ACCURATE

D5.3 Sovereign Data Sharing Integrated components for secure, privacy-preserving, and compliant data sharing among project participants

Actual Submission Date: **31/01/2026**
Produced by: DAO: deltaDAO AG

Accurate

https://accurateproject.eu/

HORIZON-CL4-2023-TWIN-TRANSITION-01

Grant Agreement no.: 101138269

Start date of project: 01 12 2023 - Duration: 36 months

Funded by the European Union

The ACCURATE project is funded by the European Union, under Grant Agreement number 101138269. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

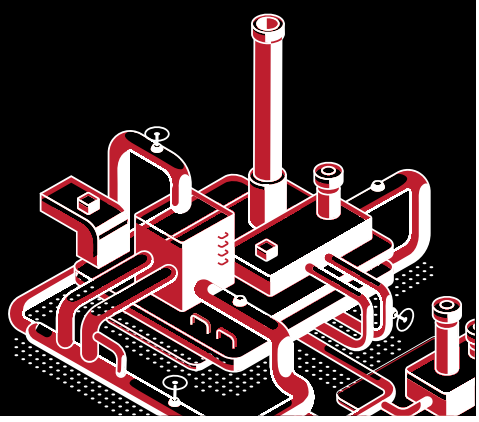

DELIVERABLE FACTSHEET

Deliverable D5.3	
Nature of the Deliverable:	Demonstrator
Due date of the Deliverable:	M26 – 31/01/2026
Actual Submission Date:	M26 – 31/01-2026
Produced by:	DAO: deltaDAO AG
Contributors:	Thomas Komenda
Work Package Leader Responsible:	DAO: deltaDAO AG
Reviewed by:	Shortname: AU, IAO

Dissemination level	
PU	PU = Public
	PP = Restricted to other programme participants (including the EC)
	RE = Restricted to a group of the consortium (including the EC)
	CO = Confidential, only members of the consortium (including the EC)

Contents

Figures

Terms and abbreviations

Term / Abbreviation	Definition
ACCURATE	The European Horizon project pioneering a decentralized data space infrastructure for the manufacturing sector.
Aquarius	The metadata cache component that indexes data assets to make them discoverable without exposing the raw data itself.
CLI	Command Line Interface. A text-based interface used to interact with software and operating systems (e.g., kubectl).
CtD	Compute-to-Data. A privacy-preserving mechanism where algorithms are executed on data within the data owner's secure infrastructure, preventing raw data exposure.
Data Space Authority	A democratically elected governance body responsible for manually verifying the legitimacy of participants before they are admitted to the ecosystem.
DDO	Decentralized DID Document. A document containing metadata about a data asset or participant, stored on the distributed ledger.
DID	Decentralized Identifier. A globally unique identifier that allows participants to prove ownership and identity without a central registry.
D5.3	Deliverable 5.3. The project deliverable focusing on Sovereign Data Sharing, including component deployment and governance integration.
EUROe	A European Union-regulated digital Euro used for financial settlements within the ACCURATE ecosystem.
Gaia-X	A European initiative establishing a federated and secure data infrastructure based on values of openness, transparency, and sovereignty.
GXDCH	Gaia-X Digital Clearing House. A service that automatically validates participant credentials against the Gaia-X Trust Framework.
Identity Wallet	A software or hardware tool that manages a participant's public-private key pairs, enabling them to sign transactions and manage digital assets locally.
K8s	Kubernetes. An open-source system for automating the deployment, scaling, and management of containerized applications, used as the foundation for Sovereign Nodes.
MaaS	Manufacturing as a Service. A business model where manufacturing capabilities are offered as a service, enabled by secure data sharing in the ACCURATE ecosystem.
Ocean Provider	The technical component acting as the API gateway for a Sovereign Node, enforcing access control and managing data connections.
Operator Engine	The orchestration component within the Compute-to-Data environment that executes compute jobs in isolated containers.
Operator Service	The microservice that manages the queue of compute jobs and coordinates between the Ocean Provider and the Operator Engine.
Pontus-X Identity	A verifiable credential issued to a participant after successful validation, serving as their digital passport within the ecosystem.
Pontus-X Registry	A decentralized smart contract acting as the "single source of truth" that maps wallet addresses to verified legal identities.
Private Key	A cryptographic key generated locally that allows a participant to sign transactions. It must never leave the participant's security perimeter.
Sovereign Node	A self-contained infrastructure unit hosted by a participant, consisting of the Ocean Provider and Compute-to-Data environment, ensuring full control over data.
SSI	Self-Sovereign Identity. An identity model where individuals and organizations own and control their digital identities without intervening administrative authorities.
VC	Verifiable Credential. A tamper-evident digital credential that proves specific attributes (e.g., identity, certification) of a participant.

Public Summary

The ACCURATE project is pioneering a decentralized data space infrastructure to enhance the resilience, competitiveness, and sustainability of the European manufacturing sector [1]. Following the establishment of the network architecture (D5.1) and the privacy-preserving computation framework (D5.2), Deliverable D5.3 "Sovereign Data Sharing" provides the operational blueprint for participants to transition from passive users to active, sovereign operators within the ecosystem.

This deliverable details the integrated components and processes required for technical data sovereignty, enabling manufacturing partners to share data and services while retaining absolute physical and legal control over their assets. Central to this approach is the deployment of sovereign components—specifically the Ocean Provider and Compute-to-Data environments—within the participant's own infrastructure. This decentralized architecture ensures that sensitive data is processed within the data owner's security perimeter, eliminating reliance on central intermediaries.

To establish a trusted environment for Manufacturing as a Service (MaaS), D5.3 introduces a robust end-to-end integration process based on the Gaia-X Trust Framework [2]. This includes:

- Self-Sovereign Onboarding: Participants generate their own cryptographic keys locally, ensuring that their digital identity remains under their exclusive control.
- Democratic Verification: A "Human-in-the-Loop" trust model where a democratically elected Data Space Authority manually verifies the legitimacy of participants before issuing a Pontus-X Identity.
- Transparent Discovery: The decentralized Pontus-X Registry allows ecosystem members to transparently verify the identity and compliance status of potential partners before engaging in collaboration.

Furthermore, the deliverable operationalizes data governance through Standard Data Usage Agreements. These legally vetted templates are integrated directly into the technical workflow and enforced by smart contracts, significantly reducing the administrative friction of compliant data sharing. By combining secure local infrastructure with transparent, community-governed trust mechanisms, the ACCURATE project empowers European manufacturers to collaborate confidently, complying with the European Data Act while protecting their trade secrets and intellectual property.

1 Introduction

1.1 About this Deliverable

This deliverable**,** D5.3 Sovereign Data Sharing, marks the transition in the ACCURATE project from architectural design to operational implementation and onboarding. While previous deliverables established the network infrastructure (D5.1) and the conceptual framework for privacy-preserving computation (D5.2), this document serves as the operational handbook for technical data sovereignty and participant integration. It provides the necessary guidelines, technical specifications, and governance frameworks for consortium partners to securely onboard, deploy, and operate their own independent access points to the data space.

