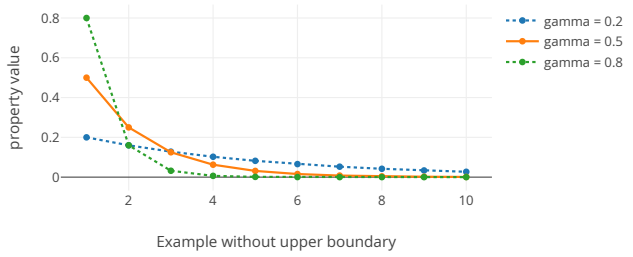

Example without upper boundary

Example

For the graph below, various transparency indexes T_T are calculated for different values of T_{T_n}.

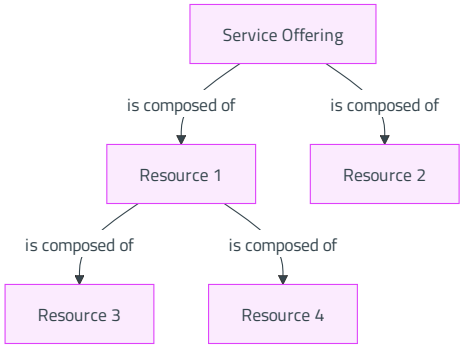

	Example 1	Example 2	Example 3	Example 4	Example 5
Service Offering (d = 1)	Tn=0.2	Tn=0.2	Tn=0.5	Tn=0.5	Tn=0.5
Resource 1 (d = 2)	Tn=0	Tn=0.1	Tn=0.1	Tn=0.9	Tn=0.9
Resource 2 (d = 2)	Tn=0	Tn=0.1	Tn=0.1	Tn=0.8	Tn=0.8
Resource 3 (d = 3)	Tn=0	Tn=0.3	Tn=0.3	Tn=0.3	Tn=0.3
Resource 4 (d = 3)	Tn=0	Tn=0.9	Tn=0.9	Tn=0.1	Tn=0.9
T_T	0.100	0.200	0.350	0.488	0.537

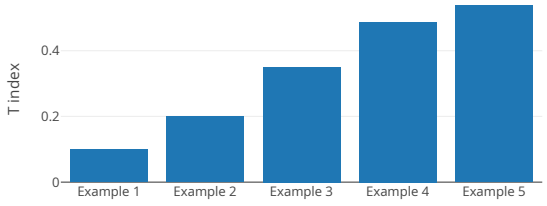

8.1.3 Composability

The Composability index T_C is a function of two or more service descriptions, computing:

- the capacity of those services to be technically composed together, e.g., software stack, compute, network and storage characteristics, plugin and extension configurations, ...

This index is computed, for example, by instances of the Federated Catalogues, by analysing and comparing the characteristics of several service descriptions.

This index can be decomposed for each service description into several sub-functions:

- the level of detail of the Software Bill of Material
- the list of known vulnerabilities, such as the ones listed on the National Vulnerability Database.
- Intellectual Property rights via the analysis of the licenses and copyrights.
- the level of stickiness or adherence to specific known closed services
- the level of readiness for lift-and-shift migration.

Check Fundamental/Linked Data for an example.

8.1.4 Semantic Match

The Semantic Match index T_{SM} is a function of:

- the use of recommended vocabularies (Data Privacy Vocabulary, ODRL, ⋯)
- the unsupervised classification of objects based on their properties
- Natural Language Processing and Large Language Model analysis from and to Domain Specific Language (DSL)
- Structured metadata information embedded in unstructured data containers (PDF/A-3a, ⋯)

1. PDP https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml ↩

9 Changelog

9.1 Release 3.1 (2026 June Release 26.06)

The Architecture Working Group meeting 09 June 2026 unanimously decided to switch to semantic versioning for the 26.06 (old versioning) = v3.1 (new semantic version) release of the Architecture Document.

It was also decided to provide a retroactive comparison table of the old versions of the Architecture Document (like 21.03) to an assumed semantic versioning in line with having a v3.1 for a 26.06 (or so) release. This is for information, comparison, and historic purposes only. References to pre-v3.1 versions of the Architecture Document can (and should) be made using the old `YY.MM` versioning system.

