Co-funded by the European Union

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Call for proposals	DIGITAL-2023-CLOUD-AI-04	**Type of action**	DIGITAL-SIMPLE
Grant Agreement No.	101167948	**Start date**	1 December 2024
Project duration	30 months	**End date**	31 May 2027
Project consortium – Coordinator: LIBELIUM LAB SL - LIBE			
AYUNTAMIENTO DE CARTAGENA		BEN - CART	
DELTADAO AG		BEN - DDAO	
DREES & SOMMER SE		BEN - D&S	
LANDESHAUPTSTADT KIEL		BEN – KIEL	
OPEN & AGILE SMART CITIES		BEN - OASC	
MOBILITY AS A SERVICE ALLIANCE		BEN - MAAS	
UNIVERSITY OF GALWAY		BEN – GALW	
KIEL-MARKETING EV		BEN – KIELM	
DUNAVNET LIMITED		BEN – DNET	

Contact: projects@libelium.com

Co-funded by the European Union

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

D4.3 Technical Guidelines

Document Identification			
Status	Draft (final once approved by the European Commission)	**Due Date**	30/09/2026
Version	1.0	**Submission Date**	18/10/2026

Related WP	WP4	**Document Reference**	D4.3
Related Deliverable(s)	All	**Dissemination Level**	PU
Lead Partner	LIBE	**Lead Author**	Juan Francisco Inglés
Contributors	DDAO DNET Imascono (incoming partner)	**Reviewers**	Inma Molina

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Document Information

List of Contributors	
Name	Partner
Kai Meinke	deltaDAO AG
Thomas Komenda	deltaDAO AG
Sebastian Abromeit	deltaDAO AG
Dejan Drajic	DunavNET Limited
WP4 front-end team	Imascono (incoming partner)
Juan Francisco Inglés	Libelium

Document History			
Version	Date	Change editors	Changes
0.1	07.05.26	Kai Meinke, Thomas Komenda, Sebastian Abromeit	Inputs from WP3
0.2	21.07.26.	Dejan Drajic (DNET)	Inputs from WP2
0.3	07.08.26	Imascono	Inputs from WP4
0.4	05.10.26	Juan Francisco Inglés	Completed missing sections and reviewed existing content for consistency and completeness.
1.0	16.10.26	Juan Francisco Inglés	Finalized the document for submission.

Quality Control		
Role	Who (Partner short name)	Approval Date
Reviewer/Quality Manager	Inma Molina	16/10/2026 (version 1.0)

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Table of Contents

Document Information ... 3
Table of Contents ... 4
List of Figures ... 6
Executive Summary ... 7
1. Introduction ... 8
1.1 Purpose of the Document ... 8
1.2 Scope and Relationship to Other SENSE Deliverables ... 8
1.3 Structure of the Document ... 8
2. Interoperability Baseline: MIMs and Local Digital Twin Reuse ... 9
2.1 MIM8 Local Digital Twin as the SENSE Reuse Baseline ... 10
2.2 Relationship to Gaia-X, FIWARE and the LDT Toolbox ... 11
3. Data Guidelines for Integration with the SENSE CitiVerse ... 12
3.1 Data Source Classification and Integration Patterns ... 12
3.2 Minimum Metadata for a Publishable Data Product ... 12
3.3 Semantic, Geospatial and Identifier Rules ... 13
3.4 Publication and Consumption Lifecycle ... 13
3.5 Data Integration Acceptance Checklist ... 14
4. Identity and Trust Guidelines for Integration with the SENSE CitiVerse Data Space ... 15
4.1 Organizational Identities and the SENSE Trust Framework ... 15
4.1.1 Basic Information about a Legal Entity ... 16
4.1.2 Verification against a Trust Service Provider ... 16
4.2 Publication and Consumption in the SENSE CitiVerse Data Space ... 18
4.2.1 Requirements for Publication through the Portal ... 20
4.2.2 Requirements for Consumption through the SENSE Portal ... 22
5. Infrastructure and Service Integration Guidelines for the SENSE CitiVerse ... 24
5.1 Integration on the Application and Scenario Layer ... 24
5.1.1 Architecture and Responsibility Boundary ... 25
5.1.2 Common Application Integration Guideline ... 27
5.1.3 Option A: Reusing or Embedding an Existing SENSE Application ... 28
5.1.4 Option B: Deploying a Custom Scenario Application ... 28
5.1.5 Scenario Portability and Reuse ... 29
5.1.6 Cross-cutting Requirements ... 30

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

5.1.7 Alignment with the Minimum Interoperability Mechanisms 31
5.2 Technical Integration in the SENSE Data Space Layer 32
5.2.1 Deployment of an Own Connector / Access Controller 32
5.2.2 Deployment of an Own Policy Server 33
5.2.3 Deployment of a Participant-operated Portal 34
5.2.4 Deployment of Own Public Key Infrastructure (PKI) and Digital Identities 35
5.2.5 Deployment of an Own SSI Wallet 35
5.3 Integration in the SENSE Data Layer 36
5.3.1 Integration of a FIWARE Context Broker 36
5.3.2 Integration of Portable Applications 38
5.4 Integration in the SENSE Cloud Infrastructure Layer 40
6. Frontrunner Adoption Guidelines 41
6.1 Cartagena Adoption Profile 41
6.2 Kiel Adoption Profile 41
6.3 Common Adoption Rules for Both Frontrunners and Follower Cities 42
7. Developer Integration Checklist and Reuse Path 43
7.1 Reuse and Open-source Contribution Practice 43
8. Conclusions 45
9. Glossary 46
10. Abbreviations 48

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

List of Figures

Figure 1 - SENSE Portal (Version 1) 19
Figure 2 - SENSE Portal (Version 2) 19
Figure 3 - Basic Connector Setup for Data Services 22
Figure 4 - Successfully connected SSI wallet 23
Figure 5 - SSI Wallet with Organizational Credentials 23
Figure 6 - Simplified Version of the SENSE Architecture 24
Figure 7 – Recommended interaction pattern 25
Figure 8 – Embedding an existing SENSE application 28
Figure 9 - Custom Scenario Application 29

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Executive Summary

Deliverable D4.3 provides the public technical guidelines for developers integrating data, services and applications with the SENSE CitiVerse. The document focuses on reuse of established European and open frameworks, such as Gaia-X and FIWARE, and on practical interoperability through the Minimum Interoperability Mechanisms (MIMs) and Local Digital Twin (LDT) approaches.

This revision aligns the guidelines with the current SENSE implementation baseline described in D4.1. The pilot scope is organised around four consolidated scenarios in the frontrunner cities of Cartagena and Kiel. Three scenarios use the common browser-native WP4 front-end approach, while Kiel SEALEVEL uses a dedicated web stack. Both implementation tracks share the same logical separation between user application, scenario/content services, and governed access to data and knowledge resources through the SENSE Data Space.

For developers, the report defines: (i) the SENSE mapping to MIMs and the reuse path for LDT components; (ii) data onboarding, metadata, semantic and geospatial interoperability rules; (iii) organisational identity, trust, publication and consumption requirements in the Data Space; (iv) integration routes across the application/scenario, Data Space, data and cloud-infrastructure layers; and (v) concrete adoption rules for Cartagena and Kiel. The guidance is intentionally technology-aware but not tied to one deployment environment, so that new cities and service providers can reuse the same integration patterns.

Because D4.3 is a public deliverable, it describes reusable interfaces, responsibilities and deployment patterns rather than operational secrets, credentials or non-public configuration details. Where supporting WP2 or WP3 deliverables contain environment-specific information, this report abstracts that information into developer-facing rules. The result is a practical integration guide that supports reuse, portability and future contribution of generic improvements to the relevant open-source ecosystems and SENSE repositories.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

1. Introduction

1.1 Purpose of the Document

The purpose of D4.3 is to translate the SENSE technical architecture and pilot experience into implementable guidance for developers. It explains how a new participant, city, data provider or application team can connect to SENSE without reproducing the complete project stack, and which interoperability conditions must be respected so that data, services, 3D assets and scenario components remain reusable.

The report addresses the deliverable mandate by connecting four perspectives: open European data-space frameworks and trust mechanisms; FIWARE/NGSI and geospatial standards for data integration; MIMs and Local Digital Twin interoperability; and the concrete web-native scenario architectures used in the two frontrunner cities. It also distinguishes mandatory integration conditions from recommended practices, so developers can start with a lightweight route and move towards sovereign deployment where required.

1.2 Scope and Relationship to Other SENSE Deliverables

D4.3 consolidates implementation guidance from the project artefacts listed below. It does not replace those deliverables; it extracts the parts that are relevant to software and integration teams and keeps the Grant Agreement as the governing source in case of discrepancy.

Source	Role in D4.3
D2.1	Provides the standards baseline for digital twins, data exchange, NGSI-LD, geospatial information and related interoperability mechanisms.
D2.2-D2.4	Provide privacy/trust recommendations, the data-source catalogue and implemented data-space integration patterns. D4.3 reuses these as public technical rules without reproducing sensitive deployment details.
D3.1-D3.4	Provide relevant input and contributions related to the SENSE Data Space.
D4.1	Defines the current four-scenario baseline and the browser-native application architecture used by the frontrunner pilots.

1.3 Structure of the Document

Section 2 defines the MIM and Local Digital Twin interoperability baseline. Section 3 gives data onboarding and metadata rules. Sections 4 and 5 cover identity/trust and the application, Data Space, data and cloud integration layers. Section 6 provides adoption rules for Cartagena and Kiel, Section 7 provides the developer checklist and reuse path, and Section 8 concludes the guidance. The glossary and abbreviations follow.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

2. Interoperability Baseline: MIMs and Local Digital Twin Reuse

SENSE uses the Minimum Interoperability Mechanisms as a practical checklist for keeping city data, services and Local Digital Twin components portable across implementations. The focus is not to implement every mechanism as a separate product; it is to ensure that the interfaces crossing organisational and technical boundaries are discoverable, machine-readable, secure and reusable.

The MIM mapping used in this deliverable concentrates on the mechanisms that materially affect the SENSE architecture: access and interlinking, common representation, governed exchange, personal-data considerations, interoperable AI, security, geospatial information and Local Digital Twins. The table below summarises how these mechanisms are reflected in the SENSE layers.