The core objective of this deliverable is to empower partners to enable full data sovereignty. True data sovereignty requires that participants retain absolute authority over their identity and data assets. Consequently, this deliverable focuses on three key pillars:

1. **End-to-End User Integration and Identity Control:** Focusing on the secure onboarding process based on the Gaia-X framework. This involves the local creation of Identity Wallets and public-private key pairs, ensuring that control over digital identity remains entirely within the participant's infrastructure.
2. **Operational Sovereignty (Technical Deployment):** Providing detailed instructions for self-hosting critical infrastructure components, specifically the Ocean Provider [3] and the Compute-to-Data (CtD) [4] environment. This ensures that data access is validated and processed within the partner's own security perimeter.
3. **Transparent Governance and Compliance:** Utilizing the decentralized Pontus-X Registry to provide full transparency regarding ecosystem participants and their public addresses. This enables partners to make conscious, informed decisions about whom to collaborate with and under which conditions. Additionally, this deliverable includes the preparation of Data Usage Agreement templates to facilitate compliant and seamless data sharing.

By following the guidelines in this deliverable, ACCURATE partners will be equipped to participate in the manufacturing data space while maintaining full control over their proprietary assets, consistent with Gaia-X principles and the European Data Act [5].

1.2 Document Structure

This deliverable is structured as follows:

- **Chapter 1: Introduction:** Outlines the scope, objectives, and the document's role within the ACCURATE project.
- **Chapter 2: End-to-End User Integration:** Details the onboarding process aligned with the Gaia-X framework. It covers the local generation of public-private key pairs for Identity Wallets and the role of the decentralized Pontus-X Registry. This registry allows participants to transparently verify the identity of ecosystem members (via public addresses) and make informed collaboration decisions.
- **Chapter 3: Deployment Guidelines – Infrastructure & Operations:** A technical guide for deploying the necessary components. This includes prerequisites, Kubernetes configurations, and specific steps for installing the Ocean Provider and the secure Compute-to-Data environment.
- **Chapter 4: Data Governance and Usage Agreements:** Focuses on the organizational aspects of data sharing. It includes frameworks for defining Data Usage Rules and provides a standard Data Usage Agreement template to ensure compliant and efficient contracting between participants.

- **Chapter 5: The Sovereign Data Sharing Process:** Visualizes the complete workflow of a data transaction, from the transparent discovery of partners in the registry to the secure execution of algorithms.
- **Chapter 6: Conclusion:** Summarizes the achievements of this task and outlines the next steps for system integration.

1.3 Relation with Other Tasks and Deliverables

This deliverable serves as the operational integration layer for Work Package 5 (WP5), effectively instantiating the architectures defined in previous tasks into a deployable reality.

- **Relation to D5.1 – Decentralized Data Space Infrastructure:** D5.1 established the foundational network layer and the distributed ledger. D5.3 builds upon this by describing how a partner creates the local identity (keys) to access this infrastructure and how they register their node in the ecosystem registry defined in D5.1.
- **Relation to D5.2 – Compute-to-Data Environment:** D5.2 provided the conceptual architecture for privacy-preserving computation. D5.3 translates this into specific technical deployment steps, guiding partners on how to stand up their own secure compute environments using the defined software stack.
- **Relation to Task 5.4 – System Integration and Decentralization:** The sovereign nodes and identities established using the guidelines in this deliverable will form the distributed network that is further integrated and tested in Task 5.4. This task will ensure that the disparate nodes deployed by different partners can interoperate seamlessly.
- **Relation to Task 6.2 – Establishing Trust with SSI:** While D5.3 focuses on the operational setup of wallets and the registry for transparency, Task 6.2 will expand this into a full Self-Sovereign Identity (SSI) trust framework with advanced credential verification. The onboarding processes defined here (local key generation) form the technical prerequisite for the advanced trust mechanisms in WP6.

2 End-to-End User Integration

2.1 Onboarding based on the Gaia-X Framework

The onboarding process in the ACCURATE data space is designed to ensure compliance with the Gaia-X Trust Framework while maintaining strict data sovereignty and robust governance. The ACCURATE ecosystem implements a "Human-in-the-Loop" trust model for initial participant admission, ensuring that all ecosystem members are vetted legitimate entities.

To facilitate this, the project provides a User-Friendly Onboarding Wizard [6]. This graphical interface guides participants through the process of generating their digital identity and submitting it for governance review without requiring deep knowledge of the underlying cryptography.

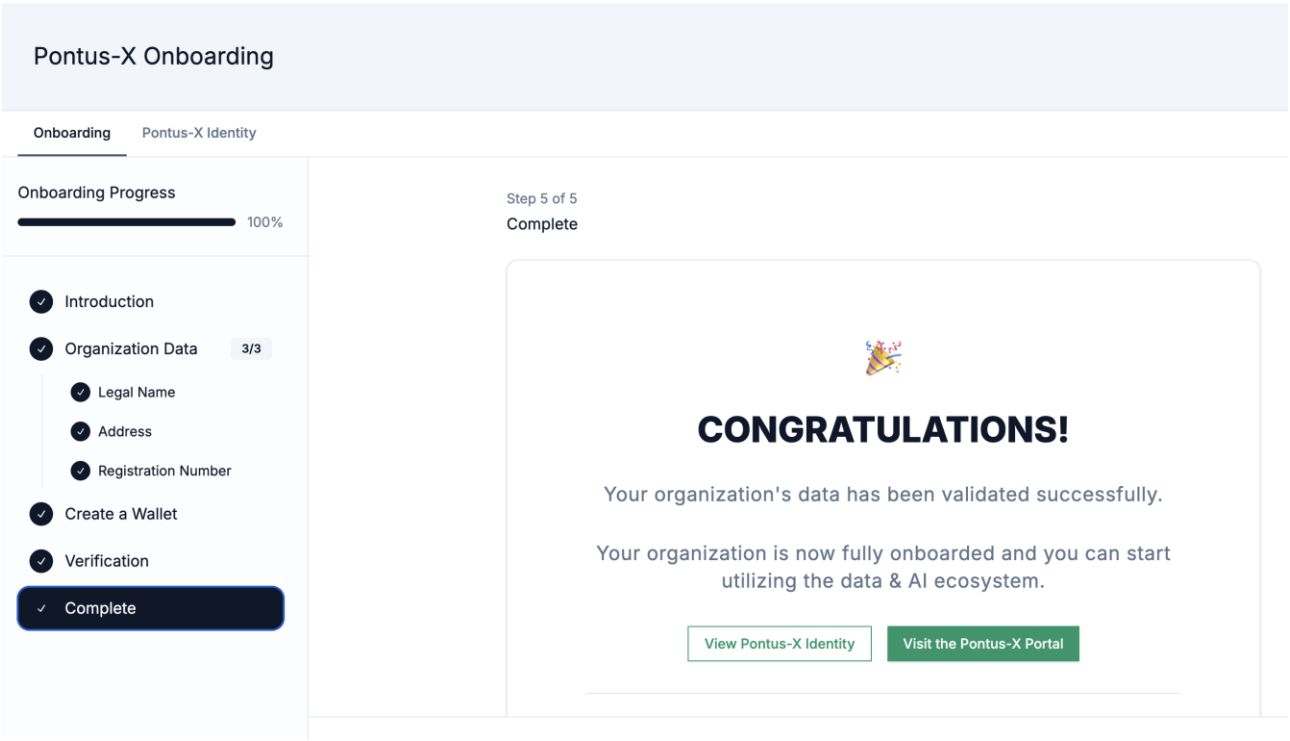

Figure 1: Onboarding Wizard and Flow (Source: deltaDAO)

