



CIRCULess

Upcycling mineral and timber-based waste from construction & manufacturing process industries through ecodesign, advanced logistics, quality control and digital solutions

D1.1

Digital Platform Alpha Version

Grant Agreement Number:	101138330
Start Date of the Project:	1 July 2024
Duration:	42 months
Deliverable lead beneficiary:	BVR
Deliverable Dissemination Level:	PU
Deliverable Nature:	R
Due date of deliverable:	30/09/2025
Actual submission date:	06/10/2025



Funded by
the European Union

Authors

Name	Partner
Jorge Almela Miralles	BVR
Peter Drahovsky	BVR
Stefan Vanya	BVR
Maria Kosarli	GCS
Gaylord Booto	GCS
Lizhen Huang	NTNU
Xiongwei Huang	NTNU
Foivos Psarommatis	ZERO
Irina Kalb	ZERO
Christoforos Kanellakis	LTU
Rocco Lagioia	TGT

Revision history

V#	Date	Name	Partner
0.1	01.08.2025	Jorge Almela Miralles	BVR
0.2	04.09.2025	Jorge Almela Miralles	BVR
0.3	12.09.2025	Ekain Cagigal Montalban	TEC
0.4	15.09.2025	Georgios Triantafyllidis	NTNU
0.5-0.7	03.10.2025	Jorge Almela Miralles	BVR
0.8	06.10.2025	Sotirios Grammatikos	NTNU

Technical references

Funding Scheme	HORIZON-CL4-2023-TWIN-TRANSITION-01-42
Project Coordinator	Professor Sotirios Grammatikos Norwegian University of Science and Technology (NTNU) sotirios.grammatikos@ntnu.no
Project Website	www.circuleless-project.eu
Service	HADEA/B/03

Disclaimer and Proprietary Rights Statement

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101138330.

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 HADEA. Neither the European Union nor the granting authority can be held responsible for them.

Table of Contents

Authors	1
Revision history	1
Technical references.....	1
Disclaimer and Proprietary Rights Statement	1
List of Figures	5
List of Tables.....	7
List of Abbreviations.....	8
Executive summary	10
1 Introduction	12
1.1 Purpose of the deliverable	14
1.2 Relation to other tasks.....	14
2 Background work and inputs	16
2.1 Standards and Data Models.....	16
2.1.1 Interoperability, APIs, and Data Models for Context Descriptors.....	16
2.1.2 Building and Geospatial Models	17
2.1.3 Communication Protocols and Security	19
2.2 State of the art and existing platforms.....	20
2.3 Business requirements.....	23
3 Requirements	25
3.1 Functional requirements	26
3.1.1 Decentralised environment requirements.....	26
3.1.2 Waste management platform requirements.....	27
3.1.3 Digital product passport requirements.....	28
3.1.4 Data requirements	28
3.2 Non-Functional requirements.....	28
3.2.1 Privacy and security.....	28
3.2.2 Availability and performance.....	29
3.3 Interface/UI/UX requirements.....	29
3.4 Public policy and regulatory requirements.....	29

4	Architecture	31
4.1	Overall design	32
4.1.1	Targeted users.....	33
4.1.2	Digital Waste Management Platform.....	33
4.2	Information model	36
4.2.1	Waste Management Platform Data Models	36
4.2.2	Digital product passport data model.....	37
4.2.3	Advanced use of semantic capabilities	37
4.3	Concept model	37
4.4	Logical view	38
4.5	Process view	38
4.6	Interfaces view	39
4.7	Information view	39
4.8	Deployment view	40
4.9	Development view	41
4.10	Security and quality considerations	41
4.10.1	Security Considerations.....	42
4.10.2	Quality Considerations	43
5	Conclusions.....	45
6	References.....	47
ANNEX I: Review of DPP tools in circular construction		48
ANNEX II: Technical annex covering the architecture of CIRCULess DWMP		49
AII.1	Overall design	49
AII.1.1	Concept and main characteristics.....	49
AII.1.2	Targeted users	50
AII.1.3	Digital Waste Management Platform	50
AII.1.4	Decentralised environment	52
AII.2	Information model	54
AII.2.1	Waste Management Platform Data Models	55
AII.2.2	Digital Product Passport data model	56

AII.2.3 Advanced use of semantic capabilities	59
AII.3 Concept model	60
AII.3.1 Actors	61
AII.4 Logical view	61
AII.4.1 CIRCULess DWMP	61
AII.4.2 Decentralised environment	62
AII.4.3 How authority and access to data works	63
AII.5 Process view	65
AII.5.1 Joining CIRCULess	65
AII.5.2 Connecting a partner infrastructure or service	69
AII.5.3 Contracts and access management	77
AII.5.4 Data Discovery and Consumption.....	85
AII.5.5 Data Ingestion and Harmonisation	89
AII.6 Interfaces view	90
AII.6.1 CIRCULess DWMP interfaces	91
AII.6.2 CIRCULess DWMP Cloud interfaces.....	92
AII.6.3 CIRCULess DWMP Node interfaces	92
AII.7 Information view	93
AII.7.1 CIRCULess Waste Management Platform information flows	93
AII.7.2 CIRCULess DWMP Cloud information flows	94
AII.7.3 CIRCULess DWMP Node information flows	95
AII.8 Deployment view	95
AII.8.1 CIRCULess DWMP Cloud	96
AII.8.2 CIRCULess DWMP Node	96
AII.8.3 CIRCULess DWMP	96
AII.9 Development view	97
AII.9.1 CIRCULess Digital Waste Management Platform	97
AII.9.2 Decentralised environment	104
AII.10 Security and quality considerations.....	108
AII.10.1 Security Considerations	109
AII.10.2 Quality Considerations.....	110

