Check for updates

METHOD ARTICLE

A federated policy-governed architecture for sovereign European data spaces

[version 1; peer review: awaiting peer review]

Paula Lamo[1], Diego Valdeolmillos[2], Marta Plaza Hernández[3], Panagiotis Dimitrakis[4], Ricardo S. Alonso[2]

[1]Tecnología Electrónica Ingeniería de Sistemas y Automática, Universidad de Cantabria, Santander, Cantabria, 39005, Spain
[2]AIR Institute, 49022, Zamora, Castile and León, Spain
[3]BISITE, Universidad de Salamanca, 37007, Salamanca, Castile and León, Spain
[4]NCSR ``Demokritos'', Institute of Quantum Computing & Quantum Technology, Paraskevi, 15341, Greece

v1 **First published:** 13 Aug 2026, **6**:293
https://doi.org/10.12688/openreseurope.24687.1
Latest published: 13 Aug 2026, **6**:293
https://doi.org/10.12688/openreseurope.24687.1

Abstract

European data spaces aim to enable trusted, sovereign and interoperable data sharing across organisational boundaries. Existing reference frameworks address parts of the problem in isolation: GAIA-X concentrates on federated trust and self-descriptions, the International Data Spaces Reference Architecture Model (IDS-RAM) emphasises usage control and the connector pattern, the Data Spaces Support Centre (DSSC) Blueprint consolidates governance vocabulary and Simpl-Open provides a vendor-neutral middleware substrate. The integration of semantic interoperability, policy-driven sovereignty, federated computation across the cloud–edge–high-performance computing (HPC) continuum and decentralised governance within a single deployable architecture remains an open design problem. This article proposes and analyses a conceptual, policy-governed federated architecture derived from the architectural work of the Horizon Europe NOUS project (Grant Agreement No. 101135927). The architecture is organised in five functional layers — data, services, compute, federation and governance — and revolves around five interacting roles: the Data Space Participant (DSP), the Data Space Infrastructure Participant (DSIP), the Data Space Common Services (DSCS) and the companion NOUS Data and Infrastructure Agents (NDA, NIA). A clear separation between Control Plane and Data Plane mediates policy negotiation, contract enforcement, federated learning and infrastructure provisioning across cloud, edge, HPC and quantum back-ends. We illustrate the model through two representative scenarios derived from NOUS use cases — crisis management with

Open Peer Review

Approval Status *AWAITING PEER REVIEW*

Any reports and responses or comments on the article can be found at the end of the article.

heterogeneous geospatial and field data, and cross-organisational federated analytics over heterogeneous compute resources — and we map architectural functions onto major European reference initiatives. The contribution is conceptual: the article does not claim empirical validation, performance benchmarks or formal certification of standards compliance. Limitations and a research agenda — prototyping, federated interoperability testing and formal verification of policy enforcement — are stated explicitly.

Keywords
federated data spaces, data sovereignty, policy enforcement, semantic interoperability, cloud–edge–HPC continuum, decentralised identity, reference architecture, European data strategy



This article is included in the Horizon Europe gateway.

Corresponding author: Paula Lamo (paula.lamo@unican.es)

Author roles: Lamo P: Conceptualization, Formal Analysis, Methodology, Visualization, Writing – Original Draft Preparation, Writing – Review & Editing; **Valdeolmillos D**: Conceptualization, Formal Analysis, Methodology, Visualization, Writing – Original Draft Preparation, Writing – Review & Editing; **Plaza Hernández M**: Formal Analysis, Funding Acquisition, Methodology, Project Administration, Writing – Original Draft Preparation, Writing – Review & Editing; **Dimitrakis P**: Conceptualization, Funding Acquisition, Project Administration, Writing – Review & Editing; **Alonso RS**: Conceptualization, Funding Acquisition, Project Administration, Visualization, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: This project has received funding from the European Union under grant agreement No 101135927 (Project name - A catalyst for EuropeaN ClOUd Services in the era of data spaces, high-performance and edge computing [NOUS]). Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union; neither the European Union nor the granting authority can be held responsible for them.The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.
The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2026 Lamo P *et al*. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Lamo P, Valdeolmillos D, Plaza Hernández M *et al.* **A federated policy-governed architecture for sovereign European data spaces [version 1; peer review: awaiting peer review]** Open Research Europe 2026, **6**:293 https://doi.org/10.12688/openreseurope.24687.1

First published: 13 Aug 2026, **6**:293 https://doi.org/10.12688/openreseurope.24687.1

1. Introduction

Data has become a structural input for science, public administration and industry, and the European Union has explicitly framed access to and governance of data as a matter of strategic autonomy. The European Strategy for Data,[1] the Data Governance Act[2] and the Data Act[3] set out a regulatory framework in which data should circulate across organisations and sectors while remaining under the control of the actors that originate it. The notion of a *data space*[12,23,24] operationalises this aspiration as a federated, voluntary and rule-based ecosystem in which participants can share data without surrendering sovereignty.

Implementing data spaces in practice is a non-trivial architectural challenge. A working data space must reconcile at least four concerns simultaneously: (i) *semantic interoperability*, so that heterogeneous data can be exchanged with unambiguous meaning; (ii) *data sovereignty and usage control*, so that policies travel with the data and are enforced beyond its initial release; (iii) *federation across heterogeneous infrastructure*, including cloud, edge, high-performance computing (HPC) and emerging quantum back-ends; and (iv) *decentralised identity, trust and auditability*, so that participants can recognise each other and verify the integrity of interactions without a single central authority. Existing reference frameworks each address subsets of these concerns. The IDS Reference Architecture Model[8,9] concentrates on data sovereignty and the IDS Connector pattern; GAIA-X[10,11] focuses on federated trust and self-descriptions; the DSSC Blueprint[12] unifies governance vocabulary; and Simpl-Open[13] provides a vendor-neutral middleware foundation. None of these initiatives, on its own, prescribes how the four concerns should be integrated in a single deployable architecture.

This article addresses that integration gap by proposing and analysing a *conceptual* federated architecture for policy-governed data spaces. The architecture is grounded in the technical work of the NOUS project, a Horizon Europe research and innovation action carried out by a consortium of twenty-one European partners (Section 3). It is structured as a layered, role-based model around five interacting roles — Data Space Participant (DSP), Data Space Infrastructure Participant (DSIP), Data Space Common Services (DSCS) and the participant-side NOUS Data Agent (NDA) and NOUS Infrastructure Agent (NIA) — with an explicit separation between a Control Plane (identity, policy, contract, governance) and a Data Plane (transfer, storage, computation). The contributions of the article are:

1. A structured conceptual architecture for federated, policy-governed data spaces that integrates semantic interoperability, sovereignty, federated computation and decentralised governance into a single layered model.

2. A consistent role-based decomposition (DSP, DSIP, DSCS, NDA, NIA) and a Control Plane/Data Plane separation that clarifies where policies are decided, where they are enforced and where data and workloads actually flow.

3. A mapping between architectural functions and the principal European data-space reference initiatives (GAIA-X, IDS-RAM, DSSC, Simpl-Open), exposing convergences and points of design discretion.

4. Two representative usage scenarios — crisis management and cross-organisational federated analytics — derived from NOUS use cases and instantiated through real project components, illustrating how the architecture supports diverse data, latency and policy requirements.

5. An explicit articulation of the conceptual scope, the limitations and the empirical and formal validation work that remains to be carried out.