2.1.1 The Onboarding Flow

As illustrated in the ecosystem's onboarding interface, the process is structured into a clear workflow designed to capture compliance data and facilitate the governance review:

1. **Introduction:** An overview of the principles, governance rules, and requirements for participation in the ecosystem.
2. **Organization Data:** The participant inputs their verifiable legal information, specifically:
    - **Legal Name:** The official registered name of the entity.
    - **Address:** The registered physical headquarters.
    - **Registration Number:** The national business identifier (e.g., VAT ID, EORI, or LEI).
3. **Create a Wallet:** The wizard facilitates the local generation of the identity wallet (see Section 2.2), ensuring keys are created on the user's device and remain under their control.
4. **Verification:** The participant signs their data with their new private key and submits a verification request. This request is routed to the Data Space Authority [7] for manual review.
5. **Completion & Pontus-X Identity:** Once the Authority has successfully verified the user manually, they are issued with a Pontus-X Identity containing a valid Gaia-X Participant Credential and a connection to the Pontus-X Identity Wallet. This serves as their digital passport in the ecosystem.

2.1.2 Governance based Verification by the Data Space Authority

A defining feature of the ACCURATE trust framework is the role of the Data Space Authority in the verification process. ACCURATE relies on a transparent governance layer to establish the root of trust.

- **Manual Verification:** When a participant submits their organization data, it is not simply accepted by an algorithm. The Data Space Authority manually reviews the submission, cross-referencing the provided Legal Name and Registration Number against official business registers (e.g., National Trade Registers, EU VIES). This ensures that no fraudulent or non-existent entities can enter the ecosystem.
- **Democratic Governance:** The Data Space Authority is not a centralized owner but a representative body. It is democratically elected by the ecosystem participants in a transparent governance process. This ensures that the power to admit or reject new members resides with the community itself, preventing vendor lock-in or monopolistic control. For now, deltaDAO AG, the leader of work package 5, is taking on the role, which can at any time be put to a vote by the ACCURATE general assembly. To facilitate democratic, immutable and transparent governance processes, a voting tool has been developed to support governance decisions [15].

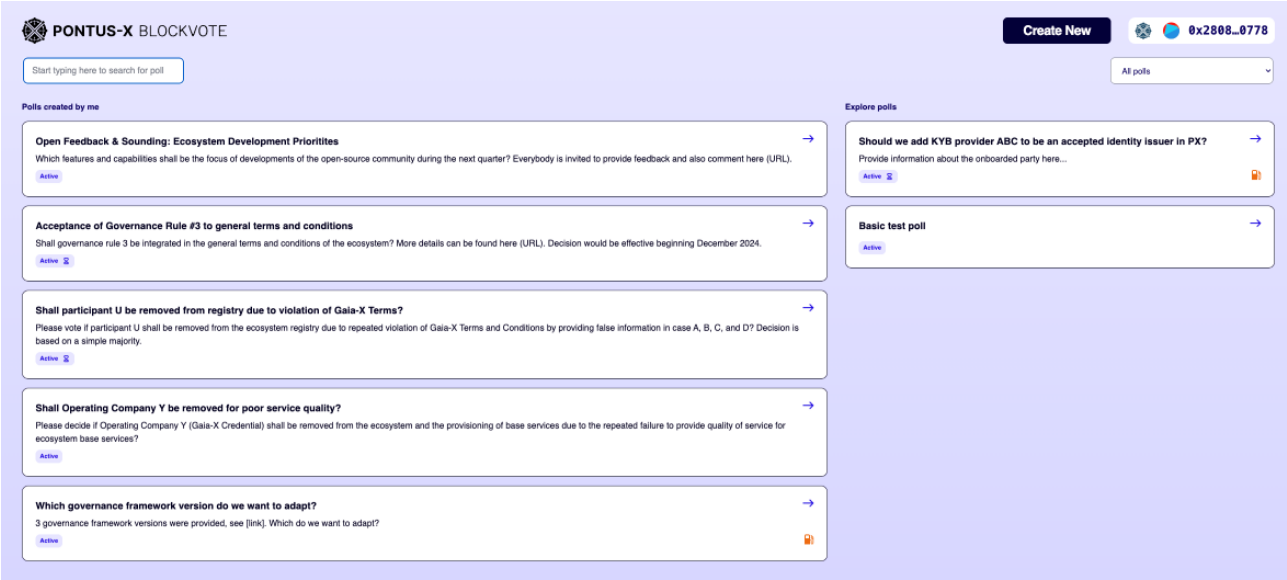

Figure 2: Data Space Governance Voting Tool (Source: deltaDAO)

- **Issuance of Credentials:** Once the Authority validates the application, and it's cryptographically signed by the GXDCH [11], they receive their Pontus-X Identity, including the Gaia-X Participant Credential. This action essentially "notarizes" the participant's digital identity, allowing them to interact with other participants with a high level of assurance.

2.2 Identity Creation: Local Key Generation

True technical sovereignty begins with the generation of cryptographic keys during the "Create a Wallet" step of the onboarding. In the ACCURATE ecosystem, a participant's identity is anchored by a Public-Private Key Pair.

To ensure that no third party, not even the Data Space Authority, has control over a participant's identity or assets, these keys must be generated locally within the participant's infrastructure.

2.2.1 The Generation process

Participants can use the onboarding interface to generate a key pair based on the *secp256k1* elliptic curve standard [8].

- The Private Key: This is the "root of control." It allows the participant to sign legally binding agreements (smart contracts), authorize data transfers, and update their profile. This key must never leave the participant's secure environment.
- The Public Address: Derived from the private key, this alphanumeric string (e.g., 0x123...) acts as the participant's public ID in the network.

By strictly enforcing local generation, ACCURATE ensures:

- Non-Repudiation: Only the holder of the private key could have signed a specific transaction.
- Censorship Resistance: While the Authority grants admission to the *trust framework*, they do not hold the keys to the participant's wallet. They cannot technically confiscate the identity or assets, only revoke the trust credential.