MIM	Objective	Relevance
0	Accessing Data	SENSE CitiVerse depends on reliable and standardised access to heterogeneous data sources, including APIs, files, databases, FIWARE Context Brokers and storage services. MIM0 provides the basic interoperability layer for making data technically reachable by applications, connectors, scenario services and Local Digital Twin components.
2	Representing Data	SENSE CitiVerse requires consistent representation of urban, environmental, mobility, energy, tourism and scenario data so that datasets can be understood and reused across cities and service providers. MIM2 supports semantic and syntactic consistency through shared data models, ontologies and mappings, including NGSI-LD and Smart Data Models where applicable.
1	Interlinking Data	CitiVerse scenarios require links between entities, datasets, services, geospatial objects, policies and decision-support outputs. MIM1 supports the ability to identify, reference and connect data across domains and systems, avoiding isolated datasets and enabling cross-domain reasoning inside Local Digital Twins.
6	Securing Data	SENSE CitiVerse includes a data space layer with identity, trust, access control, policy enforcement and secure service consumption. MIM6 is therefore central to the SENSE Trust Framework, Gaia-X-compatible identity management, connector-based access control and secure exchange of data and services.
3	Exchanging Data	SENSE CitiVerse is not only a technical platform but also a governed data-sharing ecosystem. MIM3 supports catalogue publication, service offerings, contracts, SLAs, access conditions, data usage agreements, marketplace interactions and federation between

		participants. It is a core mechanism for turning distributed city resources into reusable CitiVerse services.
4	Personal Data	SENSE CitiVerse primarily operates at the level of organisational participants and data/service offerings, but several scenarios may involve citizen interaction, mobility traces, behavioural data, participatory planning or other personal-data-related use cases. MIM4 should therefore be treated as a scenario-level requirement, especially where EUDI wallets, consent, lawful basis, data minimisation or privacy-preserving access are needed.
7	Geospatial Data	CitiVerse use cases rely on spatial representation of cities, districts, buildings, infrastructure, mobility flows and environmental conditions. MIM7 connects geospatial data, 2D/3D representations, GIS, BIM, CityGML/OGC-compatible resources and spatial layers used in Local Digital Twins and immersive visualisations.
5	Interoperable AI	SENSE CitiVerse will increasingly rely on analytics, predictive models, simulation, optimisation and potentially generative or agentic AI. MIM5 is relevant for ensuring that AI components are transparent, interoperable, reusable, and aligned with trustworthy AI requirements. In the current project phase, it should be positioned as an enabling concern rather than a fully mature implementation layer.
8	Local Digital Twin	MIM8 provides the direct interoperability frame for Local Digital Twins: connecting data, services, models, visualisations, simulations, decision-support workflows and components across city contexts. It is particularly relevant for SENSE because the CitiVerse concept depends on orchestrating LDT resources across data spaces, 3D/immersive environments and scenario-specific applications.

2.1 MIM8 Local Digital Twin as the SENSE Reuse Baseline

MIM8 is the key MIM for SENSE because the CitiVerse builds on Local Digital Twins. It provides the interoperability basis for connecting city data, models, simulations, visualisations and decision-support tools. In SENSE, this means that LDT-related resources should be published and reused through the data space as standardised services: datasets, APIs, analytical models, simulation tools, geospatial layers, 3D assets and visualisation components. Concretely, SENSE should use MIM8 to make LDT components reusable across cities and scenarios. This requires clear interfaces for accessing LDT resources, standard input and output formats, reusable simulation and analytics services, Model-as-a-Service components, workflow orchestration and links to 2D/3D visualisation environments. Without this, each city risks building a separate digital twin demonstrator that cannot easily be reused by other cities or connected to other services in the SENSE data space.

However, MIM8 is not sufficient for the CitiVerse. A Local Digital Twin mainly concerns the digital representation and analysis of a local system. A CitiVerse goes further: it combines

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Local Digital Twins with immersive environments, multi-user interaction, citizen participation, AI-supported services and cross-city federation. These elements create interoperability needs that MIM8 does not yet specify. A dedicated CitiVerse MIM is therefore needed to define how CitiVerse environments exchange and reuse: 3D scenes and assets; scenario states; avatars and user roles; interaction rules; simulation results; AI-agent services; citizen-facing interfaces; and links between immersive environments and data space resources. These are not minor presentation issues. They determine whether CitiVerse applications can be portable, trustworthy and scalable across cities. The recommendation is to use MIM8 as the technical foundation for Local Digital Twin interoperability, while launching a dedicated CitiVerse-specific MIM for the additional immersive and participatory layer. This would prevent fragmentation, reduce vendor lock-in, and give SENSE a clear contribution to the European interoperability agenda: not another isolated CitiVerse demonstrator, but a reusable framework for federated CitiVerse environments.

2.2 Relationship to Gaia-X, FIWARE and the LDT Toolbox

Gaia-X provides the trust, participant/service description and compliance concepts used by the SENSE Data Space. FIWARE provides the context-information interoperability route used for city data, in particular through NGSI/NGSI-LD interfaces and compatible data models. These two framework families are directly visible in the implementation guidance in Sections 3 to 5.

The Local Digital Twin Toolbox is reused at the level where it creates portability: MIM-aligned data access, shared data models, geospatial resources, open 3D formats, discoverable services and repeatable scenario configurations. A SENSE scenario should bind configuration to catalogue entries and persistent identifiers rather than hard-coded source endpoints. This makes a scenario transferable to a follower city by substituting equivalent local assets while preserving the application and service logic.

Where SENSE develops an improvement that is generic beyond the project - for example an adapter, data model mapping, configuration template, documentation example or reusable service component - the preferred route is to contribute it to the relevant upstream or community repository when governance and licensing permit. SENSE-specific integration code and examples can remain in the SENSE repositories. The objective is practical reuse and maintainability, not parallel forks of established frameworks.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

3. Data Guidelines for Integration with the SENSE CitiVerse

The SENSE data guidance translates the WP2 data-source catalogue and integration work into a reusable onboarding contract for developers. The same rules apply whether the underlying source is a static file, an OGC service, an NGSI-LD context service or a proprietary API that must be adapted. The goal is that the scenario layer consumes a small number of predictable, standards-oriented interfaces rather than city-specific integrations.

This public guideline deliberately separates the logical integration contract from deployment-specific credentials and endpoints. Data providers remain responsible for the authoritative source, licence and access conditions; the SENSE integration layer provides harmonisation, metadata, catalogue publication and governed access.

3.1 Data Source Classification and Integration Patterns

- Static resources: geospatial files, tabular data, documents, imagery, multimedia and 3D assets. Publish them through a controlled file/object service, record integrity and provenance, and expose them through catalogue metadata rather than hard-coded application paths.
- OGC-compliant geospatial services: prefer direct reuse of standard interfaces where they are already suitable. Preserve coordinate reference system information, identifiers and layer metadata so assets remain joinable across maps and 3D scenes.
- NGSI-LD/context services: use the context interface for dynamic city entities and time-varying observations. Reuse compatible shared data models where available and keep source-to-SENSE mappings explicit.
- Custom or legacy REST APIs: place the adapter at the data-integration boundary. Convert proprietary payloads into the agreed SENSE representation before they reach scenario applications, so application code does not accumulate one-off city integrations.

3.2 Minimum Metadata for a Publishable Data Product

Every data product must be sufficiently described for a developer to discover it, understand whether it is fit for purpose, determine the conditions of use and locate the technical access route. The following fields consolidate the metadata already captured in WP2 with the requirements used by the SENSE catalogue.

Priority	Required information
MUST	Name and description; provider/owner and rights holders; licence/terms and conditions; access restrictions; data type/nature; format and structure; update frequency or version date; technical access method; access policy; contact point; price or free-of-charge status where applicable.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Priority	Required information
SHOULD	Applicable data model/ontology and semantic context; use-case purpose; provenance; quality or known limitations; privacy information; relation to MIMs; persistent identifiers; coordinate reference system for geospatial content; expected freshness and retention where relevant.

3.3 Semantic, Geospatial and Identifier Rules

- Prefer NGSI-LD and compatible Smart Data Models for dynamic context information when the source supports them or can be mapped without loss of meaning.
- Prefer OGC-aligned interfaces and open geospatial formats for maps, features and tiles. For browser-native 3D delivery, use formats already adopted by the SENSE scenarios such as OGC 3D Tiles, glTF and CityJSON where appropriate.
- Use stable identifiers for entities, locations, datasets and service offerings. The same real-world object should be linkable across context data, 2D/3D geometry, narrative content and analytical results.
- Document transformations. An adapter that renames, aggregates, filters or converts source data must make the mapping traceable so developers can understand the relationship to the authoritative source.
- Do not make the user-facing application responsible for source harmonisation. Conversion belongs in the integration or scenario-service layer, where it can be reused by more than one application.

3.4 Publication and Consumption Lifecycle

- Describe and validate the source, including licence, ownership, update behaviour and access restrictions.
- Select the integration pattern and, where needed, implement an adapter that produces the agreed standard representation.
- Create catalogue metadata and associate the data/service offering with the provider identity and applicable policy.
- Publish the offering through the SENSE Data Space so that consumers can discover it without exposing protected credentials or internal infrastructure.
- Acquire the necessary access rights or contractual permission and consume the resource through the protected route selected by the provider.
- Validate end-to-end behaviour: freshness, identifiers, semantic mapping, timestamps, geospatial consistency, error handling and policy enforcement.
- Bind scenario configuration to the catalogue item or persistent resource identity, not to an environment-specific endpoint, so deployments remain portable.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

3.5 Data Integration Acceptance Checklist

- A developer who did not create the source can understand what the data represents and under which conditions it may be used.
- The representation is machine-readable and uses an agreed standard or a documented mapping to one.
- The access route can be changed without modifying the user-facing scenario code.
- Restricted data remains restricted after integration; publication metadata never substitutes for authorization.
- Provenance, licence and update information remain associated with derived layers and 3D/visual assets.
- Repeatable tests cover valid access, invalid access, data freshness and authorization failure.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

4. Identity and Trust Guidelines for Integration with the SENSE CitiVerse Data Space

Following the European Data Space for Smart Communities DS4SSCC Blueprint recommendation a strong identity and trust framework is applied in the SENSE CitiVerse. The DS4SSCC Reference Architecture specifically recommends the Gaia-X Trust Framework on this matter, which follows the de-facto European standards and best practices for a decentralized Identity Management based on W3C and OpenID Connect (OIDC) standards. The same Trust Framework is also recommended by the Data Spaces Support Center (DSSC) Technical Building Blocks, highlighting the relevant objectives for the inclusion of an identity and trust framework, namely:

1. Establish Trust Anchors and Trust Services to validate, issue, and verify credentials. These Trust Anchors establish the initial base and minimum level of trust in the SENSE CitiVerse (and other compatible data space ecosystems).
2. Establish Governance Rules and Conformity Assessment Processes, to operationalize trust in the SENSE CitiVerse (and other compatible data space ecosystems).

To implement and operationalize this Trust Framework, while being compatible to other ecosystems and adhering to de-facto standards, the SENSE CitiVerse has selected the Gaia-X Trust Framework and its components. Technical guidelines and steps for the necessary setup to integrate with the SENSE CitiVerse will be given in the following sections.

Note: Organizations engaging in a data space need to internally designate authorized personnel allowed to represent the organization in a data space and allowed to perform operations on behalf of the organization, i.e. handling of signatures, keys, publication of services and establishing contracts with other parties. Make sure this has been discussed and formally aligned within your organization.

Scope note: this section concerns organisational and service identities used on the Data Space plane. End-user authentication and role handling inside a scenario application are addressed at the application/scenario layer in Section 5.1 and should not be confused with participant onboarding to the Data Space.

4.1 Organizational Identities and the SENSE Trust Framework

As basis for any relationship within the SENSE CitiVerse, actors need to be identified and identities need to be verifiable. In SENSE digital identities are being utilized. The digital identity of each legal entity consists of the digital attributes and credentials. A digital identity is typically a set of verifiable credentials issued by a trusted party. These credentials contain claims about the identity of the subject, which can be verified against trusted national and international data sources in real time.

Each digital identity is anchored in a cryptographic key pair with verified ownership: the private key remains under the holder's control, while the corresponding public key enables others to verify signatures and authenticate issuers.