The work is conceptual and analytical. It does not present benchmarked implementation results, nor does it claim formal certification of compliance with any specific standard. The article is intentionally hedged on matters that depend on prototyping or external validation, in line with the editorial expectations of *Open Research Europe.*

The remainder of the article is organised as follows. Section 2 surveys the state of the art across European data spaces, sovereignty and usage control, the cloud–edge–HPC continuum with federated learning, and decentralised identity. Section 3 describes the methodology and the provenance of the architectural analysis. Section 4 presents the proposed architecture layer by layer. Section 5 develops a worked policy-as-code example. Section 6 discusses two representative scenarios. Section 7 positions the contribution, examines alignment with European initiatives and states the limitations. Section 8 concludes and outlines future work. Tables 1–6 and the legends of Figures 1–3 are collected at the end of the manuscript, and the figures themselves are supplied as separate files.

Table 1. What this article adds beyond the [1] NOUS project deliverables.

Aspect	NOUS deliverables (e.g. D6.1[35])	This article
Audience and purpose	Consortium and EU project officer; engineering reference for partners.	International peer-reviewed readership; transferable architectural contribution.
Treatment of the architecture	Exhaustive component catalogue with internal interfaces and work-package traceability.	Conceptual five-layer model with role-based decomposition and explicit Control Plane/Data Plane separation.
Relation to the literature	Compliance-oriented mapping to EIRA/EIF and selected reference frameworks.	Critical synthesis of the state of the art on data spaces, sovereignty, federated learning and decentralised identity.
Methodological framing	Iterative engineering specification driven by deliverable obligations.	Qualitative architectural analysis with explicit evaluation criteria and a requirements-to-decisions traceability matrix (Table 2).
Worked examples	Component-level instantiation diagrams for selected workflows.	Two representative scenarios and a minimal worked policy example (Section 5).
Validation claims	Document-internal review and consortium acceptance.	Explicitly conceptual; no empirical validation, benchmarks or formal certification claimed.

Table 2. Traceability between architecturally significant requirements, the evidence supporting them and the architectural decisions adopted in this work. Sections refer to the proposed architecture in Section 4.

ID	Requirement	Evidence	Architectural decision
R1	Sovereign release of data with downstream usage control.	IDS-RAM Usage Control specification[9,28]; IDS-RAM v4[8]; ODRL information model.[16]	ODRL-encoded usage policies attached to assets; DSCS PDP and participant-side PEP enforce decisions before any release (Sections 4.2 and 4.6).
R2	Cross-organisational identity without a single central authority.	W3C DID and Verifiable Credentials[17,18]; self-sovereign identity literature.[34]	DSCS Identity Provider Federation issuing and verifying VCs bound to DIDs (Section 4.6).
R3	Semantic interoperability across heterogeneous sectoral schemas.	GAIA-X self-descriptions[10]; DSSC common vocabulary[12]; W3C SHACL.[19]	RDF/SHACL-based asset descriptions and a federated catalogue; plug-and-play interoperability middleware close to the source (Sections 4.2 and 4.5).
R4	Federated computation without raw-data sharing.	Federated learning literature[29–31]; NOUS use-case requirements.[35]	Local FL agents at each NDA; global aggregator hosted by the DSCS Decentralised Learning service; contracts authorise computation, not only transfer (Section 4.5).
R5	Workload portability across the cloud–edge–HPC continuum.	Compute-continuum literature[32,33]; NOUS HPC and quantum deliverables.[35]	NIA exposes IaaS /PaaS/MLaaS abstractions; placement decisions evaluate functional and policy fit (Section 4.4).
R6	Continuous, identity-centric verification of every interaction.	NIST Zero-Trust Architecture[20]; policy-as-code practice.[22]	Zero-Trust overlay; per-flow authentication and authorisation; policy-as-code runtime in Service Gateway PEPs (Section 4.6).
R7	Tamper-evident, proportionate audit of governance events.	Regulatory and sectoral non-repudiation needs[3,5]; EBSI programme.[7]	Append-only Control Plane log; optional anchoring on a distributed ledger as a proportionate, not central, mechanism (Section 4.6).
R8	Compliance with the EU regulatory framework for data and AI.	GDPR[5]; Data Governance Act[2]; Data Act[3]; AI Act.[4]	Policy obligations explicitly model lawful-basis, purpose limitation, jurisdictional placement and AI-system role (Sections 4.2 and 5).

Table 3. Mapping between conceptual layers and architectural elements.

Conceptual layer	Principal architectural elements
Data	NDA asset, contract and policy management; federated catalogue; semantic mediation and ontology alignment; data persistence and physical-element adapters.
Service	Application and data orchestration; workflow configuration; parallel processing; data standardisation; policy-aware service exposition; domain dashboards.
Compute	NIA IaaS (cloud, edge, HPC, quantum); PaaS (blockchain, database); MLaaS (AI and analytics provisioning); admin services; runtime orchestration and monitoring.
Federation	Plug-and-play interoperability middleware; local NDA federated learning agents and DSCS Decentralised Learning coordination; cross-DS contracts; shared service and asset catalogues.
Governance	DSCS Identity & Access Management federation; certification authority; Authorisation Service (PDP); API-gateway Policy Enforcement Point (PEP); Dynamic Attribute Provisioning Service (DAPS); DID registry; verifiable credentials; ledger-anchored audit log; Zero-Trust overlay.

Table 4. Summary of governance mechanisms in the proposed architecture.

Mechanism	Realisation
Identity federation	DSCS Identity & Authentication Provider Federation; DIDs and verifiable credentials.[17,18]
Trust establishment	DSCS Certification Authority; Dynamic Attribute Provisioning Service; signed self-descriptions of participants and assets.
Policy decision	DSCS Authorisation Service acting as PDP; ODRL-based policies[16] extended with federated-computation obligations.
Policy enforcement	Service Gateway PEP at each NDA and at the DSCS edge; policy-as-code runtime[22]; default-deny semantics.
Auditability	Append-only Control Plane log; optional, proportionate anchoring on a distributed-ledger substrate (e.g. EBSI[7]) when stronger non-repudiation is required.
Network trust	Zero-Trust overlay across participants[20]; per-flow authentication and authorisation; encrypted channels managed by NIA Encryption/Key services.

Table 5. Mapping of the two representative scenarios to architectural functions. Filled bullets (●) indicate primary use; open bullets (○) indicate secondary or optional use.

Architectural function	Scenario A (crisis management)	Scenario B (federated analytics)
Semantic mediation	●	●
Policy enforcement (PEP/PDP)	●	●
Federated learning	●	●
HPC/quantum provisioning	○	●
Edge provisioning and low-latency operation	●	○
Audit logging and ledger anchoring	●	●
Identity and verifiable-credential federation	●	●

2. State of the art

2.1 European data-space initiatives

The notion of a federated data space gained traction with the European Strategy for Data,[1] which advocated common European data spaces in priority sectors such as health, mobility, energy, industry and public administration. The Data Governance Act[2] and the Data Act[3] subsequently established legal mechanisms for data altruism, intermediation services and access rights to machine-generated data. From a technical perspective, four reference efforts shape the current field.

Table 6. Functional comparison between the proposed architecture and major European data-space reference initiatives. The mapping is conceptual and does not constitute standards certification.

