Gaia-X Architecture Document

24.04 Release

Gaia-X European Association for Data and Cloud AISBL

Table of contents

1. About

1.1 1. Editorial Information

1.1.1 1.1 Publisher

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

1.1.2 1.2 Authors

Gaia-X European Association for Data and Cloud

1.1.3 1.3 Contact

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

1.1.4 1.4 Other format

For convenience a PDF version of this document is generated here.

1.1.5 1.5 Copyright notice

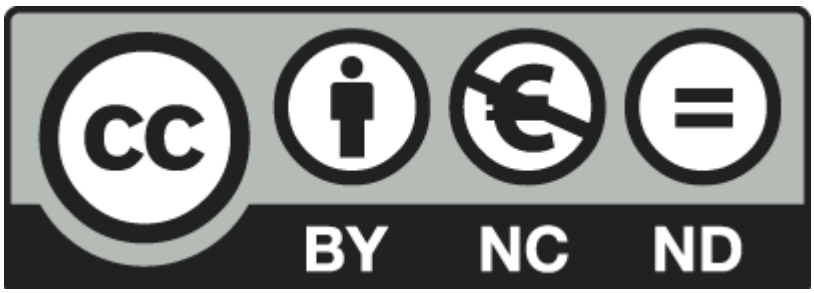

1.2 2. Context

The Gaia-X Association is working with other parties to enable the creation of infrastructure and data-driven ecosystems. This document explains the core elements that compose the Gaia-X Trust Framework and defines how they relate to each other at the functional level in the Gaia-X model. Furthermore, the document explains how the Gaia-X Trust Framework integrates with implementations of Infrastructure and Data Space-specific services.

1.2.1 2.1 Gaia-X as a member of the Data Spaces Business Alliance (“DSBA”)

Gaia-X together with the International Data Spaces Association (IDSA), the FIWARE Foundation and the Big Data Value Association has formed the Data Spaces Business Alliance to provide a comprehensive and complementary stack of specifications and OSS components that enable and create value out of distributed digital (data) ecosystems.

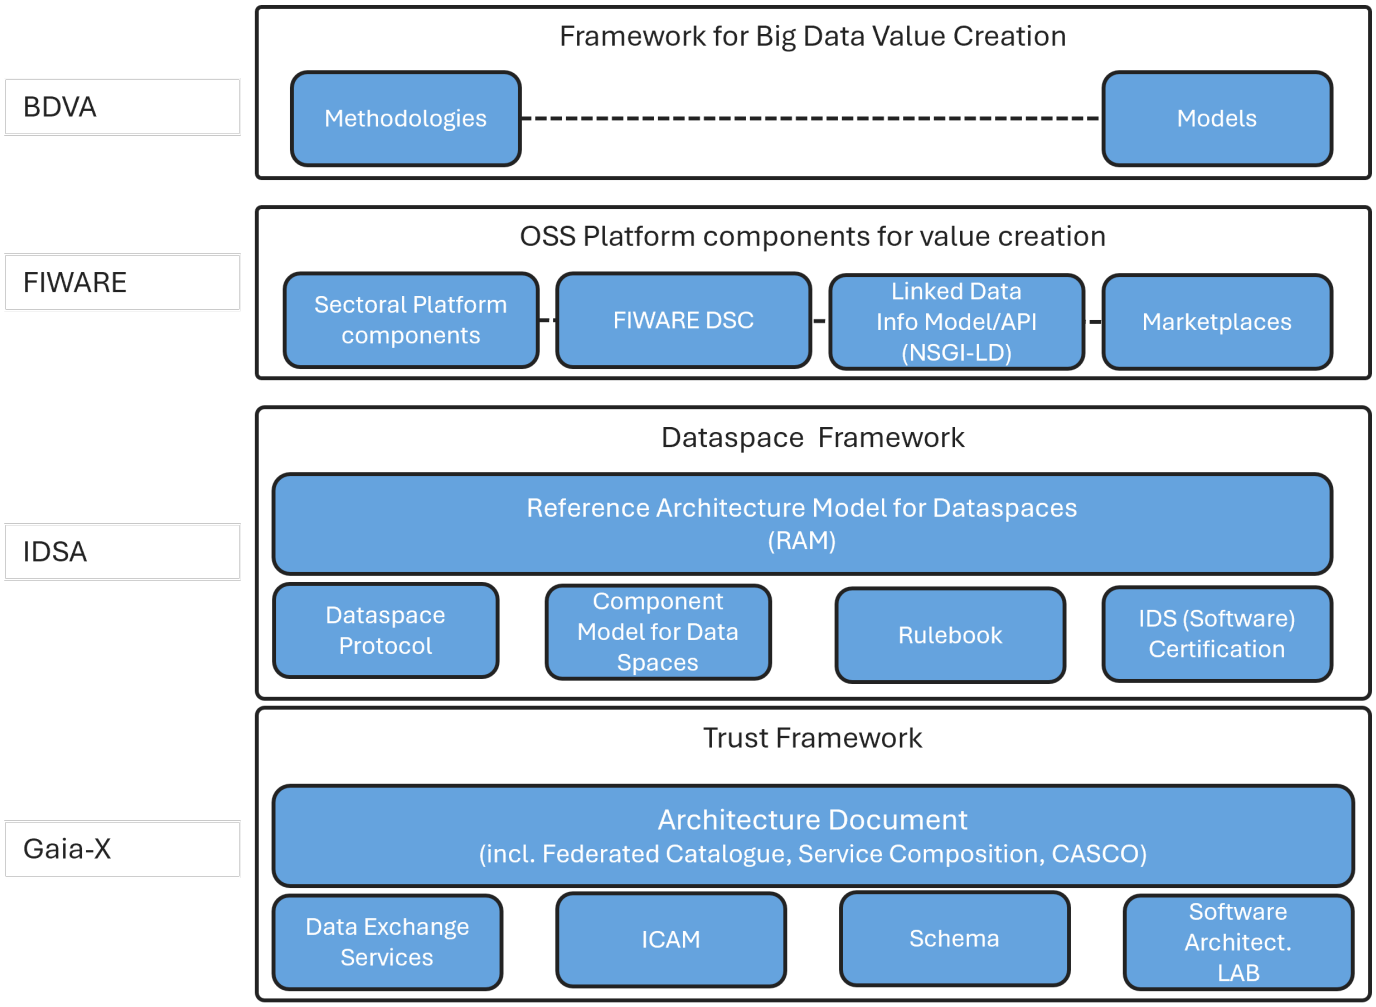

The DSBA members continue to converge and integrate their respective artefacts to provide a comprehensive, interoperable and scalable set of solutions, but they remain open to the integration of other solutions.

1.2.2 2.2 Gaia-X Trust Framework

Gaia-X provides the Trust Framework, which ensures that all participants in the Gaia-X Ecosystem are adhering to the policy rules agreed between the participants of the Ecosystem itself.

The Gaia-X Trust Framework defines the process of going through and validating the set of automatically enforceable rules to achieve the minimum level of compatibility in terms of:

- syntactic correctness;
- schema validity;
- cryptographic signature validation;
- attribute value consistency;
- attribute value verification.

The claims are all contained in verifiable credentials, independently of who is performing the assessment.

Whenever possible, the claims validation is done either by using publicly available open data and performing tests or using data from Trusted Data Sources as defined in the following sections.

The Gaia-X Trust Framework provides the interfaces to various technologies connecting data (ensuring data sovereignty) and services, which help to establish ecosystems, such as catalogues, identity solutions, and marketplaces.

The underlying trust is provided by Gaia-X Credentials.

2. Models & Components

2.1 3. Gaia-X Conceptual Model

2.1.1 3.1 Main Concepts

3.1.1 Gaia-X ecosystems

Gaia-X Credentials provide the necessary trust elements to create ecosystems that operate under a commonly defined governance. The Gaia-X credentials ensure that the policy rules agreed upon between the participants are verified and can be validated at any time.

Components of such ecosystems typically are:

- the ecosystem governance: defining the set of rules agreed upon by the parties in the ecosystem - which must be operationalised. The ecosystem governance typically defines a set of providers, who are providing services to enable basic ecosystem services (like trust services, catalogues – see chapter Enabling and Federation Services) which we define as “Federation Services”.
- infrastructure - i.e., hardware and software for computing, storage, and network services - adopting the rules defined by the governance. This includes services that support a federation of providers (“Federation Services”)
- data ecosystems and data spaces, where participants adopt the governance, using the infrastructures “to access and use data in a fair, transparent, proportionate and/ non-discriminatory manner with clear and trustworthy data governance mechanisms.” [1]. Data Spaces can span across several Infrastructure and Data Ecosystems.

The Gaia-X Ecosystem is the virtual set of Service Offerings described by Gaia-X compliant credentials, according to the Compliance schemes set by the Gaia-X Association.

3.1.2 Decentralised trust framework for ecosystems

In this challenging environment where each Data Space wants to both be interoperable and yet adapt their governance to their vertical, domain-specific needs, local market regulation, the Gaia-X Trust Framework provides a set of world-wide applicable rules and specifications usable by:

- the ecosystem governance (e.g.: Data Spaces authorities, such as Data Intermediaries from the Data Governance Act).
- ecosystems seeking interoperability and technical compatibility of their services.

The ecosystem governance defines the applicable policy rules to participate in the digital (infrastructure, data or services) ecosystem, together with the Trust Anchors and Schema Extensions policy rules which apply to operators of services in the specific ecosystem.

The interoperability in terms of governance is assessed by the Gaia-X Compliance and the Trust Indexes.

The interoperability in terms of technical compatibility is assessed by the Gaia-X Testbed.

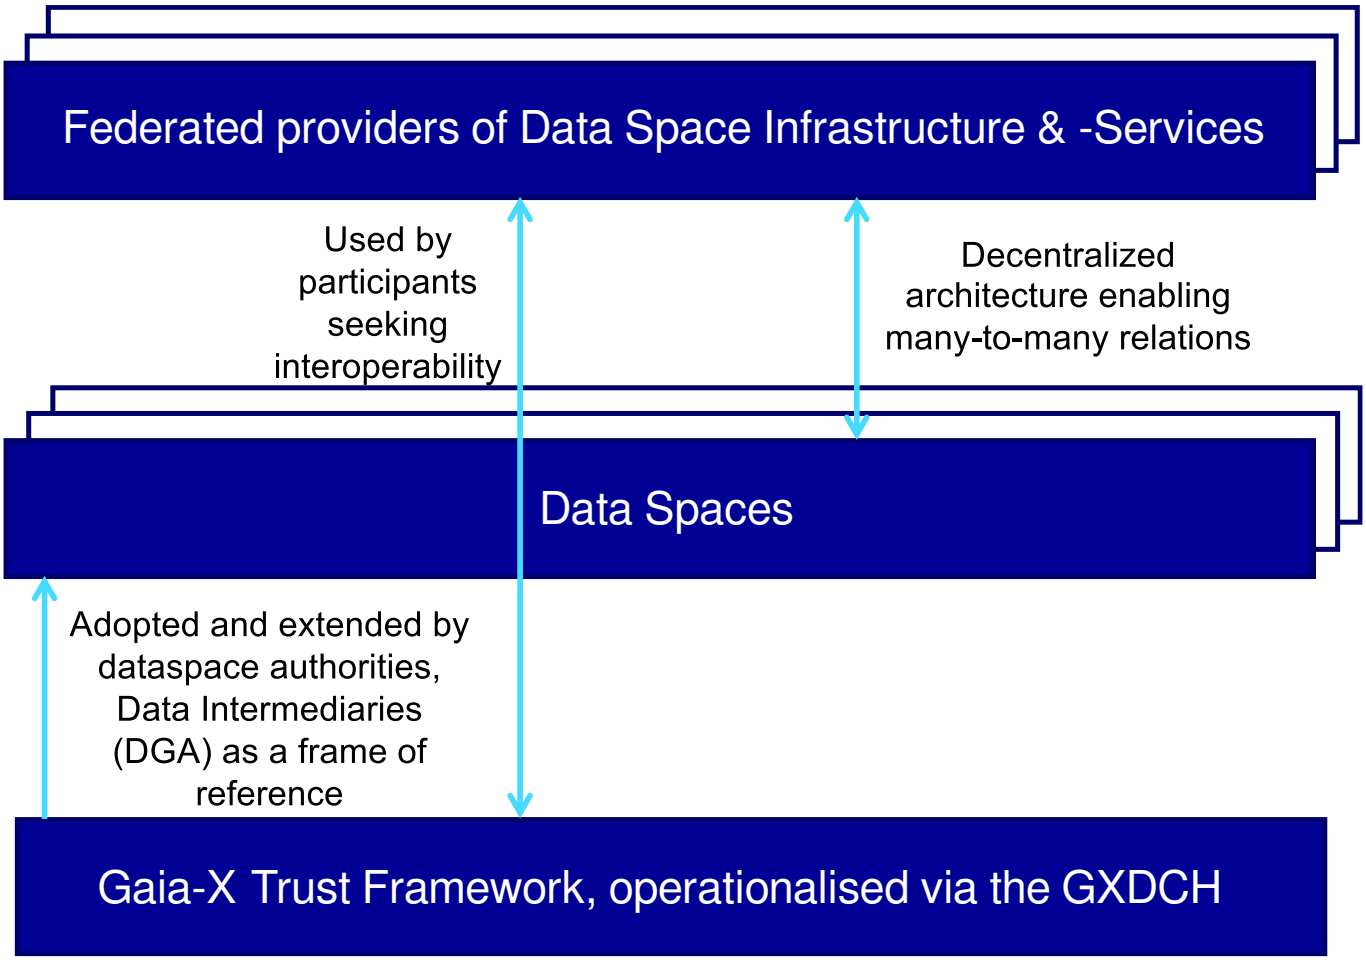

Figure 3.1 - The Gaia-X Trust Framework as a foundation for interoperable data spaces

3.1.2.1 NIST CLOUD FEDERATION REFERENCE ARCHITECTURE

The three planes from the above picture can be mapped on the three planes described in the NIST Cloud Federation Reference Architecture chapter 2:

- The Trust plane: this plane represents the global digital governance that is shared across Data Spaces and Federations.
- The Management plane: this plane represents an extension of the common digital governance to answer specific business needs.
- The Usage plane: this plane captures the technical interoperability, including the one between Service Offerings and Products.

3.1.3 The Gaia-X model building block

The building block of the Gaia-X model is based on policy expressions attached to each entity of the model. The `policies` are expressed by one or more `parties` about one or more `assets`. An `asset` can also be a `party`, making this simple model recursive and capable of handling the most complicated user scenario, with multiple providers, consumers, federators, data intermediaries, data subjects, business legal representatives, employer/employee and much more.

Kroki

The workflow above is the generic representation of the Open Digital Rights Language (ODRL) model, which is used in the rest of this document.

2.1.2 3.2 Roles

3.2.1 Trust Anchor

Gaia-X Trust Anchors are bodies, parties, i.e., Conformity Assessment Bodies or technical means accredited by the bodies of the Gaia-X Association to be parties eligible to issue attestations about specific claims.

3.2.2 Participant

A Participant is an entity, as defined in ISO/IEC 24760-1 as an “item relevant for the purpose of operation of a domain that has recognisably distinct existence”[2], which is onboarded and has a Gaia-X Participant Credential. A Participant can take on one or more of the following roles: Provider, Consumer, and Operator.

Provider and Consumer represent the core roles that are in a business-to-business relationship, while Operators are Providers that have the specific role of providing Federation Services, enabling the interaction between Providers and Consumers.

3.2.2.1 PROVIDER

A Provider operates Resources in the Gaia-X Ecosystem and offers them as services through Gaia-X Service Offering credentials. For any such service, the Gaia-X Provider defines the Service Offering including terms and conditions as well as technical policies. Furthermore, it provides the Service Instance that includes a Credential and associated policies. A Gaia-X Provider is responsible for conformity to the claims made. If third-party data, services or infrastructure are part of the service offerings the Gaia-X Provider can make those resources available (e.g., through Service Composition) but remains responsible and shall ensure coverage through appropriate back-to-back coverage.

3.2.2.2 OPERATOR

Operators are Gaia-X Providers that have been approved by the ecosystem governance to operate Federation Services and the Federation, which are independent of each other. There can be one or more Operators per type of Federation Service.

3.2.2.3 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.

3.2.3 Basic Interactions of Participants

This section describes the basic interaction of the different Participants in an ecosystem based on the Gaia-X model.

Providers and Consumers within the ecosystem are identified and well described through their valid Credentials, which are initially created before or during the onboarding process. Providers define their Service Offerings and publish them in a Catalogue. In turn, Consumers search for Service Offerings in Gaia-X Catalogues that are coordinated by Operators and the Gaia-X Registry. Once the Consumer finds a matching Service Offering in a Gaia-X Catalogue, the Contract negotiation between Provider and Consumer determines further conditions under which the Service Instance will be provided. The Gaia-X Association does not play an intermediary role during the Contract negotiations but ensures the trustworthiness of all relevant Participants and Service Offerings.

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. The technical specifications will provide more details about the different elements that are part of the concrete processes.

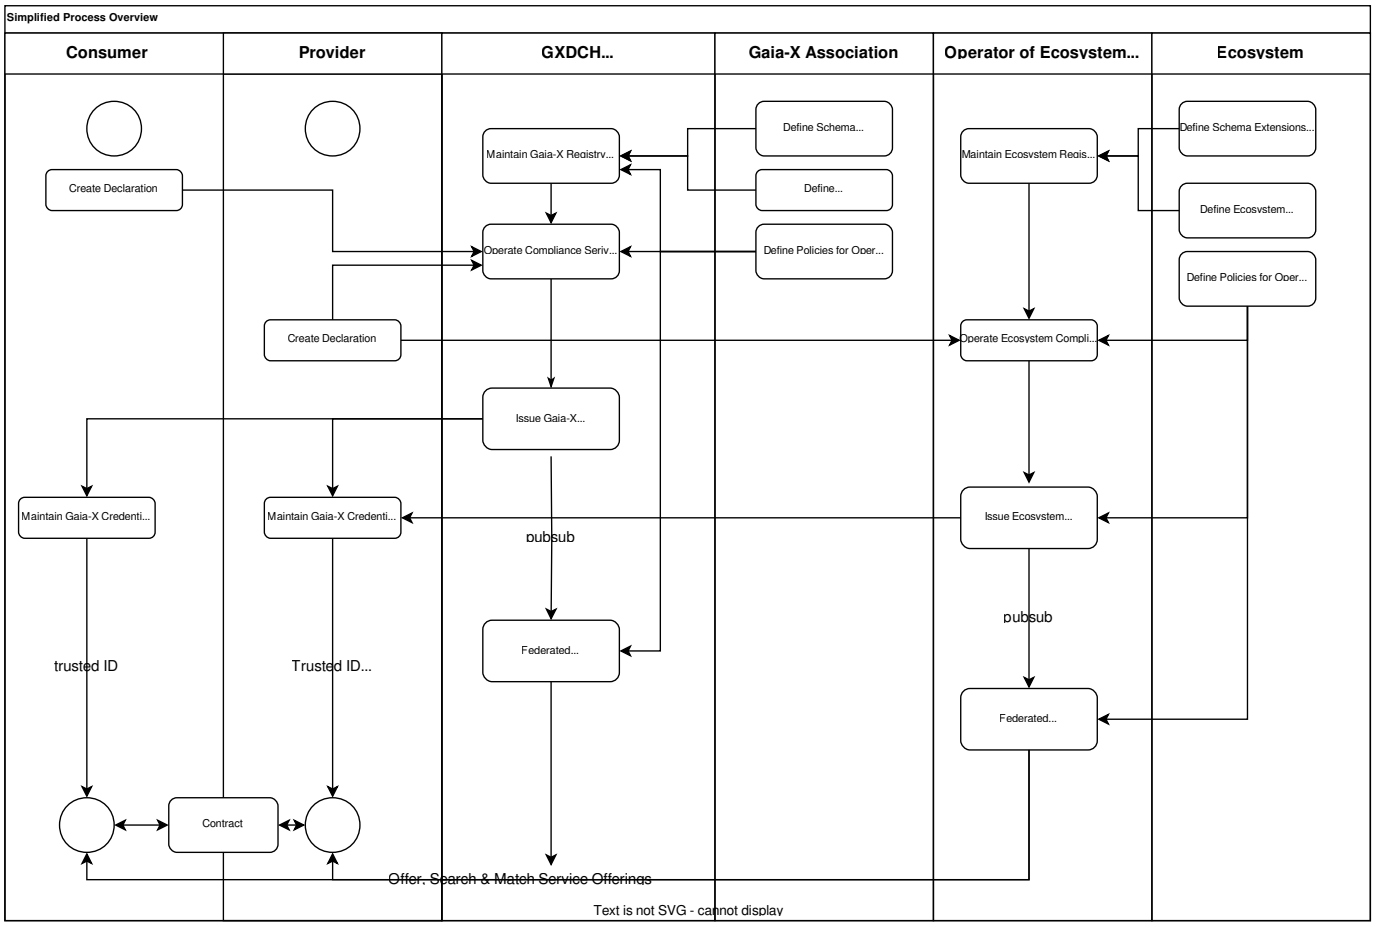

Figure 3.2 - Basic processes for provisioning and consumption of Gaia-X services

2.1.3 3.3 Components

3.3.1 Gaia-X Credentials

Gaia-X Credentials - formerly known as Self-Descriptions (*SD*) - are cryptographically signed attestations describing Entities from the Gaia-X Conceptual Model in a machine-interpretable format.