Digital identities based on digital signatures are important, especially when digital contracts are represented as a verifiable credential, because signing it with a qualified digital identity can provide strong evidence of authenticity and support legally valid electronic transactions, which is a prerequisite for Data Usage Agreements (DUAs).

4.1.1 Basic Information about a Legal Entity

Any organization aiming to join and interact with the SENSE CitiVerse must make the following information available, as defined by the Gaia-X Trust Framework applied here:

1. Full name of the Organization
2. Legal Address of the Organization
3. Full address of the organisation's headquarters
4. Legal Registration Number of the Organization
    a. as VAT ID (Value Added Tax ID, verifiable against the EU VIES)
    b. or EORI (Economic Operators Registration and Identification Number)
    c. and/or LEI Code (Legal Entity Identifier)

This information constitutes the basic information required to uniquely identify an organization. If an organization does not have one or more of the LRN available, it is recommended to register a LEI through the Global Legal Identifier Foundation (GLEIF).

You can check the VAT ID against the European VIES VAT number validation service[1]. EORIs can be checked against the EU EORI number validation service[2]. LEI Codes can be verified through the LEI Search of the GLEIF[3].

Note: Please always double-check your organizational data against the national and international registers to avoid errors and failed validation attempts.

4.1.2 Verification against a Trust Service Provider

To establish a digital identity in SENSE CitiVerse, your identity must be verified and confirmed by an accepted SENSE Trust Service Provider (TSP) who can provide proof of verification in the form of an X.509 certificate issued by an eIDAS TSP. The Trust Service Provider will verify your organisational identity information and issue you with digital proof of this verification. Within the SENSE Trust Framework, this requires a check through an eIDAS TSP accredited

[1] EU VIES VAT Number Validation, https://ec.europa.eu/taxation_customs/vies/#/vat-validation, last seen on May 4, 2026.
[2] EU EORI Number Validation, https://ec.europa.eu/taxation_customs/dds2/eos/eori_validation.jsp?Lang=en, last seen on May 4, 2026.
[3] LEI Search, https://search.gleif.org/#/search/expertMode=false, last seen on May 4, 2026.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

by Gaia-X from the EU Trusted List. We recommend acquiring a Qualified Website Authentication Certificate (QWAC) for your legal entity.

Note: If you only require access to the development infrastructure, you can bypass this step by using the deltaDAO Onboarding Application[4] which performs a soft verification. This is suitable for initial onboarding and testing.

The verification through an eIDAS TSP requires the following from a future SENSE participant:

1. The organizational details mentioned above
2. An authorized buyer from your organization must procure an eIDAS QWAC (or equivalent) from an eIDAS TSP
3. An authorized legal representative of your organization must be available for identification and authorization
4. Technical control over the domain you wish to register and its DNS entry associate with your organization
5. The technical ability to sign a Certificate Signing Request (CSR) with your own private key.

Once you have received your X.509 certificate you have successfully been verified and have proof of your successful verification. Your private key has now been officially associated with your organization, verified by an independent TSP.

Note: Self-signed certificates are NOT accepted, as they lack the verification by an accredited and regulated entity which is liable for the correctness of the information.

Next the participant should familiarize themselves with the concept of Digital Identities and their standards. We recommend the overview of “Adopted Standards and Protocols”[5] by Gaia-X AISBL, which recommendations are shared by the DSSC, DS4SSCC, the EUBW and EUDI communities.

In the SENSE CitiVerse data space participants SHOULD control Gaia-X compliant identities, based on Decentralized Identifiers (DIDs), namely the DID:WEB Method[6] and the W3C Verifiable Credential Data Model 2.0[7] (VDCM 2.0) in combination with the Gaia-X Ontologies[8].

The process on how to create a compatible DID:WEB Public Key Infrastructure (PKI) entry, based on your organizational certificate and private key can be found in the “Gaia-X Standard

[4] deltaDAO Fast Track Onboarding Application, https://onboarding.delta-dao.com/, last seen May 4, 2026.
[5] Standards and Protocol, Gaia-X ICAM Document, https://docs.gaia-x.eu/technical-committee/identity-credential-access-management/25.11/standards/, last seen May 4, 2026.
[6] Decentralized Identifiers 1.1, https://www.w3.org/TR/did-1.1/, last seen May 4, 2026.
[7] VCDM 2.0, https://www.w3.org/TR/vc-data-model-2.0/, last seen May 4, 2026.
[8] Gaia-X Legal Person Ontology, https://docs.gaia-x.eu/ontology/development/classes/LegalPerson/, last seen May 4, 2026.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Compliance for Participants" Chapter[9] of the latest Gaia-X Compliance Document. The SENSE CitiVerse team will also be glad to support any new participant with the setup procedure.

Once you completed the setup of your organizational DID:WEB and the issuance of the corresponding credentials you are ready to store your Verifiable Credentials in your own wallet or a wallet provided by a SENSE member.

SENSE does offer a suitable wallet for self-sovereign deployment, as free open-source code, via the SENSE EU Repository[10]. The wallet supports the storage of Legal Person, Gaia-X Terms & Conditions and Legal Registration Number (LRN) Credentials as well as the DID:WEB Method and the Credential Exchange via OpenID for Verifiable Presentations (OIDC4VP) as proposed in the DS4SSCC Reference Architecture[11].

As basis SENSE selected the walt.id Community Wallet[12] and forked it to match the Gaia-X Trust Framework requirements. It supports all relevant EUDI Credential Standards (W3C VC, SD-JWT VC, ISO/IEC 18013-5), the required protocols (OIDC4VCI, OIDC4VP), crypto algorithms and Status Lists.

Note: As the SENSE CitiVerse Data Space itself deals only with legal participants and thus is not directly integrated with EUDI, the applications in the ecosystem are recommended to integrate directly with EUDI wallets for Identity and Access Management as national EUDI Wallet implementations become available in accordance with the European Digital Identity framework. For the SENSE CitiVerse Data Space we are following the approach of the European Business Wallet (EUBW) with VC-JWT Credentials and the Verifiable Credential Data Model (VCDM 2.0) as recommended by the Trust Framework.

4.2 Publication and Consumption in the SENSE CitiVerse Data Space

The SENSE Data Space allows not only the publication of data services, but also the publication of (portable) applications and infrastructure services. The underlying open-source data space software framework utilized here is Ocean Enterprise[13] which has been combined with the recommended Trust Framework components to enable the SENSE CitiVerse data space. The main access point for the consumption and publication of services is the SENSE Portal: https://portal.senseverse.eu/

[9] Gaia-X Standard Compliance for Participants, Gaia-X Compliance Document, https://docs.gaia-x.eu/policy-rules-committee/compliance-document/2.5.0/Process/, last seen May 4, 2026.
[10] SENSE Repository, https://github.com/SENSE-EU, last seen May 4, 2026.
[11] DS4SSCC Reference Architecture, https://inventory.ds4sscc.eu/reference-architecture, last seen May 4, 2026.
[12] Walt.id Community Edition, https://walt.id/community-stack, last seen May 4, 2026.
[13] Ocean Enterprise Homepage, https://www.oceanenterprise.io/, last seen May 4, 2026.



Figure 1 - SENSE Portal (Version 1)

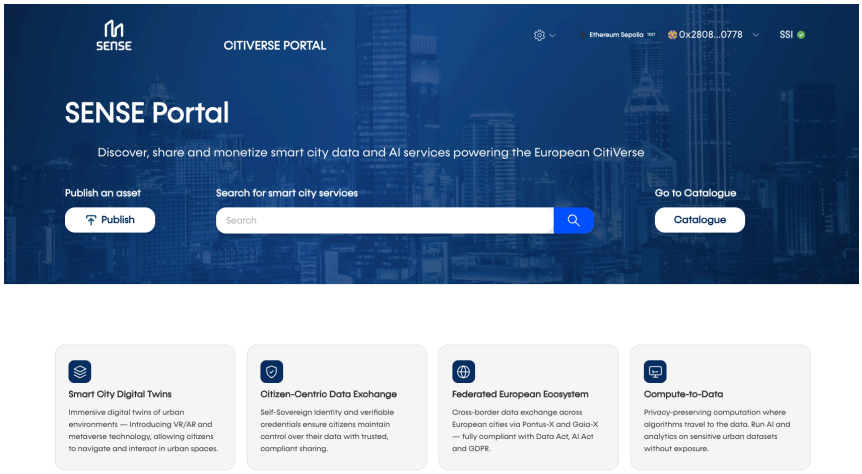


Figure 2 - SENSE Portal (Version 2)

Once connected to the SENSE Portal and your SSI wallet a participant can select to publish, manage, and consume service offerings in the SENSE data space.

A prerequisite for this is a registration with the SENSE consortium and a proper onboarding to the SENSE data space, which can be either completed through the Fast Track Application[14] of deltaDAO and/or by bringing a self-sovereign identity according to the Trust Framework.

The interaction with the SENSE data space requires the participant to be able to sign transactions with a private SECP256K1 key in a separate wallet file, in addition to the SSI

[14] deltaDAO Fast Track Onboarding, https://onboarding.delta-dao.com/, last seen May 4, 2026.

wallet. This key can be either generated through the Fast Track Onboarding and/or through a wallet and SECP256K1 key generator[15].

If the onboarding has been completed, your organization and its identity credentials are added to the data space registry. This is necessary since the SENSE data space only shows services of identified and registered users to ensure transparency and trust in the ecosystem.

4.2.1 Requirements for Publication through the Portal

As a participant agent, the SENSE Portal allows users to publish services via the integrated Metadata Wizard into the SENSE Data Space Catalogue. This makes the compilation of all required and optional information easier and a streamlined experience.

Note: The publication will require the user to be logged in with its organizational identity and connected to its organizational SSI wallet since all catalogue entries will be stored as Verifiable Credentials and signed with the organizations signature to make the entries tamper-proof and verifiable.

Before you start the publication procedure, guided by the Wizard, the necessary information from Chapter 2 needs to be compiled, such as:

- Title for the Service Offering
- Description of the Service Offering
    - Information about the data service
    - Information about all rights holders
    - Information about the use case
    - Information about relevant standards, semantics and ontologies applied
    - Information about the age and update frequency
    - Information about the publisher organization and how to contact it
- Tags for the service offering (Always include the Tag “SENSE” for reference to the SENSE data space catalogue)
- name of the publisher organization (Never include any personal data into the Metadata)
- the license under which the service is being offered (Templates are available through SENSE)
- the privacy policy under which the service is being offered.
- A list of eligible user organizations allowed to consume the service
- Access rules that shall be technically enforced, e.g. the requirement of a Trust Framework compliant identity
- A subscription duration for the data service (until the contract needs to be renewed)

[15] Free wallet generator, https://github.com/deltaDAO/wallet-generator, last seen May 4, 2026.

- A price for the usage of the service offering, i.e., free service vs. paid service.