Dimension	IDS-RAM	GAIA-X	DSSC Blueprint	Simpl-Open
Identity & trust	Connector identity, trusted parties	Self-descriptions, conformity assessment	Trust framework principles	Federated identity middleware
Policy & sovereignty	Usage Control, contract pattern	Policy rules in self-descriptions	Sovereignty as design principle	Policy hooks in middleware
Federation	Connector federation and protocol bindings (EDC and the Dataspace Protocol[14,15])	Federation Services catalogue	Cross-data-space governance	Common middleware substrate
Semantic interoperability	Information Model	Self-description vocabulary	Common vocabulary	Pluggable adapters
Infrastructure abstraction	Limited	Federated infrastructure	Conceptual	Vendor-neutral provisioning

The International Data Spaces Association proposes the IDS Reference Architecture Model, recently released in version 4,[8] together with an explicit Usage Control specification[9,28] that operationalises sovereign exchange through the IDS Connector pattern and contract-bound transfers. The Eclipse Dataspace Components (EDC)[14] provide an open, extensible reference implementation aligned with that model. GAIA-X[10,11] emphasises federated trust through self-descriptions, conformity assessment and a Federation Services catalogue. The Data Spaces Support Centre[12] consolidates a shared vocabulary and publishes a Blueprint that integrates governance, business, technical, legal and human aspects. At the European infrastructure level, Simpl-Open[13] provides an open-source middleware intended as a substrate for cloud-to-edge federations across data spaces. Recent edited volumes[23,24] synthesise these strands, while peer-reviewed surveys characterise the maturity of implementations and the open challenges that remain.[25,26]

2.2 Data sovereignty and usage control

Data sovereignty extends classical access control by attaching to data, at the moment of release, conditions of use that must be honoured downstream. Within IDS-RAM, sovereignty is operationalised by the Connector and by Usage Control

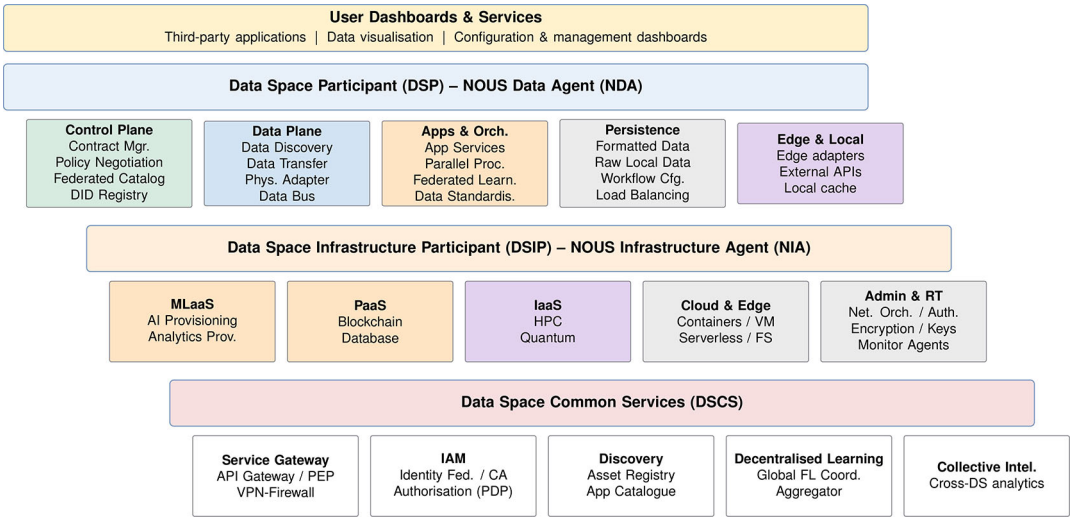

Figure 1. Overall layered architecture. The DSP layer hosts the NOUS Data Agent with an explicit Control Plane/Data Plane split; the DSIP layer hosts the NOUS Infrastructure Agent abstracting cloud, edge, HPC and quantum back-ends; the DSCS layer provides cross-cutting identity, discovery, learning-coordination and policy services. Abbreviations: DSP, Data Space Participant; DSIP, Data Space Infrastructure Participant; DSCS, Data Space Common Services; NDA, NOUS Data Agent; NIA, NOUS Infrastructure Agent; PDP, Policy Decision Point; PEP, Policy Enforcement Point; IaaS, Infrastructure-as-a-Service; PaaS, Platform-as-a-Service; MLaaS, Machine-Learning-as-a-Service; DID, decentralised identifier; FL, federated learning; HPC, high-performance computing.

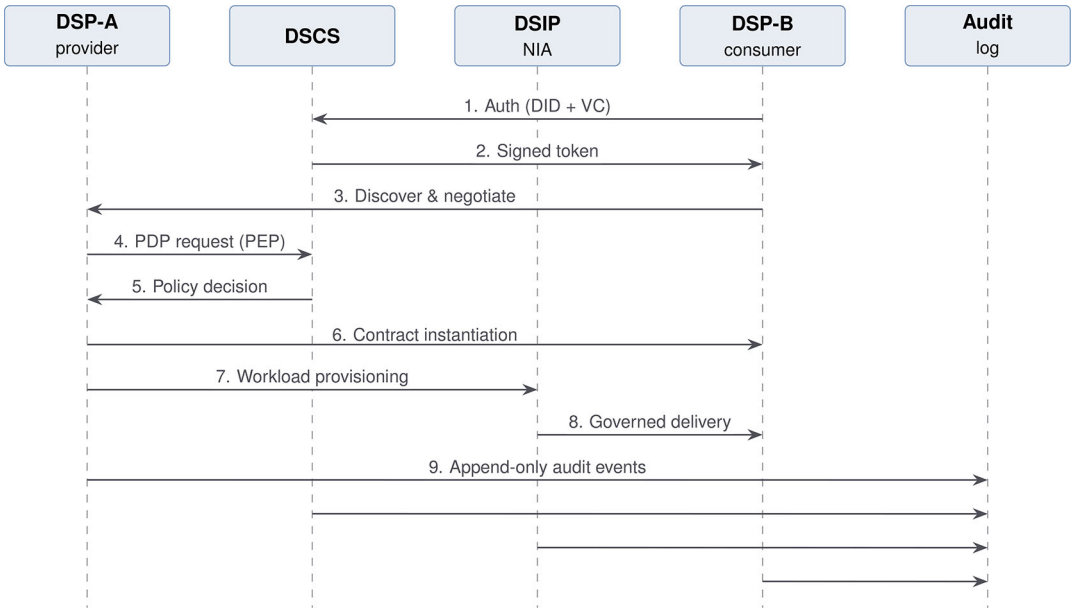

Figure 2. Policy-governed data exchange across the principal architectural roles. The lanes are, from left to right, DSP-A (a data provider), the DSCS, the DSIP with its NIA, DSP-B (a data consumer) and the append-only audit log. Numbered arrows follow the temporal order of a single exchange: authentication with a DID and VCs and issue of a signed token (steps 1 and 2), discovery and policy negotiation (step 3), evaluation by the PDP through the provider's PEP and the resulting decision (steps 4 and 5), contract instantiation (step 6), workload provisioning (step 7), governed delivery (step 8) and the appending of audit events by every role (step 9). All governance steps precede any movement of data. Abbreviations: DSP, Data Space Participant; DSIP, Data Space Infrastructure Participant; DSCS, Data Space Common Services; NIA, NOUS Infrastructure Agent; PDP, Policy Decision Point; PEP, Policy Enforcement Point; DID, decentralised identifier; VC, verifiable credential.

policies expressed at design time and enforced at run time.[9,28] A recent review[27] clarifies the multiple meanings of *data sovereignty* as a normative claim, an architectural property and a regulatory objective; the architecture proposed here addresses the second of these senses while remaining compatible with the third. At the language level, the W3C Open Digital Rights Language (ODRL)[16] provides an information model for permissions, prohibitions and obligations that has been adopted as a lingua franca for usage policies. Policy enforcement practice has been further influenced by the rise of policy-as-code approaches such as the Open Policy Agent,[22] and by the Zero-Trust paradigm of NIST SP 800–207,[20] which advocates continuous verification of identity and authorisation rather than perimeter-based assumptions of trust.