Old Version	Semantic Version	Remark
21.03	0.5.0	
21.06	0.6.0	
21.09	0.7.0	
21.12	0.8.0	
22.04	0.9.0	Elbe
22.10	1.0.0.	Tagus
23.10	1.5.0	Tagus
24.04	2.0.0	Loire
25.05	2.5.0	Loire
25.11	3.0.0	Danube
(26.06)	3.1.0	Danube; first release with native semantic versioning

Additional changes in release 3.1:

- Section 2 (Introduction): [new] (Forward) reference to Gaia-X Trust Protocol
- Section 3.5 (Gaia-X Alignment): [updates] DSSC, Simpl, CEN/CENELEC. Added: Data Spaces Adoption Forum (DSAF). Removed DSBA.
- Section 4.4 (Gaia-X Trust Framework Architecture): [new] Added the Gaia-X Trust Protocol with accompanying changes to cross- and inter-ecosystem interoperability (sub-)sections
- Section 4.11: [new] Architectural characterization of how the Gaia-X Trust Protocol/Gaia-X Trust Framework supports the creation of Trusted Artificial Intelligence Environments. This heavily makes use of the `Party Credential` concept of the ICAM document which also allows the *chaining* of credentials back to a *parent*, e.g., the company responsible for an AI agent
- Section 5 (Trusted Data Transactions) [updates] Aligned with latest version of prEN 18235 (Trusted Data Transactions, Parts 1-3) and with latest version of the *Data Exchange Document*. This includes linking the conceptual model with the three phases of prEN 1825, including the *Data Access Contract* besides the Data Usage Contract as second centerpiece of sovereign and trusted data transactions in an expanded conceptual model
- Section 6 (Gaia-X Technical Compatibility Specification) [updates] Several editorial refinements across the whole document (e.g., orthographic corrections, improved sequencing and readability, updates of links); improved section on identifiers by adopting the definition of prEN 18235-3
- Section 7 (Supported Credentials): [updates] Added OID4VC explainer text for Gaia-X 3.0 “Danube” platform. Highlighted Danube procedures for other credential exchange protocols (which would need to go into the Gaia-X Compliance Dispatcher)
- Section 8 (Trust Indexes): [updates] Added link to a new mathematical white paper on calculating two of the three trust indexes rigorously.

9.2 2025 November Release (25.11)

- In chapter 3. Gaia-X Context:
  - A new sub-chapter, *3.1 Basic Principles* is added, which explains the main characteristics of digital ecosystems and the value they provide to participants, as well as the role of Gaia-X in automating compliance verification and onboarding processes.
  - A new sub-chapter, *3.4 Complementarity with additional TSP and support in Gaia-X OSS releases* with sub-chapters *3.4.1 Integrating with external Trust Frameworks* and *3.4.2 Gaia-X 3.0 Danube Software Release* are added to explain the integration with external trust frameworks and Gaia-X 3.0 Danube release for specialized TSPs for any ecosystems.
  - A new sub-section, *3.5.1.3 - Ocean Enterprise* is added to the list of 3.5.1 Aligning with Other Associations and Foundations chapter.
- In chapter 4. Gaia-X Trust Framework Architecture, the sub-chapter *4.2.1 Performing automated onboarding and offboarding* section is updated with additional content.
- In chapter *6.2 Understanding Identity and Identifier* - conditions to trust an entity at the time of negotiation are updated.
- A new chapter *7. Supported Credential Formats and -Exchange Protocols, Wallets and DID*, is added.

9.3 2025 May Release (25.05)

The Architecture document is refactored completely. Below are the ToCs from previous version and the newly refactored Architecture document to give a clarity on what has been removed or added/updated.

Previous version ToC	Current New ToC
1. About	About
1.1 Editorial Information	1. Editorial Information
1.1.1 Publisher	1.1 Publisher
1.1.2 Authors	1.2 Authors
1.1.3 Contact	1.3 Contact
1.1.4 Other Format	1.4 Other Format
1.1.5 Copyright Notice	1.5 Copyright Notice
2. Context	2. Introduction
2.1 Gaia-X as a member of the Data Spaces Business Alliance (DSBA)	3. Gaia-X Context
2.2 Gaia-X Trust Framework	3.1 Understanding Ecosystems and Data Spaces
Models & Components	3.1.1 Federation of Ecosystems
3. Gaia-X Conceptual Model	3.2 Gaia-X high-level Positioning
3.1 Main Concepts	3.2.1 Trust Plane
3.1.1 Gaia-X ecosystems	3.2.2 Management Plane
3.1.2 Decentralized trust framework for ecosystems	3.2.3 Usage Plane
3.1.3 The Gaia-X model building block	3.2.4 Complementarity of Technical Compatibility and Compliance
3.2 Roles	3.3 Gaia-X Alignment
3.2.1 Trust Anchor	3.3.1 Aligning with Other Associations and Foundations
3.2.2 Participant	3.3.2 Aligning with Other Initiatives
3.2.3 Basic Interactions of Participants	3.3.3 Aligning with External Projects
3.3 Components	3.3.4 Aligning with Standards and Regulations
3.3.1 Gaia-X Credentails	4. Gaia-X Trust Framework Architecture
3.3.2 Policies	4.1 Elements of a Trust Framework
3.3.3 External	4.2 Using the Trust Framework
3.4 Overview Picture	4.2.1 Performing Automated Onboarding and Offboarding
4. Component Details	4.2.2 Performing Credential Verification
4.1 Resources & Service Offerings	4.2.3 Identifying Ecosystem Trust services
4.2 Policies	4.3 GXDCH
4.2.1 Policy definition	4.4 Gaia-X Conceptual Model
4.2.2 Policy description	4.4.1 Terminology Sources
4.3 Service Composition	4.4.2 Definitions
4.3.1 Assumptions	4.4.3 Model Core
4.3.2 Generic Service Cpmposition Model	4.4.4 Model for Federated Ecosystems
4.3.3 Conceptual Service Composition Model	4.5 Cross-Ecosystem Interoperability