2.3 The Pontus-X Registry: Decentralized Transparency

Once an identity is created locally and verified by the Authority, it must be made visible to potential partners. This is achieved through the Pontus-X Registry [9].

2.3.1 Function of the Registry

The Pontus-X Registry is a decentralized smart contract and database that serves as the "single source of truth" for participant information in the ACCURATE ecosystem. It maps Public Addresses to the verified Pontus-X Identity.

When the Data Space Authority approves an onboarding request, the validation status is updated in the registry. This action publicly associates the participant's localized wallet address with their real-world legal entity in a tamper-proof ledger.

2.3.2 Transparency and Discovery

The registry allows all ecosystem members to transparently view who is part of the network. It provides a searchable ledger containing:

- **Participant Name:** The legal name of the entity (e.g., "deltaDAO AG").
- **Public Address:** The wallet address used for technical interactions.

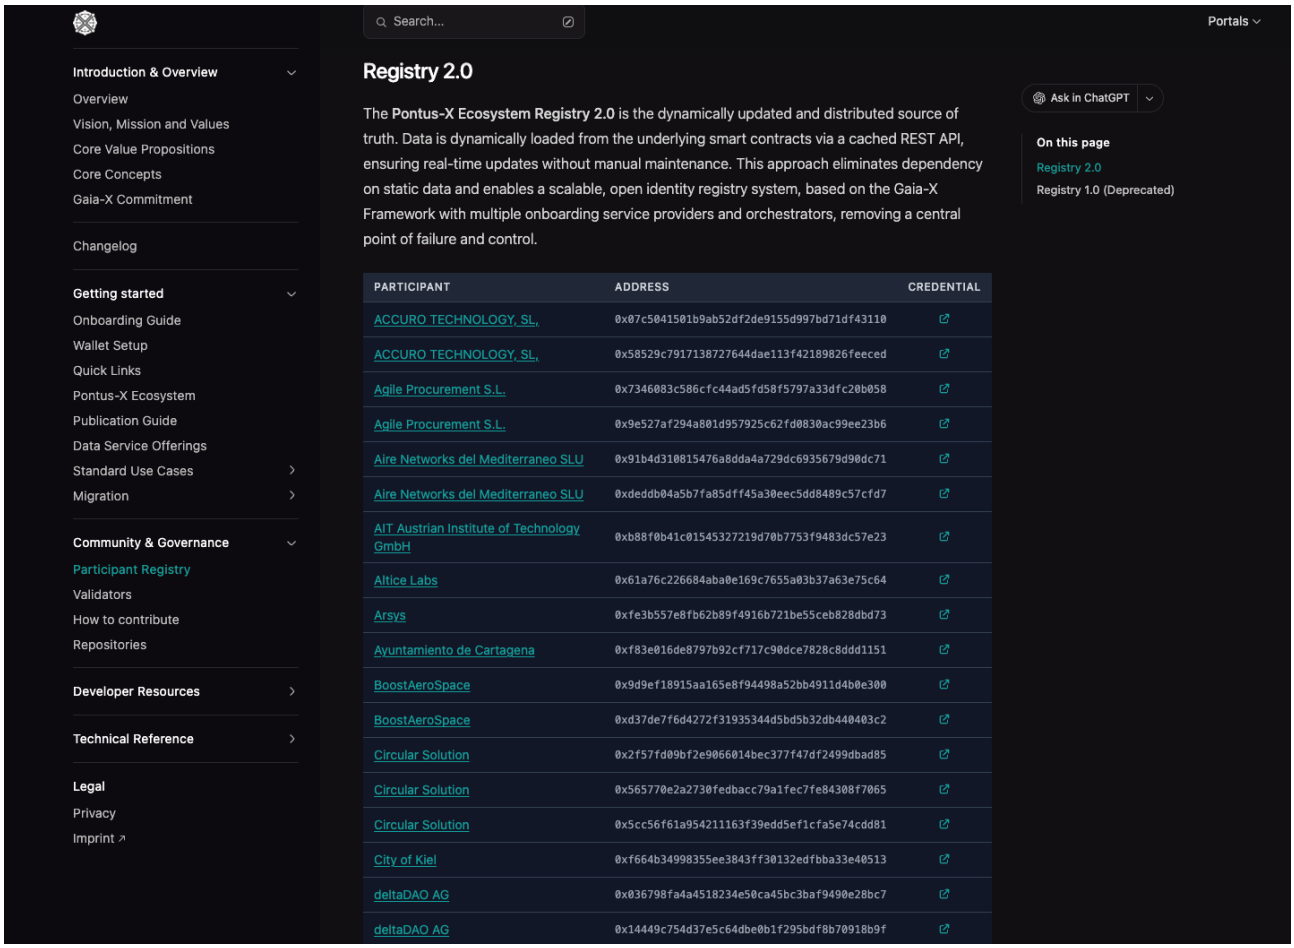

Figure 3: Pontus-X Registry (Source: deltaDAO)

The Pontus-X Registry advances decentralized trust architectures by replacing centralized Public Key Infrastructures with a transparent, smart contract-based Trust Anchor that cryptographically binds anonymous wallet addresses to verified Gaia-X credentials.

3 Deployment Guidelines – Infrastructure & Operations

To establish functional sovereign componentd within the ACCURATE ecosystem, consortium partners must deploy a specific set of decentralized software components that transform their existing infrastructure into an active, self-governed access point. Building upon the architectural foundations of D5.1 and D5.2, this deployment consists of two primary layers: the Data Access Layer, anchored by the Ocean Provider , which serves as the secure gateway for authentication and access control, and the Secure Compute Layer, enabled by the Compute-to-Data environment, which acts as an isolated "safe room" for privacy-preserving algorithms. By self-hosting these components on Kubernetes, partners achieve operational data sovereignty, ensuring that raw data remains strictly within their own security perimeter and is only processed or shared according to the specific cryptographic permissions they define.

3.1 Technical Prerequisites

To establish functional sovereign components, partners must first prepare their local infrastructure to interact with the ACCURATE decentralized data space network. The deployment utilizes Kubernetes [10] to ensure scalability and environment isolation.

3.1.1 Command Line Interface (CLI) Setup

The primary tool for managing the Sovereign Node deployment is kubectl. kubectl is the essential command-line interface for managing the complete Kubernetes infrastructure, facilitating communication between the user and the cluster's resources by transmitting HTTP requests directly to the Kubernetes API server. Partners must ensure this tool is installed and configured on the machine used for deployment management.

- **Installation:** kubectl should be installed via standard package managers (e.g., apt-get install -y kubectl for Ubuntu or brew install kubectl for Mac).
- **Verification:** The installation must be verified by running kubectl version --client to ensure the correct version is active.

3.1.2 Cluster Access Configuration

Access to the partner's Kubernetes cluster is managed via the kubeconfig file.