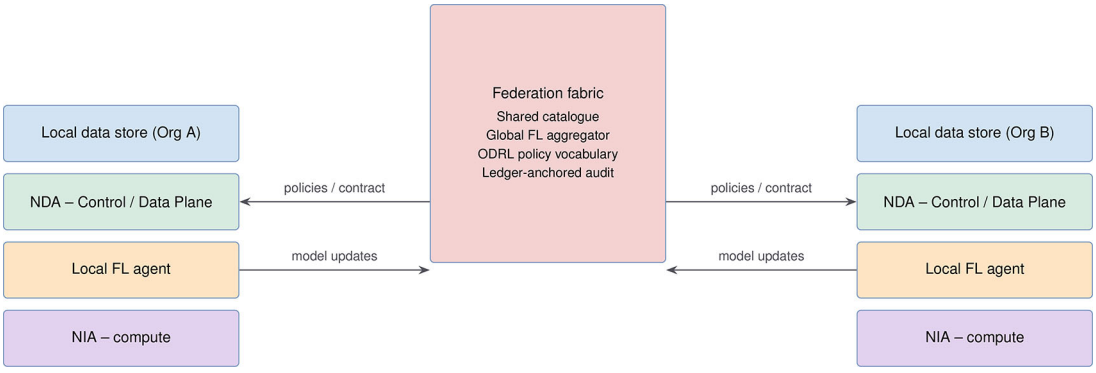

Figure 3. Policy-governed federated workflow across two participating organisations. Each organisation retains its local data store, its NDA, its local FL agent and its NIA on its own infrastructure; the central federation fabric provides the shared catalogue, the global FL aggregator, the common policy vocabulary expressed in ODRL and the ledger-anchored audit channel. Only model updates and governed views, not raw data, cross the federation fabric; policies and contracts travel in the opposite direction. Abbreviations: NDA, NOUS Data Agent; NIA, NOUS Infrastructure Agent; FL, federated learning; ODRL, Open Digital Rights Language.

Data sovereignty also intersects with regulatory regimes, in particular the General Data Protection Regulation (GDPR)[5] for personal data and the AI Act[4] for AI systems trained on shared data.

2.3 Cloud–edge–HPC continuum and federated learning

Modern data-intensive applications increasingly span heterogeneous infrastructures, from edge devices through cloud regions to HPC facilities. The compute continuum[32,33] requires abstractions that can describe and provision resources of widely different cost, latency and capability profiles. Federated learning[29–31] is a key technique for collaborative model training without raw-data sharing, and maps naturally to data-space settings in which participants are willing to share aggregate insights but not the underlying records. Open challenges include statistical and systems heterogeneity, privacy guarantees beyond data minimisation, and the orchestration of training workloads across heterogeneous compute back-ends.[29,30] The architectural question that follows from these results — how to reconcile federated learning with declarative usage control and with infrastructure-level placement decisions — has so far received less systematic attention.

2.4 Decentralised identity and verifiable credentials

Decentralised approaches to identity are central to the trust assumptions of federated data spaces. The W3C Decentralized Identifiers (DID)[17] and Verifiable Credentials (VC)[18] specifications define identifiers and signed attestations that participants can present without relying on a single central identity provider. These primitives underpin self-sovereign identity[34] and continue to evolve through recent W3C recommendations, including Verifiable Credentials Data Model v2.0,[18] and have been integrated into IDS and GAIA-X trust frameworks. Distributed ledger technology can serve as an anchor for verifiable identifiers and as an append-only audit substrate,[7] although the architecture proposed in this article uses ledger anchoring only as a proportionate audit mechanism rather than as a substrate for primary data storage.

2.5 Synthesis

Each strand offers important pieces of the puzzle, but the literature provides limited guidance on how to integrate them in a single deployable reference architecture that simultaneously honours sovereignty, semantic interoperability, federation across the compute continuum and decentralised governance. The architecture proposed in Section 4 addresses precisely this integration gap, with explicit attention to its instantiation in concrete project components.

3. Methodology and architectural analysis approach

3.1 Provenance and positioning relative to project deliverables

The architecture analysed in this article was elaborated within the NOUS project (*A catalyst for EuropeaN ClOUd Services in the era of data spaces, high-performance and edge computing*), a Horizon Europe research and innovation action (G.A. No. 101135927; Call/Topic HORIZON-CL4–2023-DATA-01-02; HORIZON-RIA) carried out by a consortium of twenty-one partners — the coordinating organisation and twenty participating organisations — from eleven European countries. The reference architecture, the use-case requirements and the open-source component catalogue of the project[35] constitute primary inputs to this work, but the article is not a deliverable rewrite. Project deliverables aim at consortium-internal completeness, traceability to contractual obligations and engineering reproducibility; an academic article must instead foreground a transferable architectural argument and situate it in the wider scholarly debate. Table 1 makes this distinction explicit.

3.2 Method

The methodology followed in this article is best characterised as *qualitative architectural analysis* in the design-research tradition. From the inputs listed below we identified a set of *architecturally significant requirements* (R1–R8) that the proposed model must satisfy; we then traced each requirement to the specific evidence that motivates it and to the architectural decisions that respond to it. The result is the traceability matrix presented in Table 2, which makes the chain *requirement → evidence → decision* auditable.

The inputs analysed were the publicly accessible NOUS architectural deliverable[35]; the reference architectures and trust frameworks of IDS-RAM,[8] GAIA-X,[10,11] the DSSC Blueprint[12] and Simpl-Open[13]; the European interoperability reference architecture[6]; the open standards relied upon by the architecture (W3C ODRL, DID, VC; NIST SP 800–207); and the peer-reviewed literature surveyed in Section 2. The analysis was structured around four evaluation criteria: (C1) *functional coverage* of the four integration concerns identified in the Introduction; (C2) *alignment* with the principles of European data-space initiatives; (C3) *separation of concerns* between policy decision, policy enforcement and data movement; and (C4) *sovereignty preservation.*

The output of the analysis is the conceptual architecture described in Section 4, the worked policy example of Section 5, the two representative scenarios of Section 6 and the discussion of Section 7. The article does not report a controlled

experiment, a benchmark or a deployed prototype, and does not claim empirical or formal verification of the architecture; such validation is part of the future work outlined in Section 8.