List of Figures

Figure 1 Structure of Deliverable D1.1.....	13
Figure 2 Task relations.....	15
Figure 3 Architecture of the CIRCULess DWMP.....	33
Figure 4 Components of the CIRCULess waste management platform.....	35
Figure 5 Entity Relationships within the CIRCULess DWMP.	38
Figure 6 CIRCULess Cloud, Platform, and Node Deployment Architecture.	41
Figure 7 Architecture of the CIRCULess DWMP.....	49
Figure 8 Components of the CIRCULess waste management platform.....	51
Figure 9 Decentralised Environment and CIRCULess Node Architecture.....	53
Figure 10 Entity Relationships within the CIRCULess DWMP.....	60
Figure 11 User and Organisation Registration Process.....	66
Figure 12 User Invitation Process.	68
Figure 13 User Removal Process.....	69
Figure 14 Node Registration Process.	71
Figure 15 Node Removal Process.....	72
Figure 16 Asset Registration Process.....	73
Figure 17 Asset Removal Process.....	75
Figure 18 Service Credential Acquisition Process.....	76
Figure 19 Asset Privacy Management Process.....	78
Figure 20 Asset policy Assignment Process.	79
Figure 21 Publishing Offering Process (DCAT).....	81
Figure 22 Offering Removal Process.....	82
Figure 23 Access Request Process.....	83
Figure 24 Access Revocation Process.....	85
Figure 25 Platform Data Discovery Process (user).....	86
Figure 26 Federated discovery request diagram.....	88
Figure 27 Data Ingestion Process (pushing data to the platform).	90
Figure 28 Interfaces View of the CIRCULess Architecture.	91
Figure 29 Information Flows.	93
Figure 30 CIRCULess Cloud, Platform, and Node Deployment Architecture.....	96

Figure 31 CIRCULess Waste Management Detail View..... 97

Figure 32 Identity Management Service - Login Screen.....100

Figure 33 Identity Management Service - Register Screen.....100

Figure 34 Management UI - List nodes.101

Figure 35 Management UI - List items.....101

Figure 36 Management UI - Configure organisation.....102

Figure 37 Management UI - Manage organisation users.103

Figure 38 CIRCULess Cloud detail view.....104

Figure 39 CIRCULess Node detail view.....106

List of Tables

Table 1 Overview of Circular Economy Platforms.....	21
Table 2 European and national circular economy projects.....	22
Table 3 DPP tools	48
Table 4 CIRCULess waste management platform interfaces.	92
Table 5 CIRCULess Cloud Interfaces.....	92
Table 6 CIRCULess Node Interfaces.	93
Table 7 CIRCULess waste management platform information flows.....	94
Table 8 CIRCULess Cloud information flows.....	94
Table 9 CIRCULess Node information flows.	95

List of Abbreviations

AEC - architecture, engineering, and construction
API - Application Programming Interface
BIM - Building Information Modelling
CA – Consortium Agreement
D – Deliverable
DCAT - Data Catalogue Vocabulary
DoA – Description of Action
DPP - Digital Product Passport
DWMP – Digital Waste Management Platform
EC – European Commission
GA – Grand Agreement
HADEA - European Health and Digital Executive Agency
IDM - Identity Management service
IFC - Industry Foundation Classes
IPR – Intellectual Property Rights
JWT - JSON Web Tokens
M – Month
M2M - Machine-to-Machine (M2M)
mTLS - Mutual TLS
OAS3 - OpenAPI Specification 3
OASC - Open & Agile Smart Cities
ODRL - Open Digital Rights Language
PC – Project Coordinator
PM – Person Month
PO – Project officer
PU – Public
SME – Small and Medium Enterprise
T – Task

UML - Unified Modelling Language

UUID – Universally unique identifier

WoT - Web of Things

WP – Work package

Executive summary

The following document is the deliverable D1.1 “Digital Platform Alpha Version”, of the European project “CIRCULess - UpCycling mIneral and timbeR-based waste from COnstruction & manUfacturing process industries through ecodesign, advanced logistics, quality control and digital solutions”, funded by the European Commission (EC) European Health and Digital Executive Agency (HADEA), under its Horizon Research and Innovation Programme. Deliverable D1.1 is the output of Work Package (WP) 1 “Digital platform development building blocks creation and integration”, Tasks T1.1 “Baseline framework for technology development”, T1.4 “Back-end alpha version development” and T1.5 “Front-end alpha version development”.

The primary objective of these tasks within the CIRCULess project is to establish the foundation for a digital waste management platform dedicated to upcycling, with a particular focus on supporting circularity in the construction and manufacturing sectors. In the Alpha version of the platform, efforts will concentrate on gathering requirements and defining the system architecture. The initial implementation will be limited to core components such as data integration, model development, and a preliminary version of the business logic. Subsequent work will continue in WP2, “Digital platform development building blocks creation and integration”, which will prioritise enhancing the platform’s value by introducing micro-services and addressing security aspects, thereby supporting the broader goal of efficient and sustainable waste management through digital innovation. Within the CIRCULess Digital Waste Management Platform (DWMP), a decentralised environment will be established to facilitate the connectivity of partner resources. This component is also integral to the platform’s architecture and is therefore part of this deliverable. However, the Digital Product Passport, which is also a key component of CIRCULess, has its own dedicated deliverable (Deliverable 1.3 “Product passport tool”). Therefore, this deliverable only includes selected sections that reference it, chiefly to describe the interconnections.

By design, the deliverable is dual layered (Chapter 4 and Annex II): Chapter 4 offers a public, accessible synopsis of the architecture and its operational logic, while Annex II presents the full technical architecture for implementers, including detailed models, process diagrams, interfaces, information flows, deployment patterns, and development views. This structure strikes a balance between accessibility for non-technical stakeholders and the engineering precision required by software developers, integrators, and operators to build, test, and deploy the platform. For more comprehensive and detailed technical information, readers are encouraged to consult the source code repositories and their accompanying documentation, which will be periodically updated, throughout the project.

This deliverable, scheduled for release in M15 (September 2025), serves as the verification mechanism for Milestone 1, “Definition of Platform Architecture.” Its contents are intended to guide the implementation of the CIRCULess DWMP by providing a clear architectural framework for development. Upon completion of the platform in M30 (December 2026), CIRCULess will be launched online, and documentation to support the release of D2.1 “Platform Beta Version” will be published. This documentation will include the final release notes, detailing the platform's components, highlighting changes or additions to the architecture from the original design, and presenting the results of validation tests.