Previous version ToC	Current New ToC
4.4 Identity and Access Management	4.6 Inter-Ecosystem Interoperability
4.4.1 Dataspace / Federation onboarding and offboarding	4.7 Services and Service Composition
4.5 Data Products and Data Exchange Services	4.7.1 Resources and Sevice Offerings
4.5.1 Data Product Conceptual Model	4.8 Policies
4.5.2 Cascading agreement and right to oblivion	4.8.1 Policy Definition 6.2.1 Using Identities in Gaia-X Credentials
4.5.3 Data Intermediaries	4.8.2 Policy Description
4.6 Gaia-X Trust Anchors	4.8.3 Gaia-X Policy Reasoning Engine
4.6.1 Gaia-X Trust Data Sources and Gaia-X Notaries	4.8.4 Policy Decision Point (PDP)
Operating & Services	4.8.5 Rights Delegation
5. Operating Model	4.9 Ecosystem Trust Functions
5.1 Context	4.9.1 Trust Indexes
5.1.1 Prerequisites	4.10 Data Space Architecture using Gaia-X Trust Framework
5.1.2 Scheme Model	5. Gaia-X Implementation of Trusted Data Transactions
5.1.3 Extension by the ecosystems	5.1 Data Product Conceptual Model
5.1.4 Gaia-X Compliance Scheme and Process	5.2 Understanding Data Usage Agreement (DUA)
5.2 Gaia-X Decentralized Autonomous Ecosystem	6. Gaia-X Technical Compatibility Specifications
5.3 Data Usage operating model	6.1 Defining Technical Compatibility
5.3.1 Data Intermediary generic operating model	6.2 Understanding Identity and Identifier
5.4 Service Composition	6.3 Using Linked Data
6. Gaia-X Trust Framework Components	6.4 Using Verifiable Credentials
6.1 Gaia-X Registry	6.5 Verifying Gaia-X Credentials
6.1.1 DNS TXT Records and Naming Convention	6.6 Gaia-X Credential Format
6.1.2 Adopting the Gaia-X Model in Other ecosystems	6.7 OpenID Connect for Verifiable Credentials
6.1.3 Basic Protocol and Interface Specification for GXDCH Registry	6.7.1 OpenID Connect for Verifiable Credentials Issuance
6.2 Gaia-X Compliance	6.7.2 OpenID COnnect for Verifiable Presentations
6.2.1 Basic Protocol and Interface Specification for GXDCH Compliance	6.7.3 Usage
6.3 Gaia-X Notary - LRN	6.7.4 Cloud/Enterprise Wallet
6.3.1 Basic Protocol and Interface Specification for GXDCH Notary	6.8 Understanding Ontologies
6.4 Gaia-X Credential Event Service (CES)	6.8.1 Versioning
6.4.1 Basic Protocol and Interface Specification for GXDCH CES	6.8.2 DCAT
6.5 Graphical overview of the Trust Framework components	6.8.3 ODRL
7. Enabling and Federation Services	6.8.4 CAP
7.1 Wizard Service	6.8.5 Gaia-X Schema
7.2 Credential Manager	6.9 Managing Trust Services