Technically the user needs to be aware of the following details necessary to publish a data service offering:

- An identity wallet and identity for the publication MUST be selected and available for the publication.
- A wallet file / SECP256K1 private key MUST be available, generated through the onboarding procedure and associated with the organization.
- A connector/node MUST be selected as Policy Enforcement Point (PEP), i.e. the public connector https://node.demo.pontus-x.eu OR an own sovereign node.
- In case the real-time verification shall be enforced (optional), the service provider MUST specify the Gaia-X Digital Clearing House (GXDCH) Compliance Service API for the verification of participant credentials. All APIs can be found in the official GXDCH overview[16]. This is not mandatory in SENSE, due to the additional effort it entails.
- The service offering source, e.g. a dataset, MUST be made available so that it can be reached by the selected Connector/PEP which acts as a proxy between provider and consumers.
- The service offering source SHOULD be protected against unauthorized access or side-channels through firewall rules or authentication headers. In case of open data or non-confidential data this is less relevant, leading to a SHOULD priority here. For sensitive information this is a MUST criteria.
- The service offering MUST be operated in accordance with service level agreements (SLAs), if they are mentioned in the respective terms and conditions.
- A regular testing and monitoring of the service offering SHOULD be in place.
- The wallet file account in the data space MUST be funded with enough fee tokens to cover the cost of writing to the catalogue. This will be automatically done with the registration and can also be requested from the SENSE governance authority.

In addition to the technical requirements, the full scope of Gaia-X compliance requirements for digital and cloud services SHOULD be considered. This is especially important if a high-level certification for security, compliance, and sovereignty is required or aspired. An overview of the requirements for Data Exchange Services[17] and Cloud Services[18] can be found in the official Gaia-X Compliance Document.

[16] Gaia-X Digital Clearing Houses Overview, https://docs.gaia-x.eu/#/gxdch, last seen May 4, 2026.
[17] Compliance Criteria for Data Exchange Services, https://docs.gaia-x.eu/policy-rules-committee/compliance-document/2.5.0/criteria_data_exchange_services/, last seen May 4, 2026.
[18] Compliance Criteria for Cloud Services, https://docs.gaia-x.eu/policy-rules-committee/compliance-document/2.5.0/criteria_cloud_services/, last seen May 4, 2026.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

SENSE does not make the full compliance mandatory to lower the entry barriers for testing and development. However, the minimal requirements SHOULD be considered in any service offering description[19].

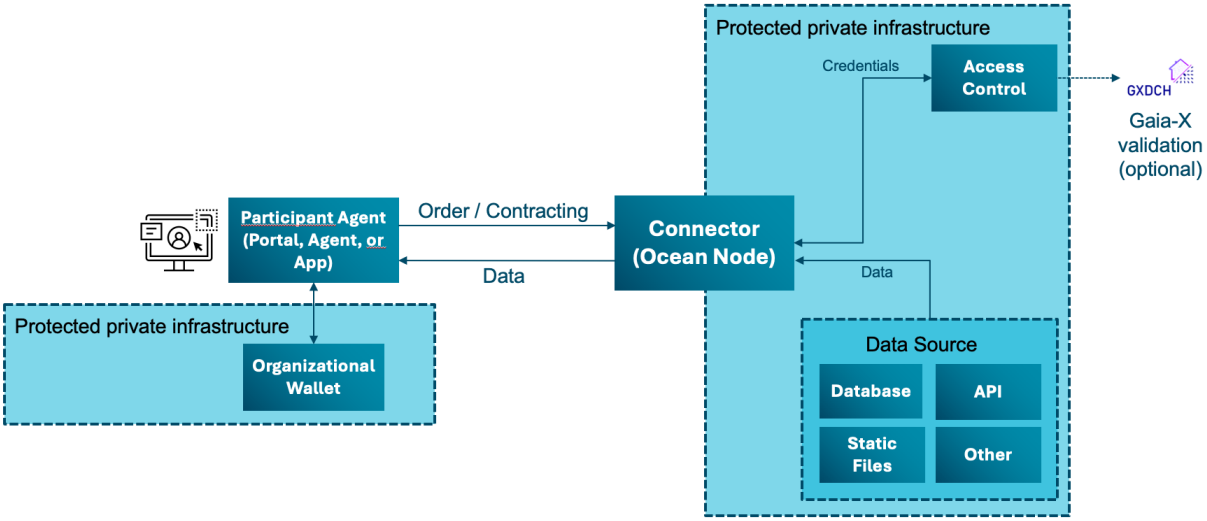


Figure 3 - Basic Connector Setup for Data Services

Data services can be shared from the following sources:

- API / URL incl. Custom Parameters and Auth Headers, e.g. Fiware Context Brokers
- S3-compatible storage
- IPFS
- Arweave
- GraphQL
- FTP/FTPS

4.2.2 Requirements for Consumption through the SENSE Portal

The technical requirements for consuming services through the SENSE data space portal are lower than the requirements for publishing services, and depend on the access controls and policies of the service providers whose services the consumer wishes to access.

The SENSE data space portal is recommended as participant agent to interact with the SENSE data space as it offers an intuitive graphical user interface (GUI).

A proper onboarding and registration with the SENSE data space is necessary which can be done in minutes using the Fast Track Onboarding App and informing the SENSE project after the successful onboarding.

In case the service providers enforce access controls based on the PEP and identity credentials, which is OPTIONAL in SENSE, but MANDATORY for full compliance with the

[19] Service Offering Criteria, https://docs.gaia-x.eu/ontology/development/classes/ServiceOffering/, last seen May 4, 2026.

Trust Framework, the consumer MUST have access to an SSI wallet and the underlying keys and identity credentials.

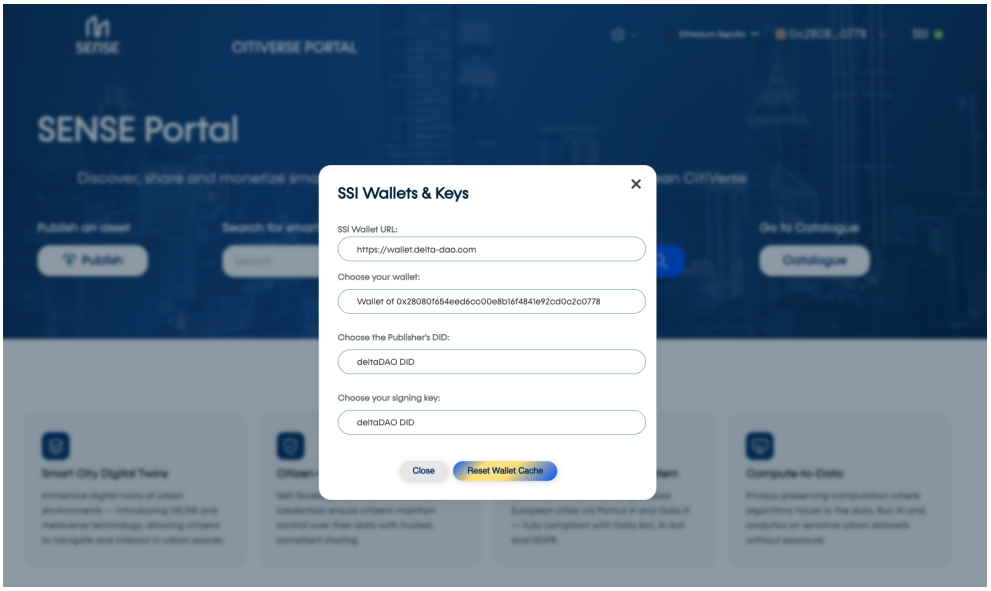


Figure 4 - Successfully connected SSI wallet

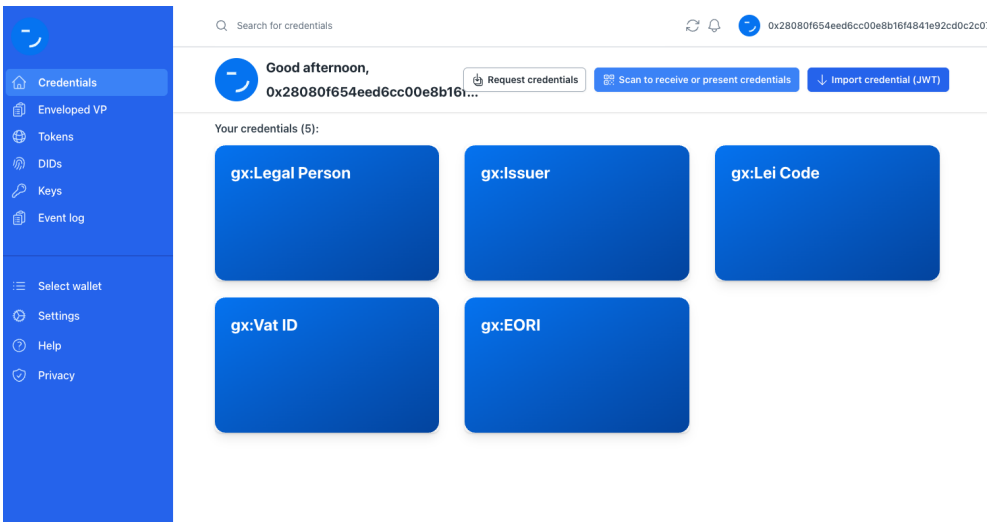


Figure 5 - SSI Wallet with Organizational Credentials

For the advanced method of consumption without a GUI, e.g., if the connection to the SENSE data space shall be facilitated directly through an application, we recommend that developers use the Nautilus[20] library which enables native Typescript-based integration.

Advanced integrations also can make use of a Model Context Protocol (MCP) server or AI Agent Skills. These advanced participant agent integrations become available through the SENSE GitHub Repository[21].

[20] Nautilus Toolkit, https://nautilus.delta-dao.com/, last seen May 4, 2026.
[21] SENSE GitHub Repository, https://github.com/SENSE-EU, last seen May 5, 2026.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

5. Infrastructure and Service Integration Guidelines for the SENSE CitiVerse

SENSE is open for integration on all four levels of the SENSE Architecture.

1. Application / Scenario Level
2. SENSE Data Space Level
3. Data Level
4. Cloud Infrastructure Level

This allows participants to operate all services in a fully sovereign manner to retain full control over their data and applications, while allowing the flexibility to rely on the publicly available services of SENSE if a sovereign deployment is not desired.

Furthermore, we describe how other participants and data spaces can integrate with the SENSE data space on different layers by deploying own instances or using the provided SENSE data space services.