1. **Configuration Directory:** Create a .kube directory in the user's home folder (e.g., mkdir -p ~/.kube).
2. **Config File Placement:** The cluster configuration file (e.g., ACCURATE.yaml for the reference implementation) must be copied to this directory and renamed to config.
3. **Connection Test:** Connectivity to the cluster is verified by listing the active nodes using kubectl get nodes.

3.2 Deployment of the Data Access Layer

The Data Access Layer serves as the gateway to the sovereign components. It consists of the **Metadata Cache (Aquarius)** and the **Ocean Provider**. This layer is deployed in a dedicated namespace to ensure logical separation from the compute environment.

3.2.1 Namespace Initialization

All components of the Data Access Layer are deployed into the pontusx-1 namespace. This namespace must be created explicitly prior to deployment:

Bash

```
kubectl create namespace pontusx-1
```

3.2.2 Credential and Secret Management

Security is paramount for the Sovereign Node. Credentials, including the node's **Private Key** and database passwords, are injected into the cluster as Kubernetes Secrets. These secrets must be deployed *before* the application services to ensure successful startup.

- **ConfigMaps:** Deploy general configuration variables using pontusx-ocean-cm.yaml.
- **Private Keys:** Deploy the private-key-pontusx-devnet.yaml secret. This contains the cryptographic key generated in Chapter 2, enabling the node to sign transactions.
- **Database Credentials:** Deploy db-credentials.yaml to secure the metadata database.

3.2.3 Deploying the Metadata Cache (Aquarius)

Aquarius indexes metadata to make the node's assets discoverable without exposing the data itself.

1. **Deployment:** Apply the configuration files for Aquarius and its event monitor (aquarius-ponutsx-devnet.yaml, aquarius-events-pontusx-devnet.yaml).
2. **Database:** Deploy Elasticsearch (elasticsearch.yaml) to serve as the backend for metadata queries.
3. **Verification:** Ensure all pods in the pontusx-1 namespace are in the Running state using kubectl get pods -n pontusx-1.

3.2.4 Deploying the Ocean Provider

The Ocean Provider is the API gateway that enforces access control. It relies on the secrets deployed in step 3.2.2.

1. **Service Deployment:** Deploy the provider service (pontusx-devnet-provider-service.yaml) to expose the application.
2. **Application Deployment:** Deploy the provider application (pontusx-devnet-provider.yaml).
3. **Log Monitoring:** Verify the provider startup sequence by monitoring the logs: kubectl logs <pod-name> -n pontusx-1 -f.

3.3 Deployment of the Secure Compute Environment

Compute-to-Data enables a secure "Code-to-Data" workflow where algorithms are executed locally on sensitive datasets, ensuring that proprietary information never leaves the data owner's infrastructure. For example, in a supply chain context, a lead manufacturer can run a carbon footprint calculation model across a supplier's private logistics data to verify sustainability compliance without the supplier ever needing to disclose their confidential shipment routes or volumes. The Compute-to-Data (CtD) environment acts as the "Safe Room" for data processing. It is deployed in a separate namespace (accurate-ctd-1) with strict network policies to prevent data leakage.

3.3.1 Isolation and Network Policies

To guarantee that raw data cannot be exfiltrated during computation, the environment is secured using Network Policies.

1. **Namespace Creation:** Create the accurate-ctd-1 namespace.
2. **Network Policies:** Apply network-policy.yaml to restrict traffic flow. This ensures that algorithms running inside the compute pods cannot initiate unauthorized outbound connections.
3. **Service Accounts:** Deploy the specific Service Account (sa.yaml) required for the Operator to manage Kubernetes resources.

3.3.2 State Management (PostgreSQL)

The Compute Environment requires a database to track the status of compute jobs (e.g., "Running", "Completed", "Failed").

1. **Persistence:** Deploy the Persistent Volume (postgres-persistent-volume.yaml) to ensure job history is retained across restarts.
2. **Database Deployment:** Deploy the PostgreSQL instance using postgres.yaml.

3.3.3 Deploying the Operator Components

The Operator consists of two parts: the **Engine** (which executes jobs) and the **Service** (which manages the API).

1. **Operator Engine:** Deploy compute-operator-pontusx-devnet.yaml. This component orchestrates the creation of algorithm pods.
2. **Operator Service:** Deploy operator-api-pontusx-devnet.yaml and its associated service.

3.3.4 Environment Initialization

Once deployed, the Operator's database must be initialized to define the required tables and schemas. This is a one-time setup step performed via the API.

1. **Port Forwarding:** Temporarily expose the Operator API to the local machine:

   Bash

   ```
   kubectl port-forward -n accurate-ctd-1 deploy/operator-api-pontusx-devnet 8050:8050
   ```

2. **Initialization Command:** Send a POST request to trigger the database initialization:

   Bash

   ```
   curl -X POST "http://0.0.0.0:8050/api/v1/operator/pgsqlinit" \
   -H "accept: application/json" \
   -H "Admin: ACCURATE"
   ```

3. **Validation:** Check the Operator logs for the message Valid admin to confirm successful initialization.

3.4 System Verification and Handover

After deployment, partners must verify that their Sovereign Node is correctly registered and operational.

3.4.1 Verifying the Provider Status

Partners can confirm their Provider is active by accessing its public endpoint.

1. **Access:** Navigate to the Provider's URL (or localhost if port-forwarded).
2. **Status Check:** The response should return a JSON object containing chainIds and providerAddresses. This confirms the Provider is successfully connected to the ledger network.

3.4.2 Verifying the Compute Environment

To ensure the Compute Environment is ready to accept jobs, check the Operator Engine logs:

Bash

```
kubectl logs -n accurate-ctd-1 deploy/compute-operator-pontusx-devnet
```

A functioning environment will display logs indicating it is polling for jobs, such as Doing SELECT FROM announce. This confirms that the "Safe Room" is active and monitoring the network for authorized compute requests.

4 Data Governance and Usage Agreements

4.1 The Governance-by-Design Approach

In the ACCURATE ecosystem, Data Governance is an integral part of the technical architecture. The sovereign components deployment described in Chapter 3 provides the enforcement mechanism, but the rules enforced by that node must be defined through an internal governance framework.

This chapter outlines how consortium partners can translate their internal business requirements into enforceable Data Usage Rules and details the **Standard Data Usage Agreement**, a legally vetted template provided to all partners to minimize administrative friction.

4.2 Internal Data Classification

Before any data is published to the catalogue, partners must classify their assets to determine the appropriate sharing mechanism. We recommend a three-tier classification model to guide these decisions:

1. **Public / Low-Sensitivity Data:**
    - *Definition:* Data that poses no risk if exposed (e.g., product catalogs, public sustainability reports).
    - *Governance Rule:* Open access or simple registration required.
2. **Protected / Business-Critical Data:**
    - *Definition:* Proprietary data that is valuable but sharable under strict conditions (e.g., supply chain schedules, quality metrics).
    - *Governance Rule:* Access restricted to specific verified partners (Allow-listing); often requires payment or a signed usage agreement.
3. **Confidential / Trade Secret Data:**
    - *Definition:* Highly sensitive data that must never leave the owner's control (e.g., raw machine logs, proprietary production recipes).
    - *Governance Rule:* **Compute-to-Data (CtD) Only.** Raw data access is strictly prohibited. Only the *result* of a specific, approved algorithm may be shared.