4. Proposed federated architecture

4.1 Overall overview

The architecture is organised as a layered, role-based model. Five interacting roles structure the system. The *Data Space Participant* (DSP) is any organisation that publishes or consumes data assets within the data space. Its technical embodiment is the *NOUS Data Agent* (NDA), a middleware component that exposes data assets, negotiates contracts and orchestrates application-level workflows. The *Data Space Infrastructure Participant* (DSIP) provides the computational and storage substrate over which DSP workloads execute. Its technical embodiment is the *NOUS Infrastructure Agent* (NIA), which abstracts cloud, edge, HPC and quantum back-ends behind common Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Machine-Learning-as-a-Service (MLaaS) interfaces. The *Data Space Common Services* (DSCS) provide cross-cutting functions that no single participant should own: identity federation, certification, policy decision and enforcement, application and data discovery, decentralised learning coordination and a shared service catalogue. A *User Dashboards & Services* layer hosts third-party applications, data visualisation tools and configuration-management dashboards on top of these roles.

A defining feature of the architecture is the explicit separation between a *Control Plane* and a *Data Plane*. The Control Plane is responsible for everything that defines and authorises an interaction: identity assertions, capability descriptions, policy negotiation, contract instantiation and audit. The Data Plane is responsible for the actual movement, storage and processing of data once an interaction is authorised. This separation makes it possible to reason about governance independently from data flow, and to deploy the two planes with different scaling, latency and trust assumptions.

Figure 1 sketches the overall architecture; it spans the four bands of *User Dashboards & Services*, DSP, DSIP and DSCS, with the Control Plane/Data Plane distinction running across DSP. The way the two planes cooperate during a single governed exchange is shown in Figure 2 and described step by step in Section 4.6. Table 3 summarises the mapping between conceptual layers and the principal architectural elements.

4.2 Data layer

The data layer is the locus at which assets become first-class citizens of the data space. Each NDA exposes data assets as governed resources described by self-descriptions that include semantic types, lineage, quality indicators and the policy templates under which the asset can be used. Asset descriptions are advertised through a federated catalogue[8,10]: each participant retains authority over its own descriptions, and the catalogue federates queries across participants without centralising the data itself. The DSCS Data Asset Registry, Metadata Repository and Application Catalogue provide the cross-DSP discovery layer.

Semantic interoperability is treated as a first-order concern. The data layer accommodates heterogeneous source formats through domain-extensible vocabularies expressed in the Resource Description Framework (RDF) and validated against Shapes Constraint Language (SHACL) shape constraints,[19] which make the conformance of exchanged data inspectable rather than implicit. A plug-and-play interoperability middleware performs syntactic and semantic translation as close as possible to the source, so that the governance layer can reason in terms of stable, agreed views rather than over the idiosyncrasies of each producer.[35] The architectural commitment is that mediation is a shared, declarative artefact (a translation specification) rather than a per-pair piece of glue code: this makes the cost of admitting a new participant linear rather than quadratic in the size of the data space.

Data sovereignty is encoded at the asset level by attaching usage policies expressed in a controlled subset of ODRL,[16] extended with obligations relevant to federated computation — for example aggregation thresholds, retention horizons, jurisdictional constraints and pre-export anonymisation requirements (a worked example is given in Section 5). These policies are not advisory: they are evaluated by the DSCS Policy Decision Point (Section 4.6) and enforced by participant-side Policy Enforcement Points (PEPs) hosted in the Service Gateway before any release of data. The decision to keep the policy language and the catalogue of self-descriptions inside the Control Plane, rather than embedded in ad-hoc service contracts, is what allows policies to travel with the data and to be re-evaluated at later stages of a workflow.

4.3 Service layer

The service layer hosts the application-level logic that operates over governed data. It is organised around two responsibilities. First, *application and data orchestration* composes data exchange, processing and analytics steps into workflows that respect the contractual context within which data has been released. Second, *policy-aware service*

exposition ensures that any service exposed to a third party publishes its self-description and the conditions under which it can be invoked, so that consumers can negotiate or refuse to proceed before any data is in motion.

The service layer is deliberately thin: it does not duplicate infrastructure-level orchestration, which is the responsibility of the NIA, nor does it assume any particular sectoral data model. Sectoral adaptation is achieved through configurable workflows, semantic profiles and domain dashboards. The architectural rationale is that the service layer should be the place where business logic is composed but not the place where governance is re-implemented: every service it exposes inherits its policies from the data layer rather than redefining them.

4.4 Compute layer

The compute layer abstracts the substrate on which data and workflows ultimately execute. The NIA exposes three layered abstractions: IaaS for virtual machines, containers, serverless functions, block and file storage; PaaS for managed components such as databases and ledger nodes; and MLaaS, which composes AI provisioning and analytics provisioning into reusable building blocks for federated and non-federated workloads.

A distinctive feature is the integrated treatment of the cloud–edge–HPC continuum.[32,33] Heterogeneous back-ends — commodity cloud, edge nodes, HPC clusters, and where available, quantum simulators reached through hybrid quantum–classical interfaces[35] — are described by uniform metadata descriptors that expose capability, locality and policy attributes. The NIA selects suitable back-ends for a workload by evaluating both functional fit (e.g. GPU, latency, memory) and policy fit (e.g. jurisdictional placement, data minimisation requirements), so that infrastructure consumption itself becomes a policy-bearing activity rather than a purely operational concern. The conceptual move here is to treat *where* a workload runs as part of the contract attached to the data, on the same footing as *whether* the workload is allowed to run at all: the same policy vocabulary that governs data release also governs placement and back-end selection.

4.5 Federation components

Federation is what makes the architecture more than a collection of bilateral exchanges. Three federation mechanisms are prominent. First, a plug-and-play interoperability middleware mediates heterogeneous data and protocol conventions through a shared interoperability catalogue, supporting protocol-agnostic transport and semantic alignment close to the point of origin. Second, federated learning is realised by local federated learning (FL) agents associated with each NDA and by a global FL coordinator that resides logically in the DSCS Decentralised Learning service; only model parameters or aggregated statistics cross participant boundaries.[29–31] The architectural commitment is that *contracts authorise computation, not only transfer*: a contract can permit the joint training of a model over an asset while forbidding any export of that asset, which is more expressive than the classical IDS contract for data transfer.[8,9] Third, the cross-data-space protocol surface is aligned with the Eclipse Dataspace Components[14] and with the IDS-RAM connector pattern, so that participants of disjoint data spaces can interoperate without a bespoke integration layer.

Figure 3 shows the resulting topology for the federated-learning mechanism: each participating organisation retains its local data store, NDA and local FL agent on its own infrastructure, while a federation fabric — shared catalogue, global FL aggregator, common ODRL policy vocabulary and ledger-anchored audit channel — mediates the exchange. Only model updates and governed views, and never raw records, cross the boundary between organisations; the contracts and policies that authorise each round travel in the opposite direction, from the federation fabric to the Control Plane of each NDA.

4.6 Decentralised governance, identity, trust and auditability

The governance layer brings the architecture together by ensuring that every interaction is authenticated, authorised, observable and revocable. Identity is federated through the DSCS Identity Provider Federation: participants are recognised through verifiable credentials[18] bound to decentralised identifiers,[17] allowing trust relationships to be established across organisational and national boundaries without requiring a single central authority. A certification authority issues digital certificates that participants use to authenticate channels and sign self-descriptions; a Dynamic Attribute Provisioning Service (DAPS) supports run-time attestation of participant capabilities.