Previous version ToC	Current New ToC
7.3 Federated Catalogues	6.9.1 Trusted Service Operators
7.3.1 Catalogue Service	6.9.2 Using Compliance Engine
7.3.2 Retention	6.9.3 Using the Registry
7.3.3 Snapshot	6.9.4 Using the legal Registration Number (LRN) Notary
7.3.4 Trust Indexes	Appendices
7.4 Notarization Services	7. Trust Indexes
7.5 Data Exchange Services	7.1 Sub-Indexes
7.5.1 Auditing	7.1.1 Veracity
8. Other Concepts	7.1.2 Transparency
8.1 Gaia-X and Data Meshes	7.1.3 Composability
8.1.1 Data Mesh Definition	7.1.4 Semantic Match
8.1.2 Gaia-X as (Higher Order) Service Mesh	8. Changelog
8.2 Computational Contracts	
8.2.1 Concept: Computable Contracts as service	
8.3 Trusted Execution of Services	
Annexes	
9. Changelog	
10. Glossary & References	
11. Annex	
12. Trust Indexes	

9.4 2024 April release (24.04)

- Provide a clear positioning of Gaia-X Trust Framework in the context of other (especially DSBA) data space initiatives
- New view on how the Trust Framework can be used to create domain and ecosystem specific extensions
- Generic Trust Framework process flow and roles for CABS fully supporting the Label definitions in the PRCD
- Introducing Signature and Party credentials (which are defined in the ICAM document)
- Add Protocol, Standards and API into the Gaia-X Services chapter
- Link to the new Technical Specification “Software Architecture”
- Move Credential Event Service to the Gaia-X Services chapter and provide functional specification details
- Clarifications and more precise wording in the Data Exchanges Services chapter based on reader feedback
- Detailed description on the use of policy engines and the link between ODRL and VC
- Added Annex for updated Trust Indexes
- Clarification on Data Usage Agreements as generalisation of Consent (GDPR) and Permission (Data Act)

9.5 2023 October release (23.10)

- Replace “Overview”, which included general concepts, with “Context”, which defines the specific scope of the Gaia-X Architecture
- Update and generalization of the “Conceptual Model”
- Generic description of “Policy Management”
- Mapping of Gaia-X Architecture to the CASCO model
- Updates based on changes in the Data Exchange Services and Identity, Credential and Access Management Document
- Introduction of the “Party-Credential” as extension of Identity and Service Offering credentials with Membership and Service credentials
- Contains the technical implementation details for the Trust Framework (from the 22.10 Trust Framework document)
- Annex includes a section on “Gaia-X Digital Clearing House”

- Chapter “Gaia-X Trust Framework components” including “Gaia-X Compliance”, “Gaia-X Registry”, and “Gaia-X Notary - LRN”
- Inclusion of the “Publication/Subscription service” (ref. Federated Catalogues) and “Trust Indexes” sections in the new chapter “Enabling and Federation Services”

9.6 2022 October release (22.10)

- Update the chapter on Service Composition
- Add a section on Data Mesh

9.7 2022 September release (22.09)

- Move Self-Description technical specs to the next Identity, Credential and Access Management document
- Update Self-Description lifecycle status
- Introduction of Interconnection Point Identifiers
- Introduction of new language terms for Data Exchange
- Update of the Gaia-X schemas and diagrams aligned with the Gaia-X Framework

9.8 2022 April release (22.04)

- Link to Trust Framework document (where the Self-Description mandatory attributes now are)
- Aligning Gaia-X architecture with NIST Cloud Federation Reference Architecture (CFRA)
- Updated definition of Data Exchange Services
- Updated Service Composition and Resource model
- Updated Self-Description Lifecycle
- Consistency and alignment with other officially published Gaia-X documents, streamlining and de-duplication of text to ease reading

9.9 2021 December release (21.12)

- Adding `Contract` and `Computable Contract` definitions in the Conceptual Model
- Update on the Self-Description lifecycle management
- Update on the Federated Trust Model

9.10 2021 September release (21.09)

- Rewrite the Operating model chapter introducing Trust Anchors, Gaia-X Compliance, Gaia-X Labels and Gaia-X Registry.
- Update of Self-Description mandatory attributes in the Appendix.
- Update of `Interconnection`, `Resource` and `Resource template` definitions.
- Gitlab automation improvement and speed-up
- Source available in the 21.09 branch.

9.11 2021 June release (21.06)

- Adding a new Operating model section introducing the first principle for Gaia-X governance.
- Adding preview of Self-Description mandatory attributes in the Appendix.
- Improvement of the Policy rules.
- Improvement of the `Asset` and `Resource` definitions.
- Complete release automation from Gitlab.
- Source available under the 21.06 tag.

9.12 2021 March release (21.03)

- First release of the Architecture document by the Gaia-X Association AISBL
- Complete rework of the Gaia-X Conceptual Model with new entities’ definition.
- Adding a Glossary section.
- Source available under the 21.03-markdown tag.

9.13 2020 June release (20.06)

- First release of the Technical Architecture document by the BMWi