Gaia-x - Architecture Document - 22.04 Release

Table of Contents

1. Overview 4
1.1 Introduction 5
1.2 Objectives 5
1.3 Scope 5
2. Gaia-X Ecosystems 6
2.1 Gaia-X as Enabler for Ecosystems 7
2.2 Goals of Federation Services 8
2.3 Goals of the Gaia-X Trust Framework 9
2.4 Gaia-X Ecosystem 10
3. Gaia-X Conceptual Model 11
3.1 Participants 13
3.2 Service Composition 13
3.3 Resources 14
3.4 Federation Services 15
3.5 Service Offerings, Service Instances and Service Contract 15
3.6 Contract 16
3.7 Additional Concepts 17
3.8 Examples 18
4. Self-Description Definition 19
4.1 Self-Description Structure 20
4.2 Self-Description Schema 23
4.3 Cryptograph Signatures in Self-Descriptions 23
4.4 Self-Description Graph 24
4.5 Self-Description Lifecycle 26
4.6 Self-Description creation 27
5. Gaia-X Operating Model 29
5.1 Gaia-X Ecosystem 30
5.2 Trust Anchors 31
5.3 Gaia-X Trust Framework 31
5.4 Gaia-X Labels 32
5.5 Gaia-X Self-Description 33
5.6 Gaia-X Decentralized Autonomous Ecosystem 37
6. Federation Services 40
6.1 The Role of Federation Services for Ecosystems 41
6.2 Interoperability and Portability for Infrastructure and Data 44
6.3 Infrastructure and Interconnection 44
6.4 Federated Catalogue 45
6.5 Identity and Access Management 45
6.6 Data Exchange services 48
6.7 Gaia-X Federation Servies for Notarization and Credential storage 51
6.8 Portals and APIs 51
7. Example Gaia-X Participant Use Cases 53
7.1 Provider/Consumer Use Cases 54

7.2 Federator Use Cases 54
7.3 Basic Interactions of Participants 54
8. Changelog **56**
8.1 2022 April release 57
8.2 2021 December release 57
8.3 2021 September release 57
8.4 2021 June release 57
8.5 2021 March release 57
8.6 2020 June release 57
9. References **58**

1. Overview

1.1 Introduction

Gaia-X aims to create a federated open data infrastructure based on European values regarding data and cloud sovereignty. The mission of Gaia-X is to design and implement a data sharing architecture that consists of common standards for data sharing, best practices, tools, and governance mechanisms. It also constitutes an EU-anchored federation of cloud infrastructure and data services, to which all 27 EU member states have committed themselves[1]. This overall mission drives the Gaia-X Architecture.[2]

This version of the Gaia-X Architecture document replaces previous version of the document.

1.2 Objectives

This document describes the top-level Gaia-X Architecture model. It focuses on conceptual modelling and key considerations of an operating model and is agnostic regarding technology and vendor. In doing so, it aims to represent the unambiguous understanding of the various Gaia-X stakeholder groups about the fundamental concepts and terms of the Gaia-X Architecture in a consistent form at a certain point in time.

It forms the foundation for further elaboration, specification, and implementation of the Gaia-X Architecture. Thus, it creates in particular an authoritative reference for the Gaia-X Federation Services specification.

The Gaia-X Architecture Document is subject to continuous updates reflecting the evolution of business requirements (e.g., from data space activities in Europe), relevant changes in regulatory frameworks, and advancements in the technological state of the art.

More can be found in the * Vision & Strategy document * Gaia-X Labeling Framework * Gaia-X Trust Framework

Additional information can be found at the publication section of the Gaia-X web-site: Gaia-X List of Publications

1.3 Scope

The Gaia-X Architecture document describes the concepts required to establish the Gaia-X Data and Infrastructure Ecosystem. It integrates the Providers, Consumers, and Services essential for this interaction. These Services comprise ensuring identities, implementing trust mechanisms, and providing usage control over data exchange and Compliance -- without the need for individual agreements.

The Gaia-X Architecture Document describes both the static decomposition and dynamic behaviour of the Gaia-X core concepts and Federation Services.

2. Gaia-X Ecosystems

2.1 Gaia-X as Enabler for Ecosystems

The Gaia-X Architecture enables Data and Infrastructure Ecosystems using the elements explained in the Gaia-X Conceptual Model, the Gaia-X Operational Model and the Federation Services together with the Gaia-X Trust Framework.

An Ecosystem is an organizing principle describing the interaction of different actors and their environment as an integrated whole, like in a biological ecosystem. In a technical context, it refers to a set of loosely coupled actors who jointly create an economic community and its associated benefits.

Gaia-X proposes to structure a Data Ecosystem and an Infrastructure Ecosystem, each with a different focus on exchanged goods and services. Despite each of them having a separate focus, they cannot be viewed separately as they build upon each other, i.e., they are complementary.

The Gaia-X Ecosystem consists of the entirety of all individual Ecosystems that use the Architecture and conform to Gaia-X requirements. Several individual Ecosystems may exist (e.g., Catena-X in the automotive sector) that orchestrate themselves, use the Architecture and may or may not use the Federation Services open source software.

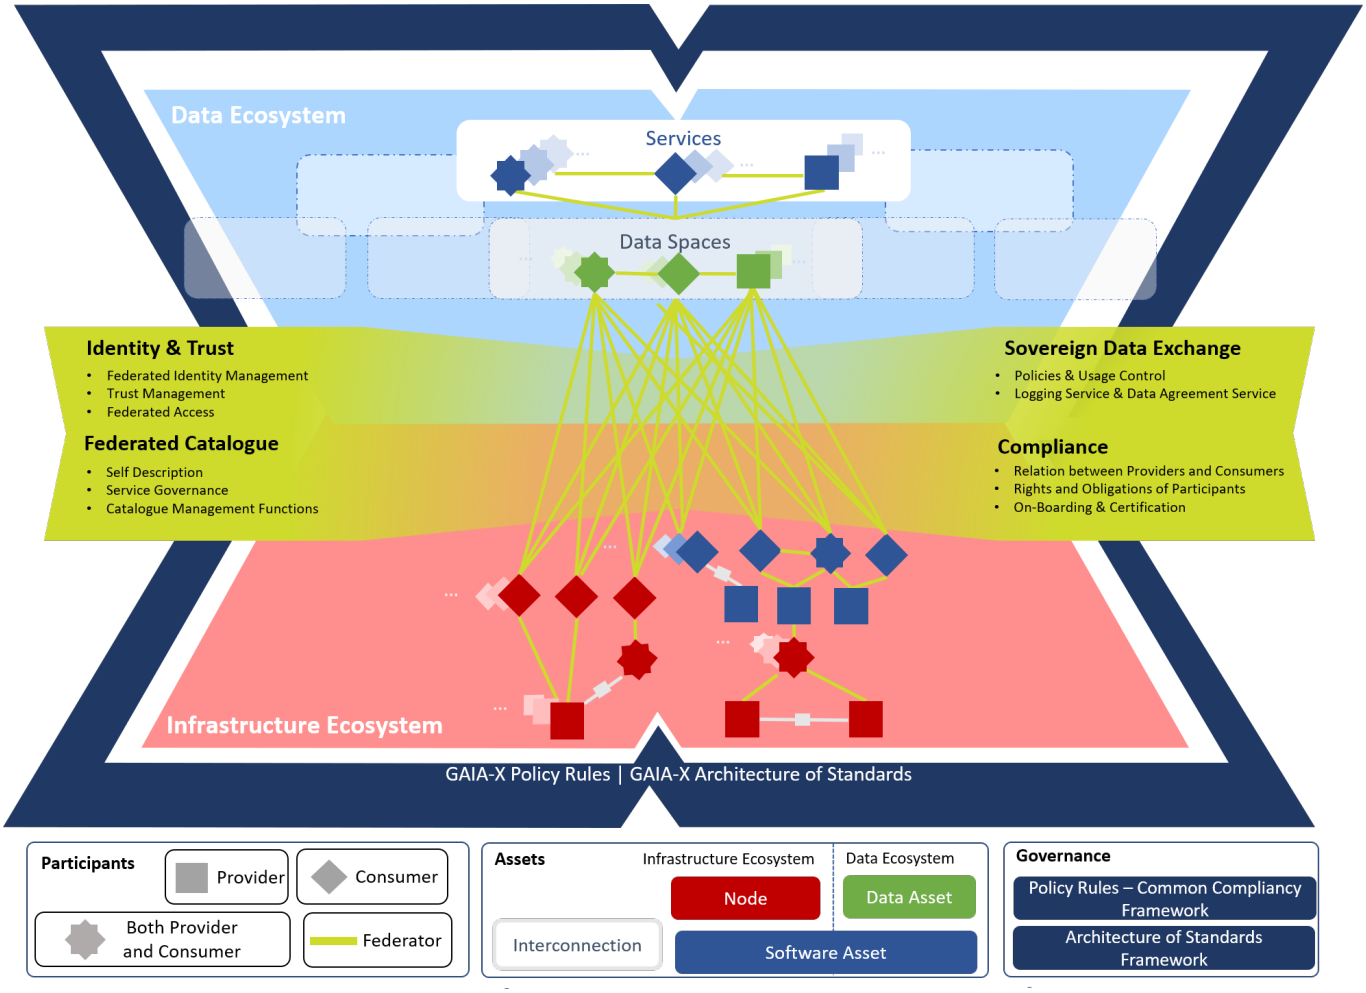

Gaia-X Ecosystem Visualization

The basic roles of Consumer and Provider are visualized as different squares, while the Federator appears as a connecting layer, offering diverse core Federation Services. Federation Services provide connections between and among the different elements as well as between or among the different Ecosystems. The star-shaped

element visualizes that Consumers can act also as Providers by offering composed services or processed data via Catalogues. Governance includes the Policy Rules, which are statements of objectives, rules, practices or regulations governing the activities of Participants within the Ecosystem. Additionally, the Architecture of Standards defines a target for Gaia-X by analysing and integrating already existing standards for data, sovereignty and infrastructure components.

2.2 Goals of Federation Services

Federation Services aim to enable and facilitate interoperability and portability of Resources within and across Gaia-X-based Ecosystems and to provide Data Sovereignty. They ensure trust between or among Participants, make Resources searchable, discoverable and consumable, and provide means for Data Sovereignty in a distributed Ecosystem environment.

They do not interfere with the business models of other members in the Gaia-X Ecosystem, especially Providers and Consumers. Federation Services are centrally defined while being federated themselves, so that they are set up in a federated manner. In this way, they can be used within individual Ecosystems and communities and, through their federation, enable the sharing of data and services across Ecosystems or communities as well as enable the interoperability and portability of data. The set of Ecosystems that use the Federation Services form the Ecosystem.

2.2.1.1 Avoiding Silos

There may be Ecosystems that use the open source Federation Services but do not go through the Compliance and testing required by the Gaia-X Association AISBL. This does not affect the functionality of the Federation Services within specific Ecosystems but would hinder their interaction.

To enable open Ecosystems and avoid "siloed" use of Federation Services, only those that are compliant, interoperable (and tested) are designated as Ecosystems. Therefore, the Federation Services act as a connecting element not only between different Participants, commodities, but also between Ecosystems (see above).

The following table presents how the Federation Services contribute to the Architecture Requirements that are mentioned in section Architecture Requirements.

Requirement	Relation to the Federation Services
Interoperability	◦ The Federated Catalogues ensure that Providers offer services through the whole technology stack. The common Self-Description schema also enables interoperability. ◦ A shared Compliance Framework and the use of existing standards supports the combination and interaction between different Resources. ◦ The Identity and Trust mechanisms enable unique identification in a federated, distributed setting.

Requirement	Relation to the Federation Services
	◦ The possibility to exchange data with full control and enforcement of policies as well as logging options encourages Participants to do so. Semantic interoperability enables that data exchange.
Portability	◦ The Federated Catalogues encourage Providers to offer Resources with transparent Self-Descriptions and make it possible to find the right kind of service that is "fit for purpose" and makes the interaction possible. ◦ The open source implementations of the Federation Services provide a common technical basis and enable movement of Resources in ecosystems and across different ecosystems. ◦ Common compliance levels and the re-use of existing standards support portability of data and services.
Sovereignty	◦ Identity and Trust provide the foundation for privacy considerations as well as access and usage rights. Standards for sovereign data exchange enable logging functions and Usage Policies. The Self-Descriptions offer the opportunity to specify and attach Usage Policies for Data Resources.
Security and Trust	◦ The Architecture and Federation Services provide definitions for trust mechanisms that can be enabled by different entities and enable transparency. ◦ Sovereign Data Exchange, as well as Compliance concerns address security considerations. The identity and trust mechanisms provide the basis. The Federated Catalogues present Self-Descriptions and provide transparency over Service Offerings.

Federation Services match the Architecture Requirements

2.3 Goals of the Gaia-X Trust Framework

Gaia-X AISBL defines a Trust Framework that manifests itself in the form of two services:

- the Gaia-X Registry, detailed in the Operating model chapter
- the Gaia-X Compliance service, as the service implementing the set of rules described in the upcoming Gaia-X Trust Framework document.

2.4 Gaia-X Ecosystem

The Gaia-X Ecosystem is the virtual set of Participants, Service Offerings, Resources fulfilling the requirements of the Gaia-X Trust Framework.
There is one Gaia-X Ecosystem federating independent autonomous existing and future ecosystems.

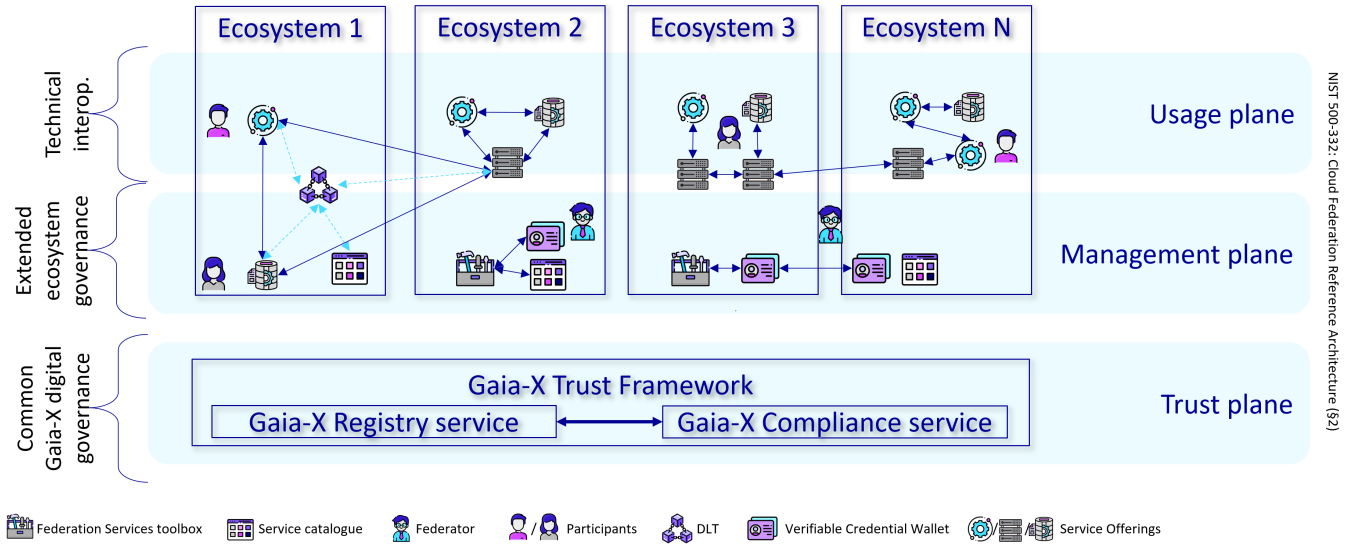

Gaia-X planes

The three planes represent three levels of interoperability and match the planes as described in the NIST Cloud Federation Reference Architecture chapter 2.

2.4.1 The Trust plane

The Trust plane represents the global digital governance that is shared across ecosystem. The rules of this common governance are captured by the Trust Framework and operationalized by two services:

- the Gaia-X Compliance service, described in the Gaia-X Trust Framework
- the Gaia-X Registry service, describe in the Operational Model chapter of this document

2.4.2 The Management plane

The Management plane represents an extension of the common digital governance provided by the Federators of the relevant ecosystems.
It includes potential contract templates specific to a vertical market.
For example, a finance or a health ecosystem will have additional rules.

Specific ecosystem governance rules are out of scope for Gaia-X.

2.4.3 The Usage plane

The Usage plane is the one capturing technical interoperability, including the one between Service Offerings.

3. Gaia-X Conceptual Model

The Gaia-X Conceptual Model, shown in the figure below, describes all concepts in the scope of Gaia-X and relations among them. Supplementary, more detailed models may be created in the future to specify further aspects. Minimum versions of important core concepts in the form of mandatory attributes for Self-Descriptions are specified in the Gaia-X Trust Framework. The general interaction pattern is further described in the section Basic Interactions of Participants.

The Gaia-X core concepts are represented in classes. An entity highlighted in blue shows that an element is part of Gaia-X and therefore described by a Gaia-X Self-Description. The upper part of the model shows different actors of Gaia-X, while the lower part shows elements of commercial trade and the relationship to actors outside Gaia-X.