Authorisation follows a clean separation between decision and enforcement. The Policy Decision Point (PDP), implemented by the DSCS Authorisation Service, evaluates access requests against the contracted policies; participant-side Policy Enforcement Points, embedded in the Service Gateway and in service-facing API gateways, intercept requests and implement the decisions. This is the architectural realisation of the policy-as-code discipline[22] in the data-space setting and aligns with the Zero-Trust paradigm of NIST SP 800–207[20]: every access is verified, no implicit perimeter is assumed, and lateral movement within a participant does not imply lateral access across participants. The architectural payoff of this separation is twofold. It allows the policy language and the decision engine to evolve

independently of the diverse runtimes that host enforcement points; and it makes policy decisions auditable as first-class events, because the PDP becomes the natural anchor for the audit log (Section 5 returns to this point with a worked example).

Figure 2 traces the resulting sequence of interactions for a single policy-governed exchange between a data provider (DSP-A) and a data consumer (DSP-B). The consumer first authenticates against the DSCS with its decentralised identifier and verifiable credentials and obtains a signed token; discovery and policy negotiation then take place between consumer and provider; the provider's Policy Enforcement Point queries the DSCS Policy Decision Point and, only if the decision is positive, a contract is instantiated; workload and transfer provisioning is delegated to the DSIP, which performs the governed delivery; and every role appends its events to the shared append-only audit log. The figure makes visible that all governance steps precede any movement of data, which is the operational meaning of the Control Plane/Data Plane separation.

Auditability is provided by an append-only audit log to which Control Plane events are written. Where stronger non-repudiation is required — for example in regulated sectors[3,5] — the log can be anchored on a distributed ledger, in line with European-level efforts such as the European Blockchain Services Infrastructure.[7] Ledger anchoring is, however, a proportionate option rather than a central commitment of the architecture. The architecture deliberately decouples the abstract property (*tamper-evident audit*) from any single substrate that realises it: a permissioned ledger, a public ledger, a hash-chain maintained by a trusted operator, or a transparency-log style substrate are all admissible, and the choice is best made on sectoral risk grounds rather than at the architectural level.

A summary of governance mechanisms is given in Table 4.

5. Policy-as-code: a worked example for federated learning

To make the role of the Control Plane concrete, we present a minimal ODRL policy that authorises a federated learning campaign over a governed asset while forbidding raw export. The policy is written in ODRL 2.2[16] as a JSON-LD document; it is deliberately illustrative and abstracts from full vocabulary issues that domain communities are expected to settle (R3, Table 2). The policy expresses three architecturally significant ideas: (i) the *permission* is to participate in a federated learning protocol, not to receive the asset; (ii) the *prohibition* forbids any operation that would export the asset or copy it verbatim; and (iii) a set of *duties* obliges the consumer to operate within an EU jurisdiction, to apply differential-privacy noise to aggregated updates, and to log every training round to the shared audit channel.

The policy is evaluated by the DSCS PDP at three points in the lifecycle of a campaign. *At contract instantiation*, the PDP checks that the consumer's verifiable credentials satisfy the `purpose` and `legalBasis` constraints; if they do not, no contract is created. *At workload provisioning*, the NIA queries the PDP with the proposed back-end placement; the `operateInJurisdiction` duty filters out placements that would violate the spatial constraint, regardless of which back-ends would be cheaper or faster. *At each round of the federated learning protocol*, the participant-side PEP intercepts model updates: it verifies that the differential-privacy obligation has been honoured by the local trainer (e.g. that the gradient has been clipped and noised under a budget $\varepsilon \leq 1.0$), drops any payload that resembles a raw record, and emits an audit event referencing the policy and the round identifier. The `prohibition` clause acts as a default deny: any action not explicitly permitted, and in particular any attempt to export raw records, is refused at the PEP.

This example illustrates two architectural points that apply more generally. First, the policy carries enough structure to be *enforced*, not only *displayed*: each duty is bound to an action that the PEP can recognise and evaluate. Second, the policy is the binding artefact between regulatory requirements (R8) and architectural mechanisms (R1, R4, R7): the same document is the unit of governance, the unit of compliance argument and the unit of audit. We make no claim that this minimal policy is sufficient for any specific clinical or industrial deployment; concrete profiles must be elaborated by the relevant sectoral communities.

6. Representative usage scenarios

This section illustrates how the architecture supports concrete problem settings. The scenarios are explicitly *representative* and *illustrative*: they describe how the architectural functions described in Section 4 would be exercised, but they do not report deployed measurements or empirical results. Their purpose is to make the architecture concrete and to expose the design choices that follow from different combinations of data, latency and policy requirements. Table 5 maps the two scenarios to the architectural functions they primarily exercise.

6.1 Scenario A: Crisis management

Crisis management combines two qualitatively different data regimes in a single operational picture. On the one hand, command and coordination benefit from large, relatively static geospatial assets — typically satellite imagery, base cartography and pre-event analyses — which are data-heavy but tolerate higher latencies. On the other hand, field operations generate low-latency dynamic streams — alerts, observations, voice and video — that must reach decision makers and peer responders within seconds. Both regimes are subject to strict governance requirements: provenance, traceability and access control are not optional, and the regulatory regime[4,5] continues to apply during emergencies. The discussion that follows is conceptual: it considers what the architecture would do in such a setting, without making claims about a specific deployment or about real-world incident data (see also the Ethics and reuse statement).

The architecture supports this scenario as follows. Geospatial assets are exposed by the corresponding DSP through its NDA, with self-descriptions that include resolution, footprint, sensor lineage and usage policies, and registered in the DSCS Data Asset Registry. Heavy assets are stored once at the publisher and made available to authorised consumers through the federated catalogue, with content-addressable references where appropriate; the Data Plane is responsible for the actual transfer, while the Control Plane records contract instantiation, policy decisions and access events on the audit log, following the interaction sequence of Figure 2. Field-side applications interact with a local NDA close to the edge and rely on the NIA's edge provisioning to host inference and pre-processing services with low latency. Federated learning components allow models for tasks such as object detection in drone or CCTV streams, or text-from-speech extraction from radio exchanges, to be improved with field experience without exporting raw recordings. Sensitive transformations, such as the anonymisation of bystanders in field imagery, are expressed as policy obligations of the kind shown in Listing 1 and enforced by the PEP before any data leaves the edge.

The combination of policy-anchored audit logging, fine-grained authorisation through DIDs and verifiable credentials, and Zero-Trust network assumptions means that the operational picture can be shared across participating agencies without dissolving accountability. The scenario also illustrates the importance of the Control Plane/Data Plane distinction: governance overhead is concentrated in the Control Plane, while latency-critical operations on the Data Plane are kept close to the edge.

6.2 Scenario B: Cross-organisational federated analytics over heterogeneous compute resources

The second scenario captures settings in which several organisations wish to collaboratively derive insight from data that none of them is willing or allowed to share in raw form. Plausible illustrations include energy lifecycle modelling combining historical consumption and weather data across producers, distributors and large consumers, and scientific analytics in which complementary cohorts contribute to shared predictive models that depend on HPC for training.

In the proposed architecture, each participating organisation operates its own NDA and exposes governed views of its data, together with policies of the form shown in Listing 1 that authorise federated computation but forbid raw export. A joint analytical campaign is initiated by negotiating a contract that references those policies and specifies the federated learning protocol.[29,30] Local FL agents train models on local data using compute resources provisioned by the local NIA; only model updates or aggregated statistics are transmitted, through the federation fabric depicted in Figure 3, to a global FL coordinator hosted by the DSCS Decentralised Learning service. Where the workload demands more capacity than a participant can supply locally — as in scientific analytics involving simulation, large-model training or hybrid quantum–classical exploration — the NIA can map workloads to HPC and quantum back-ends, while the policy attached to the data ensures that placement decisions remain jurisdictionally and contractually valid.