The Gaia-X Credentials are the building blocks of a decentralized machine-readable knowledge graph of claims, each credential carrying a tamper-proof and authenticable part of the information of that graph.

The knowledge graph can be completed and extended by Federations and Data Spaces, which always keep control of what and how much information is being shared by implementing access control at the credentials level.

Kroki

Gaia-X Credentials refer to W3C Verifiable Credentials containing claims using the Gaia-X schema in their context.

Other types of credentials or credentials not using the Gaia-X schema are out of the scope of Gaia-X.

Gaia-X Credentials are W3C Verifiable Credentials with claims expressed in RDF.

3.3.1.1 GAIA-X SCHEMA

The Gaia-X members define the Schema for Gaia-X Credentials. It is used as the vocabulary of the claims about credential subjects and must be available in the form of SHACL shapes (cf. the W3C Shapes Constraint Language SHACL[3]).

At any point where Credentials are created or received, a certain set of SHACL shapes is known, which forms a shapes graph. A Credential forms a data graph. For compliance with Gaia-X and/or specific ecosystem extensions, this data graph must be validated against the given shapes graph according to the SHACL specification.

The defined version of the Gaia-X Schema is maintained and made available through the Gaia-X Registry Service.

3.3.2 Policies

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

- Policy representation: interoperable access and usage policies that are specified in a human- and machine-readable format. Policies generally express three possible restrictions: prohibitions, obligations, and permissions. Constraints defining a rule can be combined into more complex rules, which then form the applicable policy.
- Decision-taking/policy engine: during the execution of a data transaction, the policies need to be evaluated. This decision typically requires context information. With the decision context, the policy engine will decide whether the request or usage is permitted. The evaluation process is handled by the policy engine, which is instantiated by the data product provider and the data consumer or a trusted third party.
- Enforcement and execution of policies: the enforcement and execution of policies is a key capability which needs to be implemented for both access and usage policies.

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

3.3.2.1 GAIA-X POLICY REASONING ENGINE

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

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

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

3.3.2.2 POLICY DECISION POINT (PDP)

3.3.2.2.1 Reference implementation

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

3.3.2.3 POLICY ENFORCEMENT

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

The enforcement phase starts when the actual data transaction is being executed and it takes place throughout the data transaction. The goal of the enforcement phase is to evaluate the relevant rules of the policy and decide whether or not the data transaction is allowed to proceed unless an agreement or contract has already been negotiated, in which case the only need is to verify the contract's validity.

For the different types of rules, the evaluation might occur at different stages of the data transaction: - Access rules are primarily evaluated by the data provider before any data is exchanged. A data consumer may also check these rules when receiving data to ensure it is still according to the access rules. - Usage rules are evaluated by the data consumer when the data is used. Depending on the usage rule, evaluation might be required not just when starting to use the data but constantly throughout the lifecycle of the data usage. - Consent rules require the evaluation of consent of third parties, which might be revoked during the data transaction or data use. Therefore, evaluation should occur both at the data product provider and consumer sides.

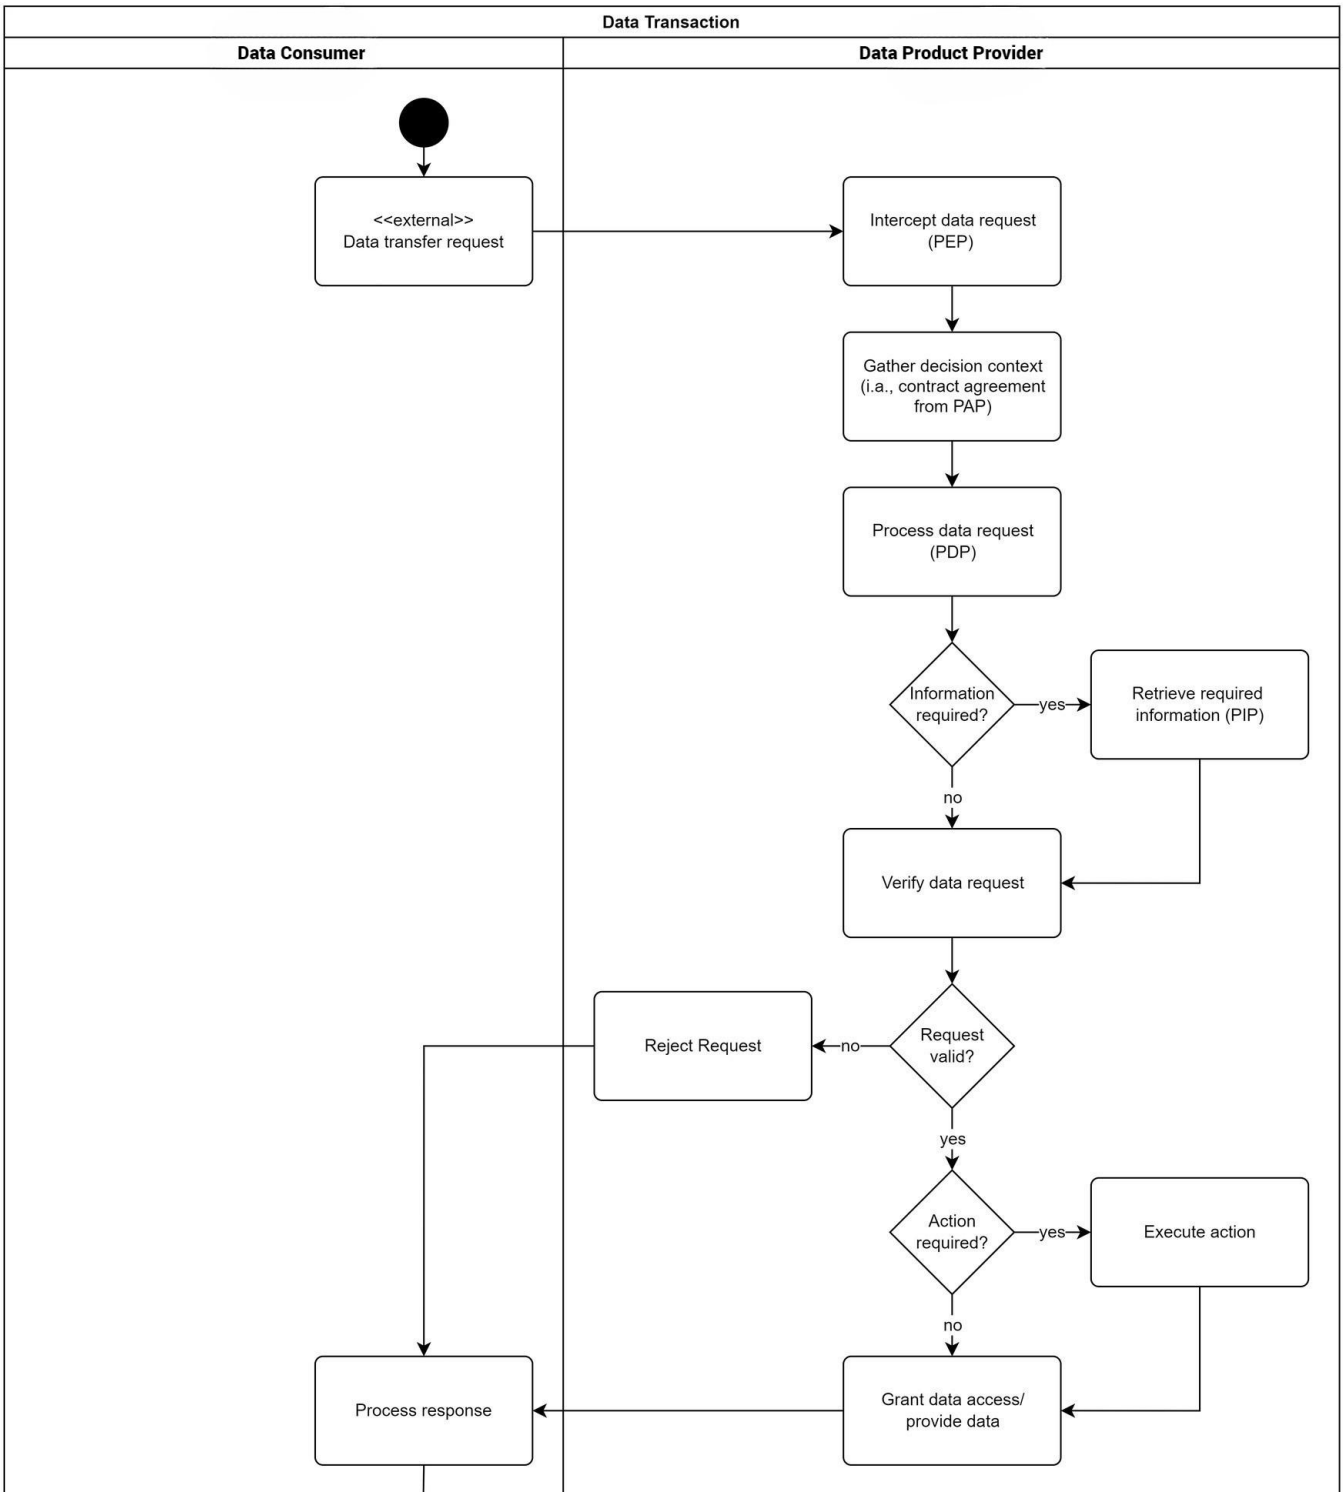

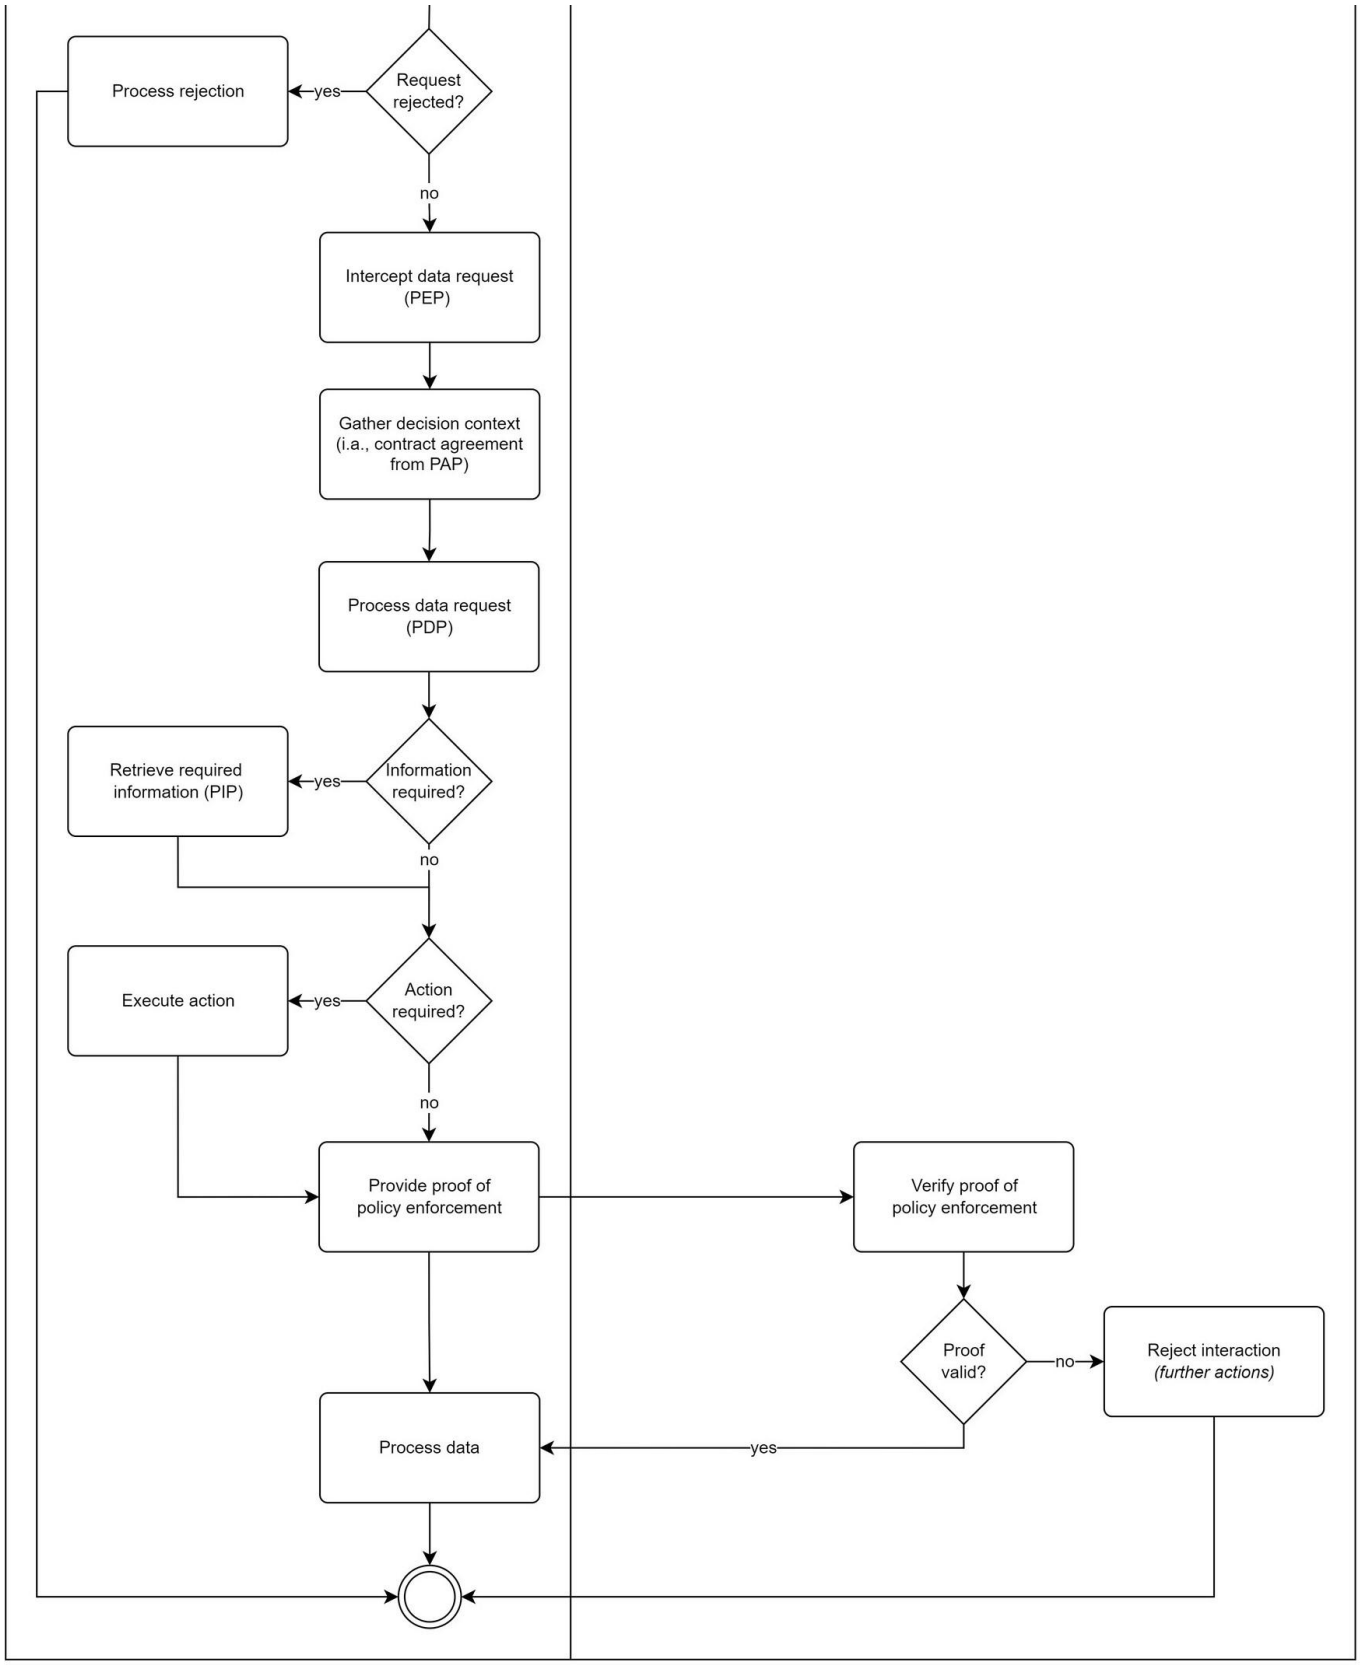

Figure 3.3 - Policy Enforcement Flow

3.3.3 External

3.3.3.1 IDENTITIES

An Identity is composed of a unique Identifier and an attribute or set of attributes that uniquely describe an entity within a given context. Gaia-X uses existing Identities and does not maintain them directly.

Identities uniquely describe participants (natural persons, companies) and resources (e.g., machines, interconnection or data endpoints). Personal and Corporate Identities are validated by the Gaia-X Compliance Service and Gaia-X Participant credentials issued. Identities are represented using **Party Credential** specializations.

3.3.3.2 SERVICES

All offerings by Gaia-X providers are considered “Services” (these can be composed of different types of physical or virtual resources). A service description that follows the Gaia-X Schema and whose claims are validated by the Gaia-X Compliance Service becomes a Gaia-X Service-Offering Credential.

2.1.4 3.4 Overview picture

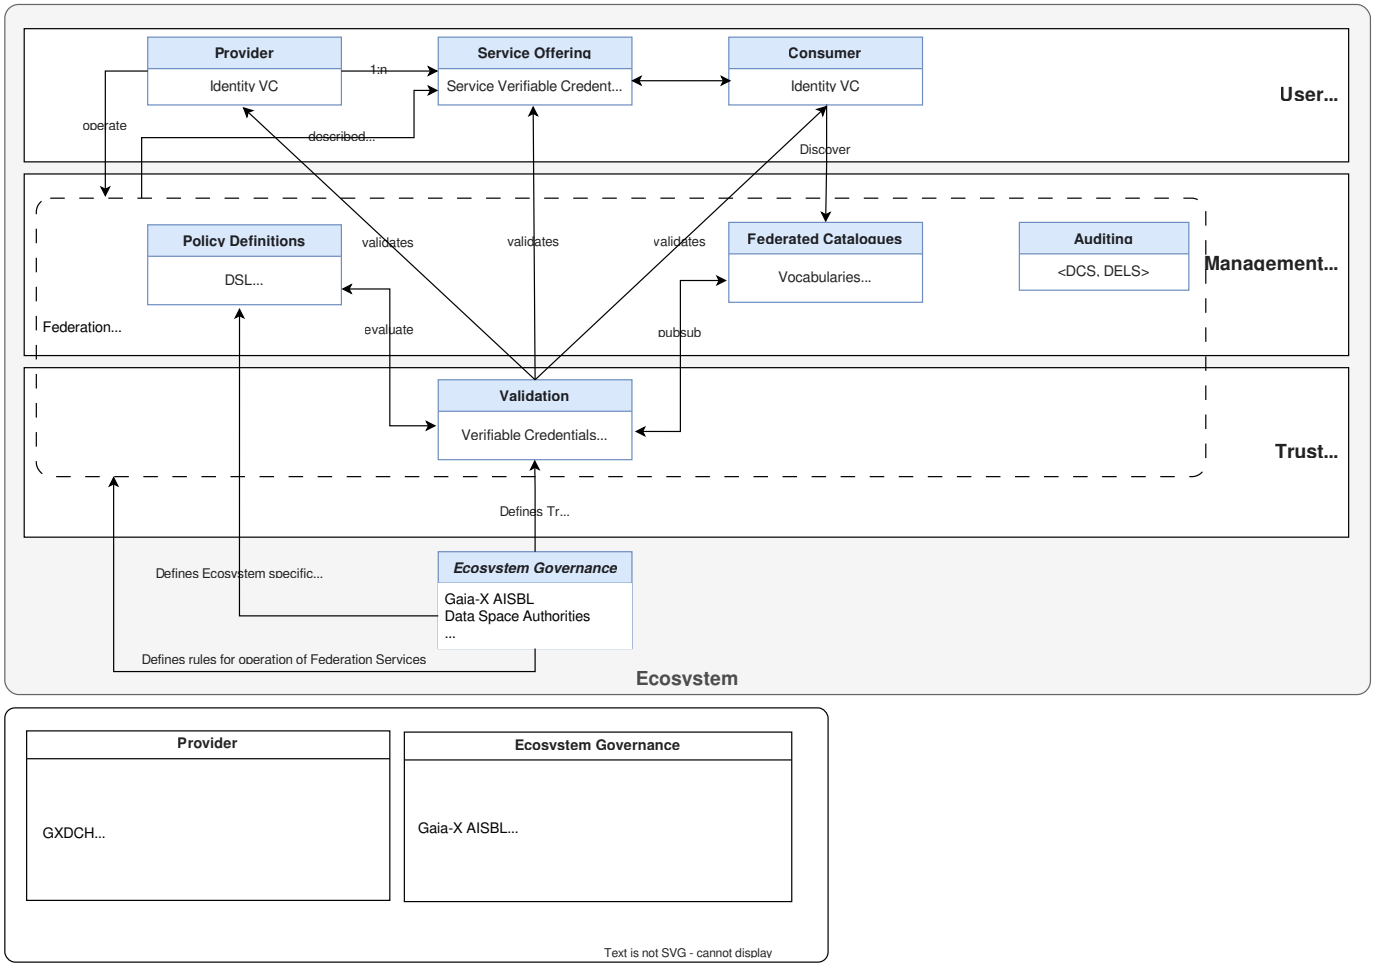

Figure 3.4 - Gaia-X conceptual model overview

1. https://joinup.ec.europa.eu/collection/semic-support-centre/data-spaces ↩
2. 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 ↩
3. https://www.w3.org/TR/shacl/ ↩

2.2 4. Component details

2.2.1 4.1 Resources and Service Offerings

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

Resource Categories

A Resource can be a:

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

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

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

Kroki

2.2.2 4.2 Policies

4.2.1 Policy definition

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[12].