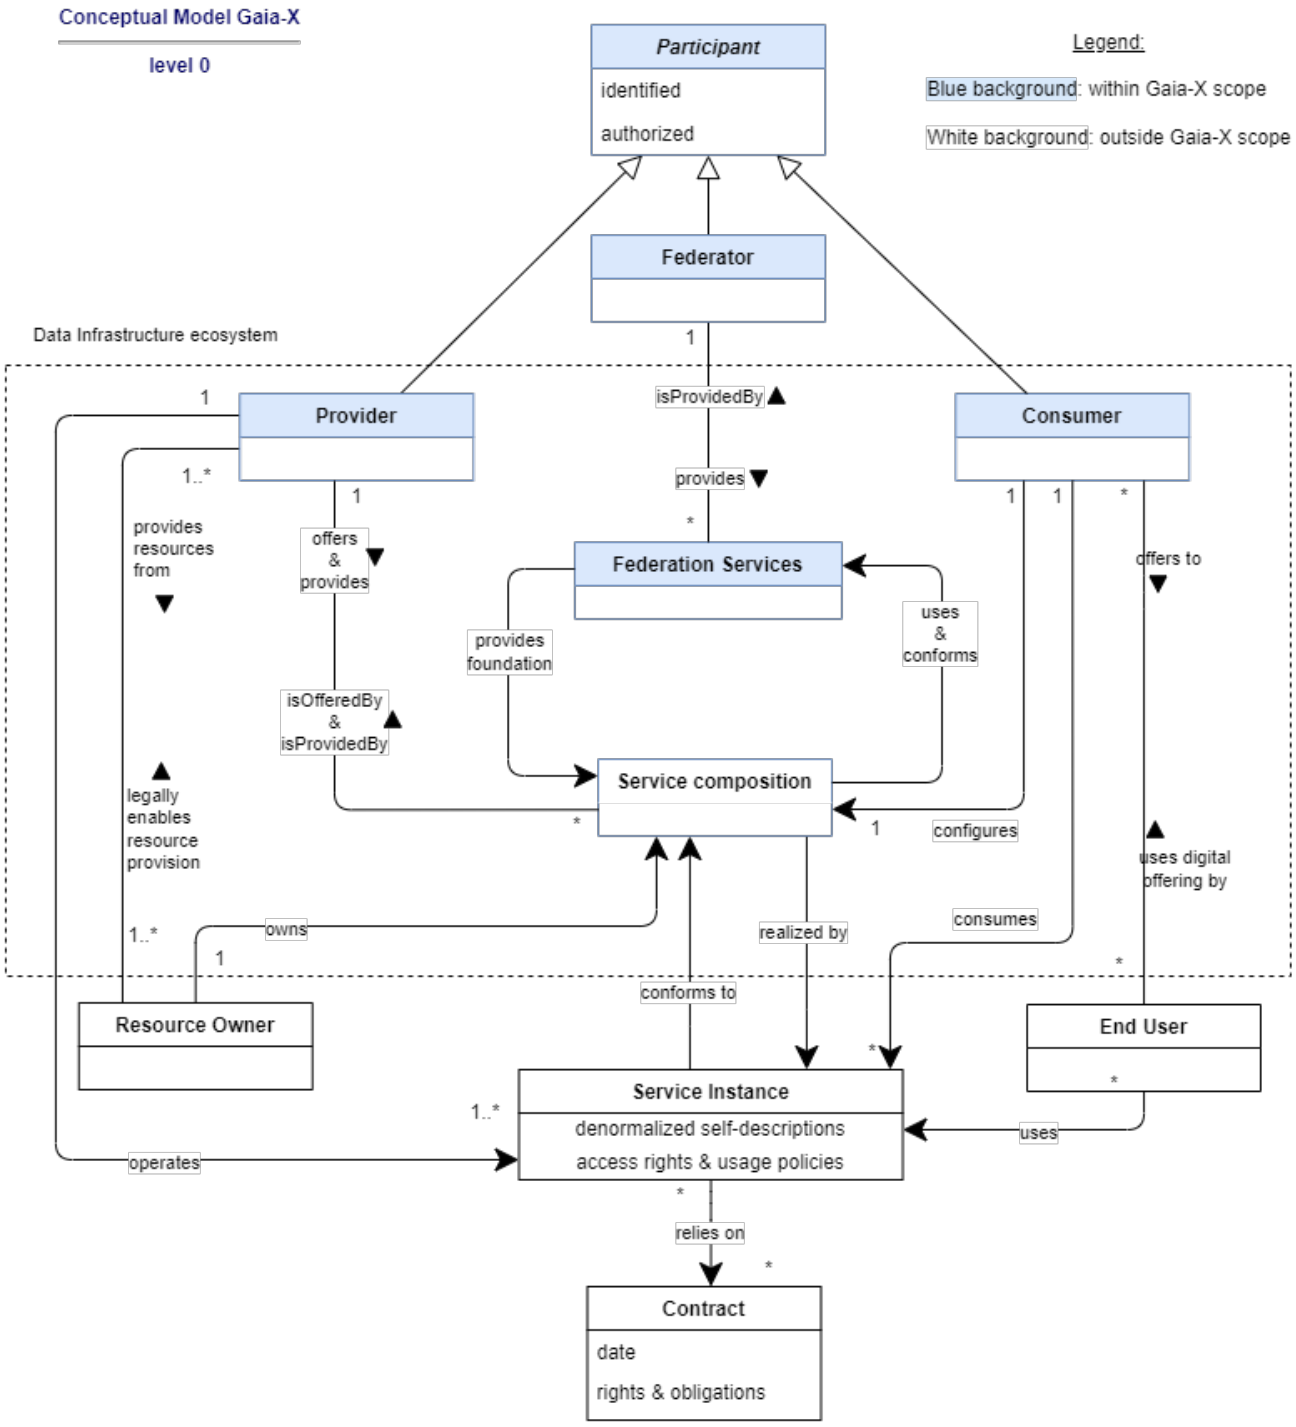

Gaia-X conceptual model

3.1 Participants

A Participant is an entity, as defined in ISO/IEC 24760-1 as "item relevant for the purpose of operation of a domain that has recognizably distinct existence"[3], which is onboarded and has a Gaia-X Self-Description. A Participant can take on one or more of the following roles: Provider, Consumer, Federator. Section Federation Services demonstrates use cases that illustrate how these roles could be filled. Provider and Consumer present the core roles that are in a business-to-business relationship while the Federator enables their interaction.

3.1.1 Provider

A Provider is a Participant who operates Resources in the Gaia-X Ecosystem and offers them as services. For any such service, the Provider defines the Service Offering including terms and conditions as well as technical Policies. Furthermore, it provides the Service Instance that includes a Self-Description and associated Policies.

3.1.2 Federator

Federators are in charge of the Federation Services and the Federation, which are independent of each other. Federators are Gaia-X Participants. There can be one or more Federators per type of Federation Service.

A Federation refers to a loose set of interacting actors that directly or indirectly consume, produce, or provide related Resources.

3.1.3 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.

3.2 Service Composition

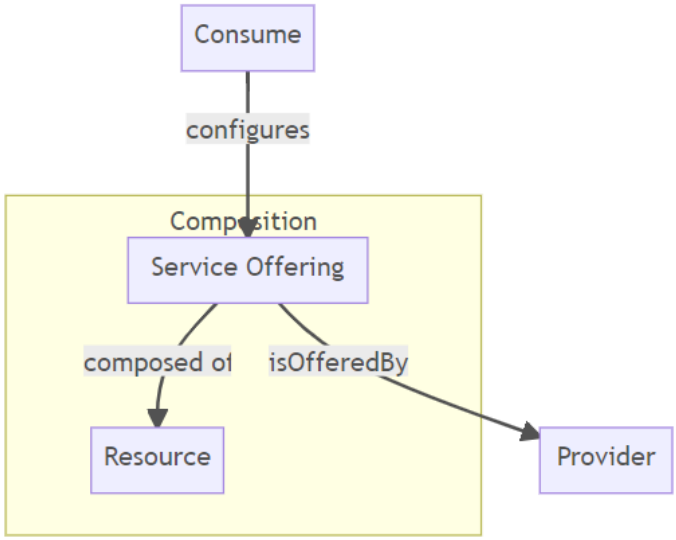

Gaia-X conceptual model

3.3 Resources

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

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

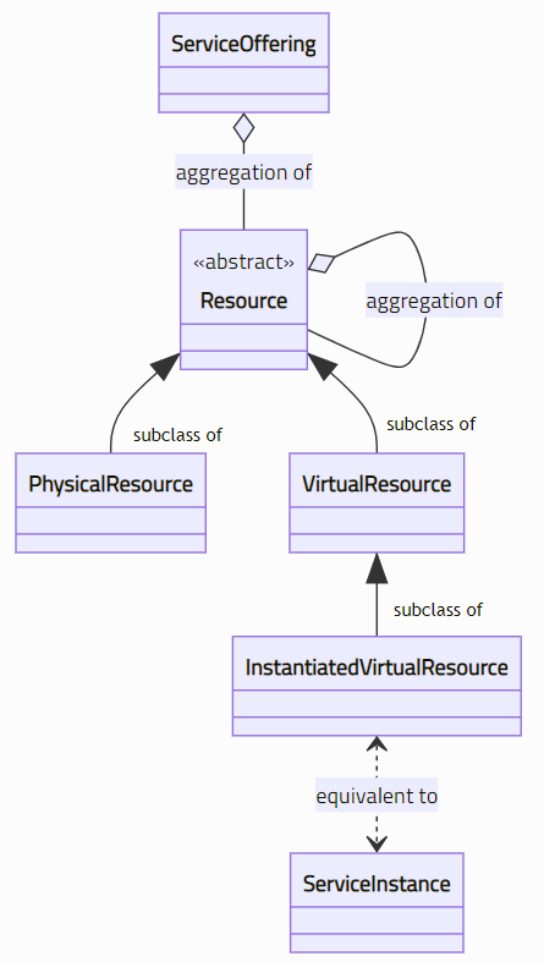

Resource Categories

A Resource can be a:

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

3.3.1 Policies

Policy is defined as a statement of objectives, rules, practices, or regulations governing the activities of Participants within Gaia-X. From a technical perspective Policies are statements, rules or assertions that specify the correct or expected behaviour of an entity[45].

The Policy Rules Document explains the general Policies defined by the Gaia-X Association for all Providers and Service Offerings. They cover, for example, privacy or cybersecurity policies and are expressed in the

conceptual model indirectly via Gaia-X Federation Service Compliance and as attributes of the Resources, Service Offerings, and Service Instances.

These general Policies 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 of 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 fulfil a certain service level.*

In the Conceptual Model, they appear as attributes in all elements related to Resources. The specific Policies have to be in line with the general Policies in the Policy Rules Document.

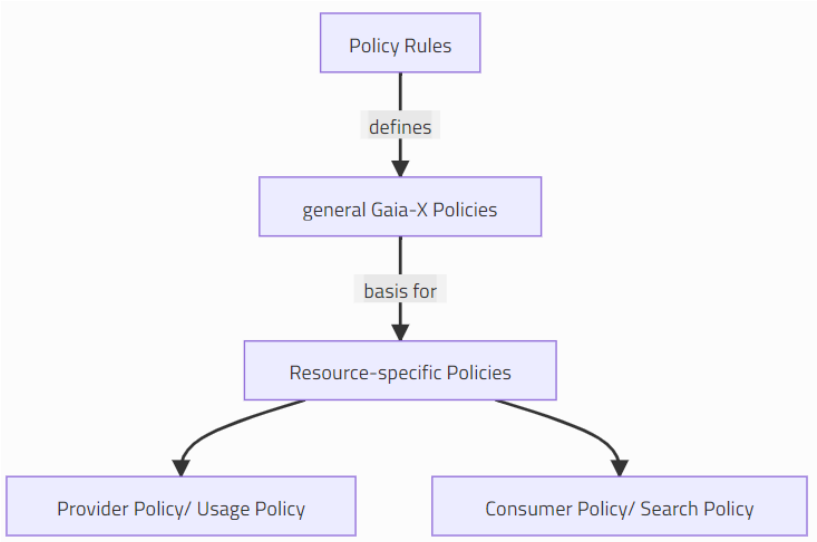

3.4 Federation Services

Federation Services are services required for the operational implementation of a Gaia-X Data Ecosystem. They are explained in greater detail in the Federation Services section.

They comprise four groups of services that are necessary to enable Federation of Resources, Participants and interactions between Ecosystems. The four service groups are Identity and Trust, Federated Catalogue, Sovereign Data Exchange and Compliance.

3.5 Service Offerings, Service Instances and Service Contract

A Service Offering is defined as a set of Resources, which a Provider aggregates and publishes as a single entry in a Catalogue. Service Offerings may themselves be aggregated realizing Service Composition.

A Service Instance is the instantiation of a Service Offering at runtime, strictly bound to a version of a Self-Description.

During the ordering phase, the Provider is invited to generated a denormalized version of a self-description for the newly instantiated Service Instance. The format and content of this self-description is out of scope for this document.

A Contract is an agreement between a Consumer and a Provider, to allow and regulate the usage of one or more Service Instances. It is related to a specific version of a Service, from which it derives the attributes of the Service Instances to be provisioned. The Contract has a distinct lifecycle from the Service Offering and additional attributes and logic, which are out of scope of Gaia-X architecture document.

3.6 Contract

The Gaia-X Association is not getting involved into the realisation of the `Contract`. However, in order to ease participants with the establishement and to enter into a contractual relationship, we are defining below a common model for `Contract`.

3.6.1 Concept: Computable Contracts as a service

- Contracts are the basis for business relationships.
- Whereas a licensor has rights with respect to a resource and is willing to (sub-)license such rights by a defined set of conditions.
- Whereas a licensee would like to get license rights with respect to a resource by a defined set of conditions.
- Licensor and licensee agree on it in form of a contract.
- Every role of the Gaia-X Conceptual Model as well as of operational model can be seen as legal persons and therefore may have a role as a licensor or licensee or both.
- In traditional centralized driven ecosystems the platform provider which is very often the ecosystem owner, defines the contractual framework and participants need to accept without any possibility for negotiation.
- In distributed and federated ecosystems individual contracting becomes much more important to support individual content of contractual relations, e.g., individual sets of conditions.
- The ability to negotiate on contracts is key for a sovereign participation. The ability to observe if all parties of a contract behave the way it is agreed, to validate their rights, to fulfill their obligations and ensure that no one can misuse information is key for a trustful relationship.
- Computable contracts aim to easy the complex processes of contract design, contract negotiation, contract signing, contract termination as well as to observe the fulfillment of contractual obligations and compliance with national law.

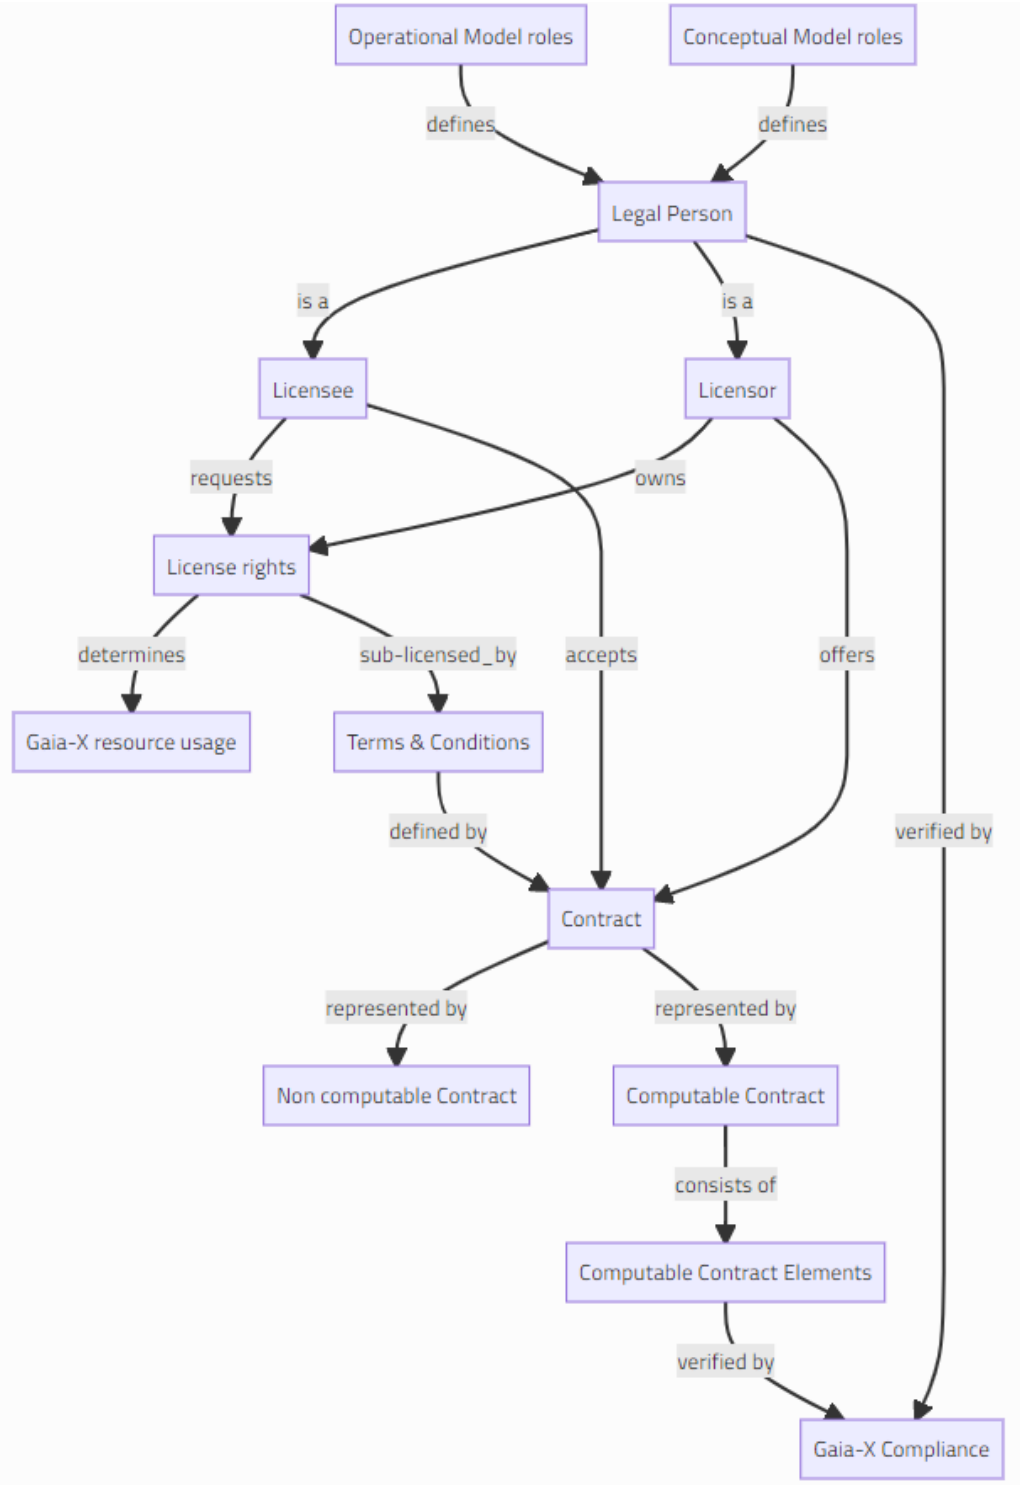

3.7 Additional Concepts

In addition to those concepts and their relations mentioned above, further ones exist in the conceptual model that are not directly governed by Gaia-X. These concepts do not need to undergo any procedures directly related to Gaia-X, e.g., do not create or maintain a Gaia-X Self-Description.