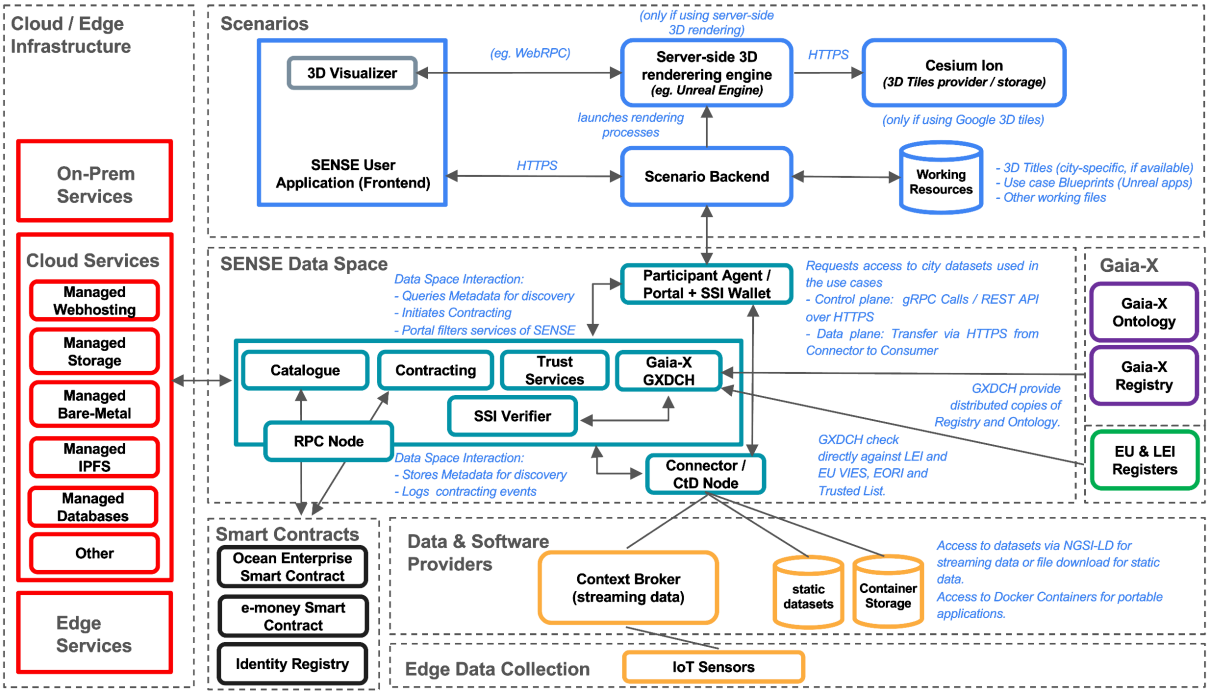


Figure 6 - Simplified Version of the SENSE Architecture

5.1 Integration on the Application and Scenario Layer

The Application and Scenario Layer is the point at which end-user applications interact with the SENSE CitiVerse. It comprises the browser-facing applications used by citizens, visitors and professional users, together with the scenario-specific services that implement the logic required by each use case.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

The objective of integration at this layer is to allow cities and application providers to reuse SENSE data, services and scenario components without reproducing the complete SENSE technical stack or creating unnecessary dependencies on city-specific source systems. Figure 7 shows the recommended interaction pattern.



Figure 7 – Recommended interaction pattern

This represents the main reusable integration path. The User Application is responsible for interaction and visualisation, while the Scenario Service Layer is responsible for scenario logic, configuration, data composition, access mediation and orchestration of analytical or simulation services.

Reusable datasets, APIs, models, geospatial resources, 3D assets and other services that may be consumed by more than one scenario or city should, where appropriate, be published and described as governed resources in the SENSE Data Space. Scenario services can then discover and consume these resources under the applicable identity, licensing, access and usage conditions.

This does not imply that every source used by a scenario must be physically routed through the Data Space. Highly scenario-specific, legacy or transient sources may be connected through dedicated adapters behind the Scenario Service Layer. The important architectural principle is that these source-specific integrations remain outside the User Application and expose a stable interface to the scenario logic.

The resulting pattern is therefore: **User Application → Scenario Service Layer →**

- **SENSE Data Space / reusable resources**, through the appropriate Participant Agent or Data Space Connector; or
- **Scenario-specific adapter → specific source**, where direct integration is justified.

Where a source initially integrated through a scenario-specific adapter later becomes useful across multiple scenarios or cities, it should be considered for harmonisation, metadata enrichment and publication as a reusable Data Space asset or service.

5.1.1 Architecture and Responsibility Boundary

The Application and Scenario Layer consists of two principal components.

SENSE User Application

The User Application is the main interaction point for public and professional users. It runs in a standard web browser and is responsible for presenting the CitiVerse experience, including, depending on the scenario:

- navigation through two- and three-dimensional city representations;
- activation and deactivation of thematic layers;

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

- selection of locations and points of interest;
- consultation of indicators and contextual information;
- presentation of multimedia and narrative content; and
- professional or role-specific interaction where applicable.

Geometry and visual data are rendered on the client. The application should use progressive loading, suitable levels of detail and appropriate caching so that the same scenario can operate on mainstream mobile, tablet and desktop devices. The User Application should not contain city-specific data adapters, privileged credentials or direct integrations with municipal databases, Context Brokers or proprietary provider APIs.

Scenario Service Layer

The Scenario Service Layer contains the use-case-specific logic that should not be embedded in the user interface. Its responsibilities may include:

- resolving scenario configuration;
- determining available layers, viewpoints, indicators and content;
- composing information from multiple sources;
- applying scenario-specific rules;
- mediating role-specific access;
- orchestrating analytical, simulation or AI services;
- invoking Data Space resources;
- invoking scenario-specific adapters where required;
- transforming results into application-oriented responses; and
- exposing stable APIs to the User Application.

Analytical, simulation and AI services should return structured outputs that can be interpreted and visualised by the application, rather than requiring server-side rendering of the complete user experience.

Integration boundary

The Scenario Service Layer defines the boundary between scenario-specific application logic and external data, services and assets.

Reusable resources intended to be shared across scenarios or cities should preferably be published as governed offerings in the SENSE Data Space and accessed through the corresponding Participant Agent or Data Space Connector. This provides a common mechanism for discovery, identity, access control, licensing, usage conditions and reuse.

The Data Space does not necessarily need to carry the complete data payload itself. Depending on the service offering, the Data Space may provide discovery, contracting, policy and access-control functions while the actual data exchange takes place through the provider's technical endpoint.

Where a scenario depends on a highly specific, legacy or transient source, the Scenario Service Layer may access it through a dedicated adapter. Such adapters should remain behind the scenario-service boundary and should expose a stable, documented interface so that provider-specific complexity does not propagate to the User Application.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

This distinction enables SENSE to combine two goals: maximise reuse through the Data Space while retaining enough flexibility to integrate sources that are meaningful only within a particular use case.

5.1.2 Common Application Integration Guideline

Regardless of how a scenario is deployed, applications integrated with SENSE should follow a common integration guideline.

Concern	Developer guideline
Application interfaces	Communication between the User Application and Scenario Service Layer should use documented web APIs over HTTPS.
Dynamic context information	NGSI-LD and compatible shared data models should be preferred where applicable.
Source-specific integration	Proprietary and legacy integrations should be encapsulated behind the Scenario Service Layer.
Reusable resources	Data, services and assets with cross-scenario or cross-city reuse value should be described and published through the SENSE Data Space where appropriate.
Scenario-specific sources	Sources with limited reuse value may be integrated through dedicated adapters.
Geospatial information	OGC-compatible interfaces and open geospatial formats should be reused where available.
3D resources	OGC 3D Tiles, glTF and CityJSON should be used where appropriate to the scenario and resource.
Identifiers	Stable identifiers should link the same real-world object across context data, geometry, content and analytical results.
Authentication	Secrets and privileged credentials must not be embedded in client-side code.
Configuration	Scenario configuration should be externalised from application code wherever possible.
Data-source binding	Applications should use scenario-service abstractions, catalogue identifiers or configuration bindings rather than hard-coded provider endpoints.
Rendering	Client-side rendering should use progressive loading, appropriate levels of detail and defined performance targets.

This guideline separates reusable application logic from the technical characteristics of the underlying city systems and makes it possible to substitute equivalent resources when a scenario is deployed in another city.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

5.1.3 Option A: Reusing or Embedding an Existing SENSE Application

This is the lightest application-integration option and is suitable for cities or organisations that already operate a municipal, tourism, cultural or other public web environment.

Under this option, an existing SENSE application or scenario is made accessible through the host environment while the scenario logic continues to be provided by the corresponding Scenario Service Layer.

The host does not need to:

- reproduce the underlying data;
- implement its own Data Space integration;
- maintain an independent scenario configuration; or
- operate dedicated server-side rendering infrastructure.

The host environment should provide HTTPS delivery, a responsive presentation context and the browser-security configuration required to access the SENSE application.

Where authenticated professional functionality is provided, authentication should be delegated to the appropriate scenario or identity service. The host application should not store privileged Data Space credentials.

Scenario configuration remains external to the host. Changes to available layers, viewpoints, indicators or content can therefore be propagated without requiring the host environment to implement the same changes independently.

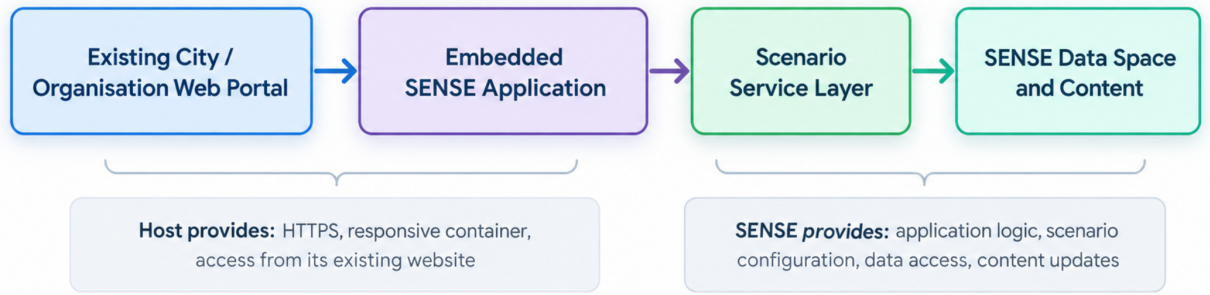


Figure 8 – Embedding an existing SENSE application

5.1.4 Option B: Deploying a Custom Scenario Application

A participant that requires greater control over branding, user interaction, functionality or deployment may develop and operate its own User Application.

The custom application may reuse an existing Scenario Service Layer or the participant may deploy its own scenario-service implementation.

Where an existing Scenario Service Layer is reused, the custom User Application remains primarily responsible for presentation and interaction.

Where the participant also operates its own Scenario Service Layer, it becomes responsible for:

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

- resolving scenario configuration;
- mediating access to protected resources;
- consuming reusable Data Space assets and services;
- implementing or invoking scenario-specific adapters;
- maintaining mappings between source representations and scenario representations;
- orchestrating analytical or simulation services; and
- exposing stable interfaces to the User Application.

The application itself should remain independent of the underlying data-provider infrastructure.

Client-side rendering remains the preferred SENSE approach. Applications should therefore implement appropriate levels of detail, progressive loading and caching and validate performance on representative target devices. Content that cannot reasonably be generated or processed at runtime may be prepared in advance and delivered as static or pre-processed assets.

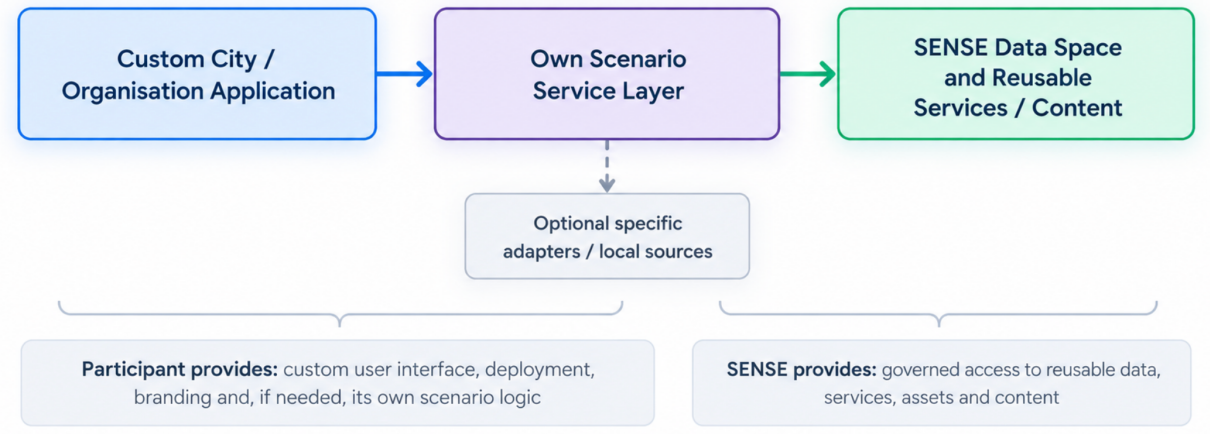


Figure 9 - Custom Scenario Application

5.1.5 Scenario Portability and Reuse

Application deployment and scenario portability are related but distinct concerns. A deployment option defines where and by whom an application is operated, while scenario portability defines how the same scenario logic can be reproduced in another city using different local data and assets.