The Gaia-X Policy Rules Compliance Document defines the Gaia-X Policies for Service Offerings. They cover, for example, privacy or cybersecurity policies and are expressed in the Conceptual Model indirectly as attributes of the Resources, Service Offerings, and Service Instances.

4.2.2 Policy description

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

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

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 defined by Gaia-X.

Kroki

4.2.2.1 POLICY ENFORCEMENT

Policy enforcement is executed by software engines (e.g. the Gaia-X Compliance Service or engines within a Data Exchange Service).

2.2.3 4.3 Service composition

4.3.1 Assumptions

The generic Gaia-X service composition model is derived assuming the availability of key related functions and systems within the Gaia-X Ecosystem. The first assumption is the availability of a Catalogue containing service offerings compliant with Gaia-X. These correspond to provider services with self-attested credentials and credentials verified by the Gaia-X Compliance Service. In addition to these services, there can be external and independent services offered by multiple third parties. A requirement on these third-party service offerings is to provide independent services compatible with Gaia-X compliant services. It is thus possible to combine them to produce composite services. Means to check for compatibility and composability are provided and potentially built into the respective credentials. Service composition

consequently assumes that searching for compatible and interoperable services in multiple Catalogues is possible. Whenever such capability is not available, service offerings nevertheless come along with service templates and service descriptors and descriptions that readily embed this key information. Hence, credentials have to contain key characteristics and must have the appropriate structure. They would embed the following TOSCA-like or similar descriptors:

- Capabilities
- Requirements
- Properties
- Operations
- Artifacts
- Compatibility, interoperability, composability, substitutability attributes (information, list, possibly revealing if service components are bundles)
- Interfaces

Furthermore, we assume this embedded additional information in service descriptions in the Gaia-X service catalog. That is, providing information on how to chain services to ensure successful and failure free composition as well as guaranteed operation at instantiation and run time. The service composition model shall also allow for the portability of applications among cloud Providers, adhering to the high-level objectives described in the Franco German position on Gaia-X.

4.3.2 Generic Service Composition Model

Considering an open European and international context involving multiple stakeholders and participants, the Gaia-X service composition model has to be abstract at the start of any initial service composition. The focus of the composition model is consequently on the service functional behaviour, with no constraints, localization, or preset values of non-functional attributes. Values are set only once the end users, tenants or consumers, have expressed their requirements, constraints and preferences. These values are gradually set as service composition moves from initial provisioning to the life cycle management of services at run time. Figure 1 depicts a high-level class diagram view of the service composition model.

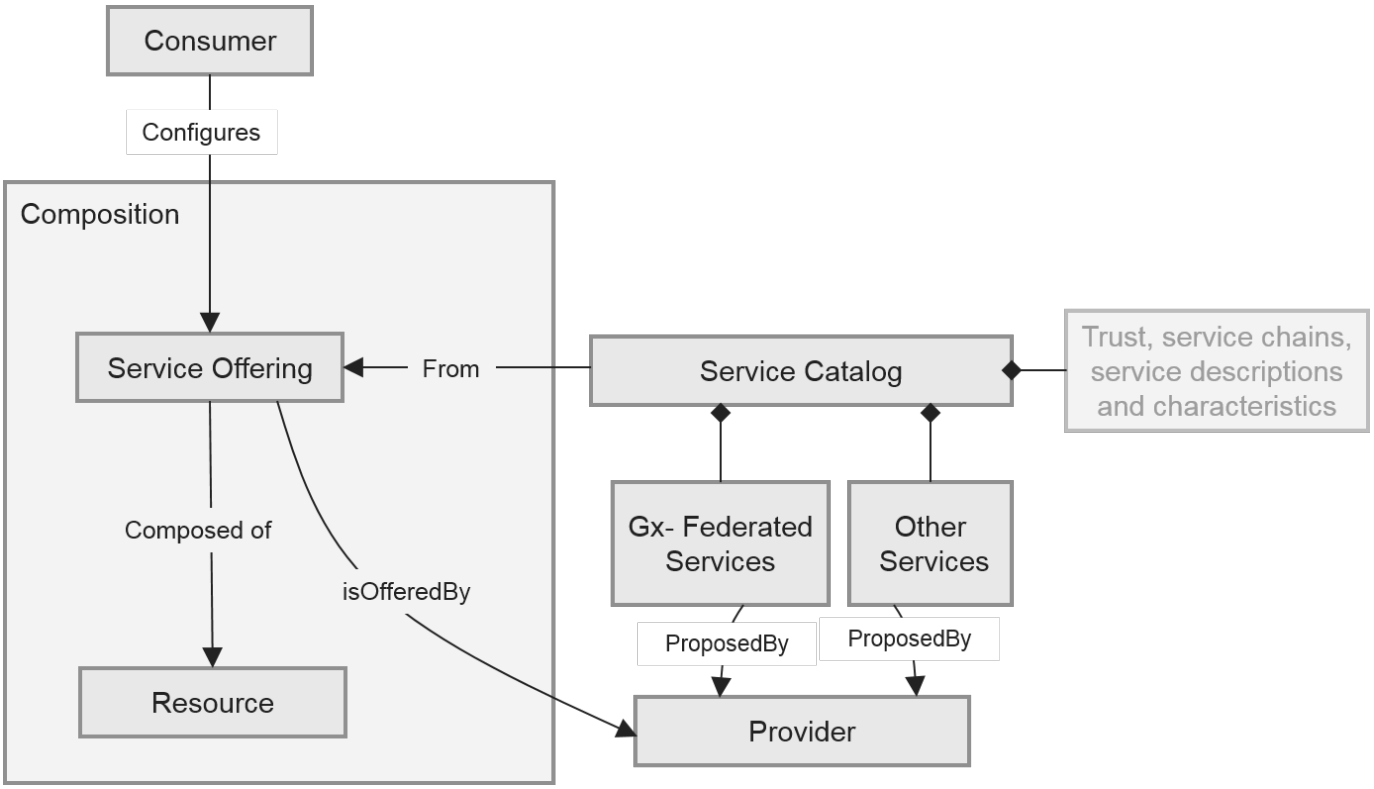

Figure 4.1 - Simplified and abstract conceptual service composition model

This process shall by no means prevent or hamper interoperability, portability, or substitution of services and applications. This should require no or only minor changes. Figure 4.1 depicts the abstract and simplified service composition model. Beyond this simplified and abstract service composition view and class diagram, Figure 4.2 illustrates a more elaborate class diagram representation of the service composition model. The figure includes details and related functions and modules involved in service composition starting from a user service request to the instantiation of the services by providers. It also includes and extends the Gaia-X conceptual model with an expanded view of the service composition model.

4.3.3 Conceptual Service Composition Model

We start with a generic and abstract view of the conceptual service composition model and illustrate where it fits in the Gaia-X conceptual model. As far as the Gaia-X service composition model is concerned, both API and template-based services have to be included in the modelling framework. The role “Operator” in the diagram below indicates a provider which has been designated by the ecosystem governance to operate “Federation Services”

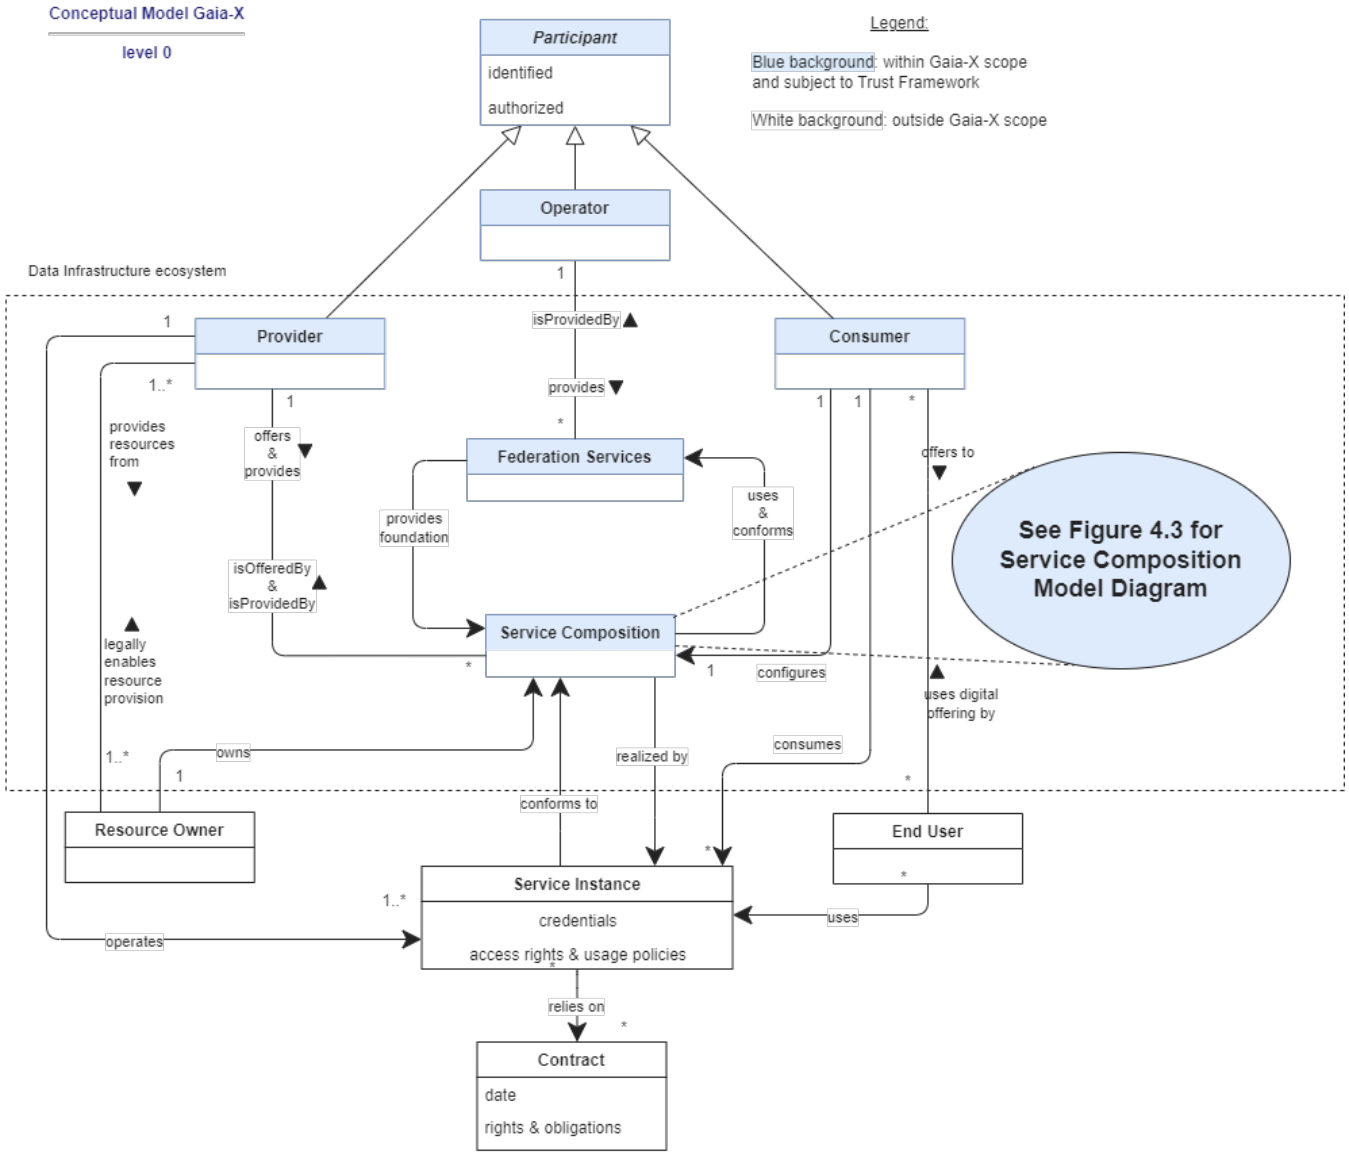

Figure 4.2 - Gaia-X Conceptual Architecture Model with Service Composition focus

Figure 4.3 depicts the need for a service discovery service that can search, match and select from the Gaia-X service catalog. The search for matching services may concern other service catalogs and offers depending on the will of the users to combine service offerings from different sources with verified interoperability and compatibility characteristics available in the catalog itself. Note that the discovery can be basic and limited to consumers selecting services by themselves or more elaborate via intelligent service discovery based on demands expressed using natural language or intent based paradigms. Such service discovery services as well as composition and orchestration engines are out of the scope of Gaia-X but found in the Gaia-X federated services catalog offerings proposed by providers, digital clearinghouses and third parties.

End users or tenants express theis initial service request and set both their requirements and constraints in their demand.

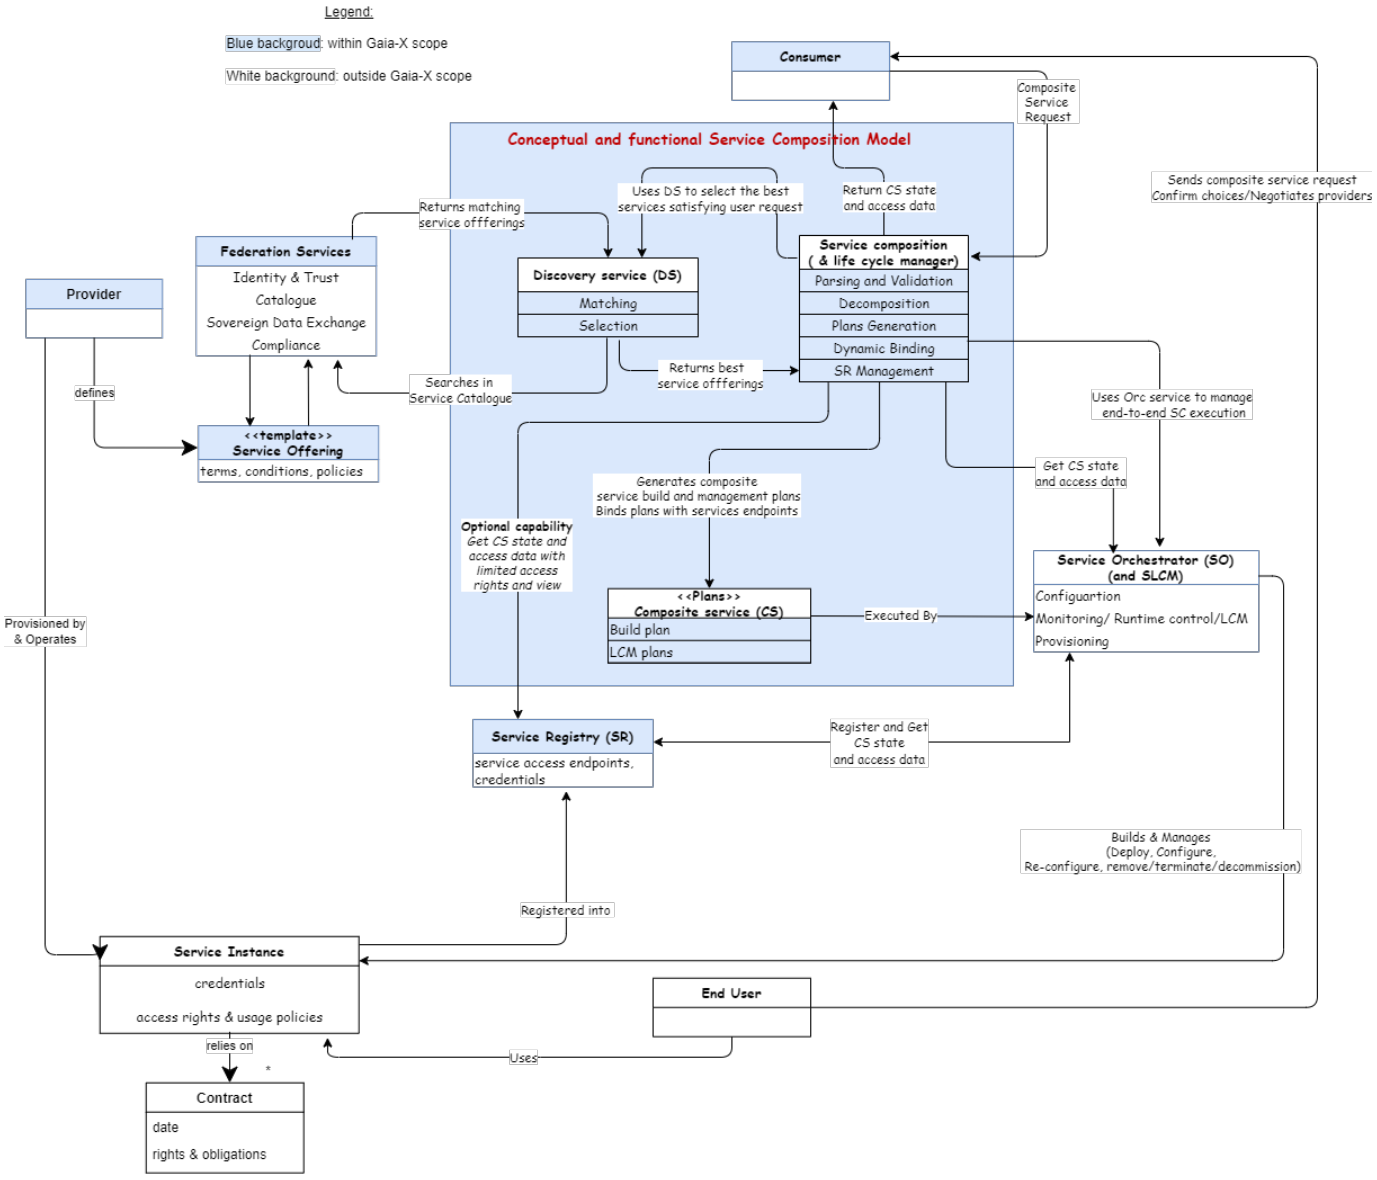

Figure 4.3 - Generic and abstract conceptual service composition model

The service composition module first “analyzes, parses and decomposes” the service request. It then prepares a set of service discovery queries to find and select the candidate services among the service offerings via the Federation Services.

Once the user has negotiated and approved the selected services and providers, candidate services are returned to service composition in order to build service deployment and instantiation workflows and plans. A contract is also established accordingly. The service composition module realizes the functional version of the solution and initiates service binding with the providers’ provisioned resources.

Starting from this abstract implementation plan for the Gaia-X Services Composition process, real implementations by suppliers will be developed. The actual deployment, binding and connectivity occur via an orchestrator, responsible for the coordination and organization of service instances, as well as the initial deployment, configuration and activation of such created service instances.

The orchestrator interacts with deployment and configuration tools ensuring the actual configuration, control and dynamic adaptation of service instances. A service life cycle manager is also associated with the orchestrator to manage services at run time. The orchestrator and service life cycle manager are typically external and act as the interface between service composition plans and providers. Optionally, the consumer acting as a broker can realize service composition. In this case, the consumer actively drives the orchestrators and partakes in the orchestration and service life cycle management according to compatible service deployment and management tools specified by the providers.

These tools and management systems can be proprietary or readily available in Catalogues. Popular configuration and deployment tools and systems such as Terraform, Ansible, Heat and Salt Stack are examples used for deployment and configuration purposes in cloud infrastructures. If service instances and resources are hosted in

containers rather than in virtual machines, an intermediate container orchestration and management system such as Kubernetes is then used in an intermediate step, which is out of the scope of this document.

2.2.4 4.4 Identity and Access Management

The identities of participants in the Gaia-X Ecosystem rely on signed attributes, which can be requested and exchanged to gain individual trust from other participants.

Each Participant provides a unique identifier that is associated with the attributes and the Participant can cryptographically prove that this Identifier is under his control.

The attributes are evaluated by the Participants to implement, enforce, and monitor permissions, prohibitions and duties when negotiating a contract for services or resources.