First, the Service Instance realizes a Service Offering and can be used by End-Users while relying on a contractual basis.

Second, Contracts are not in scope of Gaia-X but present the legal basis for the Services Instances and include specified Policies. Contract means the binding legal agreement describing a Service Instance and includes all

rights and obligations. This comes in addition to the automated digital rights management embedded in every entity's Self-Description.

Further relevant actors exist outside of the Gaia-X scope in terms of End-Users and Resource Owners.

Resource Owners describe a natural or legal person, who holds the rights to Resources that will be provided according to Gaia-X regulations by a Provider and legally enable its provision. As Resources are bundled into a Service Offering and nested Resource compositions can be possible, there is no separate resource owner either. Resources can only be realized together in a Service Offering and Service Instance by a Provider, which presents no need to model a separate legal holder of ownership rights.

End-Users use digital offerings of a Gaia-X Consumer that are enabled by Gaia-X. The End-User uses the Service Instances containing Self-Descriptions and Policies.

3.8 Examples

3.8.1 Personal Finance Management example

This example describes the various Gaia-X concepts using the Open Banking scenario of a Personal Finance Management service (PFM) in SaaS mode.

Suppose that the service is proposed by a company called `MyPFM` to an end user `Jane` who has bank accounts in two banks: $Bank_1$ and $Bank_2$.
`MyPFM` is using services provided by $Bank_1$ and $Bank_2$ to get the banking transactions of `Jane` and then aggregates these bank statements to create Jane's financial dashboard.

`Jane` is the **End-User**.

$Bank_1$ and $Bank_2$ are **Providers** defining the **Service Offerings** delivering the banking transactions and operating the corresponding **Service Instances**. They are also **Resource Owners** for the bank statements, which are **Resources** composing the **Service Offerings** (`Jane` is the data subject as per GDPR[6]).
The associated **Resource Policies** are in fact predefined by the PSD2[7] directive from the European Parliament.

`MyPFM` is the **Consumer** which consumes the **Service Instances** provided by $Bank_1$ and $Bank_2$ in order to create a financial dashboard and to offer it to `Jane`.
`MyPFM` is also likely consuming **Service Instances** from a PaaS **Provider** in order to run its own code, such as dashboard creation.

4. Self-Description Definition

Gaia-X Self-Descriptions (*SD*) describe Entities from the Gaia-X Conceptual Model in a machine interpretable format. This includes Self-Descriptions for the Participants themselves, as well as the Resources and Service Offerings from the Providers. Well-defined Self-Description Schemas (which can be extended by the Federations for their domain) enable ensuring a unified representation of the Self-Descriptions. The Self-Description allows to find and compare Entities inside Gaia-X.

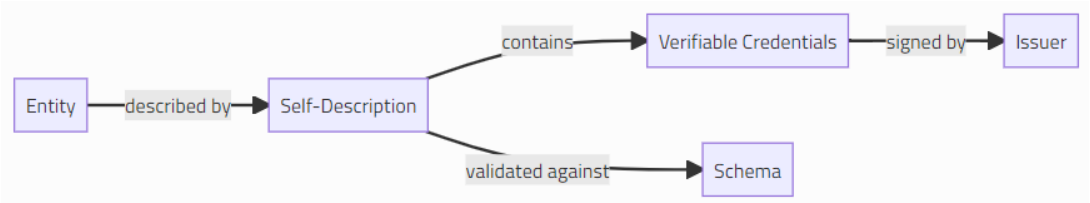

Overview on Self-Descriptions. The terms Verifiable Credential, Schema and Issuer are explained in more detail in the following sections.

Self-Descriptions in combination with trustworthy verification mechanisms empower Participants in their decision-making processes. Specifically, Self-Descriptions can be used for:

- Tool-assisted evaluation, selection, composition and orchestration of Services and Resources
- Enforcement, continuous validation and trust monitoring together with usage policies
- Negotiation of contractual terms

The Participants (particularly Providers) are responsible for the creation of their Self-Descriptions. In addition to self-declared information by Participants about themselves or their offerings, a Self-Description may comprise statements issued and signed by trusted parties.

4.1 Self-Description Structure

Self-Descriptions are W3C Verifiable Presentations in the JSON-LD format. Self-Description consist of a list of Verifiable Credentials. Verifiable Credentials themselves contain a list of Claims: assertions about Entities expressed in the RDF data model. Both Verifiable Credentials and Verifiable Presentations come with cryptographic signatures to increase the level of trust. Note that the Verifiable Credentials inside a Self-Description may be signed from different (trusted) parties. For example, a certification assessment body may assert a certification result in a Verifiable Credential. This can then be included in a Self-Description for that service.

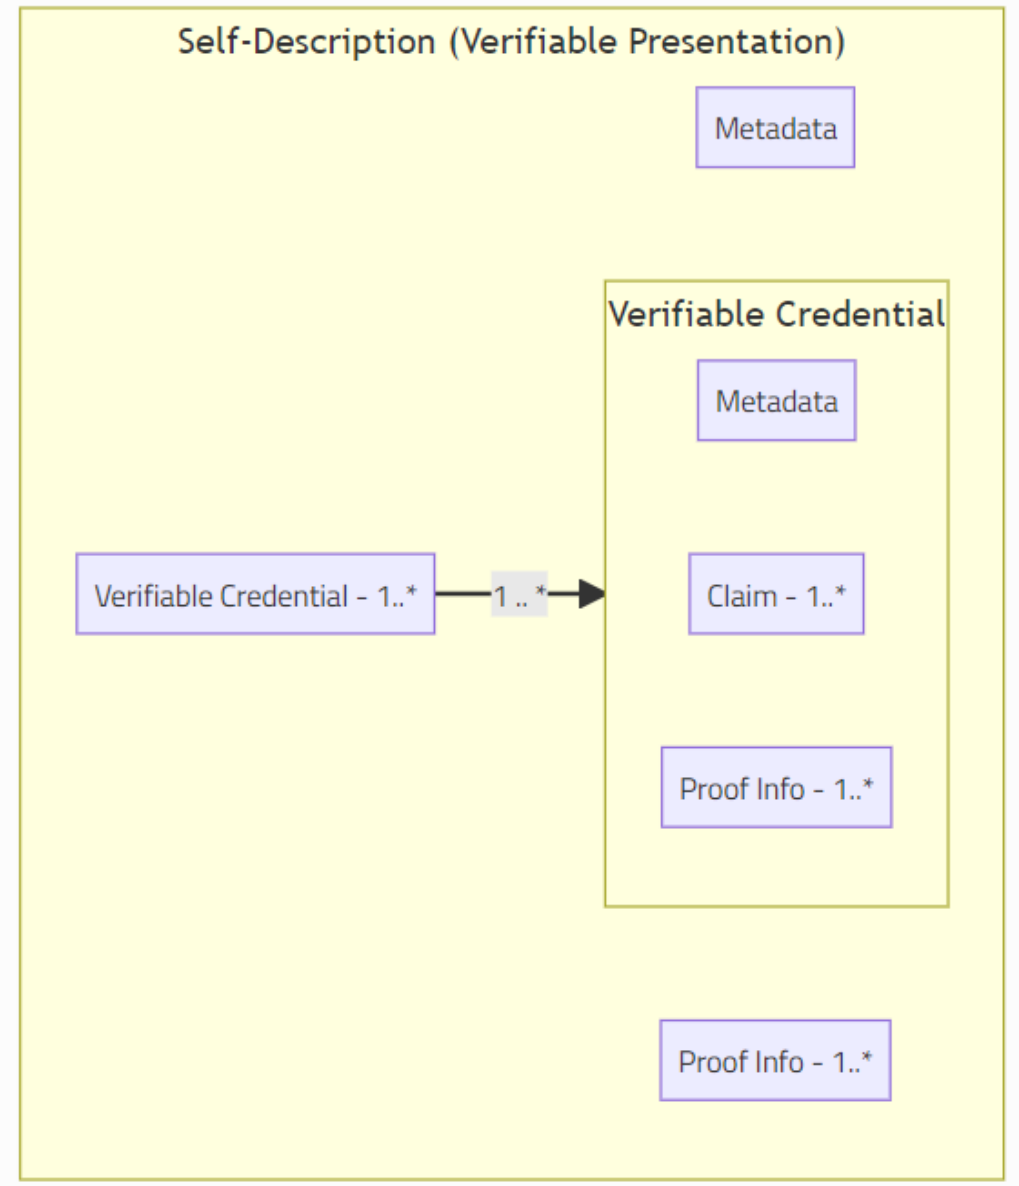

Self-Description assembly model

Verifiable Presentations and Verifiable Credentials can be expressed as graph. Below is an example for a Verifiable Presentation Graph, where a trusted party "DUV" asserts that a resource is certified according to ISO 27001.

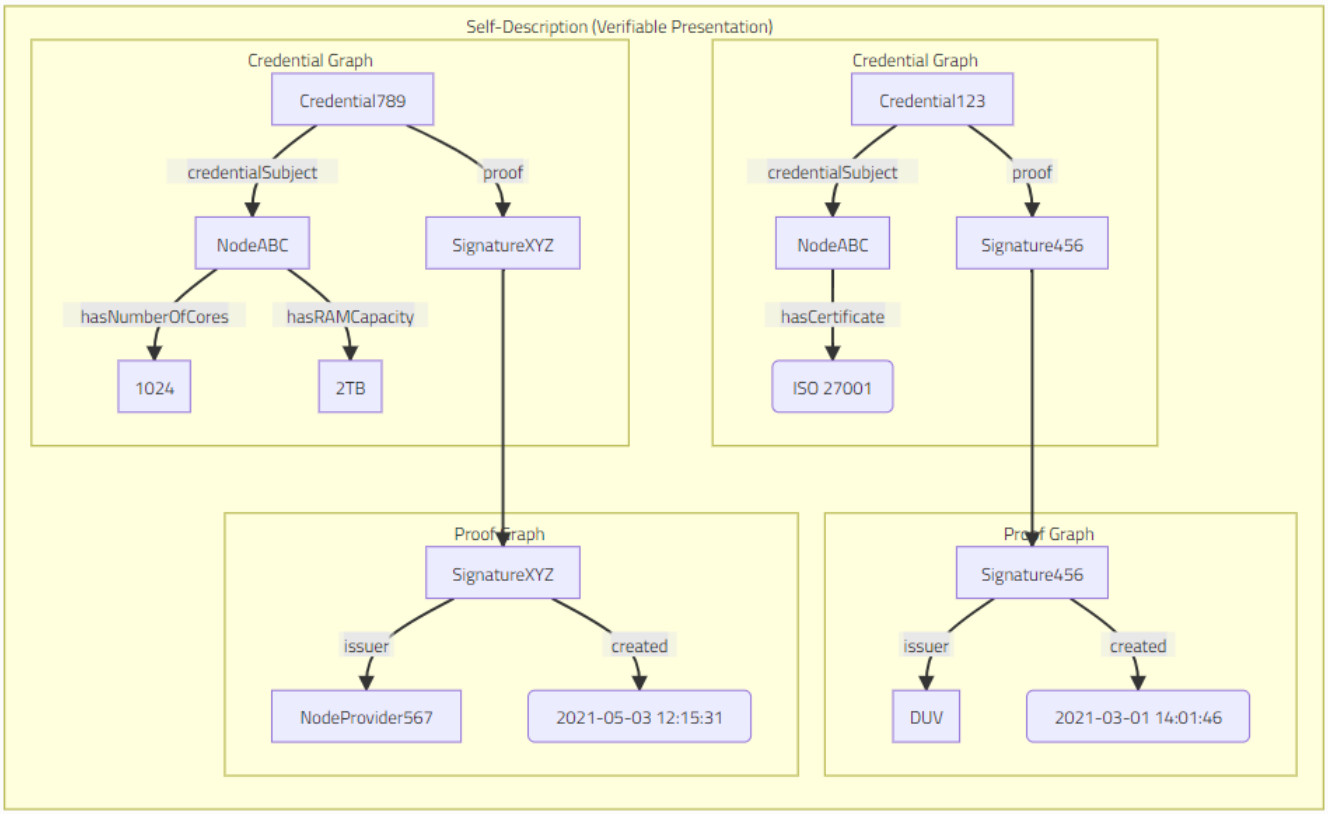

Example for Verifiable Credentials from different issuers in the same Self-Description. The `DUV` *organization asserts the certification of a Node according to ISO 27001. The provider himself provides additional technical details. The individual elements and their relation are shown as a graph (non-normative visualization).*

Self-Description contain verifiable credentials about the attributes of Entities and relations to other Entities based on subject-predicate-object triples (cf. the RDF data model). The possible attributes and relations to be used in a Self-Description come from Self-Description Schemas (see next section). Leaving out the syntactic sugar of JSON-LD, the following triples represent the payload information about the Entity NodeABC from the above figure. Typically the Provider `NodeProvider567` and the `DUV` organization have dedicated Self-Descriptions for their respective identifiers.

Each of these assertions is called a claim in the Verifiable Credentails data model.

Cross-referencing between Self-Descriptions is enabled by unique Identifiers for the Entities. Identifiers in Gaia-X are URIs and follow the specification of RFC 3986. Depending on the prefix of the URI, different technical systems are defined to ensure uniqueness of Identifiers. For example, the use of a domain-name with methods like DID:DNS as part of the Identifier, enables the domain owner to control the Identifiers by itself, eg. did:dns:example.com#z6MljvBkt9ETnxJGAFPKGgYHb33q9oNHLX7BiYSPcXqG6gZ9.

Every Self-Description has one Entity as its main topic. This Entity must have a Gaia-X compliant Identifier. Self-Descriptions can additionally describe "anonymous entities" if they are required for the Self-Description and if these Entities do not yet have their own Self-Description. The main reason for this is to give mandatory information which would go into a dedicated Self-Description, but which is not available so far. (Take as an example that the provider company of the hosting infrastructure must be described). These "anonymous entities" are defined via blank-nodes in the RDF data model and do not have identifiers. Anonymous entities from different Self-Descriptions are not merged when they are loaded into a joint Self-Description Graph (see below). So there can be duplicate anonymous entities.

4.2 Self-Description Schema

Self-Description Schema is a collection of class's data schema describing Gaia-X entities. Each data schema is part of an inheritance hierarchy having Participant, Service Offering or Resource as top-level super-class schema and defining a set of attributes available to describe the entity. Self-Descriptions must follow a common structure and well-defined semantics. Only by this it can be assured that entities can be found and compared within Gaia-X. This structure is formally described in Self-Description Schemas.

The basic set of Self-Description data schemas is defined within the Service Characteristics Working Group. Individual Gaia-X Federations can extend the schema for their application domain. Such extensions must make an explicit reference to the organization that is responsible for the development and maintenance of the extension. The Self-Description Schema defines entities that are recognized within Gaia-X. Those entities form an inheritance structure, whereas each Entity inherits from one Entity of the Conceptual Model. We call this inheritance hierarchy **Self-Description Taxonomy**. Derived classes substantiate the basic Entities of the Conceptual Model with more detailed information. For each class, properties are defined that an instance of this class can have. Those properties include *attributes*, which can have * plain values of a primitive datatype (called *datatype properties* in the W3C Web Ontology Language OWL[8] used to define terms to be used in the RDF data model), * values that are instances of auxiliary classes (e.g., a class describing an address, containing attributes like city and street), or * values reused by referencing a controlled vocabulary (well-defined terms within Gaia-X); the latter two being called *object properties* in OWL.

Properties also include *relationships* used for referencing another entity inside the Gaia-X Federation. In addition to the allowed attributes and their types, the Self-Description schema defines the cardinality of each attribute. Meaning: is the attribute mandatory, so there must be at least one value for the attribute and is it allowed to have multiple values for the attribute?

Multiple inheritance is allowed and encouraged in the Self-Description Taxonomy as well as for the instances in a Self-Description. That way, deeply nested specializations in the schema hierarchy can be avoided. For example, a REST-based database service could inherit from both *database* and *REST-based service* instead of creating a specialized class.

Gaia-X aims at building upon existing schemas, preferably those that have been standardized or at least widely adopted.

For frequently used attribute values, it is recommended that they be maintained in the same governed process as Self-Description Schemas, i.e., by giving them unambiguous identifiers maintained in Controlled Vocabularies. Examples include standards against which a Participant or a Resource has been certified, or classification schemes, e.g., for the sector in which a Provider is doing their business.

4.3 Cryptograph Signatures in Self-Descriptions

To ensure the Self-Description's integrity and authenticity the overall Self-Description (a Verifiable Presentation) must be cryptographically signed by the Participant that is issuing the Self-Description. This is done to avoid tampering and to technically allow to check the origin of the Self-Description. Inside the Self-Description each Verifiable Credential must be individually signed as well (this can be a third party different from the issuer of the overall Self-Description). A Verifiable Credential can also be signed by multiple parties to increase the trust level.

The signature mechanism used is Linked Data Signature with the `JsonWebKey2020` suite. It generates a JSON object to be included in the Verifiable Credential or Presentation. The JSON object comprises of the following fields:

- `type`: The field is set to `"JsonWebKey2020"`. See https://w3c-ccg.github.io/lds-jws2020/ for more details.
- `proofPurpose`: The field is set to `"assertionMethod"`.
- `verificationMethod`: The field identifies the party that has issued the proof. It contains either a) the Gaia-X compliant Identifier of the Participant that is signing or b) the digest (fingerprint) of an X509 certificate containing the key material. The Gaia-X compliant Identifier of the signing party can be resolved to a Self-Description that contains the public key used for the signature. An X509 certificate digest is to be matched to known certificates to resolve the public key. For example certificates from the Gaia-X Registry.
- `created`: The field contains the creation date in the ISO8601 format.

A Verifiable Credential is *Gaia-X conformant* if the Issuer of the Verifiable Credential has itself an identity coming from one of the Trust Anchors. See the section on the Trust Framework for more details on Trust Anchors and the Registry.