This scenario foregrounds three properties of the architecture. First, sovereignty is preserved by construction: raw data does not leave the publisher and policy is enforced before any release, including the release of intermediate model state. Second, federation is more than data transfer: it includes the joint training of models, mediated by Control Plane contracts and observed by audit logs. Third, the cloud–edge–HPC–quantum continuum is a first-class concern: the same architecture supports both lightweight federated training across small edge devices and heavy federated training that depends on HPC, without requiring a different governance model.

7. Discussion

7.1 Comparative analysis with related data-space approaches

The proposed architecture is best interpreted as an integration model that composes existing European data-space principles into one deployable governance and operations stack. Its comparative contribution is not to replace IDS-RAM, GAIA-X, DSSC, or Simpl, but to connect their trust and interoperability assumptions to an explicit Control Plane/Data Plane separation and to infrastructure-orchestration decisions across cloud, edge, HPC, and quantum resources.

Table 6 provides a function-level comparison with the reference initiatives most closely related to this work. The comparison is conceptual and does not constitute formal conformance certification.

In comparative terms, the architecture remains closest to IDS-RAM[8,9] for contract-bound usage control, to GAIA-X[10,11] for self-descriptions and trust services, to the DSSC Blueprint[12] for governance vocabulary, and to Simpl[13] for middleware-oriented federation practice. The main differentiator is that infrastructure placement is treated as a policy-bearing decision and not as a downstream operational concern. This matters when workloads must be shifted between sovereign domains or between edge and HPC environments without violating contractual and jurisdictional constraints.

7.2 Discussion of scenario results

The two scenarios expose complementary architectural outcomes. In Scenario A, the architecture supports low-latency operational exchange while preserving traceability through policy-linked audit events. The result is not merely secure transport, but accountable data sharing in which access decisions, obligations, and enforcement points remain inspectable after operation. In Scenario B, the architecture enables federated model training across heterogeneous compute capacities without requiring raw data centralisation, which confirms that sovereignty and distributed analytics can coexist when contract negotiation, identity, and transfer orchestration are treated as first-class governance functions.

The policy specification exercise (Listing 1) also has a methodological implication. It demonstrates that governance rules can be rendered machine-enforceable in the same lifecycle that manages data exchange and computation, reducing the gap between legal intent and runtime enforcement. This integration is particularly relevant for cross-border contexts where accountability requirements are strict and where the burden of proof cannot be deferred to post hoc procedural reports.

7.3 Limitations and future research lines

Several limitations of the present work should be made explicit. First, the analysis is conceptual: the article does not present prototype deployments, benchmarks or formal verification, and the architecture should not be interpreted as having been empirically validated. Second, the alignment with reference initiatives is at the level of architectural function; standards conformance would require formal assessment by the relevant bodies. Third, the realisation of semantic interoperability depends on the availability of domain ontologies and shape constraints typically developed by sectoral communities and outside the scope of an architectural article. Fourth, the governance model assumes that participants honour their commitments to the DSCS, including the operation of conformant PEPs; the architecture provides *technical* hooks for verification (audit logs, signed self-descriptions, ledger anchoring) but does not, by itself, replace the legal and organisational arrangements that ultimately underpin trust. Fifth, the choice of distributed ledger anchoring is presented as proportionate and optional; the trade-offs between different audit substrates — centralised signed logs, permissioned ledgers, public ledgers — depend on sectoral risk profiles and are not fully resolved here.

These limitations define a concrete research agenda. The next stage requires interoperable pilots combining NDA, NIA and DSCS services in real cross-organisational settings, with protocol-level validation based on current EDC and Dataspace Protocol implementations.[14,15] It also requires formal analysis of policy composition and enforcement chains, plus performance studies that quantify Control Plane overheads under realistic workloads. When AI-enabled services are part of the data space, evaluation criteria should explicitly connect legal compliance and risk management, including the AI Act and operational risk frameworks such as NIST AI RMF.[4,21]

8. Conclusions and future work

This article has proposed and analysed a conceptual federated, policy-governed architecture for European data spaces, derived from the architectural work of the NOUS project. The architecture is structured as a layered model around five interacting roles — DSP, DSIP, DSCS, NDA and NIA — and is organised internally by an explicit separation between a Control Plane (identity, policy, contract, audit) and a Data Plane (transfer, storage, computation). The model integrates semantic interoperability, data sovereignty, federation across the cloud–edge–HPC–quantum continuum and decentralised governance within a single coherent design, and it is positioned with respect to GAIA-X, IDS-RAM, DSSC and Simpl-Open through a function-level mapping. Two representative scenarios — crisis management with heterogeneous geospatial and field data, and cross-organisational federated analytics — illustrate how the architectural functions are exercised under different data, latency and governance regimes.

The contribution is conceptual and should therefore be treated as an architectural baseline rather than as a validated operational blueprint. The immediate next step is to execute interoperable pilots of NDAs, NIAs, and DSCS governance services across heterogeneous participants, combining open implementations such as Eclipse Dataspace Components with current Dataspace Protocol bindings to observe contract lifecycle behaviour, identity propagation, and failure

handling in realistic conditions.[14,15] In parallel, policy specifications and enforcement chains should be subjected to formal analysis, while performance studies quantify Control Plane overheads under different latency, scale, and workload profiles. Sectoral instantiations should co-design semantic profiles with domain communities and include explicit risk-management procedures for AI-enabled pipelines aligned with the AI Act and established AI risk frameworks.[4,21]

Ethics and reuse statement

The article reports a conceptual architectural analysis. No personal data, no real emergency-response data, no patient data and no operationally sensitive data were processed in the preparation of this work. The illustrative scenarios in Section 6 and the policy example in Section 5 use synthetic identifiers and are intended solely to expose architectural properties; they are not derived from, and should not be interpreted as describing, any specific real-world deployment. Because no human participants and no personal data were involved, no ethics approval and no informed consent were required for this study. Any future empirical study of the architecture will require an ethics review appropriate to the data involved (in particular, in clinical, public-safety or telecommunications contexts).

Data and software availability

Source data

The article analyses no dataset and no data have been generated or collected for it. The only external material on which the architectural analysis draws is the reference architecture of the Horizon Europe NOUS project, documented in deliverable D6.1, *The NOUS Architecture and Open-Source Development Platform.*[35] That deliverable is previously published project documentation — bibliographic material rather than a dataset: it contains no data records, no personal data and no operationally sensitive material.

No ethical, legal, commercial or security restriction applies to the material underpinning this article, and no restriction has been requested by any party. No Institutional Review Board (IRB) or research ethics committee review was required or sought, because the work involved no human participants, no personal data, no patient data and no real emergency-response data (see the Ethics and reuse statement below); consequently, no IRB has issued any statement on data sharing for this study.

Deliverable D6.1 is openly available, with no registration, embargo or other restriction, from the Zenodo repository under a Creative Commons Attribution 4.0 International licence (CC BY 4.0): https://doi.org/10.5281/zenodo.18622569. The other public deliverables of the project are reachable through the European Commission CORDIS record for Grant Agreement No. 101135927 (https://cordis.europa.eu/project/id/101135927). Should a reader or reviewer nevertheless be unable to retrieve any document cited here, a copy may be requested by e-mail from the corresponding author, Paula Lamo (paula.lamo@unican.es), or from the coordinator of the NOUS project, Ricardo S. Alonso (ralonso@air-institute.com); a request need only state the requester's name, affiliation and intended use, and copies are supplied free of charge, normally within ten working days.