A portable SENSE scenario should therefore separate: application code, scenario configuration, resource bindings, reusable assets/content, and access and policy configuration.

Scenario configuration should describe what the scenario requires without embedding unnecessary references to a specific provider infrastructure.

For example, an application should not require a hard-coded Cartagena-specific URL to obtain air-quality information. Instead, the scenario configuration should identify the required resource or service, while the local deployment binds that requirement to the corresponding available offering.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Resources intended for reuse may include:

- scenario configuration templates;
- catalogue and service identifiers;
- client-side rendering profiles;
- reusable data-model mappings;
- 2D and 3D assets;
- content and multimedia structures;
- scenario-service components;
- adapters that are generic enough to be used by more than one scenario; and
- data-source onboarding procedures.

This approach allows a follower city to retain the application and scenario logic while replacing the city-specific resource bindings with equivalent local resources.

Scenario-specific adapters should also be reviewed from a reuse perspective. If an adapter becomes useful for several scenarios, providers or cities, it should be generalised where practical and considered for inclusion in the reusable integration stack.

Licensing, attribution and provenance conditions remain attached to reused assets. Where reconstructed or interpretative content is reused, its source and validation status should remain traceable.

5.1.6 Cross-cutting Requirements

The following requirements apply regardless of the selected integration option.

Requirement	Application-layer guideline
Access control	Access restrictions attached to data products and services must remain effective when those resources are used by a scenario.
Credential handling	Privileged credentials must remain in trusted backend components and must not be embedded in browser code.
Privacy and data minimisation	Personal data should be processed only where necessary for the scenario and where appropriate legal, organisational and technical safeguards are in place.
External AI/analytics	Restricted or personal information must not be submitted to external analytical, generative or AI services unless this processing has been explicitly assessed and authorised.
Accessibility	Interfaces should follow WCAG 2.1 AA principles where applicable. Essential functionality should not rely only on colour or pointer-hover interaction.
Language and readability	Citizen-facing content should be available at least in the relevant local language and written for non-expert users.
Performance	Performance targets and scene budgets should be validated on representative devices.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Requirement	Application-layer guideline
Graceful degradation	Core scenario functionality should remain usable when optional analytical, simulation, AI or external services are temporarily unavailable.
Portability	City-specific endpoints and implementation details should remain separate from reusable application code.
Traceability	Derived information should remain traceable to the relevant source data and transformation logic.

5.1.7 Alignment with the Minimum Interoperability Mechanisms

Application and scenario integration relies on several Minimum Interoperability Mechanisms.

Integration concern	Main MIM relationship
Stable identifiers and relationships between entities and resources	MIM1 – Interlinking Data
Common representation of context information	MIM2 – Representing Data
Governed exchange and reusable service offerings	MIM3 – Exchanging Data
Personal-data-related scenario requirements	MIM4 – Personal Data
Secure and controlled access	MIM6 – Securing Data
2D/3D and geospatial interoperability	MIM7 – Geospatial Data
Reuse of Local Digital Twin data, services and components	MIM8 – Local Digital Twins

MIM7 is particularly relevant to geospatial resources exchanged with the application and scenario layers, including spatial data, two- and three-dimensional representations, persistent identifiers and coordinate-reference information.

MIM8 is particularly relevant to scenario portability because it promotes the ability of Local Digital Twins to share and reuse data, services and components across implementations.

In SENSE, these principles are complemented by the Data Space integration model. Resources with reuse value can be published as governed offerings, discovered by scenario services and consumed under defined policies. Scenario-specific resources can still be integrated through adapters where required, without forcing every local integration into the common catalogue from the outset.

This hybrid approach provides a practical path from local experimentation to reusable CitiVerse components: a source or adapter can initially serve one scenario and later be harmonised and promoted into the shared Data Space when its reuse value has been demonstrated.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

5.2 Technical Integration in the SENSE Data Space Layer

The SENSE Data Space Layer is open to participation by other interested parties, subject to the decisions made by the data space governance authority (DSGA).

This differs from other data space frameworks in that SENSE allows each component of the data space to be provided as a distributed/federated system. This is particularly important for components relating to technical data sovereignty and control, such as data storage, data processing and access control.

5.2.1 Deployment of an Own Connector / Access Controller

The Ocean Enterprise Node[22] is the core of data sovereignty. It is a REST API that functions as the technical gatekeeper for data assets and Policy Enforcement Point (PEP). The Provider acts as the PEP, intercepting all data requests. It must be configured to utilize REGO policies via the Policy Server, which acts as Policy Decision Point (PDP), to evaluate the consumer's Verifiable Presentation (VP) and must have access to the state of the smart contracts for contracting / access token verification. The Access Controller, also called connector, requires as dependency:

- an IPFS Storage Service for the storage of service Metadata that can either be procured as a managed service or an own IPFS node to be setup.
- an Elasticsearch Database for the catalogue search functionality to allow for performant queries against the catalogue ground truth.
- An RPC Node to fetch data from the smart contracts, which can either be procured as a managed service (recommended) or self-hosted. SENSE also offers an RPC Node with a proxy solution as a service.
- A policy server as Policy Decision Point (PDP) to evaluate services.

The node can be installed either on-premise, on the edge, or in a managed cloud environment. As the node acts as PEP it is a critical component when full technical data sovereignty is required. SENSE also offers a managed node for participants for faster development. Each node is able to query other nodes (if allowed) and be queried by other nodes in the data space (if allowed) to enable a catalogue federation. Entries are tamper-evident, integrity can be verified, as they are stored as credentials and checksums are anchored through the smart contract layer.

Together with the Connector / Access Controller, designated Compute-to-Data (C2D) environments can be made available. C2D brings algorithms to the governed data location so that processing can be executed in a controlled environment and only permitted outputs are returned. Resource sizing is workload-dependent: development nodes may be modest, while production C2D environments must be sized for the algorithm, concurrency, data volume and

[22] Ocean Enterprise Repository for Node, https://github.com/OceanProtocolEnterprise/ocean-node, last seen May 7, 2026.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

any GPU/HPC requirements. The operator should define resource limits and monitor actual consumption rather than relying on a fixed project-wide hardware profile.

For small/edge deployments, a single node with a local Docker daemon can be used.

The C2D environment requires a layered software stack to manage the lifecycle of compute jobs:

- Container Engine: Docker Engine is mandatory. The Ocean Node interacts with the Docker daemon to spin up isolated containers for algorithms / portable applications. It enables container lifecycle management, CPU core pinning, GPU passthrough, security hardening and image lifecycle management.
- Storage Class: Mounting a persistent Storage is important and must be configured to create temporary volumes where the algorithm can access raw data (read-only) and store its results. Compute jobs are stored in a sqlite database.

Note: As the provisioning and usage of the Compute features introduces costs for the operator of the node the compute resources can be priced to enable a sustainable operational model. This is part of the configuration of the node, where users then will be charged for the usage, depending on the usage of resources.

5.2.2 Deployment of an Own Policy Server

As part of the Identity, Credential, and Access Management (ICAM) stack, the Policy Server[23] functions as the Policy Decision Point (PDP). While the Connector / Access Controller serves as the enforcement gate (PEP), the Policy Server is the specialized "decision engine" that translates complex legal requirements into binary access outcomes.

For DevOps engineers, operationalizing the Policy Server involves the following technical features:

- Identity Verification: It inspects the requester's **Verifiable Presentation (VP)** to ensure they hold, e.g., a valid Gaia-X Participant Credential, verified via OIDC4VP. It also performs the forwarded and additional credential checks against the Gaia-X Digital Clearing Houses (GXDCH).
- Open Policy Agent: The SENSE ecosystem utilizes the OPA standard for decentralized policy evaluation. This allows the decision logic to be decoupled from the data-streaming middleware, facilitating more robust and scalable security architectures.
- REGO Policy Language: To ensure production readiness and direct compatibility with the SSI stack, policies are written in REGO, a high-level declarative language. This choice replaces the more abstract ODRL (Open Digital Rights Language) for dynamic, infrastructure-level enforcement. However, an REGO to ODRL translation can be included if required.

[23] Ocean Enterprise Repository for Policy Server, https://github.com/OceanProtocolEnterprise/policy-server, last seen May 7, 2026.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

- Machine-Readable Translation: When a service provider defines a Data Usage Agreement (DUA) and access policies as part of the service metadata, the specific clauses (e.g., "Gaia-X participant only," "specific organization only") are mapped to REGO scripts. These scripts are then stored within the service offering's metadata and referenced directly in the Service Credential, creating a cryptographic link between the asset and its rules.

As the policy server performs the identity and access credential check, it must be deployed alongside a sovereign access controller. SENSE offers a shared instance of a node for easy development without the need for an own setup, but for technical data sovereignty and sovereign access control enforcement this is a recommended component.

Policy-server and wallet resources should be sized against the expected number of participants, policy checks and concurrent authorization flows. Production deployments should include monitoring and capacity thresholds so that identity and policy evaluation do not become a bottleneck.

5.2.3 Deployment of a Participant-operated Portal

SENSE provides a shared portal/marketplace for interacting with the Data Space. Participants that require greater autonomy may also deploy and operate their own portal, for example to apply organisation-specific branding, expose tailored information, integrate additional applications or reduce dependency on shared user-facing services.

A participant-operated portal may query the SENSE catalogue and, where permitted, federated catalogues from compatible data spaces. It may also apply its own filtering and presentation rules to create a customised user experience while relying on the same underlying catalogue, contracting and access-control mechanisms.

The SENSE catalogue can also be exposed to other compatible portals, enabling service offerings to be discovered and reused across a wider CitiVerse ecosystem.

The minimum technical dependencies for a portal are:

- at least one Access Controller / Connector capable of querying the catalogue; and
- an RPC endpoint or node capable of retrieving the smart-contract state required by the portal.

A portal therefore acts as a participant-facing interaction layer rather than as the authoritative source of catalogue or policy information. Deploying an alternative portal does not change the underlying Data Space governance, identity, contracting or access-control mechanisms.

Through catalogue federation, a participant-operated portal may expose offerings from SENSE together with offerings from other compatible data spaces, subject to the corresponding federation and access policies.

5.2.4 Deployment of Own Public Key Infrastructure (PKI) and Digital Identities

The SENSE Data Space follows principles of digital sovereignty, decentralised identity and participant-controlled credentials, in line with the Gaia-X Trust Framework and the recommendations referenced by DSSC, DS4SSCC and the LDT Toolbox.