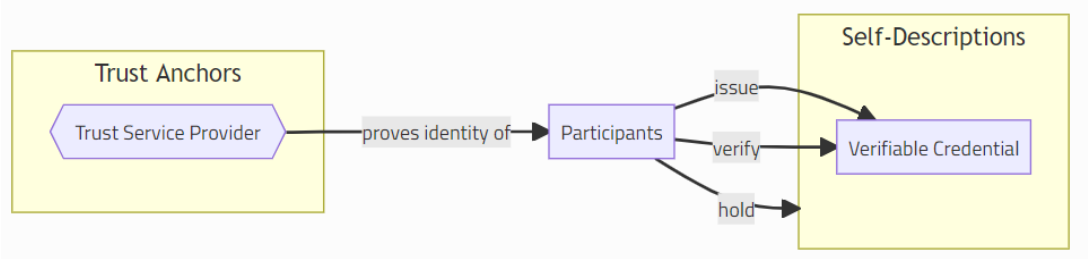

4.4 Self-Description Graph

4.4.1 Relations between Self-Descriptions

A **Self-Description of one Entity may contain typed relations to other Entities**. For instance, the Self-Description of Entity `ServiceA` may specify that this service is hosted on `DataCenterB` (another Entity) formalized as RDF triple `(ServiceA, hostedOn, DataCenterB)`. Entities (e.g., `DataCenterB`) are referred to by their respective Identifier.

In the example, the property `hostedOn` constitutes the type of the relation linking the two Self-Descriptions. Many other relation types between different Entities (not just between Services and Resources) are conceivable such as `providedBy` or `operatedBy` and are defined in the respective Self-Description data schemas. Relations between Entities may also cross organizational boundaries, for instance, when `DataCenterB` is `operatedBy` a different Participant than `ServiceA`.

The formalized (in RDF) relations between Self-Descriptions always have a **direction** (viz., from *subject* to *value*) and Self-Descriptions shall only contain one direction as defined in the respective Self-Description Schema. Reverse relations such as `(DataCenterB, hosts, ServiceA)` shall not be included in the Self-Description of `DataCenterB` but may be inferred and queried by suitable systems (e.g., a Catalogue) evaluating all Self-Descriptions and their relations. This also holds for (semantically) symmetric relations such as `collocatedTo` or `nearTo` where only one direction shall be used in the Self-Descriptions.

The set of all Self-Descriptions in the "active" lifecycle state and all their typed relations with each other are called the **Self-Description Graph** (with Self-Descriptions as the vertices of the graph and the relations between Self-Description its edges). By following the relations between Self-Descriptions in the Self-Description Graph, advanced queries across individual Self-Descriptions become possible. Such functionality will be typically implemented by Catalogues and may further involve Certification aspects and Usage Policies relating to the Self-Descriptions. This will enable, for instance, that a Consumer can use Catalogue Services to require that a Service Instance must not depend on other Service Instances deployed on Nodes outside a Consumer-specified list of acceptable countries.

4.4.2 Representing Claims in the Self-Description Graph

Self-Descriptions are Verifiable Presentations collecting Verifiable Credentials about the Entity from various issuers with metadata about the issuer itself, date and time of issuance, expiry date, and so on. In order to also include this information in the Self-Description Graph (e.g., to allow searches and queries based on this metadata), the simple graph model indicated above (`Self-Description —-property--> Self-Description`) is extended by permitting so-called **"edge properties"**. This means that the edges of the Self-Description Graph are endowed with additional attributes besides their type such as the origin of the claim, the issuer, and others.

Graphs with edge properties are supported by labelled-property graph databases (e.g., based on GQL[9] or OpenCypher) as well as in semantic triple-stores that support the RDF-Star/SPARQL-Star extension[10].

Labelled-Property Graph Representation. Labelled-property graphs allow edge properties as shown in the following (simplified) example.

```
(NodeProvider567, provides[claimedBy: NodeProvider567], NodeABC)
(NodeABC, hasCertificate[claimedBy: DUV], [certificateType:"ISO27001"])
```

RDF-Star Representation. In RDF-Star, edge properties are added by using the respective original triples (1 and 2 below) as the subject of another triple (3 and 4 below):

```
[1]   (NodeProvider567, provides, NodeABC)
[2]   (NodeABC, hasCertificate, "ISO27001")

[3]   ( (NodeProvider567, provides, NodeABC), claimedBy, NodeProvider567 )
[4]   ( (NodeABC, hasCertificate, "ISO27001"), claimedBy, DUV )
```

Such a representation allows **complex queries** mixing normal properties of Entities with metadata about the claims.

```
GQL
MATCH (provider)-[:provides]->(node),
      (node)-[rel:hasCertificate]->(certificate)
WHERE certificate.certificateType = "ISO27001",
      rel.claimedBy = DUV
RETURN provider, node
```

```
SPARQL-Star
SELECT ?provider, ?node
WHERE { ?provider provides ?node .
        ?node hasCertificate "ISO27001" .
        <<?node hasCertificate "ISO27001">> claimedBy DUV }
```

The approach has the advantage that users can both, write simple queries that do not consider claim structure information, and also complex queries that take the claims into account. Note, that simple queries are not impacted by the presence (or absence) of claim structure or metadata information.

The disadvantage is that fewer query languages support edge properties. For example, SPARQL-Star (in draft status) would have to be used instead of SPARQL. Another disadvantage is that only claims "one level deep" can be represented straightforwardly via edge properties. Deeply nested claims (such as Bob attesting that Alice has verified that a particular person owns a red car) would require a more involved representation.

4.5 Self-Description Lifecycle

Since Self-Descriptions are protected by cryptographic signatures, they are immutable and cannot be changed once published. The lifecycle state is therefore managed outside of the Self-Description itself, like for example and not limited to, Catalogues.

The lifecycle of the Self-Description depends on the lifecycle of the verifiable credentials that are contained within it. Furthermore both depend on the lifecycle status of the certificate (public/private key-pair) used to sign it. The following table shows the possible states in the lifecycle of Self-Descriptions, Credentials and certificates.

State	Self-Description	Verifiable Credential	Certificate (Key Pair)	Comment
Active	x	x	x	All Verification Rules are passed.
Partially Active	x	-	-	Some claims inside the Self-Description are not in the active state.
Expired	x	x	x	Too old and must be renewed.
Deprecated	x	-	x	Replaced by a newer version.
Revoked	x	x	x	Revoked by the issuer or a trusted party.

State	Self-Description	Verifiable Credential	Certificate (Key Pair)	Comment
Inconsistent	x	x	-	Internally inconsistent or incompatible with existing information.
Unverifiable	x	x	x	The trust level could not be verified.

*Table with the possible lifecycle states for Self-Descriptions, Verifiable Credentials and Certificates (Key Pairs). The lifecycle is the result of Verification Rules. The full list of Verification Rules and the aggregation of their result into an overall state is the subject of ongoing specification work. The `x` indicates support for the lifecycle state, a `-` indicates that the state is not supported for either Self-Description, Verifiable Credential or Certificate. *

4.6 Self-Description creation

4.6.1 Collecting claim

The first step is to collect Verifiable Credentials (= set of claims signed by its issuer). Claims can be self-signed (by using a keypair issued by a Trust Anchor) or directly signed by a Trust Anchors. Claims can be self-signed using a keypair issued to the creator of the Self-Description by one of the Trust Anchors, or it can be signed by a third-party who is one of the Trust Anchors.
The list and scope of each member of Trust Anchors are defined in the Gaia-X Trust Framework.

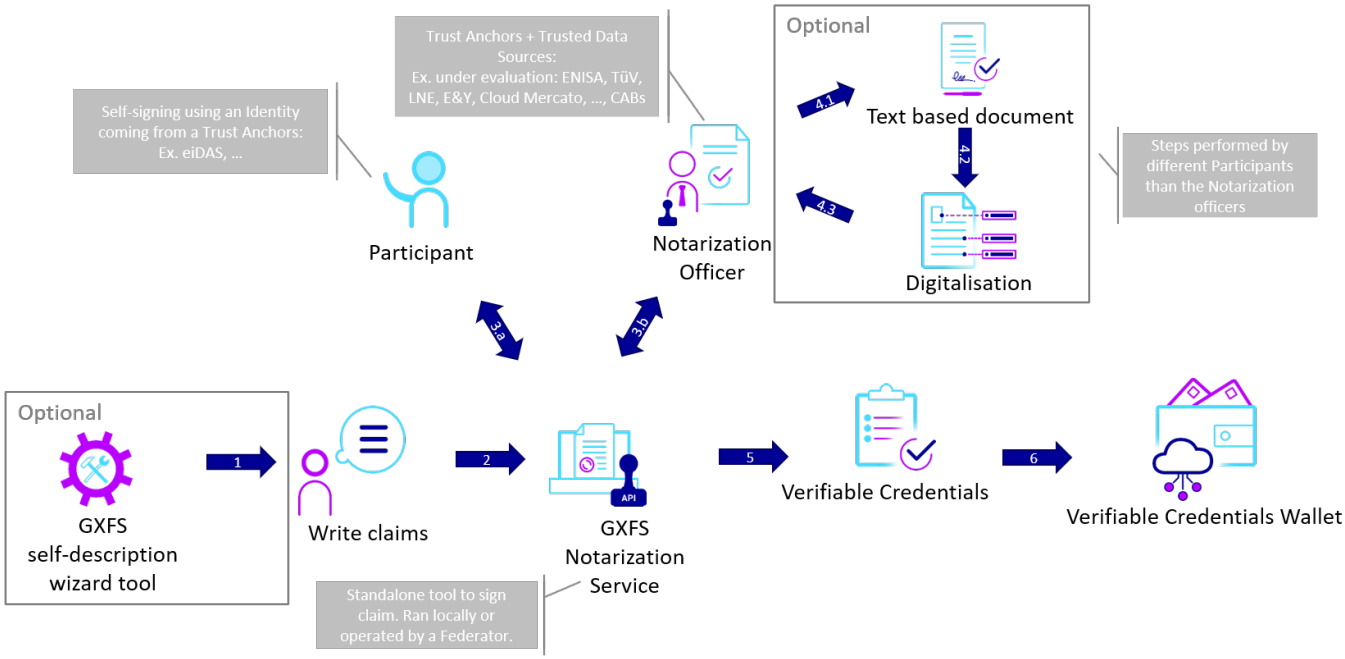

Collecting signed claims

4.6.2 Gaia-X verification

Using the collected signed claims, a participant can submit them for verification to the Gaia-X Compliance service and get in return a signed Claim with the result.
The Gaia-X Registry and Gaia-X Compliance are developed as Open-Source project inside the Gaia-X association.
The version v1.0 of the services are deployed by the Gaia-X Association. The version v2.0 of those services will be distributed. The version 3.0 of those services will be decentralized.

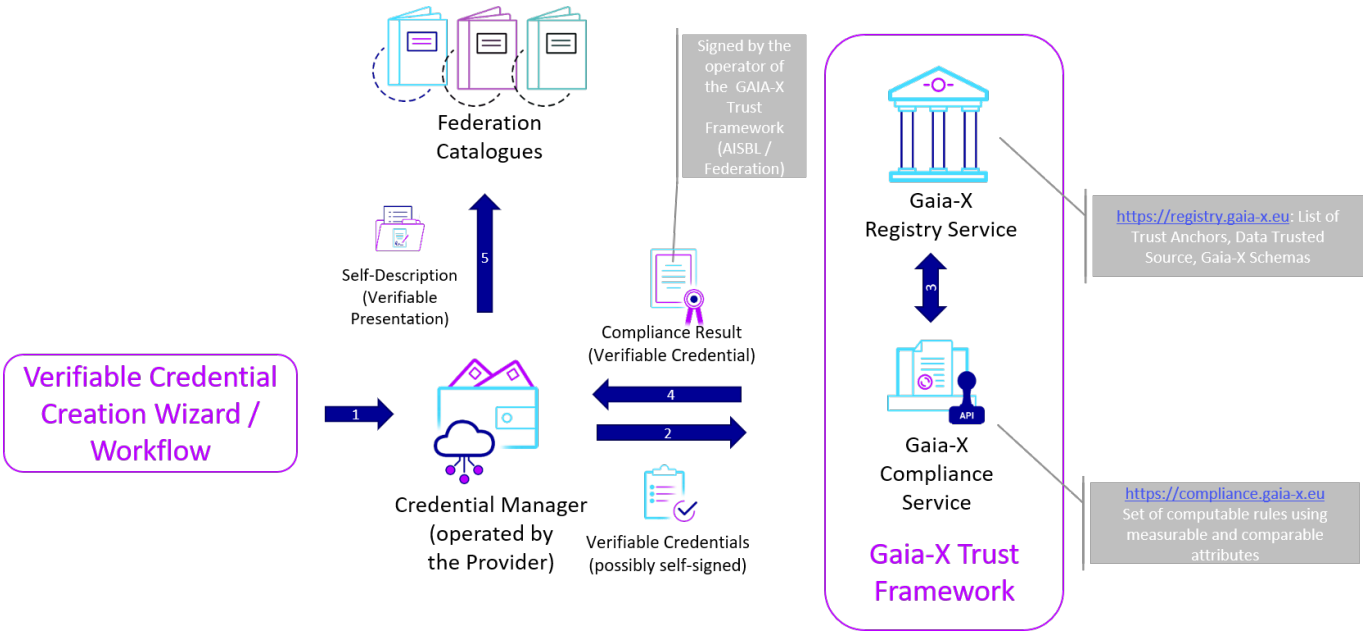

Validating signed Claims using the Gaia-X Trust Framework

4.6.3 Federation governance

Using the same general workflow of the previous step, every federation is free to extend Gaia-X governance and add custom rules and checks.

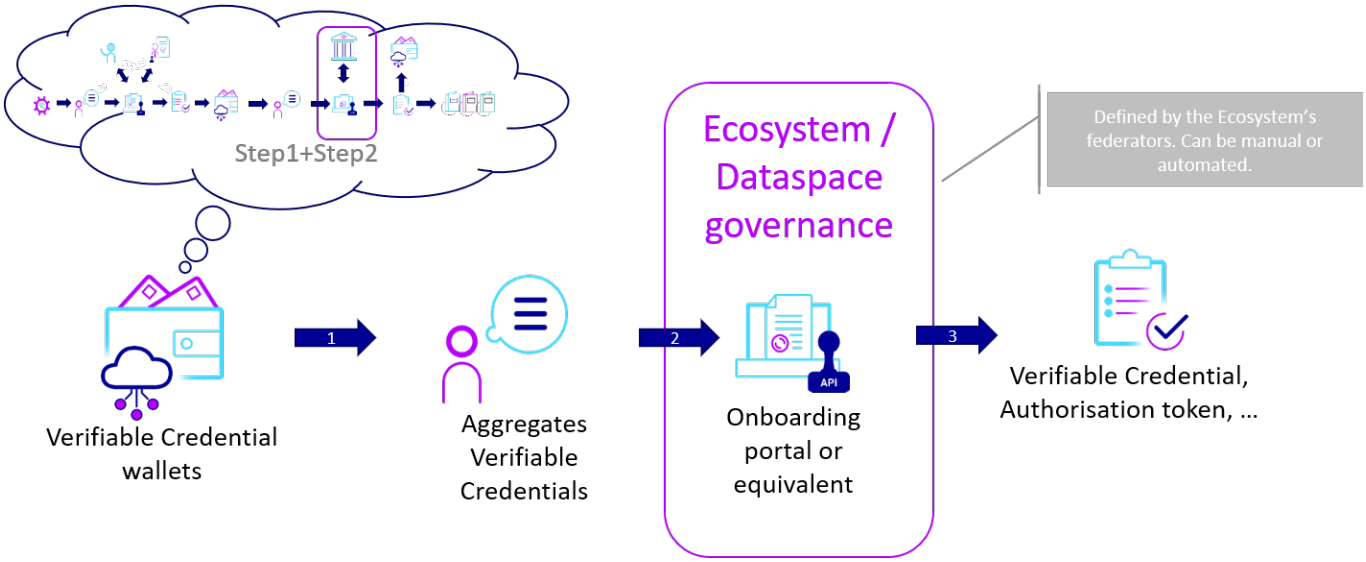

Federation extending Gaia-X governance

5. Gaia-X Operating Model

Gaia-X in its unique endeavour must have an operating model enabling a widespread adoption by small and medium-sized enterprises up to large organisations, including those in highly regulated markets, to be sustainable and scalable.

To achieve the objectives above, a non-exhaustive list of Critical Success Factors (CSFs) includes these points:

1. The operating model must provide clear and unambiguous added value to all Participants
2. The operating model must have a transparent governance and trust model with identified accountability and liability, that is clearly and fully explained to all Participants
3. The operating model must be easy to use by all Participants
4. The operating model must be financially sustainable for the Gaia-X Ecosystem
5. The operating model must be environmentally sustainable.

The first part of this chapter introduces the Gaia-X Ecosystem, as well as Trust Anchors. Trust Anchors are defined, including details about who defines them and how they will be nominated.

The second part defines Gaia-X Compliance, and how to become compliant. It introduces the Gaia-X Compliance Service as well as the usage of Gaia-X Labels.

Finally, the last section will cover the Gaia-X Self-Descriptions life-cycle and the Gaia-X Registry, which provides essential support for the Gaia-X Decentralized Autonomous Ecosystem.

5.1 Gaia-X Ecosystem

An ecosystem is an independent group of Participants that directly or indirectly consume, produce, or provide services such as data, storage, computing, network services, including combinations of them.

Those Participants autonomously decide which information to share, with whom, and if they want to do it in a Gaia-X compliant way.

Independently of the functional, business scope or interoperability levels of those ecosystems, the Gaia-X Ecosystem is the virtual set of all entities described with a Gaia-X compliant Self-Description and that conform to Gaia-X requirements.

Several individual Ecosystems may exist that orchestrate themselves, use the Architecture and may or may not use the Federation Services open source software.

Examples of Ecosystems:

- Catena-X - Automotive ecosystem
- Agdatahub - Agriculture ecosystem

Access to an ecosystem is under the full control of the participants providing the ecosystem's federation services, the ecosystem's federators.

Only the Self-Descriptions following the requirement of the Gaia-X Trust Framework are eligible for Gaia-X compliance.

Only compliant Service Offering Self-Descriptions are eligible to Gaia-X Labels.

More details are given in the Gaia-X ecosystem chapter.

5.2 Trust Anchors

For a given ecosystem, the Trust anchors are the entities considered by all Participants to be trustworthy when establishing the chain of cryptographic certificates.
Ecosystems can select their own Trust Anchors, however, cross-ecosystem trust requires the selected Trust Anchors to comply at least with the same rules that the common Gaia-X ecosystem Trust Anchors shall comply with.

The Gaia-X Association defines:

- the sets of rules to define the Trust Anchors:
  - Trust Service Providers.
  - Gaia-X Label Issuers
  - Trusted data source for Gaia-X Compliance
- the format of the Self-Descriptions and their compliance rules
- the Gaia-X Labels rulebook.