4.3 Defining Technical Usage Rules

Once classified, the data owner must define the specific conditions under which access is granted. In the ACCURATE ecosystem, these rules are configured via metadata and are directly enforced by the Ocean Provider, which acts as the automated policy enforcement point.

- **Participant Allow-listing:** Leveraging the Pontus-X Registry, data owners can restrict access to specific wallet addresses or only to participants with valid Gaia-X credentials.
- **Commercial Terms:** Partners define pricing models (e.g., EUROe digital Euro [12]) for data access or compute resource utilization.

4.4 The Standard Data Usage Agreement

To facilitate compliant data sharing without the need to create individual contracts with legal departments for every transaction, the ACCURATE project provides a Standard Data Usage Agreement. Developed with legal consultation, this template serves as the baseline contract for the ecosystem.

This agreement is designed as a hybrid contract, explicitly bridging the gap between traditional legal text and the automated execution of the distributed ledger.

4.4.1 Key Provisions of the Template

The template includes several critical clauses that protect the data owner (Licensor) and clarify the rights of the consumer (Licensee):

- **Smart Contract Acknowledgment:** The agreement explicitly recognizes that parts of the execution (payment, access delivery) are handled by "Smart Contracts". It includes a primacy clause stating that if the code deviates from the written agreement, the written agreement prevails, and the Licensor is obliged to reverse erroneous executions.
- **License Grant & Scope:** The template defines a non-exclusive, non-transferable right to use the data. Crucially, it clarifies that acquiring an "Access Token" grants *usage rights* but does not transfer ownership or intellectual property rights. The scope can be restricted to specific purposes (e.g., "internal business purposes" or "demonstration purposes") and territories (e.g., "limited to EU territory").
- **Confidentiality:** It imposes strict confidentiality obligations, defining "Confidential Information" broadly to include any data provided, regardless of explicit marking. The Licensee is prohibited from disclosing this information to third parties without consent, except where mandated by law.
- **Liability and Warranty:** The Licensor warrants they have the right to grant the license. Liability is limited to cases of intent, gross negligence, or breach of material obligations, protecting the data owner from unlimited liability for minor programming errors.
- **Governing Law:** The agreement is governed by the laws of the Federal Republic of Germany, providing a stable and predictable legal foundation for European partners.

4.4.2 Integration with Onboarding

This legal template is integrated directly into the technical publication workflow:

1. **Hash Linking:** When a partner publishes a data service, the PDF of this agreement is hashed (cryptographic fingerprint).
2. **Metadata Reference:** This hash and a link to the document are embedded in the asset's metadata (DDO [13]).
3. **Digital Signature:** When a consumer purchases the data, the transaction on the distributed ledger serves as an immutable acceptance of these terms.

4.5 Facilitating Compliant Sharing

By using this standardized, legally vetted template, ACCURATE ensures:

- **Reduced Legal Friction:** Partners do not need to draft new contracts for every exchange.
- **Clear Liability:** The "Smart Contract" clause protects partners from the risks of experimental code.
- **Data Protection:** The strict confidentiality clauses ensure that even "shared" data remains legally protected against unauthorized redistribution.

5 The Sovereign Data Sharing Process

The Sovereign Data Sharing Process is the operational realization of the ACCURATE architecture. It transforms the static components—Identity Wallet, Sovereign Node, and Data Usage Agreement—into a dynamic, secure transaction.

Unlike centralized platforms where trust is placed in an intermediary, the ACCURATE workflow relies on cryptographic verification at every step. Crucially, the technical execution of the transfer is determined by the data classification defined in Chapter 4.2:

- **Direct Data Exchange (Peer-to-Peer Download):** Used for Public or Protected data where the raw file is transferred to the Consumer.
- **Compute-to-Data (Remote Execution):** Used for Confidential data where the raw data never leaves the producer. Instead, a specific, trusted algorithm travels to the data.

5.1 Phase 1: The Producer's Journey (Publication)

The process begins within the secure perimeter of the Data Owner's infrastructure. The Producer configures the service type based on the sensitivity of the asset.

1. **Ingestion & Classification:** The Data Owner selects a dataset from their internal storage. Using the framework from Chapter 4.2, they determine the service type:
    - *Case A (Protected Data):* E.g., a quality certificate. The Owner selects an Access Service, allowing the raw file to be downloaded.
    - *Case B (Confidential Data):* E.g., machine telemetry. The Owner selects a Compute Service. Crucially, to prevent malicious code execution, the Owner configures the asset to only accept specific, trusted algorithms. This allowlist may include algorithms provided by the Producer themselves (e.g., a standard "KPI Calculator") or verified algorithms from a certified third-party software provider.
2. **Governance Configuration:**
    - **Terms of Use:** The Owner selects the appropriate Standard Data Usage Agreement.
    - **Access Controller clarification**: The owner defines the endpoint of their access controller (Ocean Provider).
    - **Access Control:** The Owner configures the Allow-list on their Ocean Provider (e.g., "Only allow wallet 0xABC... verified by the Data Space Authority").
3. **On-Chain Registration:** Using their Local Private Key, the Owner signs a transaction to publish the metadata (DDO). The asset becomes discoverable, but the actual data endpoint remains encrypted and stored only within the self-hosted Provider.