In the context of Gaia-X, the attributes are encapsulated in the **Party Credential** associated with its controlled identifier. Examples of **Party Credential** specializations are:

- Natural Person Party Credential
- Legal Person Party Credential
- Service Party Credential
- Membership Party Credential
- Custom Party Credential (domain specific party credential)

Another important type of credential is represented by the **TrustAnchorCredential** that provides a semantic description, the scope, the issuer(trust anchor), and any additional needed information by the **PartyCredential**. The **TrustAnchorCredential** mainly defines:

- The **Scope** where the **PartyCredential**s are considered valid.
- The **TrustedIssuer** entitled of issuance of the **PartyCredential** in the above **Scope**
- The **Vocabularies** (expressed in SHACL) that semantically defines the **PartyCredential** extended properties and the embedded attributes
- The **Trusted List** of issued **PartyCredential** in the scope by the trusted issuer.

Like for the **PartyCredential**, there are also several specializations of the **TrustAnchorCredential**, here are some examples:

- **Organization Trust Anchor –** a self issued credential that entitles an organization to define:
- **scope** - *Organization Credential Management* (OCM).
- **trusted issuer** - itself.
- **vocabularies** - defines the semantics of *Roles/Identity Attributes* and **Domain Specific Credentials** valid in the defined scope.
- **trusted list** - to assign and revoke *Roles/Identity Attributes* to its parties (users, natural persons, endpoint services, etc) issuing **PartyCredential** (referencing vocabularies and embedding the above *Roles/Identity Attributes*)
- **Ecosystem Trust Anchor** - a self issued credential that entitles an Ecosystem operator to define:
- **scope** - manage an *Ecosystem* (onboarding/offboarding/role assignement etc.).
- **trusted issuer** - itself.
- **vocabularies** - defines the semantic of *Roles/Identity Attributes* and **Domain Specific Credentials** valid in the defined scope.
- **trusted list** - to assign and revoke *Roles/Identity Attributes* to its members (other Participants) issuing **MembershipPartyCredential** (referencing vocabularies and embedding the above *Roles/Identity Attributes*)
- **Labelling Trust Anchor** - a Gaia-X issued credential that entitles a Participant to act as a CAB entitled to issue attestations related to services conformity labels.
- **scope** - issue Gaia-X conformity labels .
- **trusted issuer** - the Participant selected by Gaia-X to operate as CAB.
- **vocabularies** - defines the semantic of *Label*s that are valid in the defined scope.
- **trusted list** - to assign and revoke *Labels* issued to other Participants

> **Note**
>
> The reader must strongly distinguish between identifiers and identities. An identifier is a unique attribute that is part of the Participant's identity.

The key elements of the trustworthiness of an identity at the time of the negotiation are:

- The demonstration that the participant identifier is under the control of the credential holder using the verification method bound to the credential(**holder** property of **Party Credential**).
- The issued identity attributes conform with the KYB/KYC rules associated with the issuance of the identifier.
- The identifier issuers are trustworthy parties by checking the Gaia-X Registry and optionally other Verifiable Data Registries extending the Gaia-X rules.
- The check of **credentialStatus** property to be not equal to **revoked/suspended** (see Party Credential Status and W3C Verifiable Credentials Bitstring Status List v1.0 for additional information)

> **Note**
>
> did:web requires the control of the DNS or web service resolving the DID identifier. However, even with DNSSEC and an EV SSL cert, a did:web identifier offers a low level of confidence for legal participant authentication. For this reason, the following additional constraint would help to raise the level of confidence: Considering the key pair used by the legal participant to onboard to Gaia-X (to sign his own Participant Credential) as **the only one** valid to be used to sign/issue subsequent credentials, will ensure that only the holder of the private key associated with the key pair is trustworthy, making any kind of did:web hacking and/or DNS attack completely useless.

4.4.1 Dataspace / Federation onboarding and offboarding

The onboarding and offboarding of entities - participants, services, resources - within a Data Space or Federation is under the responsibility of the Data Space and Federation authorities which are autonomous from Gaia-X.

The above onboarding and offboarding processes are convenient optional ways for the participants themselves to request, demonstrate, collect, and revoke attributes/properties about themselves, their services, and their resources.

Those attributes/properties are expressed and exchanged via the **Membership Party Credential** specialization.

Gaia-X does not mandate nor enforce a particular timeline to gather Gaia-X Credentials.

Those could also be gathered at the time of contract negotiation, on the fly, depending on the policies being negotiated.

4.4.1.1 LAYERED IDENTITY AND ACCESS MANAGEMENT

The identity and access management relies on a two-tiered approach. In practice, this means that participants use a few selected identity systems for mutual identification and trust establishment, SSI being the recommended option for interoperability. After trust has been established, underlying existing technologies already in use by participants can be federated and reused, for example Open ID Connect (using OpenID4VP and SIOPv2) or domain-specific X.509-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, like multi-factor authentication. Server-to-server communication plays a crucial role in Gaia-X.

4.4.1.2 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 End User (PartyCredential holder) 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 building blocks provided by the mandatory attributes defined by Gaia-X must ensure that there is no lock-in to any implementation of identity networks and identity systems.

4.4.1.3 CHAIN OF TRUST AND IDENTITY

In the participant layer, the Gaia-X specifications indicate how to resolve and verify the participant identity of the contracting parties. The Consumer verifies the Provider's identity, the Provider verifies the Consumer's identity. Successful mutual participant verification results in a verified participant Token representing the trust between Provider and Consumer. The token payload contains one or more Gaia-X Credentials.

2.2.5 4.5 Data Products and Data Exchange Services

Gaia-X aims at promoting open innovation through sovereign data sharing based on trust between all involved actors. Trusted data sharing requires specific mechanisms because (a) it is difficult (or practically impossible) to technically un-share data when they have been shared and (b) it must comply with specific regulations (e.g. GDPR and the various European acts on data). The Gaia-X Data Product concept and Operational model provide such appropriate mechanisms.

4.5.1 Data Product Conceptual Model

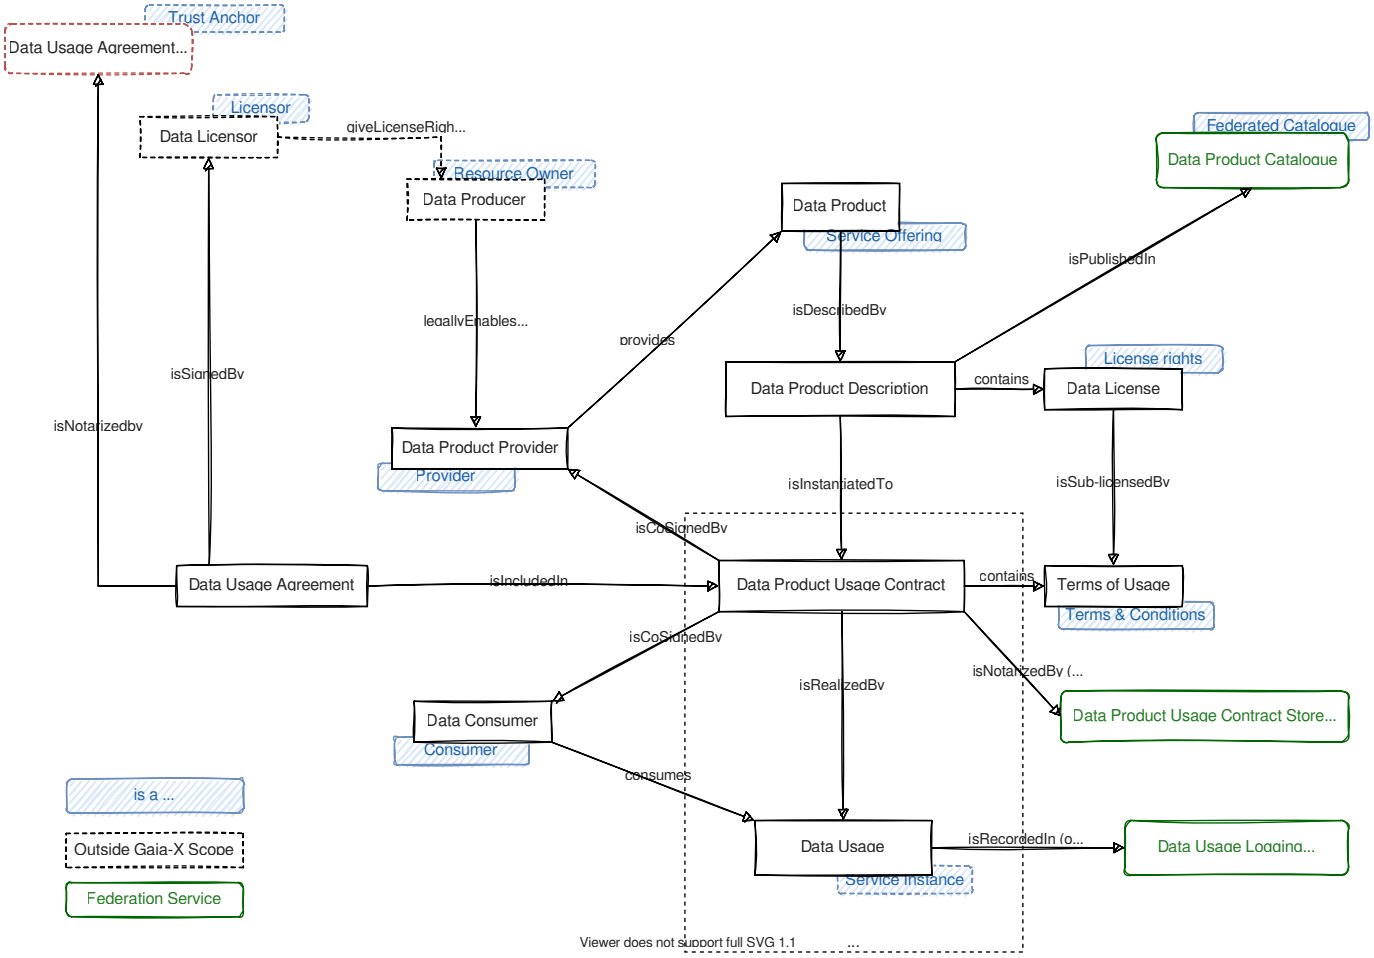

Figure 4.4 - Data Product conceptual model

Data are furnished by `Data Producers` (who are either data owners or data controllers in the GDPR sense, or data holders in the Data Act sense) to `Data Providers` who compose these data into a `Data Product` to be used by `Data Consumers`.

If a specific data license is attached to the data, then a `Data Usage Agreement` (DUA), including usage terms and conditions associated with these data, shall be signed by the `Data Licensor` (data subject or data controller as per GDPR, user as per Data Act, data owner as per copyrights laws, …) giving (a) to the Data Producer the right to distribute the data and (b) to the Data Consumer the legal authorization to use these data in accordance with the specified constraints. The *Data Usage Agreement* concept is a general concept which addresses every kind of licensed data and hence encompasses also the concepts of *Consent* from GDPR and of *Permission* from the EU Data Act (more detail in Annex C). In case of data liable to GDPR or Data Act, the Data Usage Agreement must contain all information required by the regulation (in particular the purpose of usage). Data Usage Agreements are notarized by a `Data Usage Agreement Trust Anchor` and might be revoked at any time by the data licensor.

Data Products are described by a `Data Product Description`, which must be a valid Gaia-X Credential. This description is stored in a (searchable) `Federated Data Catalogue`. Each Data Product Description contains a `Data License` defining the usage policies for all data in this Data Product – it also contains other information related to billing, technical means, service level, etc. Hence a Data Product Description constitutes a data usage contract template.

Before using a Data Product, the Data Consumer negotiates and co-signs a `Data Product Usage Contract` (DPUC) with the Data Provider. This Data Product Usage Contract is based on the Data Product Description but may differ from the original one: the Data License of the Data Product Description is sub-licensed, possibly after modification during the negotiations, by enforceable `Terms of Usage` contained in the Data Product Usage Contract. For each licensed data included in

the Data Product, the Data Product Usage Contract must include an explicit Data Usage Agreement signed by the corresponding Data Licensor (if the Data Product contains data from serveral Data Licensors, then the Data Product Usage Contract must include a Data Usage Agreement from each Data Licensor).

The Data Product Usage Contract is a Ricardian contract: a contract at law that is both human-readable and machine-readable, cryptographically signed and rendered tamper-proof, and electronically linked to the subject of the contract, i.e., the data. The parties can (optionally) request this contract to be notarized in a federated `Data Product Usage Contract Store`.

After such a contract has been agreed upon and has been signed by both parties, the Data Consumer can start accessing and using the data, realizing the Data Product Usage Contract. The Data Product Provider must check the Data Usage Agreement validity each time the data is provided to a Data Consumer, especially for recurrent data usage. The Data Consumer must also check the Data Usage Agreement validity each time it uses the data (to make sure that the agreement has not been revoked).

Such `Data Usage` with the associated Data Product Usage Contract corresponds to a *Data Transaction* as defined by the Data Spaces Support Centre (cf. DSSC Glossary).

The contract negotiation can lead to both parties agreeing on a `Data Usage Logging Service` (these logs might also include information needed for billing, inc. service level details, even if billing is outside Gaia-X perimeter).

Note: mapping between the Gaia-X Data Product concepts and the terminologies used in the various European regulations is given in Annex D.

4.5.2 Cascading agreement and right to oblivion

The above conceptual model can be applied recursively: the Data Consumer can integrate the Data into a new Data Product that can be used by other Data Consumers, who can in turn create new Data Products. It provides convenient mechanisms for data licensors or data controllers to control who is using their data and to revoke usage agreements.

Indeed, the above model blocks subsequent data transmissions between the Data Product Provider and the Data Consumer when the Data Usage Agreement is revoked and hence it provides a more robust mechanism than the usual cascading mechanism where the chain of revocations will be broken if one of the participants in the usage chain is deficient.

In order to better support the right to oblivion, it is recommended that the Data Space defines a general policy mandating each participant to check the Data Usage Agreement validity before reusing the data. How a participant implements this policy depends on its own internal data management procedures and is outside the scope of Gaia-X.

4.5.3 Data Intermediaries

Data intermediary services are explicitly mentioned as an important concept in European Commission acts around data. Their definition is quite broad: “Data intermediation services may support users or third parties in establishing a commercial relation for any lawful purpose on the basis of data of products in the scope of this Regulation e.g., by acting on behalf of a user” (cf. Regulation (EU) 2022/868).

Compared to the general Gaia-X Conceptual Model, a Data Intermediary is an actor who combines several roles: Data Product Provider acting as a proxy between the Data Producer and the Data Consumer, Federator operating a Data Product Catalogue (often named data marketplace), Data Usage Agreement Trust Anchor acting as a trusted proxy between the Licensor and the Data Consumer, etc. A generic operational model is detailed in the Operating Models section later in the document.

Several Gaia-X Lighthouse projects are implementing this concept of Data Intermediary, with different perimeters and operating models adapted to their ecosystem specificities. Accordingly, it was decided to stick here to basic data exchange concepts and to let Lighthouse projects define higher level concepts like Data Intermediary according to their needs.

2.2.6 4.6 Gaia-X Trust Anchors

Gaia-X Trust Anchors are bodies, parties, i.e., Conformity Assessment Bodies or technical means accredited by the Gaia-X Association to be trustworthy anchors in the cryptographic chain of keypairs used to digitally sign statements or claims about an object.

For each accredited Trust Anchor, a specific scope of attestation is defined.

The list of valid Trust Anchors and rules is available via the Gaia-X Registry.

There are cases where the Trust Anchor is relative to another property in a claim.

Example: In the case of a `gx:DataResource`, the `consent` property must be signed by the `gx:Participant` identified by the `dataController` property, itself signed by the `gx:Participant` identified by the `producedBy` property, itself signed by the credential’s `issuer` with a keypair from a chain of certificates with at its root a Gaia-X accredited state Trust Service Provider.

Ecosystem Trust Anchors follow the same principle, but are defined by the ecosystem-specific governance.

Kroki

In this example, the `dataController`'s participant, `producedBy`'s participant and the Trust Service Provider are Trust Anchors because they are mandatory fixed anchors in the credential chain of signatures. In this same example, all the claims could also be signed by the `issuer` if the `dataController`'s participant, `producedBy`'s participant and `issuer` are the same participant.

4.6.1 Gaia-X Trusted Data sources and Gaia-X Notaries

When an accredited Gaia-X Trust Anchor is not capable of issuing cryptographic material nor signing claims directly, then the Gaia-X Association accredits one or more Notaries, which will perform a validation based on objective evidence from a Gaia-X Trusted Data source and will issue an attestation about the previously made attestation from the Trust Anchors.

The Notaries are not performing audit nor verification of the object of conformity.

The Notaries are converting "not machine readable" proofs into "machine-readable" proofs.

ISO/IEC 17000:2020	Not machine readable	Machine readable
first-party conformity assessment activity	signing a paper, a PDF, a doc, …	signing a verifiable credential, a PDF/A-3a, …
second-party conformity assessment activity		
third-party conformity assessment activity		

Example: The European Commission provides several APIs, including one to check the validity of EORI number. Unfortunately, those APIs are not returning Verifiable Credentials. Hence Gaia-X accredited the GXDCH to act as a Notary for EORI verification using the European Commission API as the Gaia-X Trusted Data source for EORI validation.

Example: For a CAB or a set of CABs performing third-party conformity assessment activities but not capable of digitally signing the attestations, the Gaia-X Association could accredit one or more Notaries. The role of those Notaries is to digitalize an assessment previously made. The Notaries are not performing the assessment. The **Verifiable Credential** signed by such a Notary will be a second-party conformity assessment activity and have as `credentialSubject` a statement of a third-party conformity assessment activity. The overall trust level of such a **Verifiable Credential** will be lower than if the CAB was able to digitally sign the **Verifiable Credential** directly without involving a Notary.

4.6.1.1 EVIDENCE

It is expected that the credentials issued by the Notaries contain the evidence of the validation process.

The Gaia-X Trust Framework provides the definitions and means to manage Gaia-X Conformity, Gaia-X Labels and Ecosystem Trust. The Architecture Document describes the technical means to digitally notarize claims by a Gaia-X participant based on the W3C Verifiable Credentials Data Model Recommendation.

1. Singhal, A., Winograd, T., & Scarfone, K. A. (2007). Guide to secure web services: Guide to Secure Web Services - Recommendations of the National Institute of Standards and Technology. Gaithersburg, MD. NIST. https://csrc.nist.gov/publications/detail/sp/800-95/final https://doi.org/10.6028/NIST.SP.800-95 ↩
2. Oldehoeft, A. E. (1992). Foundations of a security policy for use of the National Research and Educational Network. Gaithersburg, MD. NIST. https://doi.org/10.6028/NIST.IR.4734 ↩

3. Operating & Services

3.1 5. Operating Models

This section describes the link between the documented policy rules to which participants of an Ecosystem agree (for Gaia-X defined in the Policy Rules Compliance Document, Domains or Ecosystems may extend with own policy rules) and the implementation through the software components and processes defined in the technical architecture, described in this document, and associated specifications.

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.

3.1.1 5.1 Context

Gaia-X provides a mechanism to translate Policy Rules as defined in regulatory, standards or commercial documents into criterial which can be digitally validated by a set of “Trust Anchors” which each ecosystem can define. This chapter describes the components an processed that operationalize this process.