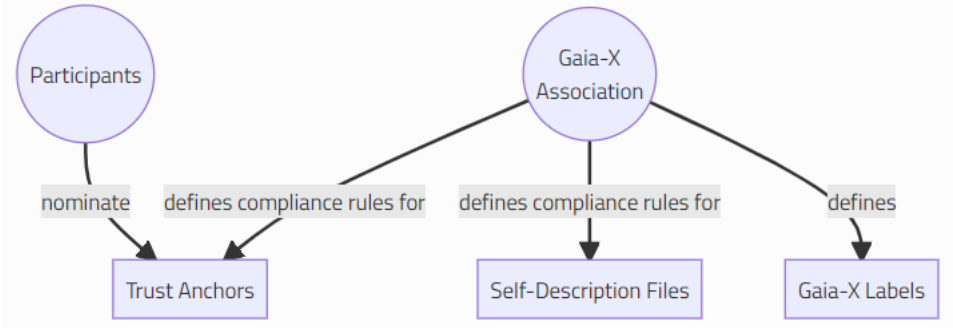

The Trust Anchors are nominated by the Participants. The validation of the nominees is done automatically by validating the rules defined by the Gaia-X Association and supervised by the Gaia-X Association.

In turn, the Trust Anchors are used by the Participants to operate the Ecosystem(s).

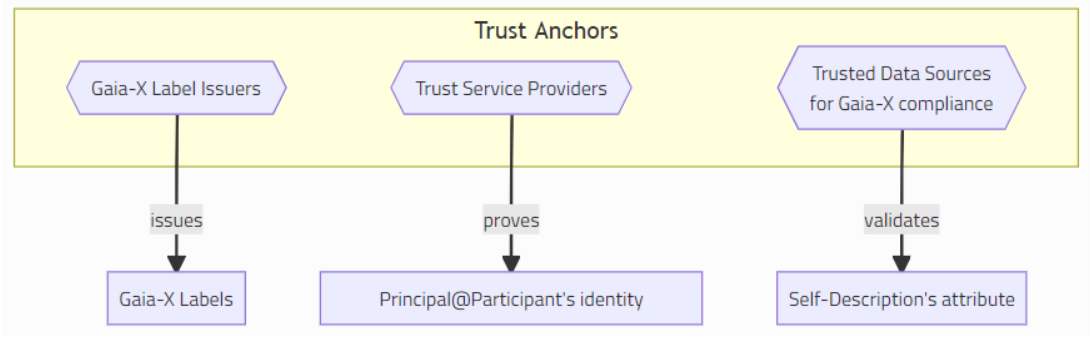

5.3 Gaia-X Trust Framework

The Gaia-X Trust Framework – formerly known as Gaia-X compliance or Regulation by Automation – is defined as the process of going through and validating the set of automatically enforceable rules to achieve the minimum level of Self-Description compatibility in terms of:

- syntactic correctness.

- schema validity.
- cryptographic signature validation.
- attribute value consistency.
- attribute value verification.

Whenever possible, the verification of Self-Descriptions' attribute values is done either by using publicly available open data, and performing tests or using data from Trusted Data Sources as defined in the previous section. This verification is captured using Verifiable Credentials issued by either: - the Gaia-X association when performing live tests (Trust Anchor) - the owner of the Trusted Data source (Trust Anchor)

However, it is expected that checking the validity of Self-Descriptions using data will introduce costs. In the context of the Gaia-X Ecosystem, a proposal to cover the operating cost is described later in this document with the introduction of a Gaia-X Decentralized Autonomous Ecosystem.

Other ecosystems are autonomous and this operating model does not cover how the operating cost of ecosystems should be handled.

The set of rules is versioned and will evolve over time to adapt to legal and market requirements. Those rules are set in a separate Gaia-X Trust Framework document.[11]

The rules will be implemented using open-source code and a service instance of that source code is called a Gaia-X Compliance Service.

One of the first Gaia-X added values is the creation of a FAIR (findable, accessible,interoperable, reusable) knowledge graph of verifiable and composable Self-Descriptions.

5.4 Gaia-X Labels

From the European Data Governance Act proposal:

> As a compulsory scheme this could generate higher costs, which could potentially have a prohibitive impact on SMEs and startups, and the market is not mature enough for a compulsory certification scheme; therefore, lower intensity regulatory intervention was identified as the preferred policy option.
> However, the higher intensity regulatory intervention in the form of a compulsory scheme was also identified as a feasible alternative, as it would bring significantly higher trust to the functioning of data intermediaries, and would establish clear rules for how these intermediaries are supposed to act in the European data market.

The decision for the Gaia-X Association is to adopt a compulsory scheme for Gaia-X compliance – see previous section – and an optional scheme for Gaia-X Labels, to ensure a common level of transparency and interoperability while limiting the regulatory burden on the market players.

Labels are issued for Service Offerings only and are the result of the combination of several Self-Description compliant attributes, that individually would be insufficient to support business or regulatory decisions.
The issued Labels must include a version number to allow continuous evolution of the set of rules and the precise set of rules in a "rulebook" defined by the Gaia-X Association, which must include a workflow for compliance re-validation.

From a technical point of view, a Label is a W3C Verifiable Credential, similar to Self-descriptions' attributes credentials that are described in the next section.

The management of the rulebook and its governance is described in the Gaia-X Labels document expected in October 2021.

5.5 Gaia-X Self-Description

Gaia-X Self-Descriptions describe in a machine interpretable format any of the entities of the Gaia-X Conceptual Model.

It means that there are Self-Descriptions for Participants in all roles: `Consumer`, `Federator`, `Provider` and all the other entities in an Ecosystem's scope such as `Resource` and `Service Offering`.

Each Gaia-X entity makes Claims, which are validated and signed by 3rd parties. Those signed Claims are defined as Verifiable Credentials and presented by the entity as Verifiable Presentations.

Technically speaking, Self-Descriptions are W3C Verifiable Presentations in the JSON-LD serialization of the RDF graph data model.

The following workflow describes how Gaia-X Self-Descriptions are created following the vocabulary of the W3C Verifiable Credentials Data Model standard.

W3C Term	Example with a car
Claim	My car is red
Verifiable Credential	The garage's attestation that my car is red
Verifiable Presentation	Me showing to my friend the garage's attestation that my car is red
Issuer	The garage
Holder	Myself
Verifier	My friends

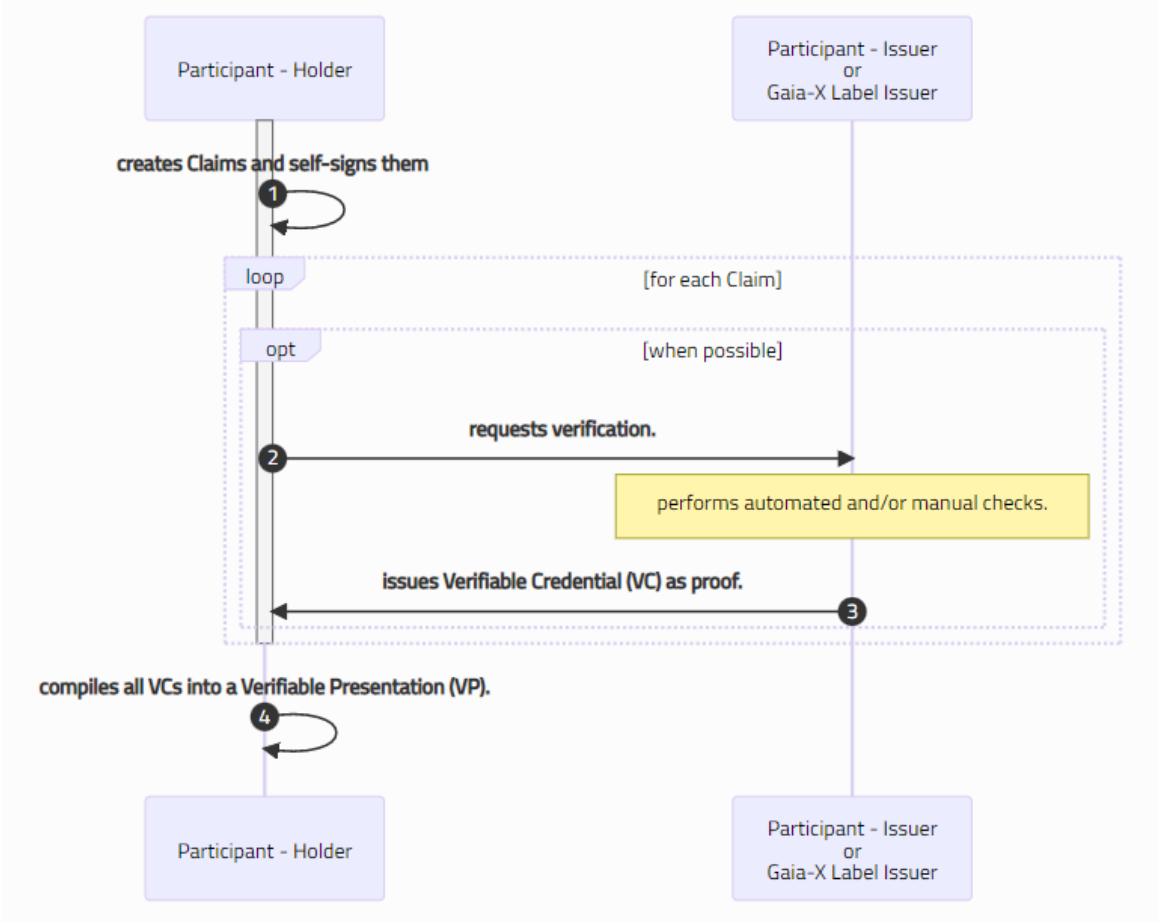

5.5.1 Difference between Self-Description's proofs (VC), Gaia-X Compliance and Gaia-X Labels.

The Verifiable Credentials are issued by other Participants, including Conformity Assessment Bodies. Verifiable Credentials can also be used to build a reputation system in the knowledge graph.

The Gaia-X Compliance ensures that the required level of information for the users to take educated decisions is available and the information is verified or verifiable.

The Gaia-X Labels set thresholds for specific industries, markets or regulated activities.

	Attribute's Verifiable Credentials	**Gaia-X Compliance**	**Gaia-X Labels**
Technical implementation	W3C Verifiable Credentials	W3C Verifiable Credentials	W3C Verifiable Credentials
Credential Issuer	Any Gaia-X Participant	Gaia-X Compliance Service	Gaia-X Label Issuer
Application scope	All Self-Descriptions	All Self-Descriptions	Service Offerings

	Attribute's Verifiable Credentials	Gaia-X Compliance	Gaia-X Labels
Assessment method(s)	Manual or Automated	Fully automated	Manual or Automated
Issuance's temporality	Frequent updates	Frequent updates	Slow updates (~yearly)

A attribute's Verifiable Credential is Gaia-X conformant if the Issuer of the Verifiable Credential has itself an identity coming from one of the Trust Anchors.

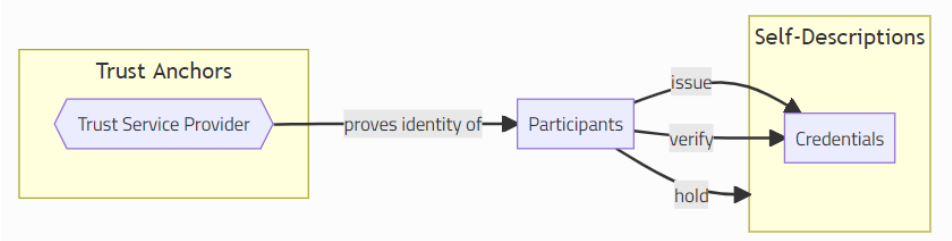

A Label is Gaia-X conformant if the Issuer of the Credential is one of the Trust Anchors' Gaia-X Label Issuers.

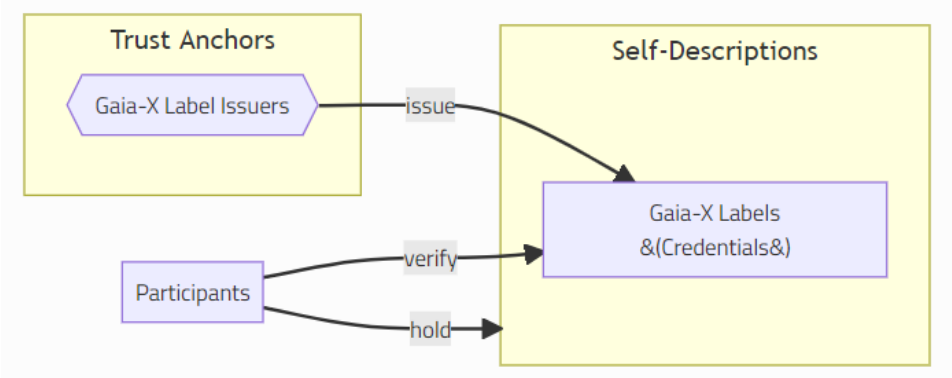

5.5.2 Self-Description compliance

A Self-Description qualified as Gaia-X compliant must be submitted to a Gaia-X Compliance Service instance as defined in the section above.

The result of that submission is captured in two ways:

- as a Verifiable Credential: If the compliance is validated, the Gaia-X Compliance Service issues a Verifiable Credential that can later be inserted into the Self-Description. This method aligns with the self sovereign principle of the holder being in charge of the information.
- If the Gaia-X Compliance Service is called via the Gaia-X Registry, the Gaia-X Registry will emit an event to synchronize Catalogues. The event contains the URL the of the Self-Description file. The Gaia-X Registry is defined in the next section.

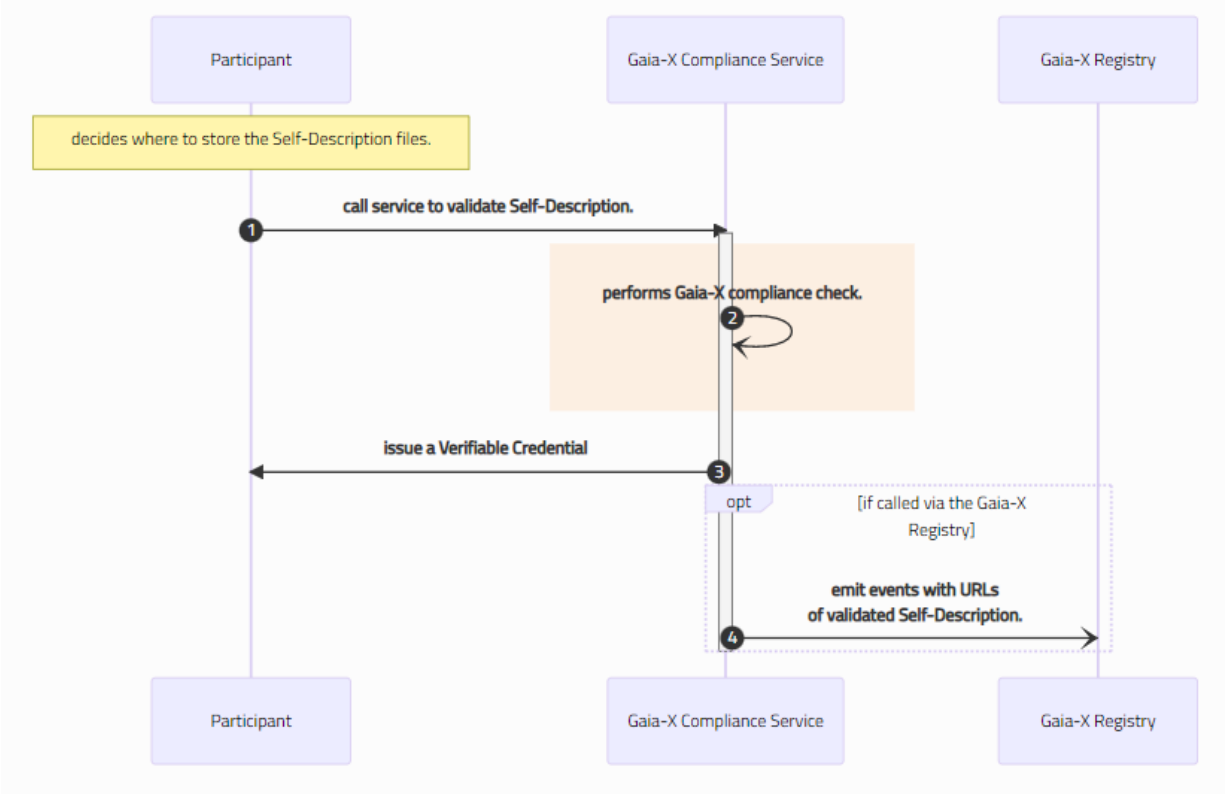

5.5.3 Self-Description Remediation

Self-Descriptions may become invalid over time. The Chapter *Federated Catalogue* section *Self-Description* describes three states declaring a Self-Description as invalid:

- Expired (after a timeout date, e.g., the expiry of a cryptographic signature)
- Deprecated (replaced by a newer Self-Description)
- Revoked (by the original issuer or a trusted party, e.g., because it contained incorrect or fraudulent information)

Expired and Deprecated can be deduced automatically based in the information already stored in the Gaia-X Registry or Gaia-X Catalogues. There are no additional processes to define. This section describes how Self-Descriptions are revoked.

The importance of Gaia-X compliance will grow over time, covering more and more Gaia-X principles such as interoperability, portability, and security. However, automation alone is not enough and the operating model must include a mechanism to demotivate malicious actors to corrupt the Registry and Catalogues.

The revocation of Self-Descriptions can be done in various ways:

- **Revocation or Deprecation by authorship**: The author of a Self-Description revokes or deprecates the Self-Description explicitly.
- **Revocation by automation**: The Gaia-X Compliance Service found at least one self-described attribute not validating the compliance rules.
- **Suspension and Revocation by manual decision**: After an audit by a compliant Gaia-X Participant, if at least one self-described attribute is found to be incorrect, the suspension of the Self-Descriptions is automatic. The revocation is submitted for approval to the Gaia-X Association with the opportunity for the Self-Description's owner to state its views in a matter of days. To minimize subjective decisions and

promote transparency, the voting results will be visible and stored on the Gaia-X Registry or in the local Ecosystem's Registry.