The Alpha deliverable establishes a coherent and implementable architecture aligned with European practice and the operational realities of construction and manufacturing value chains, thereby furnishing the stable foundation required for rapid, low risk progression to the Beta stage. By separating public exposition from technical specification and distributing responsibilities across platform, shared services, and organisational nodes, the project combines accessibility, security, and scalability to support large-scale circularity decisions and operations in WP2 and beyond.

Keywords: Digital platform, requirements, related projects, standards, architecture

1 Introduction

This deliverable addresses the definition of the CIRCULess DWMP, which focuses on Construction & Demolition Waste (CDW) management. This platform will draw on existing knowledge and best practices contributed by partners, as well as research on efforts made in various projects or commercial activities.

The CIRCULess DWMP will be connected to a decentralised data sharing environment, based on previous works in Horizon projects and tailored to the needs of CIRCULess. Therefore, the platform is conceived as a decentralised ecosystem, in which distributed nodes can standardise and facilitate access to digital assets. Building upon this foundational layer, a variety of services can be deployed to generate added value from these resources. The envisaged core functionalities include:

- (i) the establishment of fair governance mechanisms and the capacity to model stakeholder interactions;
- (ii) robust security provisions, such as authentication protocols and secure data exchange mechanisms;
- (iii) comprehensive identity management, encompassing validation and certification of participating entities;
- (iv) a service catalogue enabling the discovery, access, and utilisation of microservices, complemented by additional tools that foster collaboration among nodes;
- (v) data sharing flows between entities or with the platform.

A distinctive component of the platform will be the microservice orchestration module, responsible for accessing, navigating, and coordinating microservices through Application Programming Interfaces (APIs) or by retrieving containerised images. This capability will empower developers to design and execute workflows tailored to specific application scenarios within the construction circular value chain.

In general, the platform will be developed in two phases. The initial phase, during WP1 or designated as 'Alpha version', will primarily concentrate on defining the overall system architecture and establishing the core technical capabilities required for the platform. This platform must be capable of managing both data and user interactions. During this stage, a thorough assessment of the consortium's technical strengths and the specific requirements of the pilot projects will be conducted. Emphasis will be placed on gathering information about existing technologies and outlining relevant use cases.

The Alpha version is expected to deliver foundational functionalities, including:

- Asset registration and management
- Data ingestion
- User registration and configuration
- Search and discovery capabilities

These features will serve as the backbone for subsequent development phases, ensuring the platform is aligned with both technical objectives and pilot needs.

A second phase, corresponding to WP2 or designated as 'Beta version', will be dedicated to implementing the core waste upcycling capabilities and associated services within the CIRCULess DWMP. During this phase, the platform will introduce and continuously expand a comprehensive catalogue of microservices, including but not limited to:

- LCA (Life Cycle Assessment), LCC (Life Cycle Costing), and SLCA (Social Life Cycle Assessment) tools
- Zero-defect modelling
- Logistics optimisation at both micro and macro levels
- Materials performance simulation
- Additional microservices developed during project execution or integrated from third-party providers via API

A key objective of this second phase will be to establish ways to collaborate and utilise microservices to generate value from the data available within the consortium and public sources.

This deliverable will be carried out through the following methodological steps, shown in Figure 1. Starting from the collection of requirements and good practices, continuing with the refinement of the system requirements in a feedback loop involving the consortium, and finally defining the architecture. These steps will be reflected in Sections 2 to 4.




Figure 1 Structure of Deliverable D1.1.

1.1 Purpose of the deliverable

The contents of Deliverable D1.1 are strictly limited to the work carried out in Tasks T1.1, T1.4, and T1.5 of WP1:

- T1.1, “Baseline Framework for Technology Development”, focuses on three primary objectives:
 - Gathering relevant inputs to support the development of a digital platform for waste management (to be referred herein as CIRCULess DWMP). Key inputs include use cases defined by pilot projects and end users, widely adopted standards in the sector, and business requirements or best practices identified from similar initiatives. Described in Section 2.
 - Defining the system requirements based on the collected information, ensuring alignment with sector needs and digital transformation goals. Described in section 3.
 - Designing the technical architecture of the platform, informed by the established requirements and industry benchmarks. Described in section 4.
- T1.4, “Back-end Alpha Version Development”, implements the foundational functionality necessary for the platform’s initial operation. This includes data ingestion and storage capabilities, as well as the development of open APIs to facilitate interaction with the platform. Included in sub-section 4.9.
- T1.5, “Front-end Alpha Version Development”, delivers the first iteration of the user interfaces. These early-stage UIs enable users to register and manage assets, providing essential functionality for initial testing and feedback. Included in Sub-Section 4.9.

This deliverable thus provides a focused account of the initial technical groundwork laid for the CIRCULess DWMP, covering requirements collection, architectural design, and the development of both back-end and front-end alpha versions.

1.2 Relation to other tasks

The design and development of the CIRCULess DWMP involves a large number of related tasks, as illustrated in the diagram in Figure 2.

Specifically, this deliverable reports the work done in tasks T1.1, T1.4, and T1.5. These are connected to the rest of the tasks and inputs as described below:

- T1.1, “Baseline Framework for Technology Development,” receives inputs from T7.2, “D&C implementation and synergies with existing projects/initiatives”, and T7.3 “Standardisation activities: Landscape and strategy”, as well as from research conducted on other existing digital

platforms and websites related to circularity. Additionally, contributions are provided by project experts from various fields. The architecture is then evaluated by T2.7 “Functional and security validation of requirements of the digital platform”.

- T1.4, “Back-end alpha version development”, receives design specifications from T1.1, and the resulting implementation is continued in T2.1 “Back-end beta version development”, and tested in T2.7 “Functional and security validation of requirements of the digital platform”.
- T1.5, “Front-end alpha version development”, receives design specifications from T1.1, and the resulting implementation is continued in T2.2 “Front-end beta version development”, and functionally tested in T2.7 “Functional and security validation of requirements of the digital platform”.