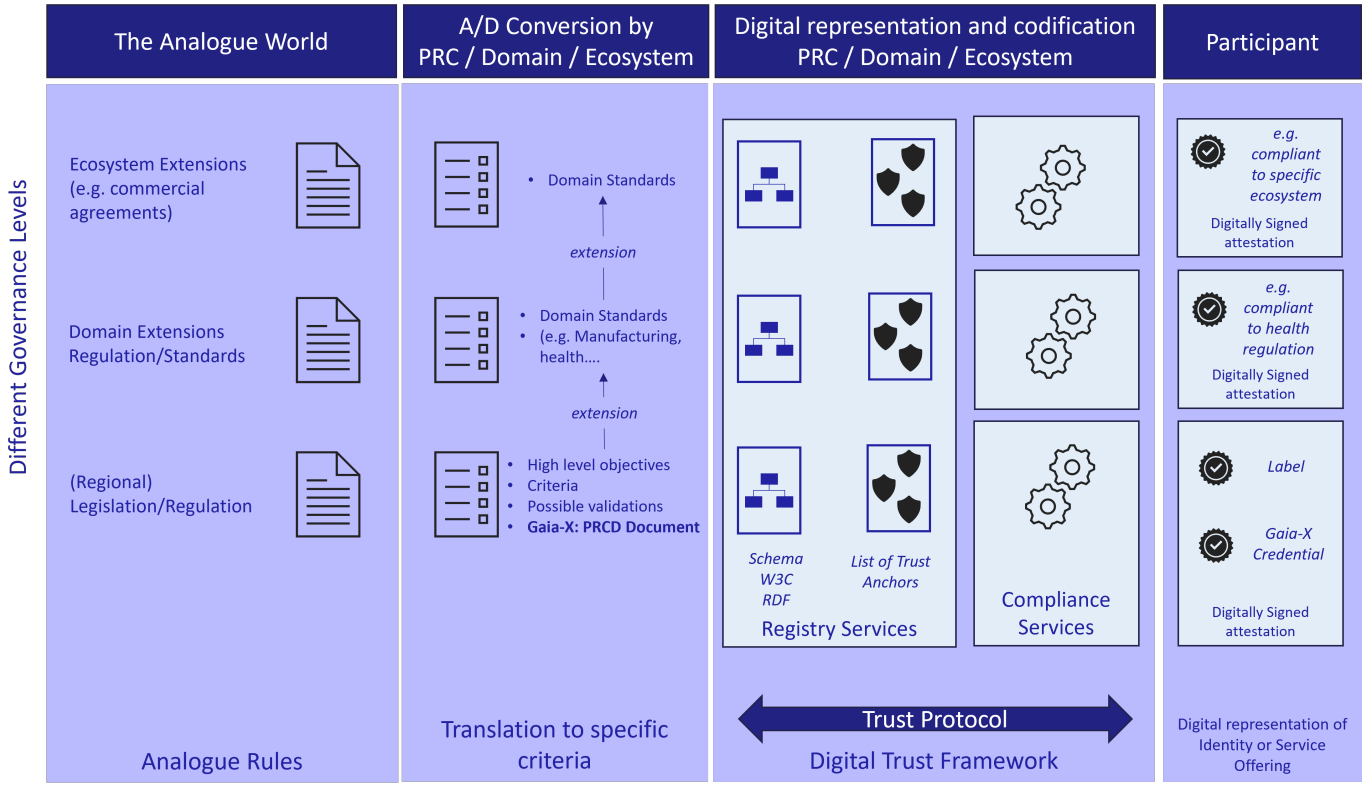

5.1.1 Prerequisites

To achieve the above goals, it is recommended for the readers to be familiar with the two documents below:

- ISO/IEC 17000:2020 Conformity assessment definitions prepared by the ISO Committee on Conformity Assessment CASCO
- Verifiable Credentials Data Model definitions from the World Wide Web Consortium (W3C)

> **Note**
>
> Versions 22.11 and below of Gaia-X documents are referring to Verifiable Credentials Data Model 1.1. It is expected to adopt Verifiable Credentials Data Model 2.0 in later releases.

The two mentioned documents address different types of readers, and the following sections will establish links between the various terms used in those documents.

5.1.1.1 CASCO WORKFLOWS

There is an overarching conformity system including one or more conformity schemes.

Kroki

5.1.2 Scheme model

Kroki

5.1.2.1 ROLES AND MAPPING BETWEEN CASCO AND VERIFIABLE CREDENTIAL DATA MODEL

The ISO/IEC 17000:2020 Conformity assessment and Verifiable Credentials Data Model documents mentioned above have different but overlapping scopes. To ease the analysis, two roles must be distinguished:

1. the body or party making a claim or attestation about an object.
2. the body or party cryptographically signing a claim or attestation.

While it is commonly expected that those two roles would be endorsed by the same identifiable legal or natural participant, it's not always the case, especially when the **CAB** is not capable of performing digital signatures or when the **credentialSubject** contains several `<subject> <predicate> <object>` triples as defined in Resource Description Framework (RDF).

> **Note**
>
> `<subject> <predicate> <object>` triples in RDF are also called `<subject>-<property>-<value>` claims in the Verifiable Credential Data Model.

Scenario	ISO/IEC 17000:2020 term	Verifiable Credentials Data Model term
1. the object being described.	object	credentialSubject (RDF object)
2. the body or party making a claim or attestation about an object.	CAB	credentialSubject (RDF subject)
3. the body or party cryptographically signing a claim or attestation.	CAB	issuer

5.1.2.2 GAIA-X CREDENTIALS AND ATTESTATIONS

As seen in the previous section, an object can be described by bodies or parties having different relations with the object itself.

Gaia-X Self-Descriptions is a legacy term still used in different Gaia-X documents and refers to a self-signed set of claims.
This is equivalent to a Verifiable Credential where the holder is the issuer. In general, the term "Gaia-X Self Description" is replaced by "Gaia-X Credential"

This simple approach is not enough and below there is a mapping between a newer vocabulary used in the Gaia-X documents and ISO/IEC 17000:2020 definitions prepared by the ISO Committee on Conformity Assessment CASCO.

ISO/IEC 17000:2020	Type of Attestion	Example
first-party conformity assessment activity	declaration	a person self-declaring itself competent
second-party conformity assessment activity		assessment of a person's knowledge and skills conducted by a trainer/instructor
third-party conformity assessment activity	certification	assessment of a person's knowledge and skills conducted with a national exam

To be noted that all the terms above can be generically referred as attestation and all **attestations** are issued by **conformity assessment body (CAB)**, including declarations.

5.1.2.3 GAIA-X COMPLIANCE SCHEMES

The Gaia-X Conformity and Gaia-X Labels are implementations of the Gaia-X governance.

To operationalise them, a software implementation must be executed and turned into up-and-running services, providing traceable evidence of correct executions.

While the governance of the Gaia-X Compliance rules and process is and will stay under the control of the Gaia-X Association, the Gaia-X Compliance Service will go through several mutually non-exclusive deployment scenarios[1].

During the pilot phase, the current, non-final deployment is described below.

Deliverables	Notes
Gaia-X Compliance Service	This service validates the shape, content and signature of Gaia-X Credentials and issues back a Gaia-X Credential attesting of the result. Required fields and consistency rules are defined in this document. Format of the input and output are defined in the Identity, Credential and Access management document.
Gaia-X Registry	This service provides a list of valid shapes and valid and revoked public keys. The Gaia-X Registry will also be used as the seeding list for the network of catalogues.

5.1.3 Extension by the ecosystems

An ecosystem can extend the definition of the Gaia-X Trust Anchors, Gaia-X Trusted Data Sources and Gaia-X Notaries as follows:

- the ecosystem governance can add more requirements on the eligibility of the Gaia-X Trust Anchors, hence selecting a subset of the Gaia-X Trust Anchors for the ecosystem domain.
- an ecosystem governance can add additional rules on top of the ones in this document by:
- adding more criteria for a credential type. Example: adding requirements for a participant to join the federation, enforcing a deeper level of transparency for a Service Offering
- selecting new domain-specific Trust Anchors for the new criteria.

! warning!: For already defined criteria, it will break Gaia-X Compliance to extend the list of Gaia-X Trust Anchors eligible to sign the criteria.

Ecosystem rules

$$\mathbb{R}_{GaiaX} \subseteq \mathbb{R}_{domain}$$

Rule property	**A domain refining Gaia-X Conformity** $\forall r \in \mathbb{R}_{GaiaX}$ and $r_{domain} \in \mathbb{R}_{domain}$	**A domain extending Gaia-X Conformity** $\forall r \in \mathbb{R}_{domain} \setminus \mathbb{R}_{GaiaX}$
Attribute name: r_{name}	$r_{name_{CamelCase}} \equiv r_{name_{snake\_case}}$	$r_{name} \notin \mathbb{R}_{GaiaX}$
Cardinality: $\|r\|$	$\|r_{domain}\| \geq \|r\|$	*no restriction*
Value formats: $r \rightarrow \texttt{VF}(r)$	$\texttt{VF}(r_{domain}) \subseteq \texttt{VF}(r)$	*no restriction*
Trust Anchors: $r \rightarrow \texttt{TA}(r)$	$\texttt{TA}(r_{domain}) \subseteq \texttt{TA}(r)$	*no restriction*
Trusted Data Sources: $r \rightarrow \texttt{TDS}(r)$	$\texttt{TDS}(r_{domain}) \subseteq \texttt{TDS}(r)$	*no restriction*

5.1.3.1 COMPLIANCE REMEDIATION

Gaia-X Compliance credentials may become invalid over time. There are three states declaring such a credential as invalid:

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

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

The revocation of Gaia-X credentials can be done in various ways:

- **Revocation or Deprecation by authorship**: The author of a Gaia-X credential revokes or deprecates the credential explicitly.
- **Revocation by automation**: The Gaia-X Compliance Service found at least one credential claim not validating the Gaia-X conformity rules.
- **Suspension and Revocation by manual decision**: After an audit by a compliant Gaia-X Participant, if at least one credential claim is found to be incorrect, the suspension of the Gaia-X credential is automatic. The revocation is submitted for approval to the Gaia-X Association with the opportunity for the credential issuer 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.

5.1.4 Gaia-X Compliance Schema and Process

The following graph describes the generic compliance schema which is applied to Gaia-X conformance, labels and can be used for domain and ecosystem specific extensions:

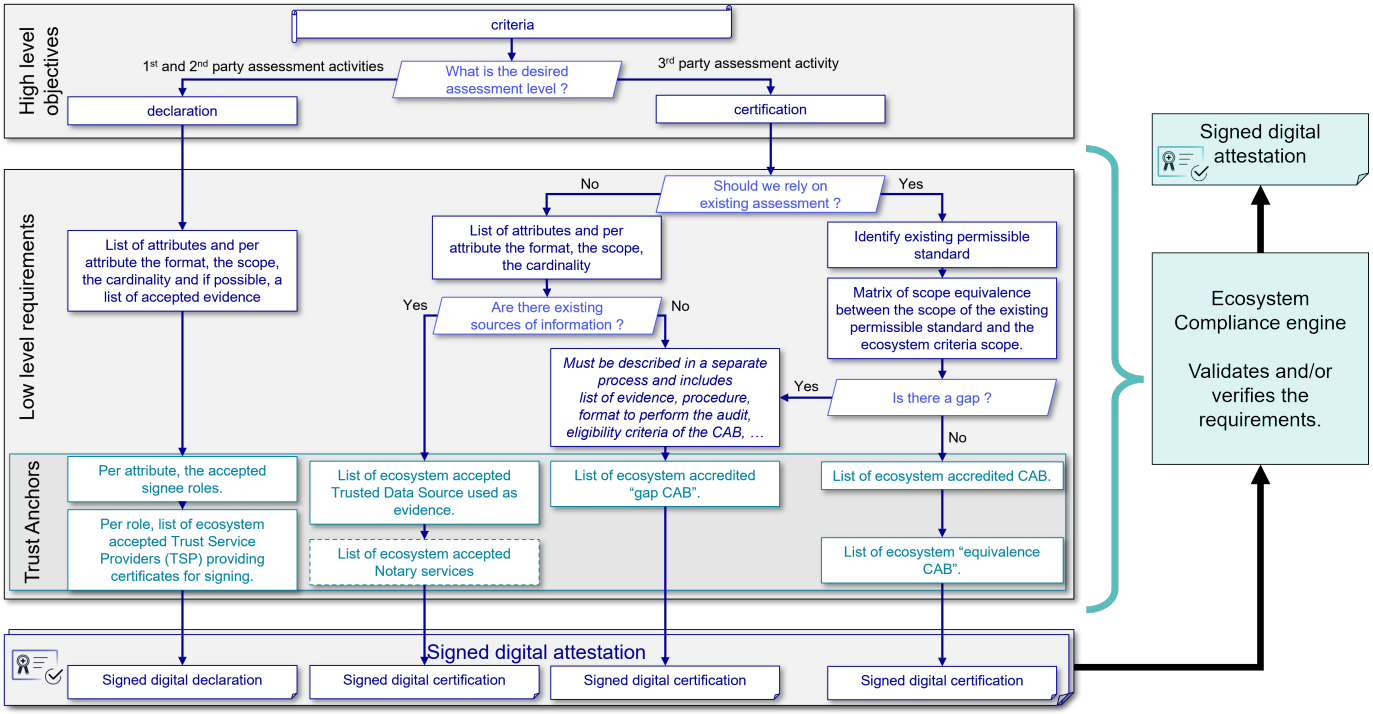

The different types of Conformity Assessment Bodies are specified in the PRCD document.

3.1.2 5.2 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[2], 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

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

3.1.3 5.3 Data Usage operating model

The generic basic operating model for data usage of un-licensed data is quite simple:

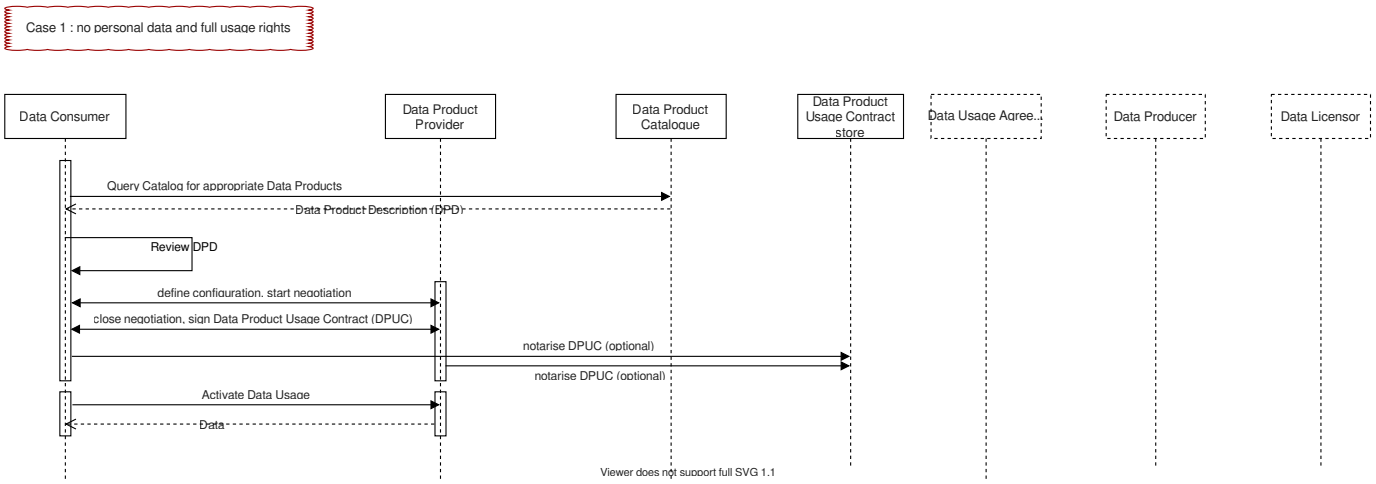

Figure 5.1 - Data usage operating model – case 1

1. The Data Consumer queries the Data Product Catalogue and reviews the Data Product Descriptions to select a Data Product that corresponds to its needs.
2. The Data Consumer configures the Data Product in terms of data scope and operational characteristics and starts negotiating with the Data Product Provider.
3. When both find an agreement, they sign a configured Data Product Description to create the Data Product Usage Contract (DPUC) and can optionally notarize it in a Federated Data Product Usage Contract Store.
4. Then the instantiated Data Product can be activated resulting in actual Data Usage.

The operating model is a bit more complex when the Data Product includes licensed data. For instance, if the Data Product includes some personal data, then the GDPR imposes that the data subjects explicitly give their consent for the data usage and that they can revoke this consent at any time. Accordingly, a Data Usage Agreement has to be signed by the data subject, who acts as a Data Licensor, before the first data usage and the Data Usage Agreement validity has to be checked before each data usage. It is recommended to establish and sign the Data Usage Agreement during the negotiation phase between the Data Product Provider and the Data Consumer because, without this signed Data Usage Agreement, the agreement between the parties (i.e., the Data Product Usage Contract) would be legally void. Note that the Data Usage Agreement may be signed by a guardian (for minor persons) or a third party through a specific or generic power of attorney or specific legislation (for instance in the case of a sick person in a hospital). The process is similar for a Data Product for which the Data Owner did not give an unconditional usage right to the Data Product Provider and wants to precisely know who will use her/his data and for which purpose.

If the Data Consumer has access to the Data Licensor, then it can directly request the Data Usage Agreement – this is usually the case when the Data Consumer is using the data to provide a service to the Data Licensor (for instance when a sports training app get historical monitoring data from a sports watch provider). This is the most convenient and simple case, and it provides better privacy (the Data Product Provider does not know what the data will be used for). Otherwise, the Data Usage Agreement has to be collected by the Data Product Provider through the Data Producer.

In the first sub-case, the operating model is:

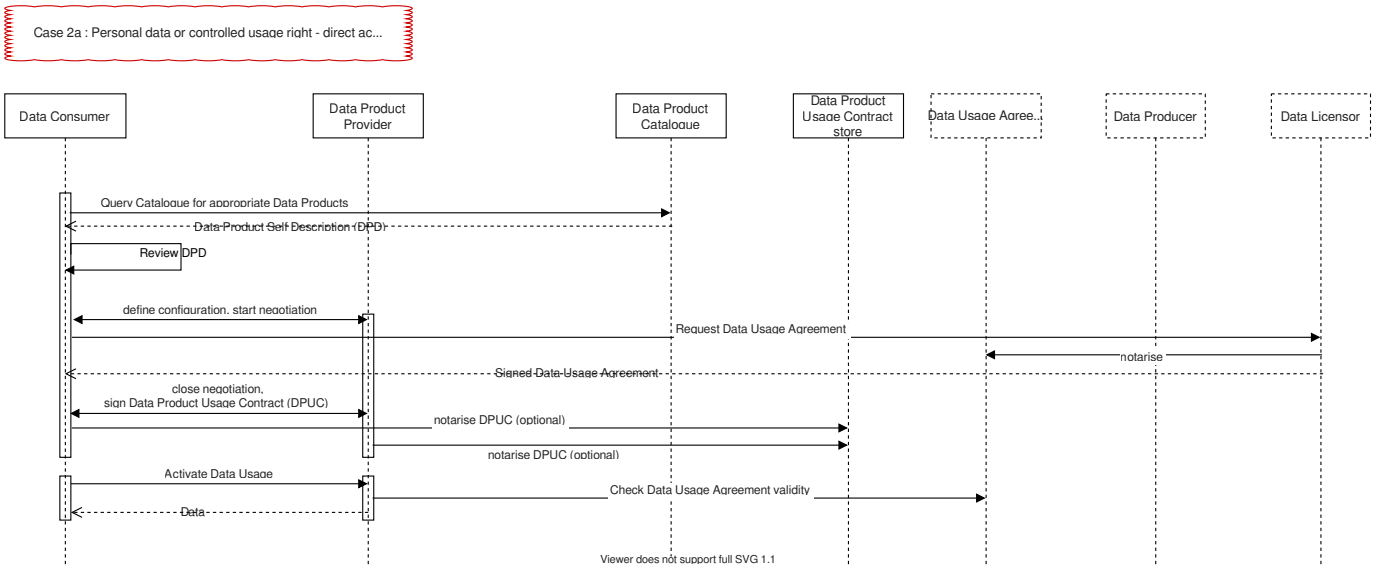

Figure 5.2 - Data usage operating model – case 2a

1. The Data Consumer queries the Data Product Catalogue and reviews the Data Product Descriptions to select a Data Product that corresponds to its needs.
2. The Data Consumer configures the Data Product in terms of data scope and of operational characteristics and starts the negotiation with the Data Product Provider.
3. The Data Consumer extracts the Data Usage Agreement template from the Data Product Description, fills it and adds its specific information (how the data will be used, the purpose, …), usually in a separate section that might not be communicated to the Data Product Provider. This Data Usage Agreement is sent to the Data Licensor who will sign it through a Data Usage Agreement Trust Anchor and send it back to the Data Consumer.
4. The Data Product Provider and the Data Consumer close the negotiation, and they sign the configured Data Product Description to create the Data Product Usage Contract (DPUC) which includes the appropriate part of the Data Usage Agreement. They can notarize this Data Product Usage Contract in a Federated DPUC Store.
5. Then the instantiated Data Product can be activated, and actual Data Usage can start. Before each Data usage, the Data Product Provider has to check the validity of the Data Usage Agreement through the Data Usage Agreement Trust Anchor.

The operating model for the second sub-case is the same except for step 3, where the Data Usage Agreement is requested by the Data Product Provider through the Data Producer. Note that the Data Consumer will have to provide the purpose of the Data usage, and this will be included in the Data Usage Agreement sent to the Data

Licensor – this is part of the service configuration phase. Note also that the Data Producer might to counter-sign the Data Usage Agreement to guarantee that the person who signed the Agreement is really the Data Licensor of the data.