For development and initial onboarding, SENSE participants may use the Fast Track Onboarding application. This provides a simplified route for establishing the credentials required to interact with the development environment.

For a fully sovereign deployment, a participant operates and controls its own organisational identity, key material and associated public identity infrastructure. This approach is intended for production environments in which the participant requires direct control over its identity lifecycle.

A sovereign identity setup requires:

- ownership and secure management of the relevant private keys;
- verification of the organisation's legal identity by an accepted Trust Service Provider, using an appropriate X.509 certificate such as a QWAC where required by the applicable trust framework; and
- control of a domain and associated web hosting through which the required public identity information can be resolved.

The participant-controlled domain is used to publish the public information required for identity and signature verification. This may include the relevant X.509 certificate chain and the DID document associated with the participant.

A DID:WEB identifier is created using the controlled domain and the participant's key pair. The resulting DID document must be publicly resolvable so that other components can resolve the participant identity and verify signatures.

Participant Verifiable Credentials can then be created using the Gaia-X ontology and the corresponding credential-generation tools referenced by SENSE. The resulting credentials are stored and presented through an SSI wallet as described in Section 5.2.5.

5.2.5 Deployment of an Own SSI Wallet

The An SSI wallet allows a participant to store and manage its identity credentials and to present the required evidence when interacting with protected services in the SENSE Data Space.

When an access policy requires identity or organisational attributes to be verified, the wallet creates a Verifiable Presentation containing the credentials authorised by the participant. The presentation is then evaluated by the Policy Decision Point (PDP), while the corresponding Policy Enforcement Point (PEP) enforces the resulting access decision.

This mechanism can be used, for example, where access depends on a participant's organisational status, Legal Registration Number (LRN), role or other attributes defined by the service provider's policy.

Credential exchange should use the mechanisms supported by the SENSE Trust Framework, including OpenID for Verifiable Presentations (OIDC4VP) where applicable.

SENSE provides an SSI wallet implementation that supports the Trust Framework and the mechanisms required for integration with the SENSE Portal and other participant agents.

A self-operated wallet requires:

- web hosting for the wallet frontend and APIs;
- persistent storage, such as PostgreSQL, and appropriate secure storage for sensitive key material, such as HashiCorp Vault where applicable;
- an established DID:WEB identity;
- participant Verifiable Credentials compliant with the Trust Framework; and
- access to the private key associated with the participant DID and used to sign Verifiable Presentations.

Once configured, the wallet can be used with the SENSE Portal or other compatible participant agents. Where a participant does not require a self-operated wallet, a managed or shared SENSE wallet instance may be used where permitted by the applicable governance and security requirements.

5.3 Integration in the SENSE Data Layer

Data storage for the integration in SENSE can come in various forms, such as file servers, databases, IPFS or S3-compatible storage, or, e.g. the FIWARE Context Broker. We will not go into the details of these standard solutions.

In any case, if data from these resources is to be shared, a publicly available endpoint must be made accessible via the access controller as a private or public protected connection, which is being encrypted during publication. The access controller is responsible for decrypting these endpoints after the (positive) policy decision has been made, without ever revealing them to the consumer, thereby ensuring the data owner's data sovereignty.

5.3.1 Integration of a FIWARE Context Broker

The FIWARE Context Broker is integrated into the SENSE Data Layer as a source of dynamic context information represented according to the NGSI-LD standard. The platform additionally uses Apache Kafka as its internal data backbone, PostgreSQL/TimescaleDB for persistent storage of time-series data, Keycloak for identity and access management, and a Kong API Gateway for controlled exposure of northbound services.

Two main integration patterns are supported:

1. **Direct integration of an NGSI-LD Context Broker**. An external FIWARE-compatible Context Broker may push context updates to the SENSE ingestion interface. This approach is suitable when the source platform already provides data in NGSI-LD format and can actively notify or forward entity updates.

2. **Integration through a data adapter.** Where the source Context Broker uses a different NGSI version, data model, authentication mechanism, or payload structure, a dedicated ingest adapter is deployed. The adapter retrieves or receives the source data, validates it, maps it to the agreed SENSE data models and forwards the harmonized records to the platform's Kafka backbone. This mechanism isolates source-specific implementation details and provides a uniform representation to the remaining SENSE components. The reference implementation uses this approach for integrating external Context Brokers, including the Smart City Platform of Cartagena.

The technical entry point for broker and external-service integration is the SENSE Brokers API:

https://api-broker.sense.dunavnet.eu/swagger/index.html

The API receives data through HTTPS requests, primarily using HTTP POST operations. Incoming payloads are validated before being transferred to the platform processing and persistence services. Depending on the configured integration flow, the data may be published to Kafka topics and subsequently stored in PostgreSQL/TimescaleDB for historical access.

An integration should therefore implement the following sequence:

1. Identify the Context Broker endpoint and the NGSI version supported by the source system.
2. Define the context entities and attributes that will be exchanged.
3. Map the source entities to the data models adopted by SENSE.
4. Configure authentication credentials and secure HTTPS communication.
5. Configure either push-based notifications from the source broker or periodic retrieval through an adapter.
6. Submit the context data to the Brokers API.
7. Verify payload validation, ingestion and storage.
8. Verify that the ingested information can be retrieved through the SENSE data-access interface.
9. Register the resulting data service as a data asset in the SENSE Data Space catalogue and associate it with the applicable access and usage policies.

For push-based integration, the external Context Broker should be configured with an NGSI-LD subscription whose notification endpoint points to the relevant Brokers API operation. The notification payload must contain the required entity identifiers, types, attributes and temporal information. When an adapter is used, the adapter is responsible for translating the source payload into the representation accepted by the SENSE platform.

Data ingested through the Context Broker integration can subsequently be accessed through the Device DataProvider API:

https://api-data.sense.dunavnet.eu/swagger/index.html

This API exposes previously ingested time-series and device data stored in PostgreSQL/TimescaleDB. Access to the endpoint must be protected using the authentication and access-control mechanism configured for the respective data product.

When the resulting API is published as an asset in the SENSE Data Space, its actual endpoint should not be disclosed directly to unauthorized consumers. Instead, it is published through the SENSE access controller, which evaluates the applicable access and usage policies and establishes a protected connection after a positive policy decision. This approach keeps the technical endpoint under the control of the data provider and preserves data sovereignty.

Before production onboarding, the provider should test at least:

- successful submission of a valid context entity;
- rejection of malformed or incomplete payloads;
- correct mapping of entity identifiers, types and attributes;
- correct handling of timestamps and updates;
- persistence of the received measurements;
- successful retrieval through the data-access API;
- authentication and authorization failures;

This integration approach enables existing FIWARE-based city platforms to participate in the SENSE CitiVerse without replacing their local context-management infrastructure. Source-specific differences are handled at the adapter boundary, while NGSI-LD data models, standard APIs and the SENSE Data Space access-control mechanisms provide interoperable and policy-controlled data exchange.

5.3.2 Integration of Portable Applications

For the SENSE CitiVerse project, developing portable applications for a Compute-to-Data (C2D) environment requires adhering to specific standards to ensure interoperability, security, and reproducibility. The C2D infrastructure is designed to bring the "algorithm to the data" rather than moving the data to the algorithm. This happens within a secure, isolated environment. To move applications to this environment the applications must be portable – dockerized - and agnostic to the underlying infrastructure. The hardware abstraction and fulfilment of dependencies is provided through docker images.

1. Isolated Execution Environment: The application runs within a restricted docker environment. By default, the environment has no internet access during execution to prevent data exfiltration. All necessary libraries and dependencies must be pre-installed in a Docker image.
2. Storage and Pathing: The infrastructure mounts data and provides specific paths that the application must use. Input Data is located at /data/inputs. Each input dataset is typically provided in a subfolder named after its DID (Decentralized Identifier). Output Data - any results, logs, or derived data must be written to /data/outputs. Only files in this directory are stored and can be returned after the job finishes.
3. Resource Constraints: Applications must operate within defined CPU, RAM, and GPU limits set by the Provider's infrastructure. The upper limits are defined by the node operator, while – within this limit – the user can select how much resources shall be used (and paid for). This requires the application provider to indicate which resources are required.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Code Availability and Requirements:

1. Publishing as a Service Offering: The code must be published as a service offering (Algorithm) in the SENSE data space. This involves providing a URL to the source code or a compiled artifact.
2. Accessibility: The provider's node must be able to fetch the code. This can be, e.g., a public URL (e.g., a raw file on GitHub or GitLab) or an encrypted URL (managed via the node) to ensure only authorized compute jobs can access the source code. Protection mechanisms like auth tokens, basic auth, and API keys are supported to protect the application code even if public docker images are being used.
3. Entry Point: The developer must specify the exact command to run the code (e.g., python3 main.py or node index.js). This entry point must correspond to the structure of the provided code and the environment set up in the Docker container.

Docker Container Provisioning:

1. Registry Availability**:** The Docker image must be hosted on a reachable container registry. Supported options include:
    a. Public Registries: Docker Hub, GitHub Container Registry (GHCR), or similar.
    b. Private Registries: If using a private registry, the credentials must be provided. Either
        i. User-provided at request time: The consumer passes an encryptedDockerRegistryAuth field in the compute start/initialize request. This is an ECIES-encrypted JSON object containing { username, password } or { auth } (base64-encoded credentials), encrypted with the node's public key.
        ii. Node operator configuration: The node operator can pre-configure registry credentials via the dockerRegistrysAuth field in OceanNodeConfig. This maps registry URLs to auth credentials and is checked as a fallback when no user-provided auth is given.
2. Image Structure:
    a. Base Image: It is recommended to use lightweight, secure base images (e.g., Alpine Linux or slim Python/Node versions).
    b. Dependencies: All requirements (e.g., via pip install or npm install) must be baked into the image during the build process because the container will not have internet access at runtime.
    c. Non-Root User: For security, it is a best practice to configure the Dockerfile to run the application as a non-root user.
3. Versioning: Use specific tags (e.g., v1.0.1) rather than latest to ensure the compute job is reproducible and does not break when the image is updated.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

If applications change, their derived checksums will also change. As security precaution this requires any data provider to re-authorize a new version of the portable application to prevent unintended data loss or processing.

5.4 Integration in the SENSE Cloud Infrastructure Layer

SENSE does not require a specific cloud provider. Components may run on premises, at the edge or in a managed European cloud, provided that the interfaces, trust and policy conditions defined above are preserved. This avoids infrastructure lock-in and lets each city or provider select an operating model that matches local sovereignty, procurement and resilience needs.

Application portability should rely on containerised services, explicit configuration and externalised secrets. Images should be versioned and reproducible; environment-specific host names, credentials and storage locations must be supplied at deployment time rather than embedded in code or scenario configuration.

Production services should use encrypted transport, least-privilege identities, controlled secret storage, health checks and auditable logs. Detailed privacy and cybersecurity controls remain governed by the corresponding SENSE recommendations; D4.3 focuses on the integration boundary rather than duplicating operational security procedures.

The browser-native WP4 architecture reduces central rendering demand because geometry and data are rendered on the user device. Backends therefore scale mainly for API traffic, scenario orchestration, content delivery and governed data access. Public/static assets may use caching and content delivery, while restricted resources must remain on the policy-controlled path.