Figure 2 Task relations.

2 Background work and inputs

This section collects information to provide a solid foundation for defining the technical requirements of CIRCULess DWMP. This is achieved by reviewing and incorporating relevant technologies, standards, studying examples from other circularity-related platforms, and considering business requirements.

2.1 Standards and Data Models

Digital platforms for waste and circularity management should leverage interoperable standards to ensure data discoverability and secure information exchange among stakeholders, aligned with EU requirements. Open & Agile Smart Cities (OASC) Minimal Interoperability Mechanisms (MIMs) provide a layered strategy that encompasses context management, harmonised data models, trust frameworks, and geospatial capabilities. This comprehensive approach serves as a reference for establishing technical specifications and guiding procurement processes.

This section summarises current best practices as well as EU-endorsed standards for interoperability, APIs, data models, geospatial infrastructure, and secure communication protocols tailored to the requirements of an EU-wide circularity data platform.

2.1.1 Interoperability, APIs, and Data Models for Context Descriptors

OASC Minimal Interoperability Mechanisms (MIMs) establish an effective, practical baseline for interoperability in digital platforms. The MIMs framework covers essential aspects, including Context Management (MIM1), Shared Data Models (MIM2), Contract Mechanisms (MIM3), Trust (MIM4), Transparency and AI (MIM5), and Places/Geospatial Foundations (MIM7). Ongoing advancements also address Security (MIM6), Indicators (MIM8), Analytics (MIM9), and Resources (MIM10). Collectively, these mechanisms provide robust guidelines for designing systems that must facilitate data collection, processing, and sharing across diverse stakeholders.

For APIs, adherence to established standards such as OAS3¹ (Open API Specification) ensures consistent implementation and documentation of service interfaces. Beyond basic API standards, advanced data models substantially enrich the interoperability and expressiveness of interactions, particularly with IoT devices.

¹ [OpenAPI Specification - Version 3.1.0 | Swagger](#)
CIRCULess

Two principal information models widely endorsed in the EU and smart city domains are:

- NGSI-LD² (by FIWARE and standardised by ETSI³): A graph-based context information API and data exchange model, NGSI-LD enables dynamic, interoperable data sharing for smart environments. It is a cornerstone in EU data space infrastructures. The Smart Data Models initiative further enhances NGSI-LD by providing canonical, reusable domain-specific models to harmonise data entities across sectors.
- W3C Web of Things⁴ (WoT) Thing Description 1.1: A JSON/JSON-LD-based specification for describing device metadata, interfaces, and security attributes. WoT Thing Description supports cross-vendor interoperability, operating seamlessly within linked data ecosystems.

Complementary semantic models such as SAREF⁵ (ETSI) offer modular ontologies and reference templates designed for interoperability, reuse, and extensibility. SAREF enables semantic harmonisation for devices, measurements, and actions, integrating with both NGSI-LD entities and WoT Thing Descriptions.

For managing catalogues of digital assets, the W3C's Data Catalogue Vocabulary (DCAT⁶) is commonly employed. DCAT supports federated, discoverable data spaces and promotes easy integration of datasets across platforms and portals.

2.1.2 Building and Geospatial Models

The integration of georeferenced buildings and transportation data within CIRCULess necessitates an examination of current best practices in spatial data management. Concretely, INSPIRE and the OGC API for geospatial (GIS⁷) data and Industry Foundation Classes (IFC⁸) open file format with the building information modelling (BIM⁹) methodology for representing three dimensional and semantic information about buildings, building elements, and components.

In the realm of geospatial data, the INSPIRE Directive, codified by the European Union, serves as a foundational framework for the seamless sharing of spatial (geographic) information across the European continent. Its primary objective is to facilitate the support of EU environmental policies and any activities that impact the environment by empowering public sector organisations and other

² [Official Website for NGSI-LD](#)

³ [ETSI](#)

⁴ [Home - Web of Things \(WoT\)](#)

⁵ [SAREF Portal](#)

⁶ [Data Catalog Vocabulary \(DCAT\) - Version 3](#)

⁷ [Geographic information system - Wikipedia](#)

⁸ [Industry Foundation Classes \(IFC\) - buildingSMART Technical](#)

⁹ [EU BIM Task Group – Aligning national efforts into a common European approach to develop a world-class digital construction sector.](#)

stakeholders to readily access and exchange high-quality, harmonised geospatial data. INSPIRE establishes universally applicable standards and protocols, thereby enabling the interoperability and effortless integration of data collected from various member states. This interoperability is of paramount importance for facilitating informed policy-making, effective disaster management, and sustainable development at both local and continental levels.

This harmonised approach eliminates significant barriers to collaboration, boosts data reliability, facilitates cross-border environmental management, and empowers both governments and citizens with transparent, easily discoverable spatial data. Ultimately, INSPIRE makes geospatial information more accessible and useful, helping to drive efficiency, innovation, and better outcomes for society and the environment across the EU.

The OGC API, which adheres to ISO 19168-1:2020, provides a modern, RESTful approach for retrieving and interacting with geospatial feature data. It supports pagination and offers content in widely used formats like JSON and HTML, with content negotiation to cater to various client needs. The API aligns with web architecture and data management best practices and relates to the W3C Data Catalogue Vocabulary (DCAT) by linking feature collections to dataset and distribution metadata, improving discoverability in federated data ecosystems.

In a practical circular economy or waste management context, the OGC API Features can deliver spatial datasets like waste container locations, collection service zones, road network segments, and waste facility footprints. The standard's support for spatial predicates and attribute-based queries ensures precise and efficient data retrieval. By referencing NGSI-LD context entities, these spatial collections can be semantically enriched and integrated into broader smart city or EU data space infrastructures.