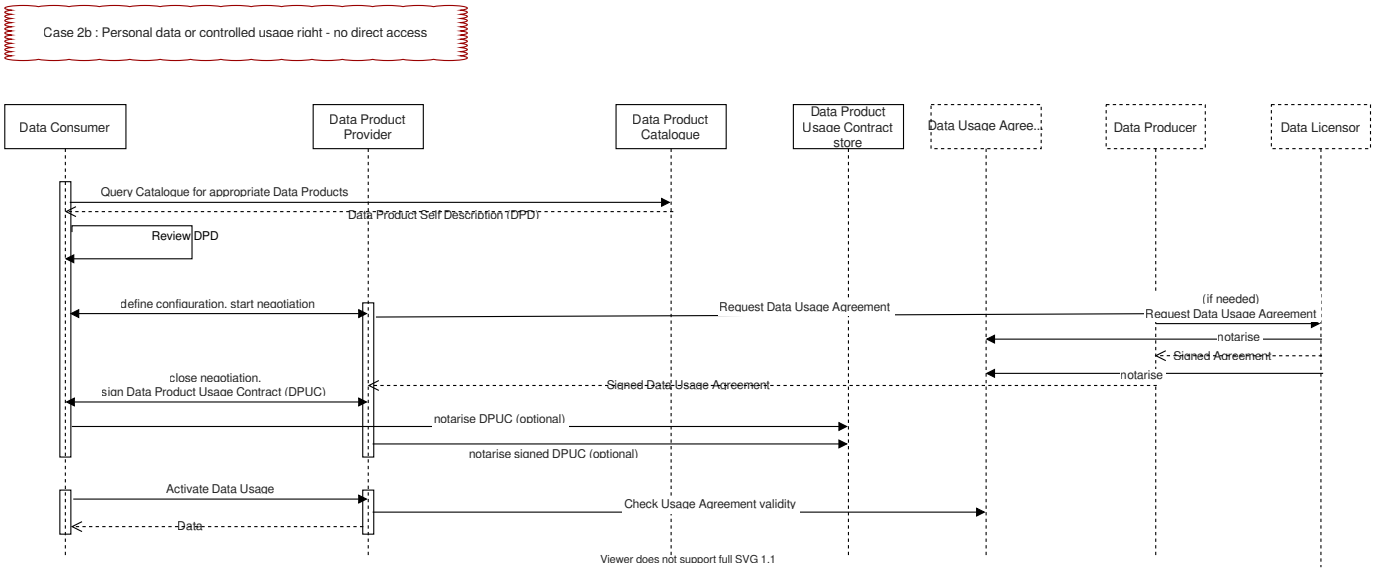

Figure 5.3 - Data usage operating model – case 2b

Note: The above operating model minimizes the functionalities of the Data Usage Agreement Trust Anchor. It is possible to imagine a slightly more complex Data Usage Agreement Trust Anchor that receives part of the Agreement from the various parties and communicates only the relevant parts to the appropriate actors. Hence in case 2b, the Data Product Provider will not know the purpose of the Data Consumer but only that the Data Licensor authorizes the data usage, while the Data Licensor will have access to the Data Consumer Purpose. We might also imagine a more complex Data Usage Agreement Trust Anchor able to compare predefined Usage Clauses and to advise the Data Licensor or even automatically grant agreement on behalf of the Data Licensor. For instance, the Data Licensor might specify that she/he allows transmission of some medical data to non-profit research laboratories which have appropriate certificates in terms of data security and data privacy – the Data Usage Agreement Trust Anchor would directly provide the general Data Usage Agreement and the Data Licensor would just receive a notification (and will always be able to revoke the agreement) - that could enable a more agile data economy. Data Intermediary services provided by some Gaia-X Lighthouse projects provide such elaborated functionality.

5.3.1 Data Intermediary generic operating model

Several Lighthouse projects implement the concept of Data Intermediary. Compared to the general Gaia-X Conceptual Model, a Data Intermediary is an actor who plays several roles: Data Product Provider acting as a proxy between the Data Producer and the Data Consumer, Provider operating a Data Product Catalogue (often named data marketplace), Data Usage Agreement Trust Anchor acting as a trusted proxy between the Licensor and the Data Consumer.

The generic operating model for the Data Intermediary role is as follows:

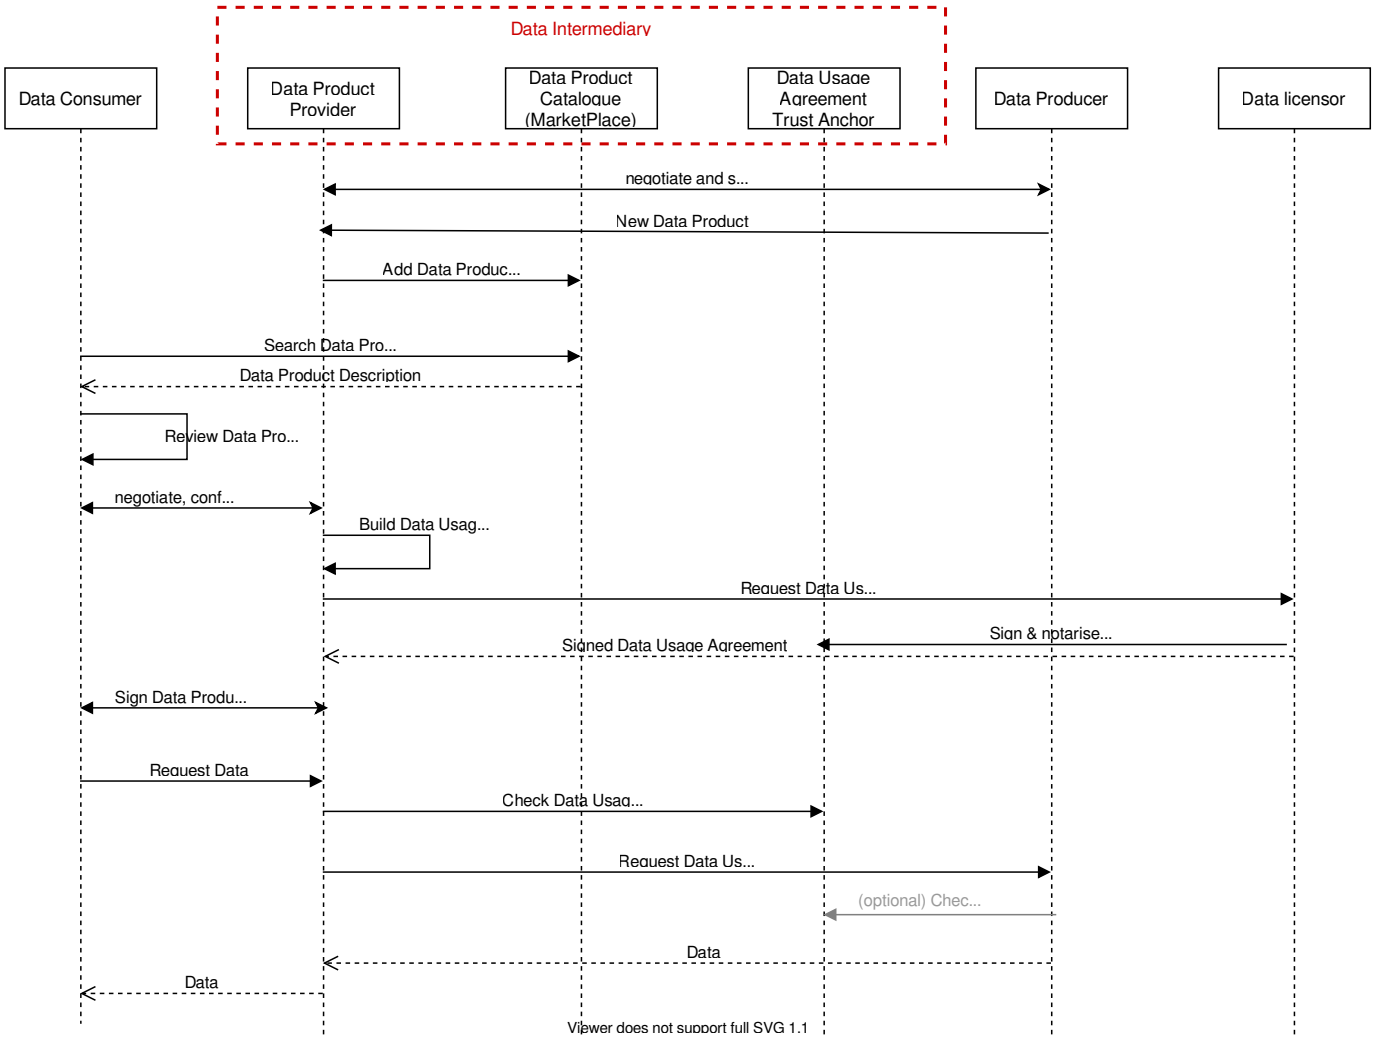

Figure 5.4 - Data Intermediary generic operating model

Many variants of this generic model can be defined. For instance, a general Data Usage Agreement may be proposed to the Data Licensor upfront when the Data Product is inserted in the Catalogue and the Licensor can trust the Data Intermediary to check the conditions and purpose before delivering data usage to a consumer.

Hence, the above model is given as a generic example. The ecosystems are free to define and implement the operating model adapted to their specificities.

3.1.4 5.4 Service Composition

To pursue the description of the service composition model, we focus on the steps required to implement service composition and describe the functions required to achieve end-to-end service composition across multiple service providers. Figure 5.5 depicts these steps and starts with a service request from the user (end user, tenant, information system manager) request.

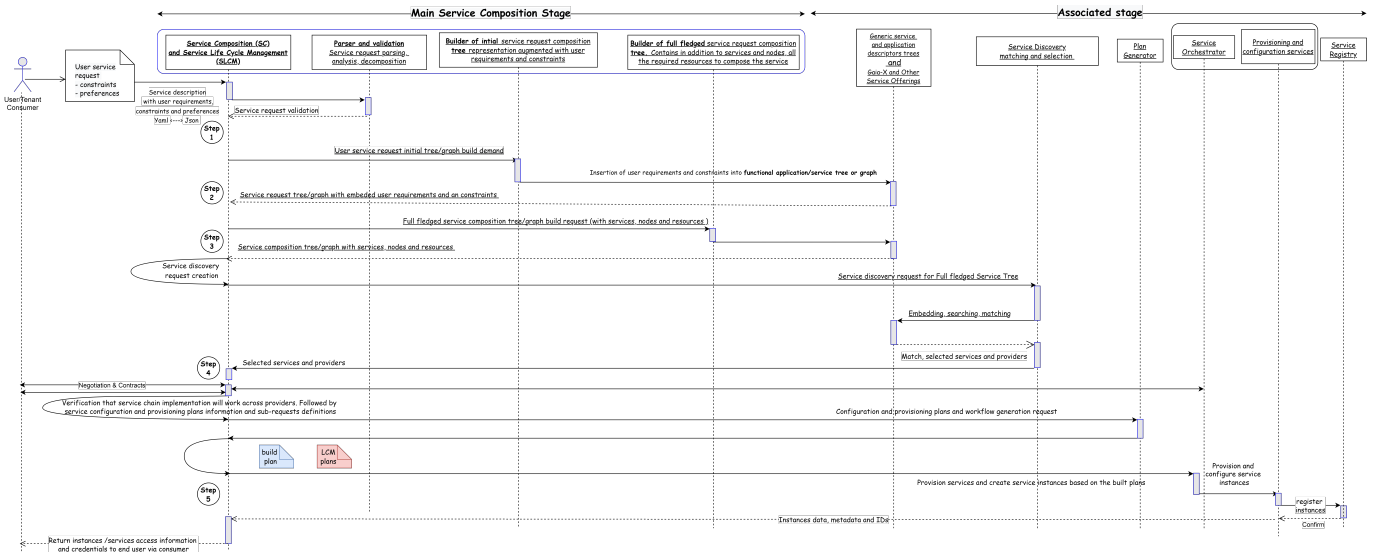

Figure 5.5. Gaia-X Compliant Service Composition Steps

The user first describes the desired service using a domain-specific description language, generally JSON and YAML-based. This request is then parsed for specific end-user requirements extraction and added (concatenated) to a generic and high-level catalogue that contains service descriptions as part of a Catalogue and holds only functional descriptions of services and applications relevant to the user request, to ensure interoperability, compatibility, and portability requirements. These services are not localized a priori; location is to be set at a later stage in the service composition process. The same holds for some additional service attribute values, in some parts of the service description. These values are empty fields or wild cards that will only take on specific values once the user service request has been parsed for preferences, constraints and specific limitations (such as provider type, restrictions on localization and geographical area, unconstrained naming and addressing spaces and more). Some attribute values are only set (or resolved) and known at later stages of the service composition (as late as steps 4 and 5).

The user-expressed constraints, in step 1, are added to the functional description in step 2 to produce the actual service request that will be used to derive in a third step an extended representation of the service request. This extended representation, usually in the form of a tree (or complex service graph) representation and description of the extended service, now contains required services and resources as well as generic hosting nodes (not yet fully specified) needed to fulfil the service request. All properties and requirements from the initial service request therefore propagate to these tree nodes and leaves. This tree contains the relationships and dependencies of all services involved in building a solution to the initial user request and has as mentioned inherited all the service properties, constraints, restrictions, and obligations stated by the user in step 1. The fourth step is service discovery which will search in multiple catalogues.

The output of the fourth step produces, with the help of orchestrators and service composition, workflows and deployment plan production engines (external or even taken from the Gaia-X service offerings) to deploy the service instances and ensure their interconnection. The result is the production of a composite service graph that fulfils the initial end user request and its concrete instantiation in multiple hosting infrastructures and providers. In Figure 5.5, the service life cycle manager handles the service management at run time by interacting with the involved providers. This manager makes recommendations about possible adaptations, substitutions and adjustments of services depending on operational conditions to maintain SLAs according to established contracts between all stakeholders and processes. These contracts are considered during stages 1 to 5 of the service composition process depicted in Figure 5.5. These steps, not shown at this stage, will receive special attention in a future dedicated document.

Annex E provides a practical “service composition model example” involving multiple cloud providers, services and infrastructures that can be accomplished using currently available cloud services.

---

1. https://gaia-x.eu/news/latest-news/gaia-x-compliance-service-deployment-scenario/ ↩
2. Example of the setup of a DAO https://blockchainhub.net/dao-decentralized-autonomous-organization/ ↩

3.2 6. Gaia-X Trust Framework components

This section focuses on the mandatory software components that are being operated only by the Gaia-X Digital Clearing Houses (GXDCHs) to enable the issuing of Gaia-X Compliance Credentials.

Note: The sections on Protocol/Interface Spec need to be completed/corrected!

3.2.1 6.1 Gaia-X Registry

The Gaia-X Registry (see https://registry.gaia-x.eu) is a public distributed, non-repudiable, immutable, permissionless database with a decentralised infrastructure and the capacity to automate code execution.

ℹ The Ecosystems may want to have their own instance of a local Registry or equivalent. Technically, this component can be part of the ecosystem's 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's 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 shapes and schemas for the Gaia-X VCs.
- the URLs of Gaia-X Catalogue's credentials.
- the text of the Terms and Conditions for Gaia-X Conformity.
- …

The Gaia-X Registry 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[1].
5. A list of Participants' identities and credentials' 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.

The Gaia-X Registry is used by the Gaia-X Compliance Engine to perform the checks needed to assess Gaia-X Compliance and can be used by 3rd parties to get correct information (about the shapes, T&C, etc). The Gaia-X Registry will be used as the seeding list for the network of Catalogues.

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.

The Gaia-X Registry leverages a combination of DNS, DNSSEC, and IPFS to ensure the integrity, availability, and non-repudiability of stored data. This approach not only secures the data but also makes it globally accessible and resistant to tampering. Data stored in IPFS includes essential governance documents, shapes, schemas, and other critical ecosystem information. Each piece of data or file added to IPFS is identified by a unique content identifier (CID), which is a cryptographic hash of the data's content.

6.1.1 DNS TXT Records and Naming Convention

To facilitate the discovery and accessibility of the stored data, the Gaia-X Registry utilizes DNS TXT records. These records are used to advertise the current URIs (in the case of the Gaia-X ecosystem, they are ipfs:// links which include CIDs) associated with the latest versions of the stored documents. The structure of these DNS TXT records follows a specific naming convention to ensure easy and systematic access:

DNS TXT records are formatted as follows: `[gxdch-version]._[ontology-version]._[type].[?subdomain.domain]`

Where: - `[gxdch-version]` Indicates the version of the Clearing House, formatted as (e.g., `vmajor.minor`, `vmajor`). - `[ontology-version]` Indicates the version or codename of the ontology (e.g., `2404`, `danube`). - `[type]` Specifies the type of content, such as `shapes`, `scheme`, or `trust_anchors`. - `[domain]` Represents the domain and optional subdomain where the records are hosted. For Gaia-X the domain is `gxdch.eu`. Other ecosystems can adopt the same design. - Here is a table illustrating the DNS TXT records format for Gaia-X:

Record Name	Record Type	TTL	Value
`v2._2404._shapes.gxdch.eu`	TXT	600	`ipfs://[CID]`
`v2._2404._scheme.gxdch.eu`	TXT	600	`ipfs://[CID]`
`v2._2404._trust_anchors.gxdch.eu`	TXT	600	`ipfs://[CID]`
`v2._2404.gxdch.eu`	TXT	600	`ipfs://[CID]`

Each TXT record contains a URI that points to the IPFS location (where `CID` is the Content Identifier for the respective document(s) or data stored on IPFS) or an external URL where the data can be accessed. This flexible system allows the registry to use `ipfs://` for decentralized storage or standard `https://` URLs for more traditional data hosting.

If you omit the `[type]`, you will obtain the ipfs URI pointing to the root folder containing all documents available for the given `[version]`.

6.1.2 Adopting the Gaia-X Model in Other Ecosystems

The Gaia-X Registry's approach to data management and distribution is adaptable. Other ecosystems can implement a similar strategy by setting up their DNS to point to the relevant data URIs. For example, an ecosystem could use a similar DNS TXT record strategy to point to different types of content:

Record Name	Record Type	TTL	Value
`v1._2404._scheme.your-ecosystem.eu`	TXT	600	`https://yourapi.com/v1/scheme`

By adopting this model, ecosystems can ensure that their data is easily accessible, verifiable, and secure. The use of versioned records also helps manage updates and maintain compatibility across different versions of the ecosystem's services. Participants and ecosystems within the Gaia-X ecosystem can query these DNS TXT records to retrieve the latest URIs and ensure that they access the most current and verifiable versions of the data, fostering a trustworthy and transparent ecosystem. This integration not only enhances the resilience and efficiency of data distribution within the Gaia-X Ecosystem but also aligns with the broader goals of decentralization and robustness in data handling.

6.1.3 Basic Protocol and Interface Specification for GXDCH Registry

Standards used:	W3C:RDF, JSON-LD, OWL, SHACL, SPARQL
Protocol:	Rest-API
API:	EBSI, OpenAPI via Swagger @ registry.gaia-x.eu
Trust Anchors:	ETSI TS 119 612

3.2.2 6.2 Gaia-X Compliance

The service takes as input the Verifiable Presentations provided by the participants, checks them against the SHACL Shapes available in the Registry and performs other consistency checks based on the Gaia-X Policy Rules.

The service returns a Verifiable Credential, the "Gaia-X Compliance Credential" with a Gaia-X signature, as a proof that the input provided has passed all the verifications.

6.2.1 Basic Protocol and Interface Specification for GXDCH Compliance

Standards used:	W3C:RDF, W3C:VC, W3C:DID [DID:WEB] JSON-LD, SHACL Shapes; OIDC4VP; OIDC4VCI
Protocol:	Rest-API
API:	OpenAPI via Swagger @ registry.gaia-x.eu

3.2.3 6.3 Gaia-X Notary - LRN (Legal Registration Number)

- Takes as input a LegalRegistrationNumber VC from the user
- Verifies that the VC contains at least one identification number as requested by the Gaia-X rules, and checks that the number is valid
- Returns a Gaia-X VC with the proof that the number has been verified

6.3.1 Basic Protocol and Interface Specification for GXDCH Notary

Standards used:	W3C RDF, JSON-LD, SHACL Shapes
Protocol:	Rest-API
API:	OpenAPI via Swagger @ registry.gaia-x.eu

3.2.4 6.4 Gaia-X Credential Event Service (CES)

For the Gaia-X Catalogues of the GXDCH to be notified about new, updated, revoked credentials, a common Publication / Subscription service (pubsub service) must be deployed via decentralized GXDCH instances. This service is called “Credential Event Service” (CES) Its Technical Specification is available in the Software Specification of the Gaia-X Lab

It is expected that the pubsub service will have different implementations over time from a distributed service during the pilot phase - Apache Kafka or similar - to a decentralized one - a Gaia-X consortium blockchain like for the Gaia-X Registry.

The deployed solutions have to accommodate for convenience from a user point of view and for security from an operator point of view.

- The service only accepts push/pull requests from participants which have a valid Gaia-X Participant Credential to ensure that Terms & Conditions of the service have been agreed. Abuse of the service can lead to revocation of the right to use the CES service.
- The service should provide synchronization within a defined time limit specified in the GXDCH Operations Handbook (tbp) to ensure consistency for users across GXDCH providers.
- Potential handling of downtimes of the CES are specfied in the GXDCH Operations Handbook (tbp). To allow users to mitigate potential lazy syncronization a timestamp for the latest synchchronization must be provided and pointers to potential alternative catalogues should be provided. The design principle follows the CAP Theorem which basicly states that the service should be available and partition tolerant but not necessarily consistent.
- The CES is limited to push/pull Gaia-X Compliance credentials only. To avoid publication of confidential data only the credential ID is going to be published. The subscribers of the pubsub services have to resolve the ID and request access to the holder for the credential. The service storing the credential - agent or wallet - is from a functional point of view a Policy Decision Point leaving full control of the credential access to the holder.
- Identification of catalogue providers consuming the data in some way.
- “Log events” on participants consuming data shall be provided to avoid and track abuse.
- The service is going to accept standard OIDC4VP API.