5.6 Gaia-X Decentralized Autonomous Ecosystem

The operating model described in this chapter motivates the creation of a Gaia-X decentralized autonomous Ecosystem following the principles of a Decentralized Autonomous Organisation[12], with the following characteristics:

- Compliance is achieved through a set of automatically enforceable rules whose goal is to incentivize its community members to achieve a shared common mission.
- Maximizing the decentralization at all levels to reduce lock-in and lock-out effects.
- Minimizing the governance and central leadership to minimize liability exposure and regulatory capture.
- The ecosystem has its own rules, including management of its own funds.
- The ecosystem is operated by the ecosystem's Participants

Other ecosystems are autonomous and this operating model does not enforce how internal ecosystem governance should be handled.

5.6.1 Gaia-X Registry

The Gaia-X Registry is a public distributed, non-repudiable, immutable, permissionless database with a decentralized infrastructure and the capacity to automate code execution.

The Ecosystems may want to have their own instance of a local Gaia-X Registry or equivalent. Technically, this component can be part of the ecosystem local Catalogues.

The Gaia-X Registry is the backbone of the ecosystem governance which stores information, similarly to the Official Journal of the European Union, such as:

- the list of the Trust Anchors – keyring.
- the result of the Trust Anchors validation processes.
- the potential revocation of Trust Anchors identity.
- the vote and results of the Gaia-X Association roll call vote, similar to the rules of the plenary of the European Parliament
- the URLs of the Self-Description Schemas defined by Gaia-X
- the URLs of Catalogue's Self-Descriptions
- ...

It also facilitates the provision of:

1. A decentralized network with smart contract functionality.
2. Voting mechanisms that ensure integrity, non-repudiation and confidentiality.
3. Access to a Gaia-X Compliance Service instance.
4. A fully operational, decentralized and easily searchable catalogue[13].
5. A list of Participants' identities and Self-Description URIs which violate Gaia-X membership rules. This list must be used by all Gaia-X Trusted Catalogue providers to filter out any inappropriate content.

6. Tokens may cover the operating cost of the Gaia-X Ecosystem. This specific point can be abstracted by 3rd party brokers wrapping token usage with fiat currency, providing opportunities for new services to be created by the Participants. Emitting tokens for the Gaia-X Association's members is also considered.

Each entry in the Gaia-X Registry is considered as a transaction. A transaction contains DIDs of all actors involved in the transaction and metadata about the transaction in a machine readable format. The basic rule for a transaction to be valid is that all DIDs have one of the Trust Anchors as root Certificate Authorities. Please also note that the Registry stores all revoked Trust Anchors.

This model enables the Participants to operate in the Gaia-X ecosystem, to autonomously register information, and to access the information which is verifiable by other Participants.

5.6.2 Ecosystem launching phase

In order to enable the 1st scenario which is:

> As a Gaia-X Provider, I want to publish the self-description of my Service Offerings and I want my Service Offerings to be made available to all Ecosystems.

and until the inter-catalogue synchronization is documented, the Registry will also be used to store, directly or indirectly via an external storage, the Self-Descriptions' URLs.

5.6.3 Verifiable Presentation Verification

The Gaia-X Registry, or a private one, independently of its implementation, is the single source of truth for the Ecosystem.
It allows any `Participant` to verify the validity of signatures.

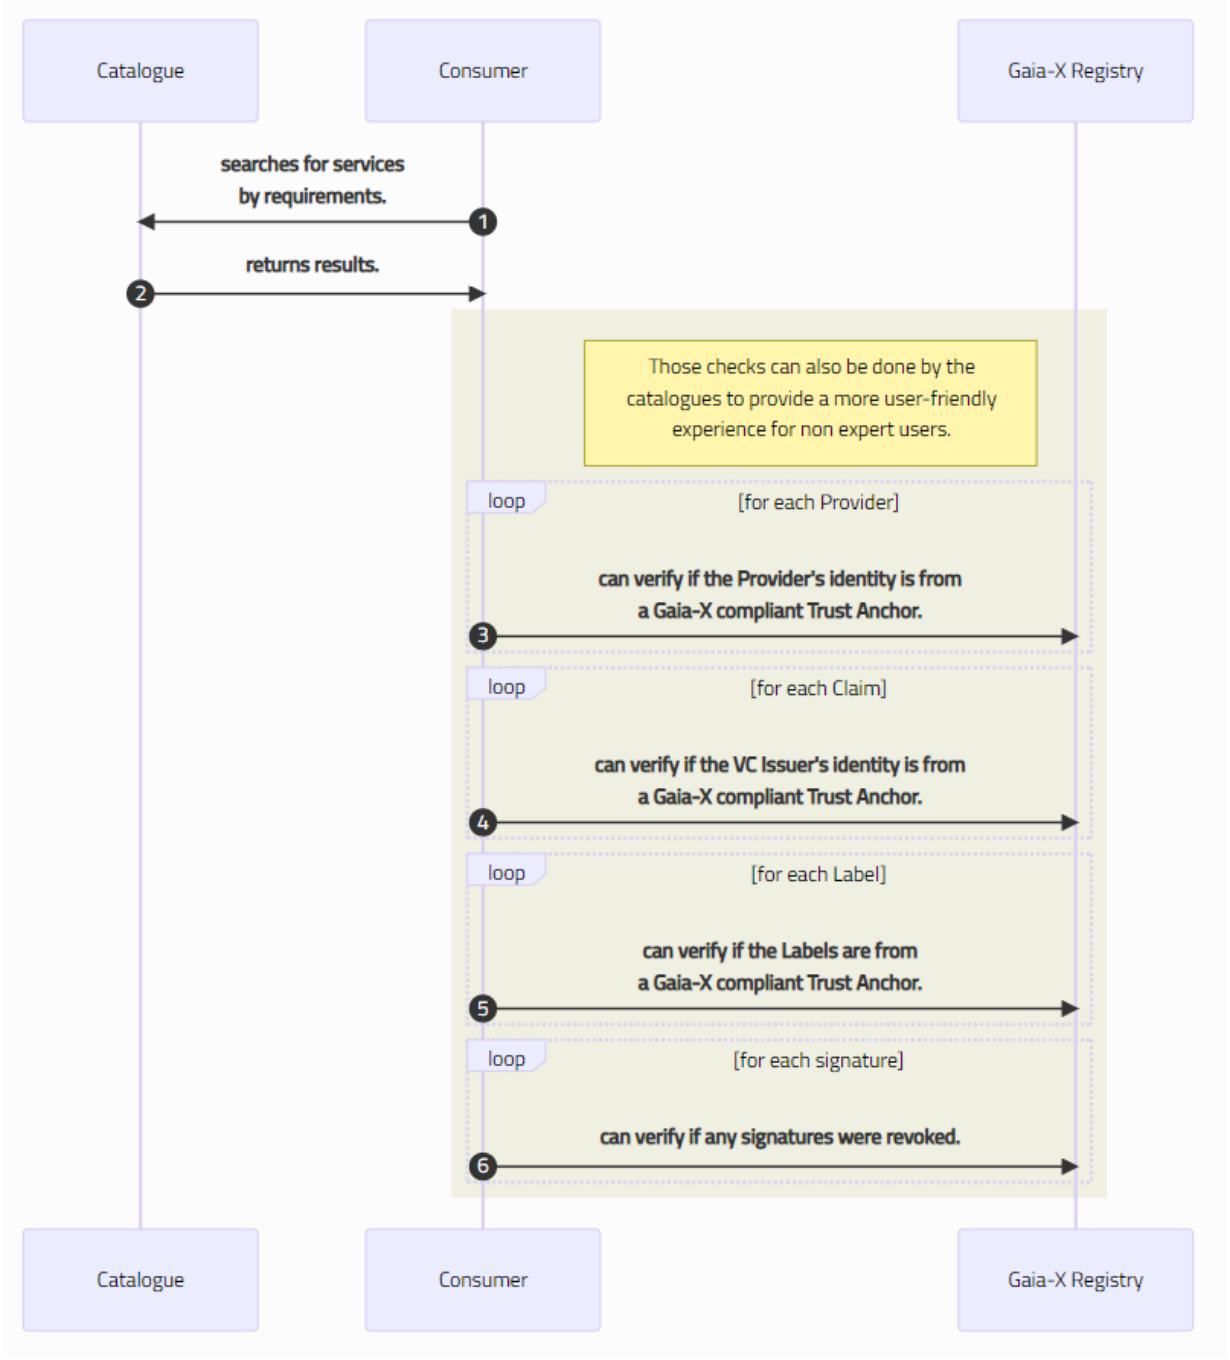

Example with `Consumer` *and main* `Gaia-X Registry`

6. Federation Services

The *Federation Services* are necessary to enable a Federation of infrastructure and data and to provide interoperability across Federations.

- The *inter-catalogue synchronisation* constitutes an indexed repository of Gaia-X Self-Descriptions to enable the discovery and selection of Providers and their Service Offerings.
- Common vocabulary for *Identity and Access Management* covers identification, authentication and authorization, credential management, decentralized Identity management as well as the verification of analogue credentials, including between existing identity federation.
- *Data Exchange services* enable data exchange of Participants by providing a Data Agreement Service and a Data Logging Service to enable the enforcement of Policies. Furthermore, usage constraints for data exchange can be expressed by Provider Policies within the Self-Descriptions.
- *Gaia-X Trust Framework* includes mechanisms to ensure that Participants adhere to the Policy Rules in areas such as security, privacy, transparency and interoperability during onboarding and service delivery.
- *Portals and APIs* will support onboarding and management of Participants, demonstrate service discovery, orchestration and provisioning of sample services within ecosystems.

6.1 The Role of Federation Services for Ecosystems

The following figure visualizes how Federation Services Instances are related to the Federator described in the conceptual model (see section Federator). The Federators enable Federation Services by obliging Federation Service Providers to provide concrete Federation Service Instances. The sum of all Federation Service Instances form the Federation Services.

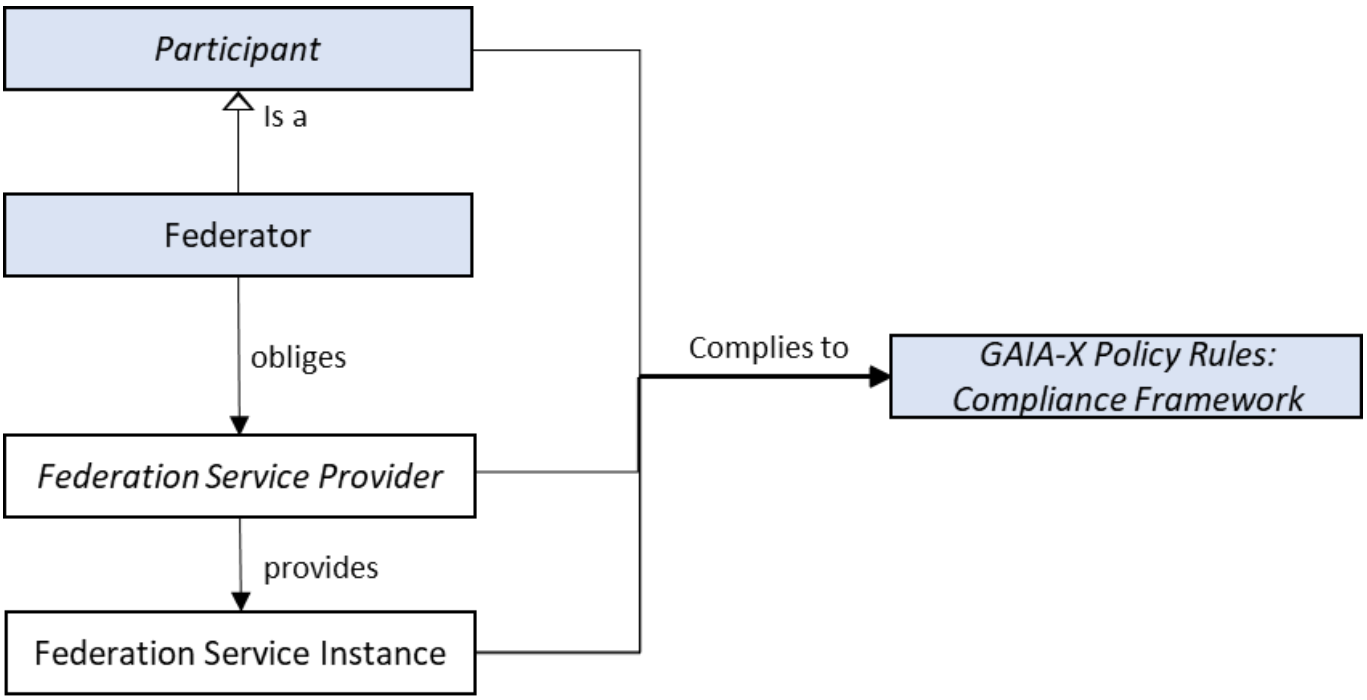

Federation Services Relations

6.1.1 Nesting and Cascading of Federation Services

Federation Services can be nested and cascaded. Cascading is needed, for example, to ensure uniqueness of identities and Catalogue entries across different individual Ecosystems / communities that use Federation Services. (Comparable to DNS servers: there are local servers, but information can be pushed up to the root servers).

Therefore, a decentralised synchronization mechanism is necessary.

6.1.2 Ecosystem Governance vs. Management Operations

To enable interoperability, portability and Data Sovereignty across different Ecosystems and communities, Federation Services need to adhere to common standards. These standards (e.g., related to service Self-Description, digital identities, logging of data sharing transactions, etc.) must be unambiguous and are therefore defined by the Gaia-X Association AISBL. The Gaia-X Association AISBL owns the Compliance Framework and related regulations or governance aspects. Different entities may take on the role of Federator and Federation Services Provider.

6.1.3 Infrastructure Ecosystem

The Infrastructure Ecosystem has a focus on computing, storage and Interconnection elements. In GAIA-X Ecosystem these elements are designated as Nodes, Interconnections and different Software Resources. They range from low-level services like bare metal computing up to highly sophisticated offerings, such as high-performance computing. Interconnection Services ensure secure and performant data exchange between the different Providers, Consumers and their services. Gaia-X enables combinations of services that range across multiple Providers of the Ecosystem. The Interconnection Services are also the key enabler for the composition of services offered by diverse and distributed providers, ensuring the performance of single-provider networks on a multi-provider "composed" network.

6.1.4 Data Ecosystem

Gaia-X facilitates Data Spaces which present a virtual data integration concept, where data are made available in a decentralised manner, for example, to combine and share data of stored in different cloud storage backends. Data Spaces form the foundation of Data Ecosystems. In general, Data Ecosystems enable Participants to leverage data as a strategic resource in an inter-organizational network without restrictions of a fixed defined partner or central keystone companies. For data to realize its full potential, it must be made available in cross-company, cross-industry Ecosystems. Therefore, Data Ecosystems not only enable significant data value chain improvements, but provide the technical means to enable Data Sovereignty. Such sovereign data sharing addresses different layers and enables a broad range of business models that would otherwise be impossible. Trust and control mechanisms encourage the acceleration of data sharing and proliferate the growth of Ecosystems.

6.1.5 Federation, Distribution, Decentralization and Sharing

The principles of federation, distribution, decentralization and sharing are emphasized in the Federation Services as they provide several benefits for the Ecosystem:

Principle	Need for Gaia-X	Implemented in Gaia-X Architecture
Decentralization	Decentralization will ensure Gaia-X is not controlled by the few and strengthens the participation of the many. It also adds key technological properties like redundancy, and therefore resilience against unavailability and exploitability. Different implementations of this architecture create a diverse Ecosystem that can reflect the respective requirements and strengths of its Participants. (example: IP address assignment)	The role of Federators may be taken by diverse actors. The open source Federation Services can be used and changed according to specific new requirements as long as they are compliant and tested.
Distribution	Distribution fosters the usage of different Resources by different Providers spread over geographical locations. (Example: Domain Name System)	Self-Description ensures that all Resources and Service Offerings are defined standardized ways, which enables them to be listed in a searchable Catalogue, each with a unique Identifier. Therefore, it facilitates the reuse and distribution of these components.
Federation	Federation technically enables connections and a web of trust between and among different parties in the Ecosystem(s). It addresses the following challenges: ◦ Decentralized processing locations ◦ Multiple actors and stakeholders ◦ Multiple technology stacks ◦ Special policy requirements or regulated markets (Example: Autonomous Systems)	Each system can interact with each other, e.g., the Catalogues could exchange information and the Identity remains unique. Furthermore, different Conformity Assessment Bodies may exist.
Sharing	Sharing of the relevant services and components contributes to the Ecosystem development. Sharing and reuse of Resources across the Gaia-X Ecosystem enables positive spillovers, leading to new and often unforeseen economic growth opportunities.	The Federated Catalogues enable the matching between Providers and Consumers. Sovereign Data Exchange lowers hurdles for data exchange and Ecosystem creation.

Summary of Federation Services as enabler

By utilizing common specifications and standards, harmonized rules and policies, Gaia-X is well aligned with specifications like NIST Cloud Federation Reference Architecture[14]:

- Security and collaboration context are not owned by a single entity
- Participants in the Gaia-X Association AISBL jointly agree upon the common goals and governance of the Gaia-X Association AISBL
- Participants can selectively make some of their Resources discoverable and accessible by other Participants in compliance with Gaia-X
- Providers can restrict their discovery and disclose certain information but could risk losing their Gaia-X compliance level

6.2 Interoperability and Portability for Infrastructure and Data

For the success of a Federated Ecosystem it is of importance that data, services and the underlying infrastructure can interact seamlessly with each other. Therefore, portability and interoperability are two key requirements for the success of Gaia-X as they are the cornerstones for a working platform and ensure a fully functional federated, multi-provider environment.

Interoperability is defined as the ability of several systems or services to exchange information and to use the exchanged information mutually. Portability refers to the enablement of data transfer and processing to increase the usefulness of data as a strategic resource. For services, portability implies that they can be migrated from one provider to another, while the migration should be possible without significant changes and adaptations and have an equivalent QoS (Quality of Service).

6.2.1 Areas of Interoperability and Portability

+The Gaia-X Ecosystem includes a huge variety of Participants and Service Offerings. Therefore, interoperability needs to be ensured on different levels (Infrastructure as a Service [IaaS] , Platform as a Service [PaaS], Software as a Service [Saas], data resources, and others) by means of Service Composition.

Regarding interoperability of data, core elements to be identified in this endeavour are API specifications and best practices for semantic data descriptions. The use of semantic data interoperability is seen as a foundation to eventually create a clear mapping between domain-specific approaches based on a community process and open source efforts.

6.3 Infrastructure and Interconnection

To best accommodate the wide variety of Service Offerings, the Gaia-X Architecture is based on the notion of a sovereign and flexible Interconnection of Infrastructure and Data Ecosystems, where data is flexibly exchanged between and among many different Participants. Therefore, Interconnection Services represent a dedicated category of Resources as described in section Gaia-X Conceptual Model.