For sensitive or high-value data, Compute-to-Data can execute approved algorithms where the provider controls the data and return only permitted outputs. It remains part of the same catalogue, identity and usage-control model.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

6. Frontrunner Adoption Guidelines

The Grant Agreement requires D4.3 to include adoption rules for the two frontrunner cities. The rules below use the current four-scenario baseline in D4.1 v1.3 and translate the common integration guidance into concrete implementation profiles. They are not separate architectures: both cities use the same data-space and interoperability principles, while the user-facing technology may differ by scenario.

6.1 Cartagena Adoption Profile

- Use the common WP4 browser-native scenario architecture for Smart Urban Mobility and Public Space, Environment and Quality of Life. The user-facing client renders locally in the browser and obtains scenario configuration, content and results from the Scenario Service Layer.
- Use open geospatial and 3D delivery formats already adopted by the common front end (including OGC 3D Tiles, glTF and CityJSON where applicable) and use MapLibre/CesiumJS/deck.gl-compatible data products rather than server-rendered imagery as the integration contract.
- Treat dynamic mobility, environmental and city-context information as governed data services. Prefer NGSI-LD/compatible models and place any custom source mapping in the data-integration boundary, not in the public browser.
- Professional management views may use named accounts and role-based functions. The lightweight public mobility view must remain usable without privileged access and must expose only information that is licensed and authorised for public consultation.
- Reuse scenario configurations across Cartagena use cases where layers, viewpoints and indicators overlap. Configuration should refer to catalogue resources or stable identifiers, allowing a source to be replaced or versioned without changing the client.

6.2 Kiel Adoption Profile

- Hidden World uses the common browser-native WP4 front-end and is optimised for a small set of curated points of interest. Historical 3D reconstructions, before/after material and panoramas should be pre-produced as lightweight assets with recorded provenance, rights and validation status.
- Where historical geometry is uncertain, the implementation should prefer clearly sourced imagery or panoramas over invented 3D detail. Source and licence information must remain attached to each published asset and derivative.
- SEALEVEL follows the same logical separation between user application, content/scenario services and governed data access but uses its dedicated web stack: Next.js for the application, Three.js/React Three Fiber for client-side 3D, Strapi for content administration and Redis/local caching where appropriate.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

- The public SEALEVEL browser should consume approved content from the scenario backend rather than connect directly to governed Data Space services. The backend is the integration point that retrieves authorised resources through the SENSE participant/data-space path and exposes only the result needed by the public experience.
- Kiel public experiences should be designed for standard browsers and short visitor interactions. Heavy assets require progressive loading and pre-production; content-management and editorial roles remain separated from the public visitor surface.

6.3 Common Adoption Rules for Both Frontrunners and Follower Cities

Rule	What the city/provider does	Why it matters
Catalogue-first binding	Bind scenario configuration to a published data/service identity rather than a fixed endpoint.	Supports portability, controlled replacement and federation.
Adapter at the boundary	Map legacy/custom APIs once in the integration layer.	Prevents city-specific logic from spreading into applications.
Client-side rendering	Stream geometry/data and render in the browser within a declared scene budget.	Reduces central rendering cost and improves replicability.
Rights and provenance	Carry licence, rights-holder and provenance metadata into derived assets and layers.	Allows lawful reuse and traceability.
Role separation	Keep public, professional and content-administration capabilities distinct.	Reduces exposure and simplifies authorization.
Graceful degradation	Core scenario functions remain usable when optional AI/simulation services are unavailable.	Improves operational resilience.
Reuse before replacement	Prefer MIM-aligned, FIWARE/OGC-compatible interfaces and LDT components already used in SENSE.	Reduces vendor lock-in and onboarding effort.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

7. Developer Integration Checklist and Reuse Path

The following sequence is the recommended minimum path for integrating a new data product, service or scenario with SENSE. A developer may use shared SENSE services during prototyping and move to sovereign components later, but the interoperability contract should remain unchanged.

Step	Developer action	Minimum evidence before integration is accepted
1	Choose the integration route	Option A – reuse/embed an existing SENSE application; Option B – deploy a custom scenario application; scenario portability/reuse; or a data/service-only integration.
2	Establish identity and responsibility	Provider/consumer organisation, authorised operator and service ownership are known; production identity requirements are planned.
3	Describe the asset/service	Catalogue metadata includes owner, licence/terms, structure, update behaviour, access conditions and contact.
4	Standardise the interface	Use an agreed NGSI-LD/OGC/open-format interface or document the adapter that creates one.
5	Apply policy-controlled access	Restricted resources are reachable only through the authorised path; the browser or client does not receive standing provider credentials.
6	Integrate the scenario/service layer	Configuration references persistent resources; results are returned as data/content rather than server-rendered imagery unless the scenario explicitly requires otherwise.
7	Validate non-functional requirements	Security, privacy, accessibility, language, provenance, scene performance, failure behaviour and logging are tested.
8	Package for reuse	Version configuration, mappings, containers and documentation; identify generic improvements that should be proposed to SENSE or upstream repositories.

7.1 Reuse and Open-source Contribution Practice

A reusable SENSE integration is more than working code. It includes the interface contract, machine-readable metadata, configuration, versioned mappings, test data where licensing permits, and deployment documentation. Project-specific credentials, city secrets and internal endpoints must never be required to reproduce the software.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Generic fixes or extensions to reused open-source frameworks should be proposed upstream when technically and legally appropriate. SENSE-specific adapters, examples and scenario configurations can be maintained in the SENSE repositories. Repository-level evidence such as issues, pull requests and release notes is best maintained with the code itself; this report therefore defines the contribution rule rather than claiming a static number of upstream contributions.

Before a contribution is considered reusable, the component owner should confirm licence compatibility, remove sensitive configuration, provide a minimal example or test, and document which MIM/LDT capability or integration problem it addresses.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

8. Conclusions

D4.3 turns the SENSE architecture into a developer-facing integration guide. It connects Gaia-X, FIWARE, MIMs and Local Digital Twin reuse with the data-space, data-integration and scenario practices already established by the project. The core principle is separation of concerns: city-specific source complexity is handled at the integration boundary, trust and policy are handled in the Data Space, scenario logic is handled in the service layer, and user experiences remain portable browser applications.

The data guidelines establish a common contract for metadata, formats, semantics, access and provenance. The identity and trust guidelines describe how organisations and services participate in the Data Space, while the infrastructure section provides multiple integration routes ranging from embedding an existing SENSE scenario to operating sovereign connector, policy, wallet and cloud components. This allows developers to start with shared services and increase autonomy without changing the basic interoperability model.

The frontrunner adoption profiles show that interoperability does not require one identical frontend stack. Cartagena and Kiel Hidden World use the common WP4 browser-native geospatial approach, whereas SEALEVEL uses a dedicated exhibition-oriented stack. What remains common is the logical architecture, client-side delivery, governed data access, catalogue-first binding, rights/provenance handling and reuse of standard interfaces.

These guidelines are intended to support the next SENSE iteration and later replication by follower cities. As implementation and validation continue, D4.4 can capture the empirically validated best practices and lessons learned. D4.3 should therefore be maintained as the stable technical baseline: new integrations should conform to the rules defined here, and generic improvements should be fed back into the SENSE and upstream open-source ecosystems where appropriate.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

9. Glossary

The Glossary follows the SENSE, DSSC, and the Gaia-X Glossaries.

Data Space Governance Authority (DSGA)	The body of a particular data space, consisting of participants that are committed to the governance framework for the data space, and is responsible for developing, maintaining, operating and enforcing the governance framework.
Issuer	A role an entity can perform by asserting claims about one or more subjects, creating a verifiable credential from these claims, and transmitting the verifiable credential to a holder.
Trust Framework (TF)	A trust framework is comprised of: • (business-related) policies, rules, and standards collected and documented in the rulebook. • procedures for automation and implementation of the business-related components.
Trust Anchor (TA)	An entity for which trust is assumed and not derived. Each Trust Anchor is accepted by the data space governance authority in relation to a specific scope of attestation.
Trust Service Provider (TSP)	Trust Service Providers (also referred to as Trusted Issuers) are legal or natural persons deriving their trust from one or more Trust Anchors and designated by the data space governance authority as parties eligible to issue attestations about specific objects.
Trusted Data Source	Source of the information used by the issuer to validate attestations. The data space defines the list of Trusted Data Sources for the Data Space Conformity Assessment Scheme/s.
Notary	Notaries are entities accredited by the Data Space, which perform validation based on objective evidence from a data space Trusted Data source, digitalising an assessment previously made.
Verifiable Credential (VC)	A verifiable credential is a tamper-evident credential that has authorship that can be cryptographically verified. Verifiable credentials can be used to build verifiable presentations, which can also be cryptographically verified
Sovereign Digital Identities	Participants use Self-Sovereign Identity (SSI) wallets to manage their digital identities. Participants prove their identity using credentials and private keys rooted in established trust service providers.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

Gaia-X Compliance	Every service and participant can be automatically verified against the Gaia-X Registry and Gaia-X Compliance Standards using a Compliance Engine embedded in the GXDCH, ensuring that only valid, legally recognized entities can operate within the CitiVerse and that the de-facto Trust Framework and semantic standards of Gaia-X are upheld.
Policy Enforcement Point (PEP)	Component that ensures that customized access control rules (e.g., "EU-only," "Actor with a specific Legal Entity Identifier") are strictly enforced at the infrastructure level before any data is exchanged or utilized.
Policy Decision Point (PDP)	Component that ensures that credentials are validated against the policies of a service provider. The outcome can be used by the PEP for access control.

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

10. Abbreviations

API	Application Programming Interface
C2D	Compute-to-Data
CI/CD	Continuous Integration / Continuous Delivery
CMS	Content Management System
CRS	Coordinate Reference System
DCAT	Data Catalogue Vocabulary
DCAT-AP	DCAT Application Profile
DID	Decentralized Identifier
DSSC	Data Spaces Support Centre
DS4SSCC	Data Space for Smart and Sustainable Cities and Communities
eIDAS	Electronic Identification, Authentication and Trust Services
FIWARE	Open-source framework for smart solutions and context data management
GIS	Geographic Information System
glTF	GL Transmission Format
LDT	Local Digital Twin
MIM	Minimum Interoperability Mechanism
NGSI-LD	Next Generation Service Interfaces - Linked Data
OASC	Open & Agile Smart Cities
OGC	Open Geospatial Consortium
OIDC	OpenID Connect
OIDC4VP	OpenID for Verifiable Presentations
OPA	Open Policy Agent
PDP	Policy Decision Point

Funded by the European Union. 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 Commission. Neither the European Union nor the granting authority can be held responsible for them.

PEP	Policy Enforcement Point
PKI	Public Key Infrastructure
QWAC	Qualified Website Authentication Certificate
RBAC	Role-Based Access Control
REGO	Policy language used by Open Policy Agent
REST	Representational State Transfer
SHACL	Shapes Constraint Language
SSI	Self-Sovereign Identity
T&C	Terms and Conditions
TSP	Trust Service Provider
URL	Uniform Resource Locator
VAT	Value Added Tax
VC	Verifiable Credential
VIES	VAT Information Exchange System
VP	Verifiable Presentation
W3C	World Wide Web Consortium
WCAG	Web Content Accessibility Guidelines