BIM is a methodology that involves the generation and management of digital representations of a building's physical and functional characteristics. It is not merely a 3D modelling tool but a comprehensive process that integrates various types of data into a single, coherent model. BIM enables stakeholders across architecture, engineering, and construction (AEC¹⁰) disciplines to collaborate more effectively by providing a shared digital environment. This environment supports parametric modelling, where building components are defined by parameters and relationships, allowing for dynamic updates and consistency across the model. BIM also incorporates time (4D), cost (5D), and operational data (6D and beyond), making it a powerful tool for lifecycle management of built assets.

¹⁰ [What Is AEC? Exploring the Future of Architecture, Engineering & Construction | ARKANCE NZ](#)

IFC is a standardised, open data schema developed by buildingSMART to facilitate interoperability between different BIM software platforms. IFC written in the EXPRESS language¹¹, which defines a comprehensive set of object classes and relationships used in the construction and facility management industries. These classes include elements like IfcWall, IfcDoor, and IfcSpace, each with attributes and relationships that describe their properties and how they interact with other elements in the model. IFC supports both geometric representations (such as boundary representation and extrusion) and semantic data, enabling the exchange of rich, structured information. These characteristics make BIM the most suitable foundation for DPP in the construction sector.

The integration of GIS and BIM - often referred to as GeoBIM - represents a convergence of detailed building-level information with broader geographic context. BIM provides highly detailed, object-based models of buildings and infrastructure, including geometrical information, and semantic information such as materials, systems, and lifecycle data (e.g., Environmental Product Declaration-EPD). However, BIM models are typically confined to a local coordinate system and lack geographic context. GIS, on the other hand, excels at managing large-scale spatial data but traditionally lacks the fine-grained detail of BIM models. Combining the two enables a more comprehensive view of the built environment.

2.1.3 Communication Protocols and Security

For any robust data platform, particularly in the context of EU waste and circularity management, communication protocols and security mechanisms are critical to ensuring safe, trusted, and compliant operations.

For data transport security, all endpoints should employ HTTPS with TLS and use modern, secure cipher suites to protect data in transit from interception or tampering.

Mutual TLS (mTLS¹²) with X.509¹³ certificates should be implemented to establish strong machine-to-machine trust between nodes or data services. This approach aligns with the stringent security requirements of EU data space regulations and the NIS2 Directive¹⁴, ensuring authenticated and encrypted communication channels.

It is essential to manage the certificate lifecycle, with renewal and revocation procedures to minimise risks from compromised or expired certificates.

¹¹ [EXPRESS Language Foundation | Don't panic, EXPRESS](#)

¹² [What is mTLS? | Mutual TLS | Cloudflare](#)

¹³ [What Is an X.509 Certificate? - SSL.com](#)

¹⁴ [NIS2 Directive: securing network and information systems | Shaping Europe's digital future](#)

The principle of least privilege should be enforced via API gateways and integration with policy enforcement points, restricting access strictly to authorised entities.

Finally, when it comes to identities, authentication and authorization, adoption of the OAuth 2.0¹⁵ / OpenID Connect¹⁶ frameworks are recommended for managing user and client authentication, leveraging JSON Web Tokens (JWTs¹⁷). Best practices to follow include issuing short-lived access tokens, implementing refresh token rotation, using Proof Key for Code Exchange (PKCE) for public clients, and defining token audiences and scopes that are tightly aligned with API resources.

2.2 State of the art and existing platforms

As part of CIRCULess, a systematic review was undertaken of digital platforms and current projects in the circular economy. Use cases and functional capabilities were catalogued and compared using a common rubric to identify gaps and transferable best practices. An overview of the review is shown in Table 1. The objective of this assessment was to determine the feasibility and strategic value of integrating such solutions into the CIRCULess DWMP, as well as to identify opportunities for future technological and operational synergies.

Some of the uses of the studied platforms include retail, search, consultancy services, and support to achieve circularity goals.

Platform	Description	Link
Madaster	Madaster simplifies the management of construction materials throughout their lifecycle, tracking and documenting them as valuable assets for reuse.	https://madaster.com
Reusewood	Resource to help salvage, recycle or reuse wood products.	https://reusewood.org
Opalis	Connect retailers selling reclaim material facilitate their reuse in construction and renovation projects	https://opalis.eu
Loopfront	Loopfront's web-based platform helps organisations save money, reduce carbon emissions and waste by allowing them to reuse	https://www.loopfront.com

¹⁵ [OAuth 2.0 — OAuth](#)

¹⁶ [How OpenID Connect Works - OpenID Foundation](#)

¹⁷ [JSON Web Token - Wikipedia](#)

	everything from furniture and fixtures to building materials.	
Restado	Obtain building materials from deconstruction projects and surplus orders. This practice conserves resources, reduces emissions, and minimises waste.	https://restado.de
Recykal	Waste management platform targeting mostly hazardous waste.	https://recykal.com
BarbourABI	Intelligence solutions for the building construction	https://barbour-abi.com
Globechain	Track ESG data and optimise waste disposal costs	https://globechain.com

Table 1 Overview of Circular Economy Platforms.

Table 2 presents relevant EU and national-level projects in the circular economy domain that have potential synergies with the CIRCULess DWMP.

Project	Link with CIRCULess	Link
RRReMaker	Reuse Reduce Recycle AI-based platform for automated and scalable Maker culture in circular economy	https://www.rrremaker.com/
REFRESH	A novel digital blockchain traceability platform covering the whole supply chain	https://refresh-project.eu/
FREE4LiB	Recycling technologies for lithium-ion batteries and deploy two open platforms: a Battery Passport (BP) and Data-driven models for process optimization	https://www.freeforlib.eu/
DigiPrime	A circular economy digital platform to create circular business models based on the data-enhanced recovery and reuse of functions and materials.	https://cordis.europa.eu/project/id/873111
DIGITAL DECONSTRUCTION	Advanced Digital Solutions Supporting Reuse and High-Quality Recycling of Building Materials	https://vb.nweurope.eu/projects/project-search/digital-deconstruction/
CINDERELA	CINDERELA is a digital platform that offers companies all the information they need to manufacture and apply SRM-based construction materials in buildings and civil engineering projects. It also has a Marketplace, secondary raw materials (SRM), reducing reliance on virgin raw materials and minimising waste.	https://www.cinderela.eu/
S.W.A.N.	A digital Solid Waste reuse platform for the Balkans	http://www.interreg-balkanmed.eu/approved-project/21/

SIRKTRE	Establishing timber value cycles in Norway	https://www.sirktre.no/en
---------	--	---

Table 2 European and national circular economy projects.