The Interconnection Whitepaper provides an overview on the current strategy to expand the Gaia-X architecture and federation services.

6.4 Federated Catalogue

The goal of Catalogues is to enable Consumers to find best-matching offerings and to monitor for relevant changes of the offerings. The Providers decide in a self-sovereign manner which information they want to make public in a Catalogue and which information they only want to share privately.

A Provider registers Self-Descriptions with their universally resolvable Identifiers in the desired Catalogue to make them public relative to the Catalogue scope. The Catalogue builds an internal representation of a knowledge graph out of the linked data from the registered and accessible seld-descriptions to provide interfaces to query, search and filter services offerings.

The system of Federated Catalogues includes an initial stateless Self-Description browser provided by the Gaia-X, European Association for Data and Cloud, AISBL. In addition, Ecosystem-specific Catalogues (e.g., for the healthcare domain) and even company-internal Catalogues (with private Self-Descriptions to be used only internally) can be linked to the system of federated Catalogues. The Catalogue federation is used to exchange relevant Self-Descriptions and updates thereof. It is not used to execute queries in a distributed fashion.

Cross-referencing is enabled by unique Identifiers as described in Identity and Access Management. While uniqueness means that Identifiers do not refer to more than one entity, there can be several Identifiers referring to the same entity. A Catalogue should not use multiple Identifiers for the same entity.

Gaia-X AISBL develops an extensible hierarchy of Schemas that define the terms used in Self-Descriptions and which must be supported by any Gaia-X Catalogue. It is possible to create additional Schemas specific to an application domain, an Ecosystem, Participants in it, or Resources offered by these Participants.

A Schema may define terms (classes, their attributes, and their relationships to other classes) in an ontology. If it does, it must also define shapes to validate instances of the ontology against.

Self-Descriptions in a Catalogue are either loaded directly into a Catalogue or exchanged from another Catalogue through inter-Catalogue synchronization functions.

Since Self-Descriptions are protected by cryptographic signatures, they are immutable and cannot be changed once published. This implies that after any changes to a Self-Description, the Participant as the Self-Description issuer has to sign the Self-Description again and release it as a new version. The lifecycle state of a Self-Description is described in a separate chapter.

6.5 Identity and Access Management

Identities, which are used to gain access to the Ecosystem, rely on unique Identifiers and a list of dependent attributes. Gaia-X uses existing Identities and does not maintain them directly. Uniqueness is ensured by a specific Identifier format relying on properties of existing protocols. The Identifiers are comparable in the raw form and should not contain more information than necessary (including Personal Identifiable Information). Trust - confidence in the Identity and capabilities of Participants or Resources - is established by cryptographically verifying Identities using the Gaia- Trust Framework, which guarantees proof of identity of the involved Participants to make sure that Gaia-X Participants are who they claim to be. In the context of Identity and Access Management, the digital representation of a natural person, acting on behalf of a Participant, is referred to as a Principal. As Participants need to trust other Participants and Service Offerings provided, it is important that the Gaia-X Trust Framework provides transparency for everyone. Therefore,

proper lifecycle management is required, covering Identity onboarding, maintaining, and offboarding within Ecosystem. The table below shows the Participant Lifecycle Process.

Lifecycle Activity	Description
Onboarding	The Gaia-X Trust Framework and optional individual ecosystems workflow validates and signs the Self-Description provided by a Visitor (the future Participant/Principal).
Maintaining	Trust related changes to the Self-Descriptions are recorded in a new version and validated and signed by the same entities involved during the Onboarding. This includes information controlled by the Participant/Principal.
Offboarding	The offboarding process of a Participant is time-constrained and involves all dependent Participants/Principals. This includes keypair revocation by the entities involved during the Onboarding.

Participant Lifecycle Process

An Identity is composed of a unique Identifier and an attribute or set of attributes that uniquely describe an entity within a given context. The lifetime of an Identifier is permanent. It may be used as a reference to an entity well beyond the lifetime of the entity it identifies or of any naming authority involved in the assignment of its name. Reuse of an Identifier for a different entity is forbidden. Attributes will be derived from existing identities as shown in the IAM Framework Document v1.2[15].

A 'Secure Digital Identity' is a unique Identity with additional data for robustly trustworthy authentication of the entity (i.e. with appropriate measures to prevent impersonation) This implies that Gaia-X Participants can self-issue Identifiers for such Identities. It is solely the responsibility of a Participant to determine the conditions under which the Identifier will be issued. Identifiers shall be derived from the native identifiers of an Identity System without any separate attribute needed. The Identifier shall provide a clear reference to the Identity System technology used. Additionally, the process of identifying an Identity Holder is transparent. It must also be possible to revoke issued Identity attributes[16].

6.5.1 Layered Identity and Access Management

The Identity and Access Management relies on a two-tiered approach. In practice, this means that Participants use a selected few Identity Systems for mutual identification and trust establishment, SSI being the recommended option for interoperability. After trust is established, underlying existing technologies already in use by Participants (on the "Principal layer") can be federated and reused, for example Open ID Connect or domain specific x509-based communication protocols.

Gaia-X Participants might need to comply with additional requirements on the type and usage of credentials management applications such as mandatory minimum-security requirements, such as Multi-factor authentication. Server-to-Server Communication plays a crucial role in Gaia-X.

This chapter describes the components required to provide an attested secure chain of trust & identities. Service implementations and the corresponding service delivery layer may include End-User services, distributed microservice architectures across multiple Participant domains, and/or cross domain data or digital service delivery

6.5.1.1.1 Architecture principles for this approach

Mutual trust based on mutually verifiable Participant identities between contracting parties, Provider and Consumer, is fundamental to federating trust and identities in the Principal layer. Heterogeneous ecosystems across multiple identity networks in the Participant layer must be supported as well as heterogeneous environments implementing multiple identity system standards. The high degree of standardization of Participant layer and Principal layer building blocks of the Gaia-X Trust Framework must ensure that there is no lock-in to any implementation of identity network and Identity System likewise.

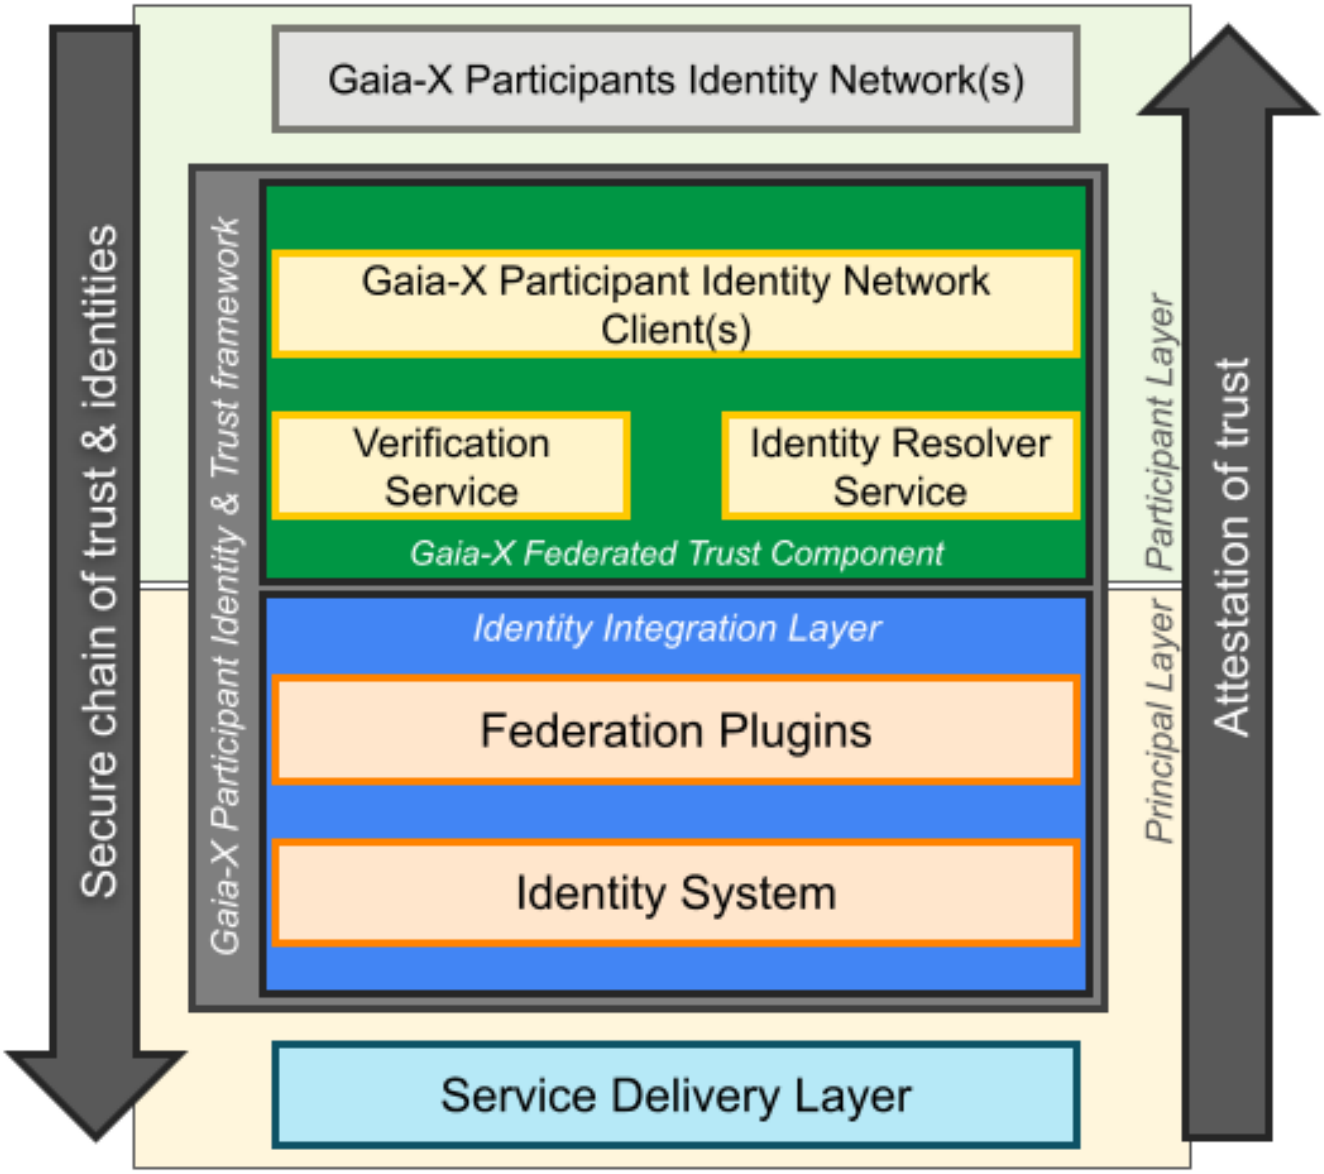

6.5.1.1.2 Chain of trust and identity

The Gaia-X Participant Identity & Trust framework delivers a secure chain of trust and identities to the service delivery layer.

Mutual participant verification

In the Participant layer, the Gaia-X Trust Framework implements the functionality to resolve and verify the Participant identity of the contracting parties. The Consumer verifies the Provider identity, the Provider verifies the Consumer identity. Successful mutual Participant verification results in a verified Participant Token representing the trust between Provider and Consumer.

Identity System federation

In the Principal layer, the Federation Plugin implements the functionality to federate trust between the Identity Systems of the contracting Participants based on the successful mutual Participant verification described above.
The federation of trust between the identity systems is based on the identity system standard implemented for the service delivery layer. Required for the federation is a secure mutual exchange of the required federation metadata. This exchange must be secured based on the Verified Participant Token. Exemplary Identity Systems standards supporting federation are: OIDC/OAuth2 (draft), SAML, SPIFFE/SPIRE. Identity System federation may also include federating the trust between certificate authorities supporting X.509 for Principals.

Identification and Authentication

Once successfully federated, the Identity Systems are enabled to identify and authenticate the Principals in the service delivery layer of the contracting parties. Federated Principal identities are mutually trusted based on the federation of the Identity Systems of the contracting Participants.

Principal Identity Integration Layer

While the interface to the Gaia-X Federated Trust Component is standardized, the federation mechanism of the Federation Plugin is specific to the implemented Identity System supporting current and future standards like OIDC/OAuth2 (draft), SPIFFE/SPIRE, PKI. Furthermore, multiple Identity Systems required for complex service offerings, like for example OIDC for user Principals, SPIRE for service Principals, are perfectly supported meaning that multiple Identity Systems on either side can be federated by corresponding plugins based on the very same mutual Participant identity verification if required for the service delivery.

6.6 Data Exchange services

The Data Exchange services represent the set of functionalities commonly used to perform controlled data exchanges.

Functional needs	Comment
Identity & Attributes	A common identity schema and vocabulary for attributes based access.
Data protocols	A common set of protocol and data format, including Enterprise Integration Patterns (EIP)
Policies negotiation	A common set of Domain Specific Language (DSL) to compute access rights and usage policies
Traceability	Means to store and trace negotiation results with the capacity to log intermediate realisations
Discoverability	Means to search and find one or more datasets based on queries and filterings, including a common means to declare and specify data ontologies, data vocabularies, and data semantics.

Each ecosystem is free to implement those functional needs as they see fit.

The Gaia-X Federation Services provides several tools to address those needs:

- a verifiable credential wallet: Organisation Credential Manager (OCM), Personal Credential Manager (PCM)
- policy negociation and logging: Data Contract Transaction (DCT), Data Exchange Logging (DEL)
- a catalogue: Catalogue (CAT)

6.6.1 Capabilities for Exchange services

The following are essential capabilities for services and products exchange in the Data-Infrastructure ecosystems:

Capability	Description
Expression of Policies in a machine-readable form	To enable transparency and control of service offering usage, it is important to have a common policy specification language to express usage restrictions in a formal and technology-independent manner that is understood and agreed by all Gaia-X Participants. Therefore, they must be formalized and expressed in a common Domain Specific Language, such as ODRL[17].

Capability	Description
Inclusion of Policies in Self-Descriptions	Informational self-determination and transparency require metadata to describe Resources as well as Providers, Consumers, and Usage Policies as provided by Self-Descriptions and the Federated Catalogues.
Interpretation of Usage Policies	For a Policy to be agreed upon, it must be understood and agreed by all Participants in a way that enables negotiation and possible technical and organizational enforcement of Policies.
Enforcement	Monitoring of services usage is a detective enforcement of service usage with subsequent (compensating) actions. In contrast, preventive enforcement[18] ensures the policy Compliance with technical means (e.g., cancellation or modification of data flows).

6.6.2 Functions of Exchange Services

Information services provide more detailed information about the general context of the services usage transactions. All information on the services exchange and services usage transactions must be traceable; therefore, agreed monitoring and logging capabilities are required for all services usage transactions. Self-determination also means that Providers can choose to apply no Usage Policies at all.

The Exchange Services in Gaia-X implement different functions for different phases of the services exchanges. Therefore, three distinct phases of service exchanges are defined:

- before transaction
- during transaction
- after transaction

Before the service exchange transaction, the Service Agreement service is triggered and both parties negotiate a service exchange agreement. This includes Usage Policies and the required measures to implement those. During transactions, a Contract Logging service receives logging-messages that are useful to trace each transaction. This includes data provided, data received, policy enforced, and policy-violating messages. During and after the transaction the information stored can be queried by the transaction partners and a third eligible party, if required.

The figure below shows the role of the aforementionentioned services to enable controlled data exchange.

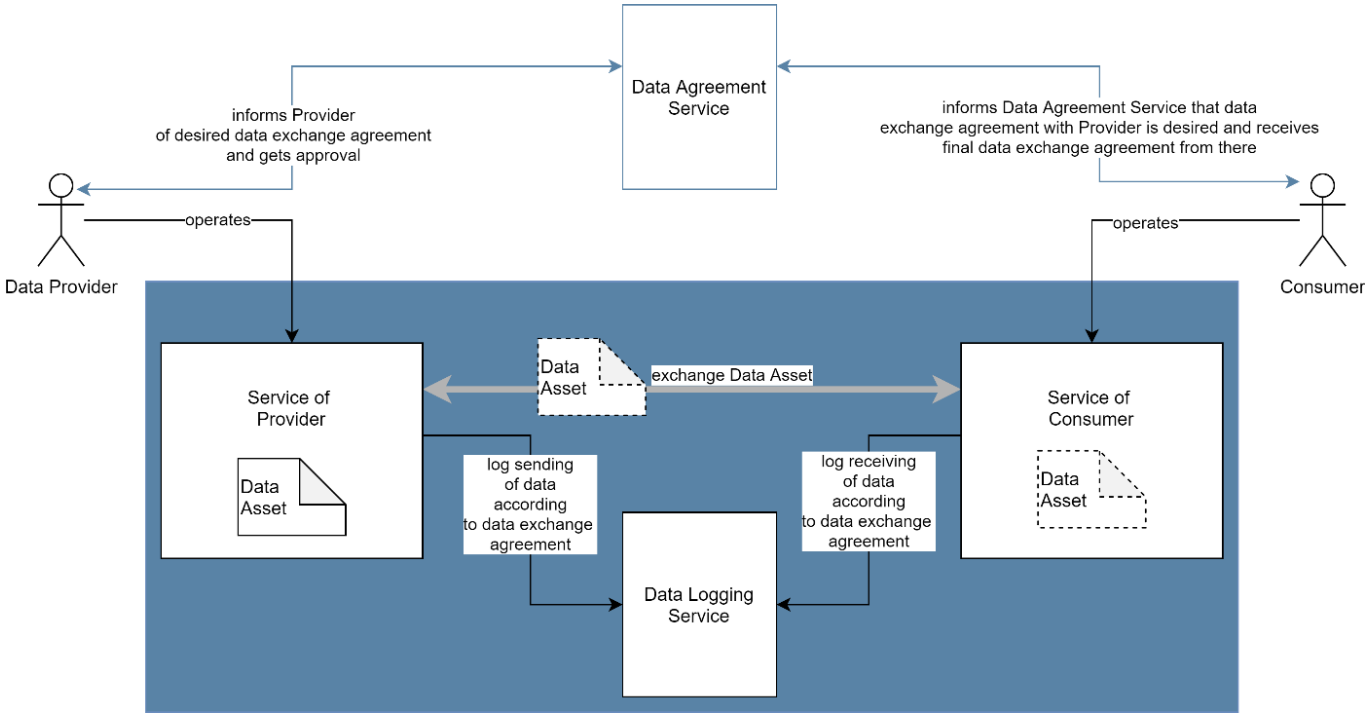

Data Services Big Picture

The Contract Agreement Service enables service transactions in a secure, trusted, and auditable way. It offers interfaces for the negotiation detailing the agreed terms for planned service exchange. The service is not meant to handle the transaction of service (which is described in the negotiated service contracts).