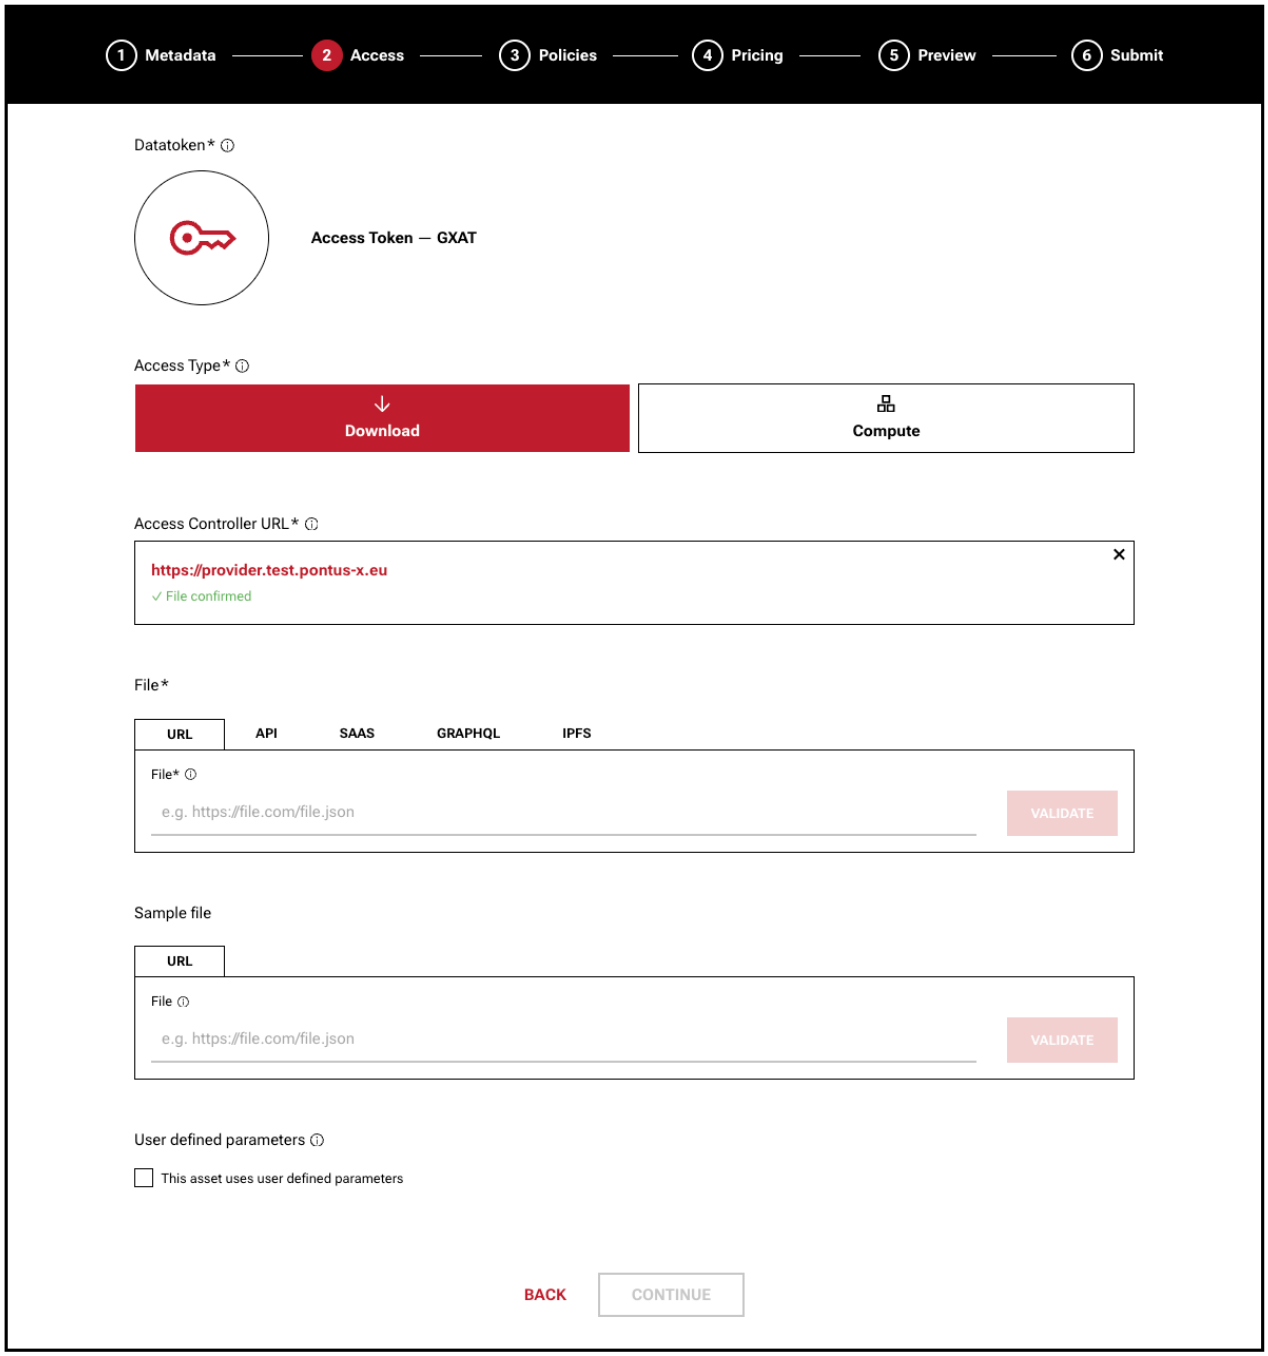

Figure 4: Ocean Provider URL and Data endpoint definition (Source: deltaDAO)

5.2 Phase 2: The Consumer's Journey (Discovery & Trust)

A potential Consumer searches the ACCURATE marketplace and identifies a relevant asset.

1. **Discovery:** The Consumer finds the asset.
2. **Identity Verification:** Before interacting, the Consumer performs a due diligence check using the **Pontus-X Registry**:
    - They resolve the Producer's wallet address visible in the metadata.
    - The Registry confirms the address belongs to a verified legal entity with a valid Pontus-X Identity.

3. **Contract Review:** The Consumer reviews the specific Data Usage Agreement attached to the asset and verifies its cryptographic hash.