Finally, research on technologies and open-source tools that might complement CIRCULess DWMP was conducted. The focus was on tools that are specific to the studied domain rather than generic data platform tools.

- IfcOpenShell is an open-source software library primarily designed to work with the IFC file format, which is widely used for BIM. It provides extensive support for reading, writing, and modifying IFC data and supports several IFC schema versions. Its capabilities include geometry processing, semantic manipulation of building models, and advanced operations relevant to 4D and 5D BIM. The tool facilitates interoperability by supporting multiple IFC serialisations (including IFC-SPF, IFC-JSON, and IFC-XML) and can be integrated into other programs like FreeCAD¹⁸ and Blender¹⁹.
- Speckle is an open, flexible, and extensible design collaboration platform tailored for AEC industries. It eliminates traditional file-based workflows by enabling live data connections between various CAD and BIM software tools, allowing seamless sharing, synchronisation, and versioning of model data. Speckle supports complex geometry, metadata, version history, and inter-element relationships without flattening data richness. Its cloud-based environment features an online 3D viewer for sharing and reviewing models, and APIs for custom integrations and automation. Speckle's open-source nature and broad interoperability enable collaborative, real-time data sharing and management in circular economy applications involving the built environment.
- xBIM Toolkit: An open-source toolkit for developing custom BIM software solutions based on IFC standards. It empowers developers to build specialised BIM middleware, aiding interoperability and data exchange in building information modelling.
- OpenTwins: An open-source platform designed to facilitate digital twin development, covering functionalities from real-time state monitoring to predictive simulation. Though still under development, OpenTwins emphasises open components for comprehensive digital twin applications.
- QGIS and PostGIS: Widely used open-source GIS tools for spatial data management and analysis, useful for mapping and geospatial analysis in waste management and circularity contexts.
- SAREF and NGSI-LD-compatible Smart Data Models: Though not software tools per se, these semantic models and data standards help achieve

¹⁸ [FreeCAD: Your own 3D parametric modeler](#)

¹⁹ [Blender Studio](#)

cross-domain interoperability in IoT and smart city contexts relevant to waste and circular economy platforms.

- Cesium is an open-source 3D geospatial platform—CesiumJS for web visualisation and Cesium ion for tiling/hosting—that streams and renders massive real-world datasets (photogrammetry, point clouds, terrain, imagery, BIM/CityGML) via open standards like 3D Tiles and glTF, enabling high-precision, time-dynamic digital twins in browsers and game engines.
- xeokit SDK: A web-first 3D engine for BIM that loads IFC (typically via compressed .xkt or glTF), renders with WebGL, and exposes rich BIM semantics (hierarchy, properties, sectioning, BCF issues) for custom apps.
- IFC.js: A Three.js-based toolkit that parses IFC in the browser (web-ifc/WASM) and renders meshes directly, enabling property queries, selection, clipping, and lightweight client-side BIM viewers.
- BlenderBIM: A Blender add-on that imports IFC via IfcOpenShell, mapping BIM elements and attributes into Blender so you can view, query, and render BIM models in the 3D viewport and produce visuals.
- three.js is a lightweight WebGL JavaScript engine that enables the visualisation of BIM in the browser by loading model geometry (typically IFC to glTF or meshes from IFC.js), rendering it interactively, and allowing for selection, property lookup, clipping, and measuring via raycasting and materials.
- Unreal Engine is a real-time 3D engine that visualises BIM by importing Revit/IFC/CAD via Datasmith (preserving hierarchy and metadata) and rendering interactively with UE5 features like Nanite/Lumen or path tracing for high-fidelity walkthroughs, simulation, VR/AR, and sequencing.

These tools and frameworks present opportunities to expand the functionality of the CIRCULess DWMP (e.g., as plugins or by extending the CIRCULess DWMP APIs), or to be integrated as part of microservices or workflows. Additionally, **Annex I** presents a review of existing tools for DPP-traceability in circular construction.

2.3 Business requirements

This section outlines the high-level requirements of the CIRCULess consortium for the DWMP.

The CIRCULess DWMP will leverage a suite of microservices to provide enhanced functionality, which will be integrated during Work Package 2 (WP2). Specific use cases for these microservices will be defined in Task 6.1. These use cases will inform how users interact with the platform's decision-making features, data storage, and search capabilities, ensuring that the system aligns with user needs and operational objectives.

1. DWMP: This is the primary digital product, which will encompass the services that end-users (see Section 4.1.2) will interact with directly. Some of its features will be:
 - a. User registration: The platform should have user identities and the capability to manage them.
 - b. Database of waste management information, good practices and resources: Users should be able to access the CIRCULess knowledge base for secondary use, although access restrictions may apply based on data licenses and ownership.
 - c. Orchestrate microservices: Users should be able to integrate their own microservices by either deploying them directly on the platform or utilising the platform as a gateway to expose their services. All services should be accessible via APIs and secured through a unified authentication mechanism.
 - d. Visualisation: The platform should support graphical user interaction through dedicated user interfaces. Some interfaces will be designed for technical users, such as those managing assets or integrating microservices, while others will enable users to visualise dashboards and service outputs.
 - e. Decision-making support based on available data and services: The various microservices available in the CIRCULess project will provide functionalities such as material and process characterisation, optimisation based on cost or emissions, and quality assessment, among others. The platform's decision-making engine will integrate these microservices into comprehensive workflows, enabling it to recommend optimal processes or the most suitable final products based on the available materials and input parameters.
 - f. Raise awareness and communicate results: This will be part of the project website (<https://circuleless-project.eu>), which will also link to the platform.
2. A decentralised data sharing environment, which will complement the waste management platform with data collection capabilities.
 - a. Federated services to support the orchestration of the different nodes connected to CIRCULess DWMP. The main service will be a catalogue of connected nodes.
 - b. Integrate assets: Different data assets should be possible to connect from partners' infrastructures over a decentralised network, keeping them uniquely identified and described with metadata.
 - c. Asset discovery mechanism: Search for assets using their metadata descriptors.
 - d. Asset management and stakeholder collaboration: Allow users to govern their assets and share them with others.
3. Digital Product Passport, which will be integrated with CIRCULess DWMP, with the following features, which are expected to cover the minimum

requirements of the different DPP models and beyond (More information about the DPP functionality contained in the Deliverable 1.3 “Product passport tool”):