The Contract Logging Service provides evidence that service has been (a) transmitted, (b) received and (c) that rules and obligations (Usage Policies) were successfully enforced or were violated. This supports the clearing of operational issues but also identifies fraudulent transactions.

The Provider can track if, how, and what data was provided, with the Consumer being notified about this. The Consumer can track if data was received or not, and, additionally, track and provide evidence on the enforcement or violation of Usage Policies.

6.7 Gaia-X Federation Servies for Notarization and Credential storage

Gaia-X AISBL defines a Trust Framework that manifests itself in the form of two services:

- the Gaia-X Registry, detailed in the Operating model chapter
- the Gaia-X Compliance service, as the service implementing the set of rules described in the upcoming Gaia-X Trust Framework document.

The Notarization Service will supprt the issuance of Verficiable Credentials to the Self Descriptions which can be stored in the Personal- or Organizational Credential Manager.

6.8 Portals and APIs

The Portals and API support Participants to interact with Federation Services functions via a user interface, which provides mechanisms to interact with core capabilities using API calls. The goal is a consistent user experience for all tasks that can be performed with a specific focus on security and compliance. The Portals provide information on Resources and Service Offerings and interaction mechanisms for tasks related to their

maintenance, including identity and access management. Each Ecosystem can deploy its own Portals to support interaction with Federation Services.

7. Example Gaia-X Participant Use Cases

The goal of this section is to illustrate how the Consumers, Federators and Providers described in the conceptual model can appear in the real word. This section focuses on the most typical kinds of actors and the list is **not exhaustive**. Examples of Gaia-X Use Cases can be found in the position paper published by the Dataspace Business Committee.

7.1 Provider/Consumer Use Cases

This section describes typical kinds of Service Offering that Provider can define, offer and provide, to be configured and consumed by Consumer.

- cloud services: Infrastructure as a Service, Platform as a Service, Software as a Service, XXX as a Service.
- datasets: data sharing, in batch, stream and event driven dataset.
- software licenses: perpetual or renewable licenses for a product without an associated online service.
- interconnection & networking services that can go beyond the capacities of the regular Internet connection and exhibit special characteristics, such as and not limited to bandwidth, latency, availability or security-related settings.

7.2 Federator Use Cases

The different Federators are not distinguished as being either domain-specific or cross-domain. Only accordance to the Policy Rules and operating according to the conditions mentioned in the Operating Model, including the respective conformity assessments and trust mechanisms, distinguish whether it is an ecosystem federated by Gaia-X policies or not.

Federators have the option to facilitate an ecosystem by using the available open source Federation Services software, such as and not limited to Catalogue, Wallet, Logging, Authorisation and Access Management

7.3 Basic Interactions of Participants

This section describes the basic interaction of the different Participants as described in the conceptual model (see section 2).

Providers and Consumers within a Ecosystem are identified and well described through their valid Self-Description, which is 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 Federators 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 AISBL does not play an intermediary role during the Contract negotiations but ensures the trustworthiness of all relevant Participants and Service Offerings.

The following diagram presents the general workflow for Gaia-X service provisioning and consumption processes. Please note that this overview represents the current situation and may be subject to changes according to the Federation Services specification. The specification will provide more details about the different elements that are part of the concrete processes.

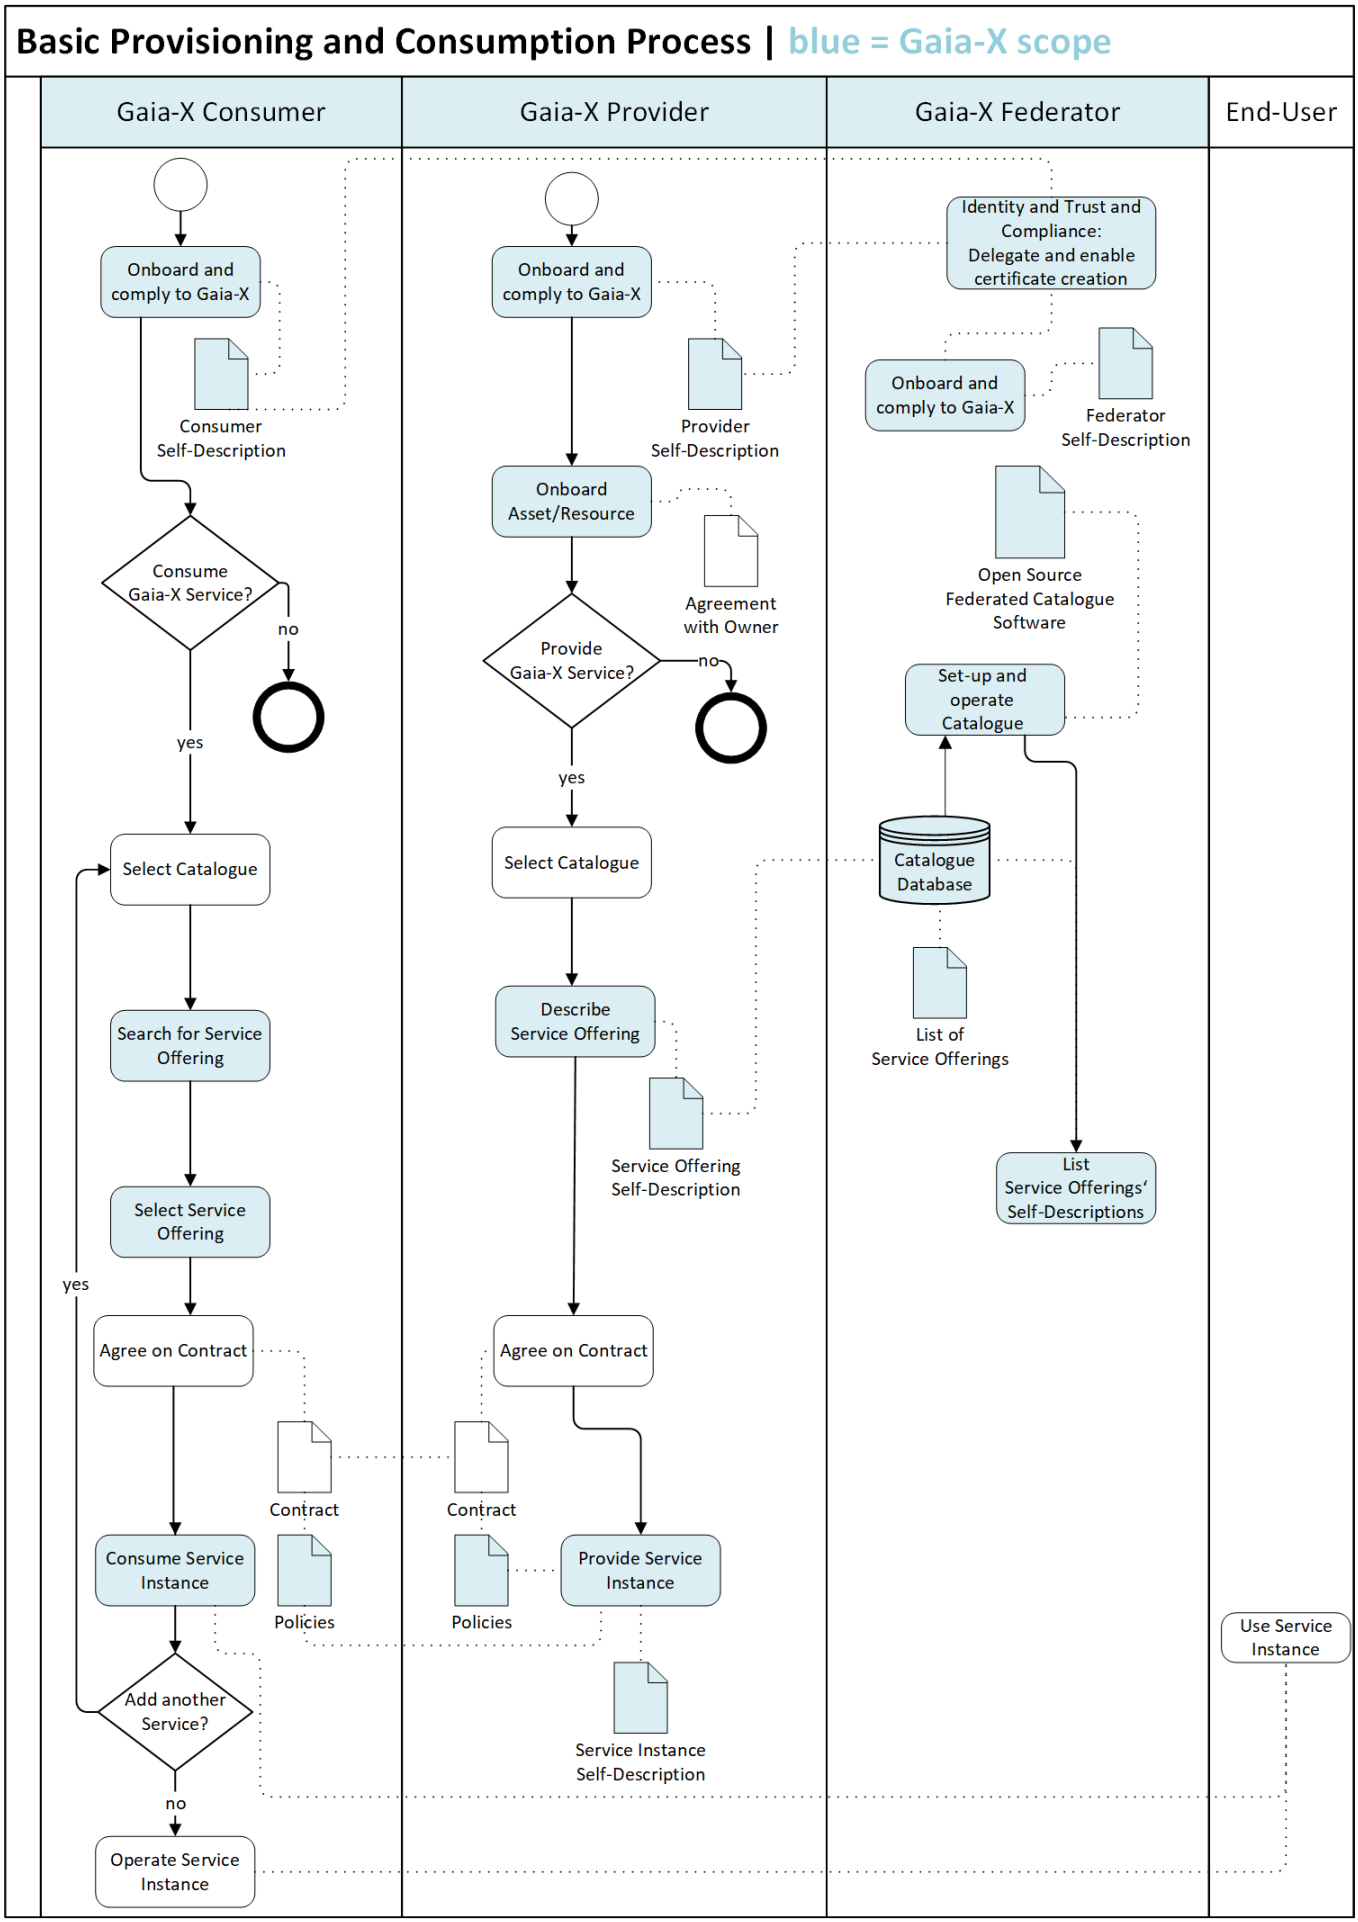

10 Basic Provisioning and Consumption Process

8. Changelog

8.1 2022 April release

- 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

8.2 2021 December release

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

8.3 2021 September release

- Rewrite of 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.

8.4 2021 June release

- 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.

8.5 2021 March release

- 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.

8.6 2020 June release

- First release of the Technical Architecture document by the BMWi

9. References

- Berners-Lee, T. (2009). Linked Data. W3C. https://www.w3.org/DesignIssues/LinkedData
- Bohn, R. B., Lee, C. A., & Michel, M. (2020). The NIST Cloud Federation Reference Architecture: Special Publication (NIST SP) - 500-332. NIST Pubs. https://doi.org/10.6028/NIST.SP.500-332
- ETSI. Network Functions Virtualisation (NFV). https://www.etsi.org/technologies/nfv
- European Commission. Trusted List Browser: Tool to browse the national eIDAS Trusted Lists and the EU List of eIDAS Trusted Lists (LOTL). https://webgate.ec.europa.eu/tl-browser/#/
- European Commission. (2020). Towards a next generation cloud for Europe. https://ec.europa.eu/digital-single-market/en/news/towards-next-generation-cloud-europe
- European Commission Semantic Interoperability Community. DCAT Application Profile for data portals in Europe. https://joinup.ec.europa.eu/collection/semantic-interoperability-community-semic/solution/dcat-application-profile-data-portals-europe
- Federal Ministry for Economic Affairs and Energy. (2019). Project Gaia-X: A Federated Data Infrastructure as the Cradle of a Vibrant European Ecosystem. https://www.bmwi.de/Redaktion/EN/Publikationen/Digitale-Welt/project-Gaia-X.htm
- Federal Ministry for Economic Affairs and Energy. (2020). Gaia-X: Technical Architecture: Release - June, 2020. https://www.data-infrastructure.eu/GAIAX/Redaktion/EN/Publications/Gaia-X-technical-architecture.html
- Gaia-X association AISBL. Architecture Decision Record (ADR) Process: GitLab Wiki. https://gitlab.com/Gaia-X/Gaia-X-technical-committee/Gaia-X-core-document-technical-concept-architecture/-/wikis/home
- ISO / IEC. Intelligent transport systems - Using web services (machine-machine delivery) for ITS service delivery (ISO / TR 24097-3:2019(en)). https://www.iso.org/obp/ui/fr/#iso:std:iso:tr:24097:-3:ed-1:v1:en
- ISO / IEC. IT Security and Privacy – A framework for identity management: Part 1: Terminology and concepts (24760-1:2019(en)). ISO / IEC. https://www.iso.org/obp/ui/#iso:std:iso-iec:24760:-1:ed-2:v1:en
- IX-API. IX-API. https://ix-api.net/
- OASIS (2013). Topology and Orchestration Specification for Cloud Applications Version 1.0. http://docs.oasis-open.org/tosca/TOSCA/v1.0/TOSCA-v1.0.html
- 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 https://doi.org/10.6028/NIST.IR.4734
- Open Source Initiative. Licenses & Standards. https://opensource.org/licenses
- Open Source Initiative. The Open Source Definition (Annotated). https://opensource.org/osd-annotated

- Plattform Industrie 4.0: Working Group on the Security of Networked Systems. (2016). Technical Overview: Secure Identities. https://www.plattform-i40.de/PI40/Redaktion/EN/Downloads/Publikation/secure-identities.pdf 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
- W3C. JSON-LD 1.1: A JSON-based Serialization for Linked Data [W3C Recommendation 16 July 2020]. https://www.w3.org/TR/json-ld11/
- W3C. ODRL Information Model 2.2 [W3C Recommendation 15 February 2018]. https://www.w3.org/TR/odrl-model/
- W3C. Verifiable Credentials Data Model 1.0: Expressing verifiable information on the Web [W3C Recommendation 19 November 2019]. https://www.w3.org/TR/vc-data-model/
- W3C. (2015). Semantic Web. https://www.w3.org/standards/semanticweb/
- W3C. (2021). Decentralized Identifiers (DIDs) v1.0. https://www.w3.org/TR/did-core/

---

1. European Commission. (2020). Towards a next generation cloud for Europe. https://ec.europa.eu/digital-single-market/en/news/towards-next-generation-cloud-europe ↩
2. Federal Ministry for Economic Affairs and Energy. (2019). Project Gaia-X: A Federated Data Infrastructure as the Cradle of a Vibrant European Ecosystem. https://www.bmwi.de/Redaktion/EN/Publikationen/Digitale-Welt/project-Gaia-X.htm ↩
3. ISO/IEC. IT Security and Privacy — A framework for identity management: Part 1: Terminology and concepts (24760-1:2019(en)). ISO/IEC. https://www.iso.org/obp/ui/#iso:std:iso-iec:24760:-1:ed-2:v1:en ↩
4. 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 ↩
5. 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 ↩
6. Rights of the data subject https://gdpr-info.eu/chapter-3/ ↩
7. Payment services (PSD 2) https://ec.europa.eu/info/law/payment-services-psd-2-directive-eu-2015-2366_en ↩
8. W3C (2012). Web Ontology Language (OWL). https://www.w3.org/OWL/ ↩
9. Graph Query Language GQL, https://www.gqlstandards.org/ ↩
10. https://w3c.github.io/rdf-star/cg-spec/editors_draft.html ↩

Gaia-X Trust Framework draft: https://gaia-x.gitlab.io/policy-rules-committee/trust-framework/ ↩
11.

12. Example of the setup of a DAO https://blockchainhub.net/dao-decentralized-autonomous-organization/ ↩

13. Example of decentralized data and algorithms marketplace https://oceanprotocol.com/ ↩

14. Bohn, R. B., Lee, C. A., & Michel, M. (2020). The NIST Cloud Federation Reference Architecture: Special Publication (NIST SP) - 500-332. NIST Pubs. https://doi.org/10.6028/NIST.SP.500-332 ↩

15. See the IAM Framwork version 1.2 for details: https://community.gaia-x.eu/s/P23ZJNLyjf7n7Zp?path=%2FReleases. ↩

16. For more details on Secure Identities, see Plattform Industrie 4.0: Working Group on the Security of Networked Systems. (2016). Technical Overview: Secure Identities. https://www.plattform-i40.de/PI40/Redaktion/EN/Downloads/Publikation/secure-identities.pdf as well as Chapter 3.4 in the IAM Framework v1.2: https://community.gaia-x.eu/s/P23ZJNLyjf7n7Zp?path=%2FReleases. ↩

17. W3C. ODRL Information Model 2.2 [W3C Recommendation 15 February 2018]. https://www.w3.org/TR/odrl-model/ ↩

18. Currently not in scope of Gaia-X Federation Services ↩