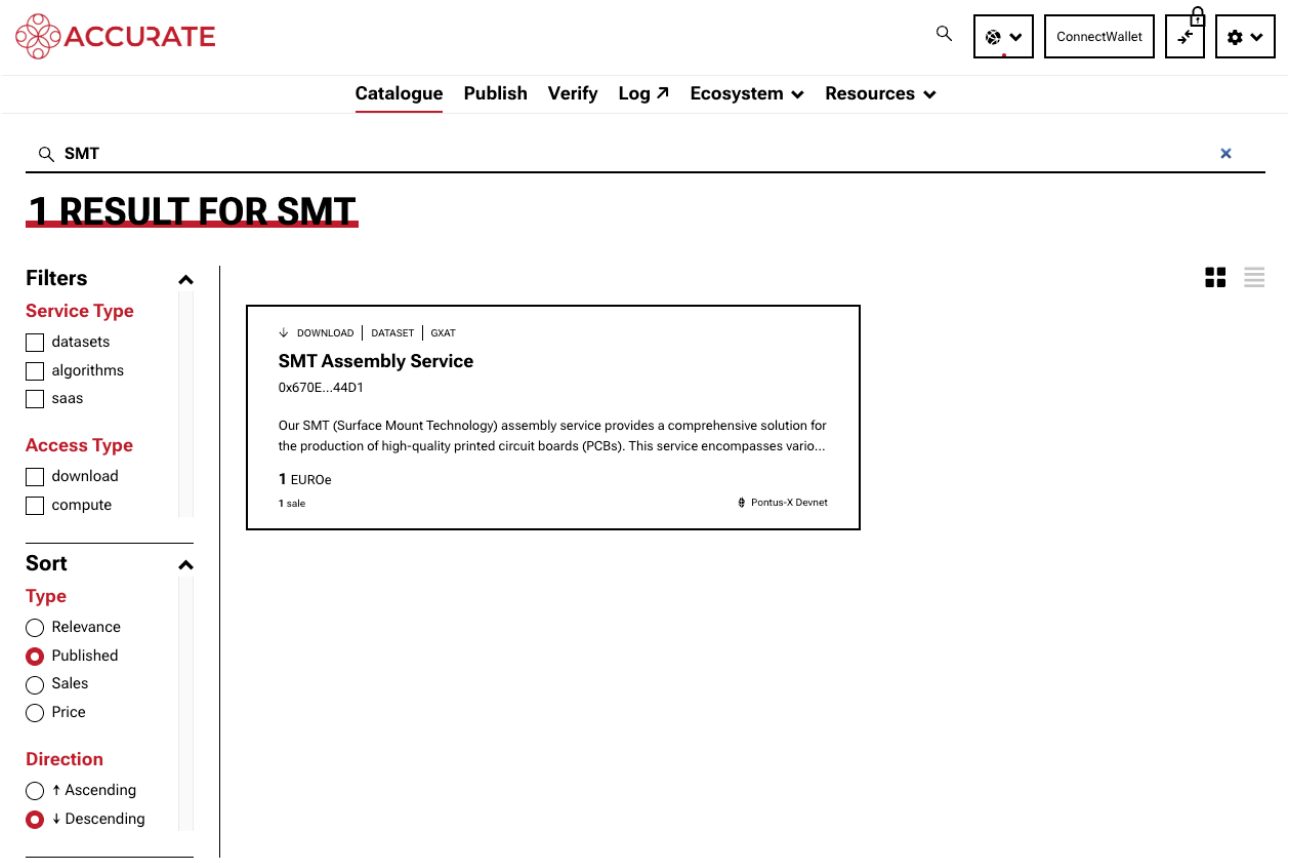

Figure 5: ACCURATE service consumer journey example (Source: deltaDAO)

5.3 Phase 3: Secure Execution (The Transaction)

Once trust is established, the Consumer initiates the transaction. The execution path branches here depending on the service type defined in Phase 1.

5.3.1 Common Step: The "Handshake"

1. **Signing & Payment:** The Consumer's Identity Wallet signs the transaction, transferring tokens (if priced) and digitally signing the Data Usage Agreement.
2. **Validation:** The Producer's Ocean Provider detects the transaction on the ledger. It validates payment, Identity Allow-listing, and contract signature.

5.3.2 Option A: Direct Data Exchange (For Protected Data)

If the asset was classified as Protected, the system executes a secure Peer-to-Peer transfer.

1. **Decryption:** Upon successful validation, the Producer's Provider decrypts the internal URL of the dataset.
2. **Streaming:** The data is streamed directly from the Producer's internal storage (e.g., private / public cloud, or server) to the Consumer via an encrypted HTTPS tunnel.

3. **Result:** The Consumer receives the raw file. The Producer retains an immutable audit log of *who* accessed the file and *when*.

5.3.3 Option B: Compute-to-Data (For Confidential Data)

If the asset was classified as **Confidential**, the system executes the "Safe Room" protocol using a trusted algorithm [14].

1. **Algorithm Selection:** The Consumer selects an algorithm to run on the data. Because the Producer locked down the asset in Phase 1, the Consumer *cannot* send arbitrary code. They must choose an algorithm that has been explicitly **allowlisted** by the Producer (e.g., a certified "Anonymization Script" or "Federated Learning Client").
2. **Execution in the "Safe Room":**
    - The **Operator Engine** pulls the selected, trusted algorithm into the Producer's isolated Kubernetes cluster.
    - The algorithm runs in a network-restricted Pod with temporary access to the local data.
3. **Result Delivery:** The algorithm generates a specific result (e.g., an optimization report). Only this result file is returned to the Consumer. The raw data never leaves the Producer's infrastructure, and the compute pod is destroyed immediately.

5.4 Result: Sovereignty in Action

This end-to-end process illustrates the practical reality of Sovereign Data Sharing. Whether transferring a file (Option A) or executing a trusted model (Option B), the Data Owner maintains absolute control.

- **The Producer** defines the mechanism (Download vs. Compute) and restricts *which* algorithms can touch their data.
- **The Consumer** interacts with a verified partner.
- **The Transaction** is fully auditable on the ledger, ensuring compliance without the need for a central intermediary.

6 Conclusion and Future Outlook

Deliverable D5.3 marks the completion of the ACCURATE data space's technical foundation, serving as the operational bridge between the architectural designs of previous tasks and real-world ecosystem deployment. By delivering the guidelines for the deployment of sovereign components, this document equips consortium partners with the blueprints to transition from passive users to active, sovereign operators who retain absolute physical and legal control over their data assets.

We have established a trust architecture that combines local identity generation with democratic verification by the Data Space Authority and transparent discovery via the Pontus-X Registry. Furthermore, the operationalization of Data Governance, through the integration of standardized, smart-contract-enforced usage agreements, ensures that legal compliance is embedded directly into the technical layer, significantly reducing the friction of secure data exchange.

Moving forward, the project focus shifts from individual component deployment to system-wide integration and validation. With the operational guidelines now established, partners can utilize the provided deployment packages to install their Ocean Provider and Compute-to-Data environments. Ultimately, this infrastructure lays the groundwork for the Work Package 6 pilots, enabling the secure, sovereign exchange of real-world manufacturing data that fulfills the ACCURATE vision of a resilient and decentralized European industrial ecosystem.

Bibliography

[1] ACCURATE Project Website, https://accurateproject.eu/

[2] Gaia-X Trust Framework 24.04, https://docs.gaia-x.eu/policy-rules-committee/trust-framework/latest/

[3] Ocean Protocol Provider Documentation, https://github.com/oceanprotocol/provider

[4] Ocean Protocol Compute-to-Data Documentation, https://docs.oceanprotocol.com/developers/compute-to-data/compute-to-data-algorithms

[5] European Data Act (Regulation 2023/2854), https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32023R2854

[6] Pontus-X Onboarding Wizard, https://onboarding.delta-dao.com/

[7] Data Space Authority (DSSC), https://dssc.eu/space/BVE2/1071253671/Organisational+Form+and+Governance+Authority

[8] Standards for Efficient Cryptography 2 (SEC 2): Recommended Elliptic Curve Domain Parameters (secp256k1), https://www.secg.org/sec2-v2.pdf

[9] Pontus-X Registry Participants, https://docs.pontus-x.eu/docs/participants-and-federators/ecosystem-participants

[10] Kubernetes Documentation - Persistent Volumes, https://kubernetes.io/docs/concepts/storage/persistent-volumes/

[11] Gaia-X Digital Clearing House (GXDCH), https://gaia-x.eu/gxdch/

[12] EUROe Stablecoin, https://www.euroe.com/

[13] Pontus-X DDO Documentation, https://docs.pontus-x.eu/docs/ddo_credential/ddo_intro

[14] Ocean Protocol Compute-to-Data Workflow Documentation, https://docs.oceanprotocol.com/developers/compute-to-data/compute-workflow

[15] Pontus-X Voting Tool, https://vote.pontus-x.eu/