6.4.1 Basic Protocol and Interface Specification for GXDCH CES

Standards used:	
Protocol:	Rest-API
API:	OpenAPI via Swagger @ registry.gaia-x.eu, OIDC4VP

3.2.5 6.5 Graphical overview of the Trust Framework components

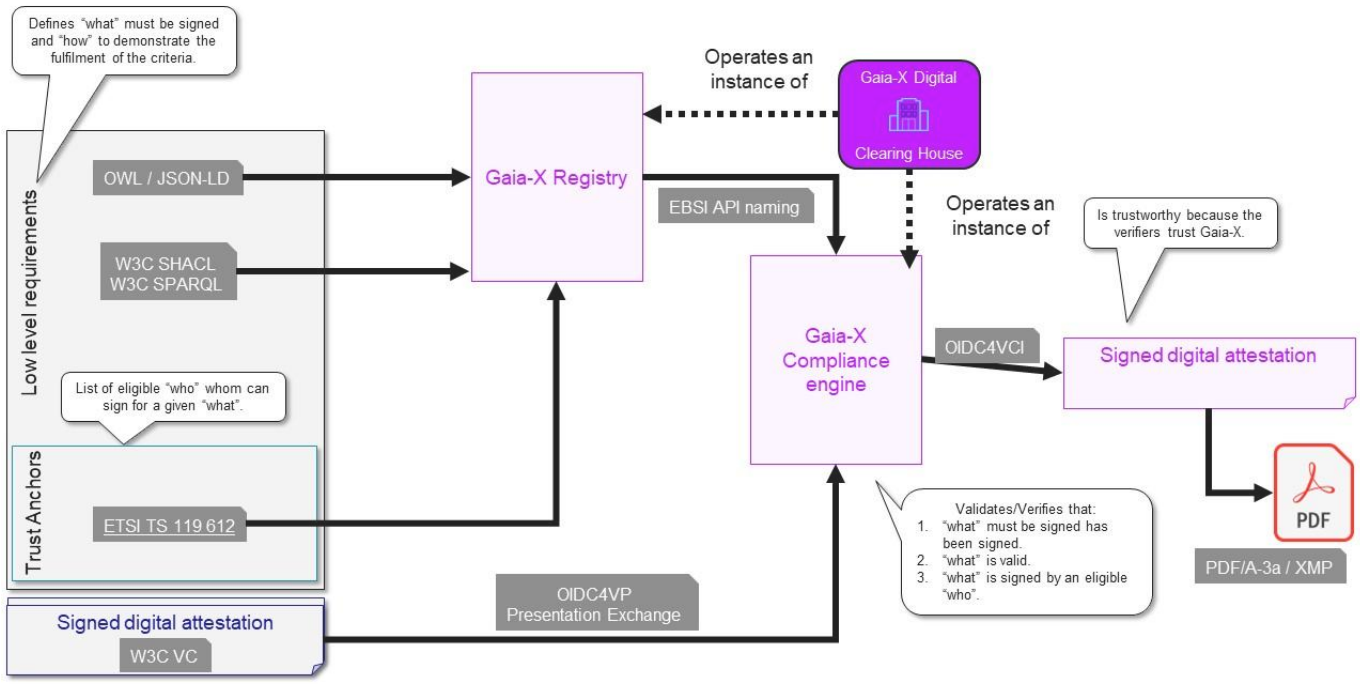

Figure 6.1 - Gaia-X Trust Framework Components

A complete software architecture is available and maintained on GitLab. It represents the current software architecture of the GXDCH components, as well as the previous versions and the upcoming ones. It's using the renowned C4 Model.

Note: see the Annex chapter for the full list of GXDCH components.

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

3.3 7. Enabling and Federation Services

Enabling Services facilitate the operation of ecosystems. There are multiple technologies, products and implementations of each of the enabling services available. This document defines the components, the interactions among them and specific requirements towards Gaia-X.

Enabling services operated under the specific rule of ecosystem governance are defined as “Federation Services”.

3.3.1 7.1 Wizard Service

UI to sign Verifiable Credentials in a Verifiable Presentation on the client side. Calling the Gaia-X Compliance Service is integrated with the Wizard, making it possible to obtain Gaia-X Compliant Credentials by directly using this tool. Depending on the implementation, a Wizard can support signing with a hardware module via PKCS#11 or other like Web eID.

3.3.2 7.2 Credential Manager

Application that allows the user to store their Credentials and present them to third parties when needed. This is also sometimes called an Agent, CredentialRegistry or IdentityHub.

3.3.3 7.3 Federated Catalogues

The goal of the Federated Catalogue is to:

- enable Consumers to find best-matching offerings and to monitor them for relevant changes in the offerings
- enable Producers to promote their offerings while keeping full control of the level of visibility and privacy of their offerings.
- enable Service Composition by including and publishing Service Descriptions, conformant to the Gaia-X Schema, that contain structured service attributes required to compose services.
- avoid a gravity effect with a lock-out and lock-in effect around a handful of catalogue instances.

7.3.1 Catalogue service

The implementation of the catalogue can vary from one catalogue owner to another and offers different user experiences.

The requirement to be listed in the Gaia-X Registry as a valid Gaia-X Catalogue is to keep the network of Gaia-X Catalogues up to date by publishing in the pubsub service the credential ID of the credential being created, updated or revoked.

A future testbed will be implemented to perform dynamic validation of the above behavior.

7.3.2 Retention

To have a basic data curation mechanism, the retention period in the pubsub service is one-year maximum.

It means that a holder, to have their credentials discoverable, should publish them at least once per year, even if the credentials have not changed.

7.3.3 Snapshot

To speed up the onboarding of new catalogue instances, it is expected that timestamped snapshots of the pubsub service data are made and stored publicly via IPFS. The URIs of the snapshots are available via the Gaia-X Registry.

7.3.4 Trust Indexes

While “Trust” is used as a building block in every Data Space and Federation related document, there are two main challenges:

- There is no unique definition of Trust; Trust is context-dependent
- The market has and will always move faster than the rules

While the Gaia-X Policy Rules provide a baseline translating European values into objectives and measurable criteria, Data Spaces and Federations need to extend this baseline for their own market, vertical domain and local regulations.

In this context, the Trust Indexes are a means to measure trust interoperability across dataspaces and federations adopting and extending the Gaia-X Policy Rules.

Four indexes are proposed:

- Veracity
- Transparency
- Composability
- Semantic Match

7.3.4.1 VERACITY

The Veracity index is a function of:

- for a given claim, the length of the signing key chain (root → intermediate → intermediate → … → leaf)

Kroki

- for a given claim, the number of signatures on that claim (root → … → leaf ← … ← root)

Kroki

7.3.4.2 TRANSPARENCY

The Transparency index is a function of:

- the number of exposed optional properties of an object versus mandatory - in terms of the Gaia-X Policy Rules - properties of the same object. This ratio is always greater or equal to 1.
- the shape of the graph formed by the linked claims, measuring its eccentricity and depth.

7.3.4.3 COMPOSABILITY

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

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

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

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

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

7.3.4.4 SEMANTIC MATCH

The Semantic Match index is a function of:

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

3.3.4 7.4 Notarization service(s)

Notarization services are built to support the validation claims to verifiable credentials by allowing to review attestations by “Trusted Data Sources” and issuance of a verifiable credential once all conditions of the policies are met (e.g. the compliance service of the GXDCH uses the “Gaia-X registryNumber notarization API” to validate the company registration number).

A more generic set of services is provided by the Eclipse XFSC project which provides components to define a complex workflow and manage the chain of attribute validations and issuance of the VC through the NOT, TSA and WFE components.

3.3.5 7.5 Data Exchange Services

Data product usage in Gaia-X is enabled by a set of Data Exchange Services that are realized by each Participant and can be supported as Federation Services. Not all Data Exchange Services are mandatory.

1. **Authentication** (mandatory) is essential to connect two Participants. Authentication is provided by The Identities and Trust Framework: Identities provide general information on the Participant, and the Trust Framework appends additional claims, like verified location, or verified application of other standards or regulations.
2. **Policy negotiation and contracting** (mandatory) includes the ability to negotiate access and usage policies between two parties. This should be a sequence between the parties, but a contracting service can support here when one or multiple parties do not have the technical abilities for this. ODRL is used to support contracting as (a) it provides interoperability (all parties must be able to understand the policies to enforce them later) and (b) it enables computerised policy enforcement during the transaction.
3. A **catalogue** (mandatory) provides mechanisms to publish Data Product Descriptions (metadata) and support search or query of the Descriptions. A catalogue may be realized as a centralized or decentralized service, but the capability can also be realized as a distributed functionality.
4. **Vocabularies** (optional) provides additional metadata to the Data Product Descriptions. The Descriptions should contain a limited amount of information as a common denominator but must be extensible with vocabularies from different (business or technical) domains.
5. **Observability** abilities (optional) - logging and audit data - are required to provide an auditable framework for transactions.
6. **Data Exchange protocols** are required to exchange data between Participants and enable Data Usage. Data exchanges are realized on a peer-to-peer basis. Gaia-X does not promote any technical protocol - the actual protocol must be agreed between the parties during the contracting phase.

Detailed descriptions of these services are provided in the *Gaia-X Data Exchange Services Specifications* document.

7.5.1 Auditing

7.5.1.1 DATA CONTRACT SERVICES

Data Contract Services act as a broker of data delivery contracts between Data Providers and Data Consumers. They facilitate and enable agreements about data deliveries.

7.5.1.2 DATA EXCHANGE LOGGING SERVICES

Data Exchange Logging Services provide evidence that data has been (a) submitted and (b) received and © rules and obligations (Data Usage Policies) were enforced or violated. This supports the clarification of operational issues, but also eventually the clarification of fraudulent transactions. The Data Provider can track that, how, and what data was provided, and the consumer can be notified. A Data Consumer can track that data was received or not received. Additionally, the Data Consumer can track and provide evidence on the enforcement of data usage policies or violation of data usage policies.

4. 8. Other Concepts

4.1 8.1 Gaia-X and Data Meshes

Data meshes have emerged recently (since 2018) as an answer to the increasing difficulties of (logically) centralized and (conceptually) monolithic cloud data warehouses, data lakes, data lakehouses and other typically cloud-based approaches to the enterprise-wide management of analytical data to scale and to provide the required flexibility in today's highly dynamic and considerably complex business environments.

This section briefly outlines how Gaia-X incorporates several foundational principles of data meshes, where Gaia-X transcends the narrow boundaries of a data mesh approach, and why Gaia-X may also be characterized as a **mesh of service meshes**.

4.1.1 8.1.1 Data Mesh Definition

The definition provided by the initial promoter of this concept, Zhamak Dehghani (Dehghani 2022), is widely accepted, accurate, and highly useful (Machado *et al.* 2022, BITKOM 2022):

```
A data mesh is a decentralized socio-technical approach to sharing and managing analytical data in complex and large-scale environments within or across organizations.
```

"Socio-technical approach" indicates that organizations need to adapt their processes and governance in addition to providing the right technology to successfully implement a data mesh. *Sharing*, for obvious reasons, also includes all required capabilities to *access* the data. Originally, data meshes almost exclusively focused on **analytical data** used to support enterprise decision-making as opposed to directly undergirding operational and transactional data management happening in the various IT systems (e.g., ERP, CRM, PDM, MES,…). Recent trends, though, including concepts such as *data hubs* and *data fabrics*, try to overcome this limitation towards a more holistic approach to enterprise data management.

Data meshes closely observe and implement the following four principles:

Principle	Explanation
domain ownership	(Analytical) data shall be *owned* by different cross-functional (as opposed to specialized!) teams organized around and within appropriate business domains similar to domain-driven design (Evans 2004).
data as a product	(Analytical) data shall be managed and treated like a true product which is *produced* by someone in order to provide a well-defined *value* to its *consumers*.
self-serve data platform	to enable the different domain teams by providing domain-agnostic capabilities often needed in a *data as a product* approach such as, for instance, data product life cycle management, pipeline workflow management, data processing frameworks, policy enforcement components, and many others.
federated computational governance	Data meshes need an appropriate data governance operating model to balance the autonomy and agility of the individual domains with the global interoperability of the overall data mesh (Dehghani 2022, 8). Typically, this will be accomplished by a cross-functional team composed of domain experts, data platform specialists, and other suitable subject matter experts (legal, security, compliance) complemented by heavy automation (e.g., *policies as code*, test and monitoring automation)

4.1.2 8.1.2 Gaia-X as (Higher Order) "Service Mesh"

Gaia-X itself incorporates several principles present in a data mesh with the marked (and defining) difference that Gaia-X relies on a **service-oriented approach** at the centre of its conceptual model as opposed to restricting itself to analytical data. This generalization notwithstanding, Gaia-X co-opts several usability characteristics of *data as a product* from the data mesh approach such as

- discoverability and addressability of services (e.g., Federated Catalogue, self-sovereign identities)
- providing trustworthy and truthful services through its Trust Framework
- interoperability and composability of services
- guaranteeing secure consumption of services through automated policies (*policies as code*), monitoring, and the Gaia-X Data Exchange Services

Gaia-X Federation Services and the Trust Framework provide capabilities similar to the *self-service data platform* of a data mesh. Also corresponding to the data mesh philosophy, there exists a dedicated domain-independent organizational unit to operate and manage this platform as well, the Federator. Yet, in contrast to the smaller focus of data meshes, the Gaia-X Ecosystem is designed *ex-ante* for federating independent autonomous ecosystems itself (cf. the three Gaia-X *planes* in section Gaia-X ecosystems).

Like with data meshes, the organizational structure of Gaia-X itself (and also Gaia-X-based ecosystems) also relies on a substantially federated governance model when considering the Gaia-X Association (or Gaia-X ecosystem federators), its Committees, Working Groups, and decision processes. In the case of Gaia-X (or Gaia-X based ecosystems), though, all actors in the various governance functions belong to different legal entities, unlike typical data mesh situations where they appertain to a single organization or enterprise.

The term “domain” in data mesh implementations characteristically denotes “bounded contexts” such as spheres of knowledge, influence, activities, or responsibilities *within* a potentially large single organization (Evans 2004, Deghgani 2022). At face value, its Gaia-X equivalent, `Participant`, possesses almost identical semantics: “A `Participant` is an entity as defined in ISO/IEC 24760-1 as an “item relevant for the purpose of operation of a domain that has recognizably distinct existence”. Practically, though, Gaia-X `Participants` will mostly consist of independent legal entities.

In summary, a high degree of similarity and overlap between the principles applied for data meshes and for Gaia-X can be recognized. On the other hand, the deviations of Gaia-X with respect to the data mesh approach identified above amount to a form of conceptual generalization. Whereas data meshes primarily concentrate on the management of (analytical and other) data for a single (albeit large and complex) organization, typically a legal entity, Gaia-X provides organizational standards and technical components for realizing whole ecosystems consisting of several independent organizations. Extending this line of thinking one may be even (rightfully) enticed to designate Gaia-X as a **mesh of service meshes** (not just a mesh of data meshes). Note, however, that this particular usage of the term “service mesh” is different from the one encountered in microservice architectures[1].

4.2 8.2 Computational Contracts

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

4.2.1 8.2.1 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 the form of a contract.
- Every role of the Gaia-X Conceptual Model as well as of the Operational model can be seen as legal persons and therefore may have a role as a licensor or licensee or both.
- In traditional centralized 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 sovereign participation. The ability to observe if all parties of a contract behave the way it is agreed, to validate their rights, and to fulfil their obligations and ensure that no one can misuse information is key for a trustful relationship.
- Computable contracts aim to ease the complex processes of contract design, contract negotiation, contract signing, and contract termination as well as to observe the fulfilment of contractual obligations and compliance with national law.

Kroki

4.3 8.3 Trusted execution of services

Trusted Execution Environments (TEEs) provide a secure enclave within a computing system where sensitive operations can be executed with a high degree of confidentiality and integrity. These environments offer a protected space that is isolated from the rest of the system, shielding it from unauthorized access or tampering. TEEs are designed to safeguard critical processes, such as cryptographic operations, secure key storage, and sensitive data processing, making them vital components in ensuring the security of modern computing platforms.

In the Gaia-X context, TEE can be used to increase the trust in the software running on a clearing house. Solutions like Intel SGX ® allows to run software with the guarantee that the source code nor binary have been tampered since they were first built. Data processed in an enclave are also invisible from the eye of an external actor, as the enclave is obfuscated from the rest of the machine.

Running components on SGX or other TEE systems allow to trust instances without having to rely solely on the Gaia-X AISBL to identify the trusted providers and communicate a list. It also allow the network of Clearing Houses provider to grow without diminishing the confidence in the issued compliance VerifiableCredentials.

1. In (typically container-based) microservices architectures a service mesh provides software components and mechanisms to separate cross-cutting concerns of service-to-service communications from the business logic of the individual microservices into a so-called “control plane”. This segregation simplifies the development, operations, and management of larger microservices environments. It is achieved by making every microservice communicate with another one (the so-called east/west traffic) over dedicated components called proxies. See, for instance, Tech Radar - Service Mesh. ↵

5. Annexes

5.1 9. Changelog

5.1.1 9.1 2024 April release (24.04)

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

5.1.2 9.2 2023 October release (23.10)

- Replace “Overview”, which included general concepts, with “Context”, which defines the specific scope of the Gaia-X Architecture
- Update and generalization of the “Conceptual Model”
- Generic description of “Policy Management”
- Mapping of Gaia-X Architecture to the CASCO model
- Updates based on changes in the Data Exchange Services and Identity, Credential and Access Management Document
- Introduction of the “Party-Credential” as extension of Identity and Service Offering credentials with Membership and Service credentials
- Contains the technical implementation details for the Trust Framework (from the 22.10 Trust Framework document)
- Annex includes a section on “Gaia-X Digital Clearing House”
- Chapter “Gaia-X Trust Framework components” including “Gaia-X Compliance”, “Gaia-X Registry”, and “Gaia-X Notary - LRN”
- Inclusion of the “Publication/Subscription service” (ref. Federated Catalogues) and “Trust Indexes” sections in the new chapter “Enabling and Federation Services”

5.1.3 9.3 2022 October release (22.10)

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

5.1.4 9.4 2022 September release (22.09)

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

5.1.5 9.5 2022 April release (22.04)

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

5.1.6 9.6 2021 December release (21.12)

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

5.1.7 9.7 2021 September release (21.09)

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

5.1.8 9.8 2021 June release (21.06)

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

5.1.9 9.9 2021 March release (21.03)

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

5.1.10 9.10 2020 June release (20.06)

- First release of the Technical Architecture document by the BMWi

5.2 10. Glossary & References

5.2.1 10.1 Glossary

The Gaia-X online glossary is at https://gaia-x.gitlab.io/glossary/

5.2.2 10.2 References

- Berners-Lee, T. (2009). Linked Data. W3C. https://www.w3.org/DesignIssues/LinkedData
- BITKOM (2022). Data Mesh – Datenpotenziale finden und nutzen. https://www.bitkom.org/sites/main/files/2022-06/220531_LF_Data_Mesh.pdf
- 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
- Dehghani, Z. (2022). Data Mesh. Delivering Data-Driven Value at Scale. Sebastopol, CA: O'Reilly
- 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
- Evans, E. (2004). Domain-Driven Design. Tackling complexity in the heart of software. Upper Saddle River, NJ: Addison-Wesley.
- 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. 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
- Machado, I., Costa, C., & Yasmina Santos, M. (2022). Data Mesh: Concepts and Principles of a Paradigm Shift in Data Architectures. Procedia Computer Science 196, 263–271
- Platform 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/

5.3 11. Annex

5.3.1 11.1 Annex A - Gaia-X Digital Clearing House (GXDCH)

The Gaia-X Digital Clearing House (GXDCH) operationalizes the Gaia-X mission. A GXDCH makes the various mechanisms and concepts applicable in practice as a ready-to-use service set. Therefore, a distinct specification and implementation of the GXDCH exists.

The Gaia-X Framework contains functional and technical specifications, the Gaia-X Compliance Service to assess Gaia-X Compliance and the testbed to validate the behaviour of software components against the Gaia-X Policy Rules and technical specifications.

The GXDCH contains both mandatory and optional components.

All the mandatory components of the GXDCH are open-source software. The development and architecture of the GXDCH is under the governance of the Gaia-X Association.

A GXDCH instance runs the engine to validate the Gaia-X rules, therefore becoming the go-to place to become Gaia-X compliant. The instances are non-exclusive, interchangeable, and operated by multiple market operators.

> **Note**
>
> It is crucial to differentiate **compliance** and **compatibility**.
> A service can be made Gaia-X compliant with Gaia-X Policy Rules (see the Policy Rules Compliance Document). Software cannot.
> However, software can be made Gaia-X compatible with Gaia-X specifications.

11.1.1 Deployment

Each GXDCH instance must be operated by a service provider according to rules defined with and approved by the Gaia-X Association.

Such providers have then the role of Operator. Gaia-X is not an operator itself.