- a. Passport creation: Stakeholders can generate a new passport for a material, product, or process, based on its origin, composition, and treatment process.
- b. Data submission: Users can add and update information at each stage of the lifecycle. The DPP²⁰ tool has two types of parameters, static and updatable. This information may include material type, history of treatment, tests of quality and certifications.
- c. Document attachment: DPP supports the upload of corresponding documents (e.g., test reports, certificates) to increase transparency of the passport.
- d. Data validation and review: DPP has the functionality to review submitted information
- e. Visualisation on dashboard: End users can explore passport data through an intuitive interface offering interactive charts, timelines and material flow diagrams.
- f. Lifecycle tracking: Shows the progress of a material or product throughout its key stages: from raw waste through processing actions to the final product, including performed transformations and the stakeholders involved.
- g. Stakeholder collaboration: DPP supports shared access to relevant data for authorised partners, improving collaboration while respecting confidentiality boundaries.
- h. Sustainability insights: DPP summarises environmental performance indicators (e.g., carbon footprint, reuse potential) based on available passport data.

3 Requirements

This section provides an overview of the technical requirements that the digital solution must meet. These requirements are generally divided into two categories:

- **Functional requirements:** Define what the system should do, focusing on its capabilities and specific behaviours or functions.
- **Non-functional requirements:** Specifies the properties the system must have, describing how the system should perform, and imposing constraints on its operation. Non-functional requirements are often referred to as system quality attributes.

Additionally, other important considerations are included in this shortlist, such as system or infrastructure requirements, user interface needs, and regulatory considerations. These aspects ensure that the solution is robust, user-friendly,

²⁰ [EU's Digital Product Passport: Advancing transparency and sustainability | data.europa.eu](https://data.europa.eu)

compliant with relevant standards, and capable of operating effectively within its intended environment.

3.1 Functional requirements

3.1.1 Decentralised environment requirements

FR1 User Accounts

- Users can register, authenticate, update profiles, and delete/unregister accounts with verified deletion workflows.
- Unique user identifiers are required for authenticated sessions, which are necessary for any access.

FR2 Organisations setup

- Create organisations, invite/join, leave/abandon, manage roles and view employee list.
- At least two roles should be available: admin and regular user, with the possibility of extending this based on the project's needs.

FR3 Assets management

- Support different asset types, such could be datasets, catalogues, data services, microservices or others. At least a type for a data producer and consumer must exist.
- Asset metadata can be seen and updated.
- Asset must be uniquely identifiable (PIDs²¹) and have rich descriptors.

FR4 Asset Lifecycle

- Assets can be registered, updated and deleted.
- Asset interactions and access are logged.

FR5 Asset access management and privacy

- Accessibility and sharing rules must be available.
- Owner can request/grant/deny/revoke usage access to other parties.
- Asset privacy or license can be set by owner.

FR6 Control plane and node fabric

- Control plan to provide governance, security policy distribution, identity/trust, node discovery, community/workspace management, and orchestration policies.

²¹ [Persistent identifier - Wikipedia](#)

- Services can run in the platform or on user premises, but access should be offered seamlessly over APIs.
- Register, verify, and list participant nodes with their basic information and public keys for identity verification.

FR7 Search & Discovery

- Search engine to query based on metadata.
- Catalogue to visualise assets.
- Federated discovery, ability to query distributed nodes and compile results.

FR8 Interoperability and open APIs

- All integrated assets can be accessed through Open APIs, which should follow a consistent structure and be clearly defined.
- At least REST/JSON APIs should be supported.
- Unified API gateway and ingress security.

FR9 Policies, consent and collaboration environment

- Usage policies should be defined to express rights/obligations/restrictions.
- The Open Digital Rights Language (ODRL²²) standard will be used to define the policies.
- Participants can agree to share assets access.

3.1.2 Waste management platform requirements

FR10 Integration of the Digital Product Passport.

- Support retrieval, creation, and updating of DPPs via standardised API calls.
- Exchange data using harmonised JSON schemas, including key metadata (materials, carbon footprint, repairability).
- Enforce secure access to all DPP interactions through JSON Web Tokens (JWT) authentication.

FR11 Data ingestion, harmonisation, and storage

- Ingestion of data via adapters.
- Process the ingested data to store metadata about them (Data types, storage path, owners and other relevant information).
- Object storage for the ingested data.

FR12 Microservices orchestration

²² [ODRL Information Model 2.2](#)

- Service store with metadata descriptions of the various tools in the platform.
- All microservices should be accessible over APIs with common authentication mechanisms.
- A more detailed definition will be provided in D2.1 “Platform beta version”, as the main work on the microservices will be done in WP2.

FR13 Decision support

- Tools and microservices provided by CIRCULess partners will be combined to build more complex processes and generate secondary data.
- Concrete and timber use cases are under definition in WP6.
- A more detailed definition will be provided in D2.1 “Platform beta version”, as the main work on the microservices will be done in WP2.

3.1.3 Digital product passport requirements

Information related to the Digital Product Passport is detailed in Deliverable 1.3 “Product passport tool”.

3.1.4 Data requirements

FR14 Data supported types

- Ingest unstructured data (Images/binaries/text/PDFs; ...)
- Ingest structured data (JSONs/CSVs/Parquet; ...)
- Support storage of structured and unstructured data.