Underlying data

No data are associated with this article. The study is a conceptual architectural analysis and reports no experiment, no measurement campaign, no prototype deployment and no processing of primary data of any kind. There are therefore no underlying datasets to deposit, and no data supporting the results have been withheld from publication. All material needed to evaluate the claims made in the article is contained in the article itself: the five-layer architecture and its role decomposition (Section 4), the requirements-to-decisions traceability matrix (Table 2), the illustrative ODRL policy (Listing 1), the mapping of scenarios to architectural functions (Table 5) and the function-level comparison with European reference initiatives (Table 6).

Extended data

No extended data are associated with this article. The illustrative ODRL policy of Listing 1 is reproduced in full in the article and uses synthetic identifiers only; it may be reused, together with the tables and figures, under the terms of the Creative Commons Attribution 4.0 International licence (CC BY 4.0) applied to this article.

Software availability

No new software was produced for this article, and no source code underlies its results. The open-source components discussed in the analysis are publicly available from their respective projects under open licences: Eclipse Dataspace Components (https://projects.eclipse.org/projects/technology.edc), the Dataspace Protocol specification (https://docs.internationaldataspaces.org/ids-knowledgebase/dataspace-protocol), Simpl-Open (https://digital-strategy.ec.europa.eu/en/policies/simpl) and Open Policy Agent (https://www.openpolicyagent.org).

Acknowledgements

The authors thank the partners of the NOUS consortium for the architectural materials and use-case discussions that underpin this analysis. The authors are responsible for any errors or omissions.

References

1. European Commission: *A European Strategy for Data.* Communication COM; 2020. 66 final, 2020.
2. European Parliament and Council: **Regulation (EU) 2022/868 on European data governance (Data Governance Act).** *Off. J. Eur. Union.* 2022; **L 152**: 1–44.
3. European Parliament and Council: **Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data (Data Act).** *Off. J. Eur. Union.* 2023; **L 2854**.
4. European Parliament and Council: **Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act).** *Off. J. Eur. Union.* 2024; **L 2024/1689**.
5. European Parliament and Council: **Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation).** *Official Journal of the European Union.* 2016; **L 119**: 1–88.
6. European Commission, ISA2 Programme: **European Interoperability Reference Architecture (EIRA), version 5.1.0.** 2022.
7. European Commission: **European Blockchain Services Infrastructure (EBSI) – programme overview.** 2024.
8. International Data Spaces Association: *Reference Architecture Model, Version 4.* IDSA; 2023.
9. Pampus J, Bader B, Steinbuß S, *et al.*: *Usage Control in the International Data Spaces, Position Paper Version 3.0.* International Data Spaces Association; 2021.
10. Gaia-X European Association for Data and Cloud AISBL; **Gaia-X Architecture Document, 25.05 Release.** 2025. (accessed 2026).
Reference Source
11. Gaia-X European Association for Data and Cloud AISBL: **Gaia-X Compliance Document/Trust Framework, 24.04-prerelease.** 2024. (accessed 2026).
Reference Source
12. Data Spaces Support Centre: **Blueprint 3.0: Europe's shared reference architecture for data spaces. Version 3.0.**(accessed 2026)
Reference Source
13. European Commission, DG CNECT: **Simpl: Cloud-to-edge federations empowering EU data spaces.**Last update: 11 October 2024. (accessed 2026).
Reference Source
14. Eclipse Foundation: **Eclipse Dataspace Components (EDC): project overview and documentation.** 2024. (accessed 2026).
Reference Source
15. International Data Spaces Association and Eclipse Foundation: **Dataspace Protocol 2024–1 (stable release candidate).** 2024. (accessed 2026).
Reference Source
16. Iannella R, Villata S, editors. *ODRL Information Model 2.2.* W3C Recommendation; 2018.
17. Sporny M, Longley A, Sabadello M, *et al.*, editors. *Decentralized Identifiers (DIDs) v1.0.* W3C Recommendation; 19 July 2022.
18. Herman I, Jones M, Sporny M, *et al.*, editors. *Verifiable Credentials Data Model v2.0.* W3C Recommendation; 15 May 2025.
19. Knublauch H, Kontokostas D, editors. *Shapes Constraint Language (SHACL).* W3C Recommendation; 20 July 2017.
20. Rose S, Borchert O, Mitchell S, *et al.*: **Zero Trust Architecture. NIST Special Publication 800-207.** 2020.
21. National Institute of Standards and Technology: *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* NIST AI 100–1; 2023.
Publisher Full Text
22. Open Policy Agent Project: *Open Policy Agent: specification and documentation.* Cloud Native Computing Foundation; 2024. (accessed 2026).
Reference Source
23. Otto B, ten Hompel M, Wrobel S, editors. *Designing Data Spaces: The Ecosystem Approach to Competitive Advantage.* Springer; 2022.
Publisher Full Text
24. Curry E, Scerri S, Tuikka T, editors. *Data Spaces: Design, Deployment and Future Directions.* Springer; 2022.
Publisher Full Text
25. Otto B: **A federated infrastructure for European data spaces.** *Commun. ACM.* 2022; **65**(4): 44–45.
Publisher Full Text
26. Solmaz G, Cirillo F, Fürst F, *et al.*: **Enabling data spaces: existing developments and challenges.** *Commun. ACM.* 2022; **65**(11): 66–75.
27. Hummel P, Braun M, Tretter M, *et al.*: **Data sovereignty: A review.** *Big Data Soc.* 2021; **8**(1): 1–17.
28. Jung C, Dörr A: **Data Usage Control.** Otto B, ten Hompel M, Wrobel S, editors. *Designing Data Spaces.* Springer; 2022; pp. 129–146.
Publisher Full Text
29. Kairouz P, McMahan HB, Avent B, *et al.*: **Advances and Open Problems in Federated Learning.** *Foundations and Trends in Machine Learning.* 2021; **14**(1–2): 1–210.
Publisher Full Text
30. Li T, Sahu AK, Talwalkar A, *et al.*: **Federated Learning: Challenges, Methods, and Future Directions.** *IEEE Signal Process. Mag.* 2020; **37**(3): 50–60.
31. Bonawitz K, Eichner H, Grieskamp W, *et al.*: **Towards Federated Learning at Scale: System Design.** *Proc. Conference on Machine Learning and Systems (MLSys).* 2019.
32. Milojicic D: **The Edge-to-Cloud Continuum.** *IEEE Computer.* 2020; **53**(11): 16–25.
Publisher Full Text
33. Beckman P, Dongarra J, Ferrier N, *et al.*: **Harnessing the Computing Continuum for Programming Our World.** *Fog Computing: Theory and Practice.* Zomaya M, Abawajy J, Buyya A, editors. Wiley; 2020; pp. 215–230.
34. Mühle A, Grüner A, Gayvoronskaya T, *et al.*: **A Survey on Essential Components of a Self-Sovereign Identity.** *Comput Sci Rev.* 2018; **30**: 80–86.
Publisher Full Text
35. Moraitis P, Zervoudakis P: *D6.1 The NOUS Architecture and Open-Source development platform (1st version).* Project deliverable, NOUS Consortium, Horizon Europe Grant Agreement No. 101135927; 31 December 2025.
Publisher Full Text