Any provider compliant with the requirements defined by the Gaia-X Association and featuring the necessary characteristics as defined by Gaia-X can become a GXDCH operator.

11.1.2 Description of the GXDCH components

The GXDCH is open for future evolution. This means that the components included in the GXDCH may be enhanced from release to release.

11.1.2.1 MANDATORY COMPONENTS

The components below are:

- technically compatible with the Gaia-X Specifications described in the functional and technical Specifications Documents of the Gaia-X Framework
- validated by the Gaia-X Testbed, which will aim to verify the behaviour of the components.

Those services are provided by GXDCH operators only.

- The Gaia-X Compliance Service
  It validates the shape and content of the Gaia-X Credentials and, if successful, issues a `GaiaXComplianceCredential`.
- The Gaia-X Registry
  This service implements the Gaia-X Registry and represents the backbone of the Gaia-X Ecosystem governance.
  It stores Trust Anchors accredited by the Gaia-X Association, the shapes, and schemas that are used for the Gaia-X Compliance validation. Beyond these core functionalities, the Gaia-X Registry Service stores further accompanying information such as the terms and conditions for the Gaia-X Compliance Service.
- The Gaia-X Notarisation Services For the business registration numbers it is a tool used to verify all registration numbers given by the participant in their Participant credentials.
- The Credential Event Service
  Provides a distributed storing solution to hold Gaia-X Compliant VCs and is used to synchronise the Credentials' IDs across the Federated Catalogues that want to publish Gaia-X conformant services.
- A InterPlanetary File System (IPFS) node
  This is used to synchronise the information across GXDCH instances, especially the information exposed by the Gaia-X Registry.
- A Logging service
  This service is not necessarily exposed to customers. This is subject to the federator depending on the federator's applicable regulations.

11.1.2.2 OPTIONAL COMPONENTS

The following services are optional and can be provided by any Provider, which may or may not be a GXDCH operator. Those services can also be operated by providers for specific Data Spaces or Federations.

- a Wizard
  UI to sign Verifiable Credentials in a Verifiable Presentation on the client side. Calling the Gaia-X Compliance Service is integrated with the Wizard, making it possible to obtain Gaia-X Compliant Credentials by directly using this tool. Depending on the implementation, a wizard can support signing with the hardware module via PKCS#11 or others like Web eID. The Wizard also provides a user-friendly way to send the Gaia-X Compliant VCs to the Credential Events Service. Finally, it provides an environment where the user can try out the Gaia-X Compliance for testing and learning purposes.
- a Wallet
  Application that allows the user to store their Credentials and present them to third parties when needed. This is also sometimes called an Agent, CredentialRegistry or IdentityHub.
- a Catalogue
  An instance of the Federated Catalogue using the Credential Event Service introduced in the previous section.
- a Key Management System
  This is provided by a provider as an additional service to help its customers to manage cryptographic material, including revocation, key rotation, key recovering, …
- Policy Decision Point (PDP)
  This service is used to perform the reasoning given one or more policy and input data. The PDP can support several policy languages. For each policy expressed in a Gaia-X Credential, the policy owner can give the list of PDP service endpoints, which can be used to compute the policies.
- Data Exchange services
  Those services are used to perform Data Exchange.

The optional services provided may be subject to fees set by the provider.

11.1.3 Gaia-X Digital Clearing House Releases

Each GXDCH release is a bundle of several components developed by the Gaia-X Lab with the support and contribution of the open-source community.

The components and their documentation are published for Operational and Lab versions at https://registry.gaia-x.eu/

11.1.3.1 ELBE RELEASE

This was an internal release only for testing purposes.

11.1.3.2 TAGUS RELEASE

This is the ongoing development cycle under `/v1`.

In order to allow several instances of different versions to be accessible in parallel, we specify paths during the deployment.
That means that `/development` uses the latest `development` image of a component, `/v1` uses the latest `v1` image, and `/main` uses the latest `main` image.

To ensure consistency, GXDCH operators are requested to use the same convention for their instances.

11.1.3.3 LOIRE RELEASE

This is the next major release developed by the Gaia-X Lab with the contribution of the open-source community.

5.3.2 11.2 Gaia-X Digital Clearing House Operations

11.2.0.1 GXDCH USAGE

For the Tagus release, the Gaia-X Association provides a load balancer in the form of some HAProxy nodes pointing to Clearing House instances.

These HAProxy nodes are managed through Load Balancer Ansible scripts.

The Gaia-X Association is responsible for these scripts and their update after a Clearing House is deployed. GXDCH operators can open merge requests on the HAProxy configuration to include their GXDCH instance.

The services are accessible behind three subdomains:

- https://registrationnumber.notary.gaia-x.eu/v1/docs/
- https://compliance.gaia-x.eu/v1/docs/
- https://registry.gaia-x.eu/v1/docs/

11.2.0.2 GXDCH SETUP

Technical details are available in the setup guide.

5.3.3 11.3 Annex B - Testbed

The testbed aims to verify the behaviour of a software component with regard to:

- Gaia-X Policy Rules
- Gaia-X Technical Specifications

The expected output of the testbed is a Gaia-X Credential containing claims about:

- the binary signature of the tested software component - like sigstore
- the version of the verified Gaia-X policy rules and technical specifications
- a validity period - not too short to have market adoption and stability, not too long to reduce software version fragmentation.

11.3.0.1 FOR LIVE SERVICE

The testbed can also be used to validate the behaviour of already deployed services.

This usage is left as an opportunity for the participants to monitor or assess the likelihood of a remote service to behave according to participants' expected rules.

5.3.4 11.4 Annex C - Mapping of Gaia-X concepts with concepts used in EU data regulation

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

European regulations concepts	Gaia-X concepts
data processor in GDPR	data product provider
data subject in GDPR / user in DxA	data licensor
consent in GDPR / permission or authorization in DxA	data usage agreement
recipient in GDPR / DxA	data consumer

The following diagram details these relationships:

Concept mapping Gaia-X vs EU regulations

Figure 11.1 - Mapping of Gaia-X concepts with EU data regulation concepts

5.3.5 11.5 Annex D - Conceptual models instantiation example: the 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 PFM service is proposed by a company called MyPFM to an end user Jane who has bank accounts in several banks: Bank(1), Bank(2)… MyPFM is using services provided by the banks Bank(i) to get the banking transactions of Jane and then aggregates these bank statements to create Jane's financial dashboard.

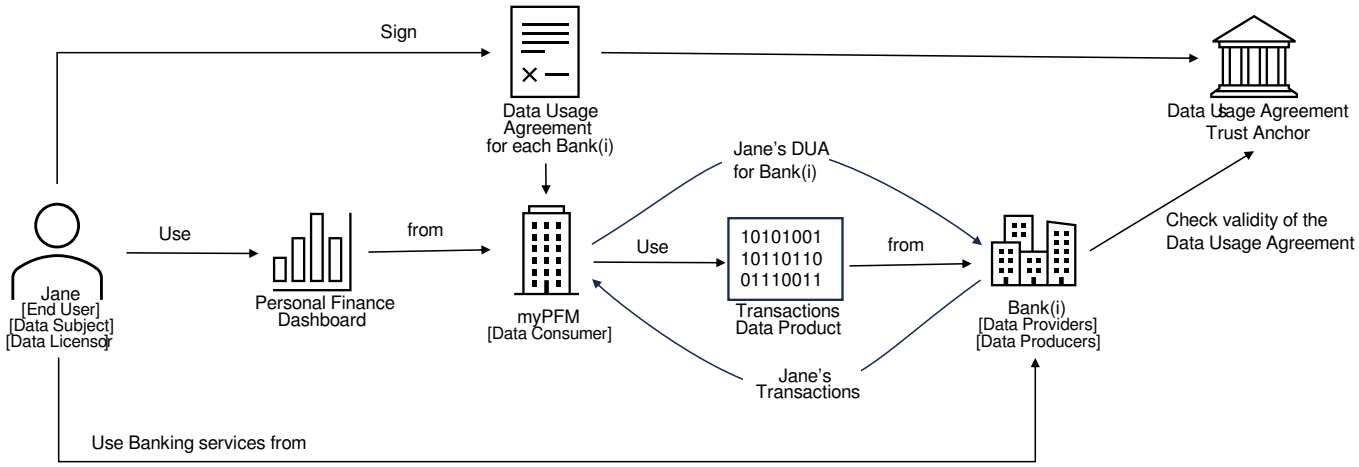

Jane is the End User and also the data licensor (as data subject per GDPR).

Bank(i) are Data Product Providers defining the Data Products (Service Offerings) delivering the banking transactions. They are also Resource Owners for the bank statements, which are Virtual Resources composing the Data Products, and as such are also the Data Producers. The associated Resource Policies are in fact predefined by the PSD directive from the European Parliament. (note: this is not the way that consent management is currently put in place for DSP2 services - the process described below is adapted to the *Financial and Insurance Data Access* (FIDA) EU regulation issued in June 2023.)

MyPFM is the Data Consumer (Service Consumer) that consumes the data (Data Usage, Service Instances) provided by Bank(i) in order to create a financial dashboard and to offer it to Jane.

MyPFM is also likely consuming other Service Instances from a PaaS Provider in order to run its own code, such as the code implementing the dashboard creation.

Because sensitive personal information is involved, she has to sign a specific "contract" stating precisely how myPFM is authorized to use Jane's data (e.g., authorized for establishing her financial dashboard), transmission to anybody else is prohibited except maybe for associating a payee and an expense category (i.e. payee X is a grocery, payee Y is a gas station, …).

Firstly, Jane will communicate her various bank account identifiers (IBAN) to myPFM. myPFM will use the Data Product catalogue to find out the appropriate services from each bank – we will suppose for simplicity's sake that they are all named GetTransactionFromIBAN. PFM will review each DataProductSelfDescription to ensure that it is compatible with their needs. As sensitive personal information is involved, Bank(i) will require a signed agreement from Jane. The agreement template is included in the DataProductSelfDescription. (Note: ecosystems will likely have predefined standard consent templates in order to enable automatic and agile processing). myPFM will fill in the template and send it to Jane for her to sign it through a digital identity and digital signature provider trusted by Jane, myPFM and Bank(i). myPFM and each Bank(i) will then configure the GetTransactionFromIBAN service for Jane, include Jane's signed Data Usage Agreement in the service contract (i.e., the updated DataProductSelfDescription) and co-sign the contract. myPFM can then get Jane's transaction data, process this data, and provide the financial dashboard to Jane.

Let us now suppose that myPFM also delivers a loan brokering service. To do so, myPFM provides a credit profile to loan institutions to receive credit proposals that they can rank and forward to Jane.

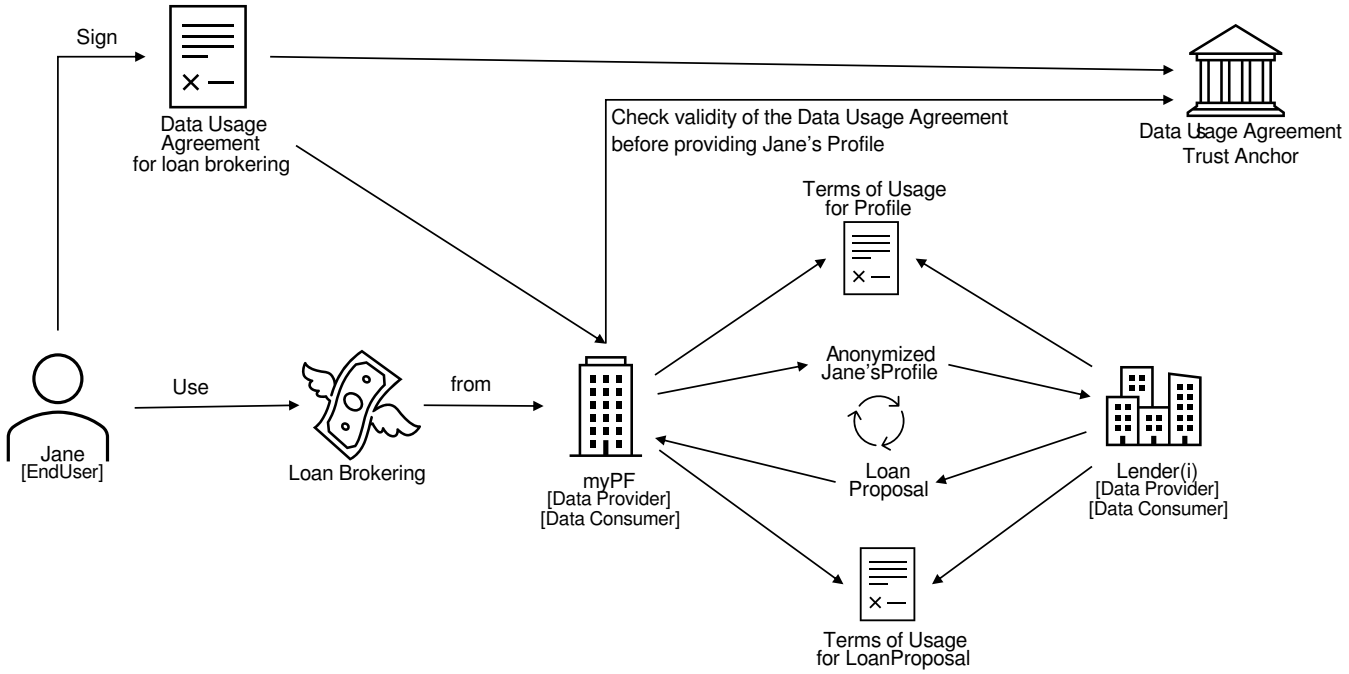

In this case, myPFM is both a Data Consumer from Bank(i) and a Data Provider to loan institutions Lender(i). If Jane wants to use this service, she will first have to sign a new data usage agreement to authorize myPFM to communicate her credit profile (total income, loan capacity, purpose of the loan, …) to some loan institutions – this credit profile is anonymous and should not enable identifying Jane. Then myPFM will query the Data Product Catalogue to find loan institutions providing such online loan services and will review the corresponding Data Product Description to check compliance with Jane's data usage agreement and with myPFM policy.

With each selected Lender(i), myPFM will negotiate, configure, and sign the service contract. The terms of usage will not include Jane's data usage agreement because the data transmitted to Lender(i) is anonymized at that stage. The terms of usage will include conditions specific to myPFM and Lender(i) business, for instance: myPFM is not authorized to communicate the credit proposal to other credit institutions for a given duration, myPFM guarantees that to their knowledge Jane is a real customer and not a data aggregator, Lender(i) commits to prepare a proposal within x hours, Lender(i) will pay myPFM some money if their offer is selected…

PFM will then call the getLoanProposal service from Lender(i). At that stage, no data is transmitted except an anonymous request identifier. In order to prepare a credit proposal, Lender(i) will have to get the credit profile associated with that identifier. For that Lender(i) will get, from the catalogue, the description of the getLoanRequestData service provided by myPFM, and Lender(i) will configure it and co-sign it. The terms of usage will still not involve Jane's agreement but should include clauses at least as strong as those in Jane's consent, for instance, that the data shall be used only for the purpose of establishing a credit proposal and shall be deleted within 30 days if the proposal is not activated. Lender(i) will then get the data from myPFM by activating the getLoanRequestData service. At that stage, PFM acts as a Data Product Provider and Lender(i) as a Data Consumer. Lender(i) will then prepare the loan proposal and make it available to myPFM.

myPFM will then collect the various credit proposals, review them, and rank them to prepare a recommendation for Jane.

The formal operating model is given below:

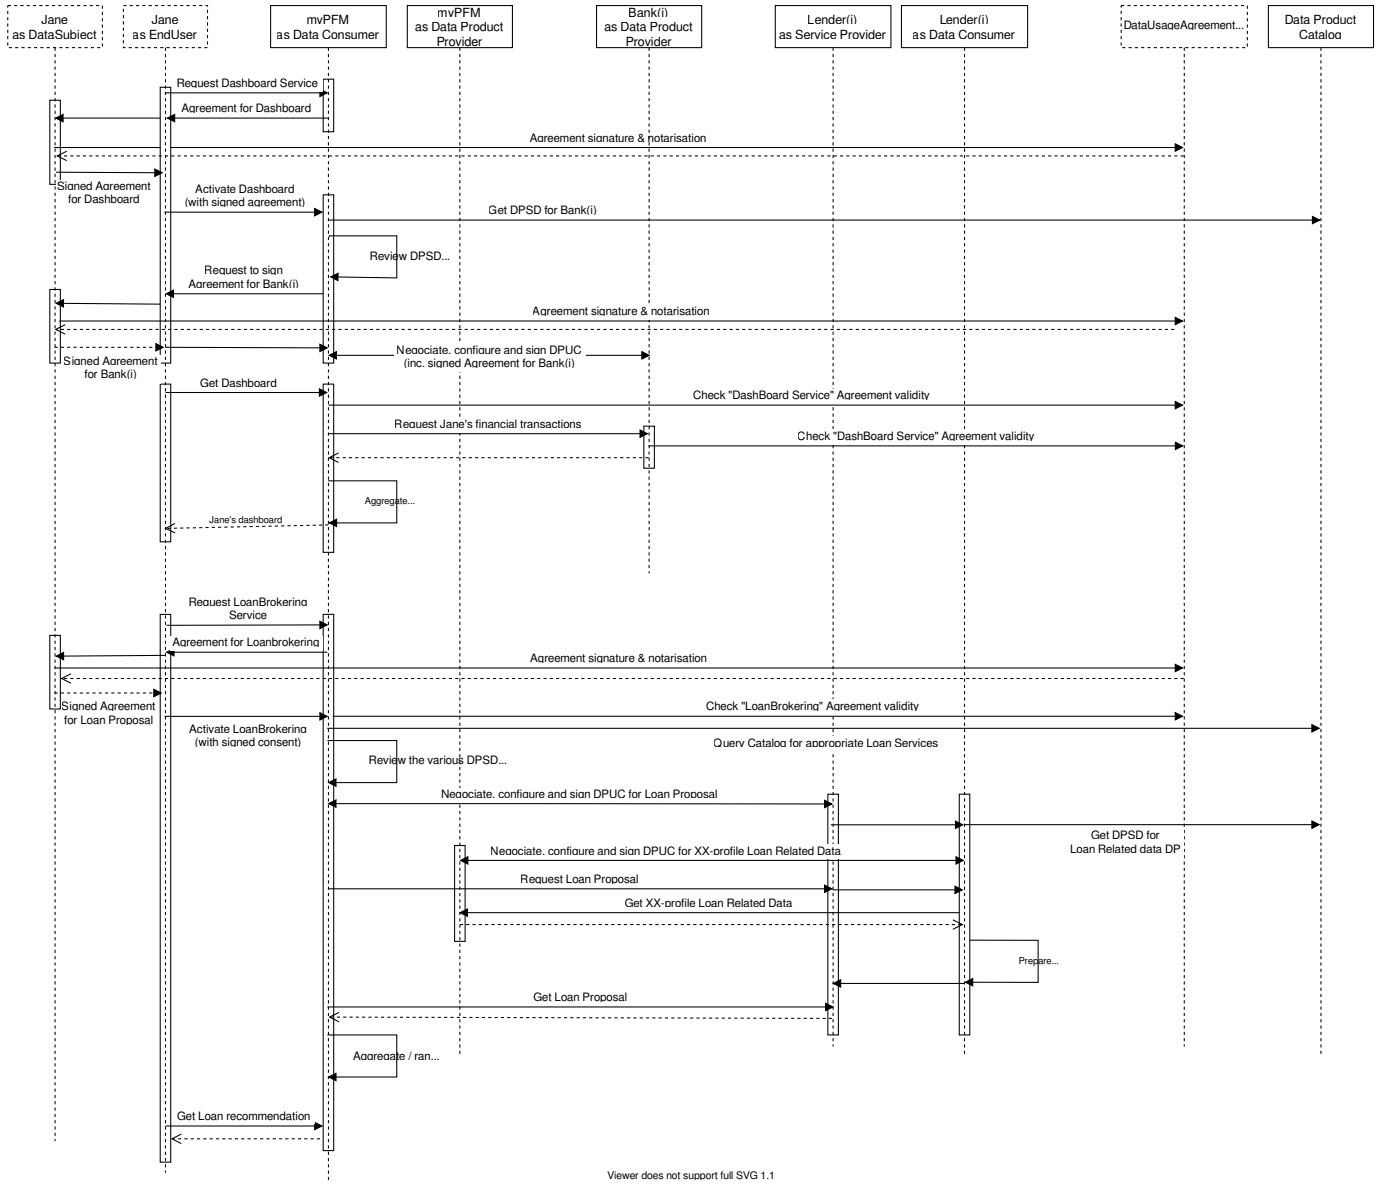

5.3.6 11.6 Annex E - Service Composition Example

11.6.0.1 PRACTICAL MULTI-CLOUD/-PROVIDER/-SERVICE COMPOSITION EXAMPLE

Figure 11.2 presents a pragmatic and practical multi-cloud, multi-provider and multi-service example (or use case) that reflects the current state of the art and practice in cloud services and computing.

Multi cloud / multiprovider / multiservice