FR15 Metadata and data models

- Common data formats and models must be used to describe assets (RDF/OWL/JSON-LD; OpenBIM/IFC where relevant).
- All data must include metadata associated with appropriate descriptors and unique identifiers

3.2 Non-Functional requirements

3.2.1 Privacy and security

NFR1 Authentication and authorisation

- All users/services authenticated (OAUTH/OIDC)
- Use of JWT (JSON Web Tokens)
- Mutual TLS for service-to-service interactions.
- Use of X509 certificates for M2M authentication

NFR2 Encryption

- In transit encryption for all communication channels.
- When possible or for sensitive information, encrypt data at rest.

NFR3 Audit and observability

- Append-only audit logs for platform operations.
- Monitor performance and logs of the servers and platform core applications.

NFR4 Privacy by design

- Default to minimal data sharing.
- Default to minimum access rights and restrictive privacy levels.
- Offer the possibility to store data on premises.

3.2.2 Availability and performance

NFR5 System must be scalable

- Horizontal scalability with the possibility to add more nodes.
- Vertical scalability

NFR6 Availability & recovery

- Aiming for availability over 99%. Should support 24x7 service availability with an acceptable downtime of 2%.
- The overall CIRCULess DWMP system should support recovery from out-of-service periods.

3.3 Interface/UI/UX requirements

NFR7 Technical and non-technical users should be able to interact through graphical interfaces

- CIRCULess DWMP should offer graphic interfaces to manage and configure assets (Technical UI)
- CIRCULess DWMP should offer (or link) graphic interfaces connected with the services integrated in the platform (User UI)

3.4 Public policy and regulatory requirements

NFR8 Consent contracts must be approved by the owner to access data.

NFR9 It should be possible to add licenses to any asset's metadata.

NFR10 It should be possible to delete user accounts. With soft-delete and audit retentions, if applicable.

NFR11 It should be possible to delete all stored metadata or data assets. With Soft-delete and audit retentions if applicable.

NFR12 Personal data: Compliance with GDPR and reduction to a minimum the storage of any personal information.

4 Architecture

Chapter 4 constitutes the conceptual and technical foundation of Deliverable D1.1, representing the culmination of the architectural design process that establishes the structural blueprint for the CIRCULess Digital Waste Management Platform. This chapter serves as the definitive reference point for the subsequent implementation phase, providing the comprehensive framework that will guide the development activities undertaken in WP2, where the Alpha version presented herein will evolve into a fully operational Beta platform with enhanced value-creation capabilities.

The architectural specifications detailed in this chapter represent the systematic translation of user requirements, technical constraints, and interoperability standards into a coherent system design that addresses the complex challenges of waste upcycling in construction and manufacturing industries. Each of the 10 subsections methodically examines distinct architectural perspectives—from overall conceptual design through information modeling, process flows, interface specifications, deployment considerations, and security frameworks—thereby establishing a holistic understanding of the platform's intended structure and operational characteristics.

Given the **public dissemination level** of this deliverable, the explanations herein ([Chapter 4](#)) have been deliberately crafted to balance technical precision with accessibility, **ensuring that stakeholders across diverse backgrounds can comprehend the platform's fundamental principles and operational logic**. This approach recognises that effective knowledge transfer in European research projects requires communication strategies that serve both technical implementers and policy-oriented decision-makers within the broader circular economy ecosystem. Consequently, while the narrative employs accessible terminology and provides contextual explanations for complex concepts, **readers seeking comprehensive technical specifications, detailed implementation guidelines, and formal system documentation are directed to [Annex II](#)**, which presents the full technical depth appropriate for software architects, system integrators, and computer science practitioners.

This dual-level documentation strategy ([Chapter 4](#) and [Annex II](#)) ensures that the deliverable fulfils its role as both a public communication instrument and a technical specification document, thereby maximising its utility across the diverse stakeholder community engaged in European circular economy initiatives.

4.1 Overall design

The technical concept of the CIRCULess DWMP is designed to facilitate the seamless integration of data and services across multiple organisations, aiming to collaboratively develop advanced workflows, visualisations, and, more broadly, added value for platform users. The architecture shown in **Error! Reference source not found.**, distinguishes three main components: (I) the Digital Waste Management Platform, responsible for value creation; (II) infrastructural components that enable partner connectivity and data sharing; and (III) the Digital Product Passport, which offers access to secure, tamper-proof digital records documenting the lifecycle of selected CIRCULess assets and processes.

This “federated” approach means data and services can be shared under clear rules without forcing everyone to copy their information into a single database, which reduces risk and respects ownership.

This architectural section focuses primarily on the first two components, while the Digital Product Passport is addressed in dedicated deliverables (i.e., D1.3 “Construction Value Chain Review”). Only the interfaces of this architecture with the Waste Management Platform relevant to this architecture are documented herein.




Figure 3 Architecture of the CIRCULess DWMP.

4.1.1 Targeted users

Two primary groups are addressed: **domain professionals** who will apply platform functions to circularity challenges, and **information technology developers** who will build and integrate microservices and workflows that operate over open interfaces. Illustrative end users include material producers, structural engineers, architects, recyclers, and construction professionals, while developer stakeholders include software providers, system integrators, small and medium enterprises, and research and technology organisations.

WP2 will refine user personas based on the final catalogue of microservices and insights from value chain analysis activities. Furthermore, in D1.3 "Construction Value Chain Review" (and through the work of partners closer to the final user, such as OMTRE and HOLCIM), CIRCULess project is investing efforts in understanding who the final user is and what it interested in.

4.1.2 Digital Waste Management Platform

CIRCULess consortium will develop a Digital Waste Management Platform that supports and enhances physical operations. This platform will serve as a central integration hub, hosting a **database**, dedicated **frontends**, and **open APIs** to allow continuous integration of external data streams, tools, and services, while

its modular microservices architecture will permit flexible and scalable expansions over time. In addition to these foundational elements, the digital platform will incorporate **decision-support intelligence**, leveraging experimental results, simulations, and integrated datasets to generate insights, guide future experimentation, and facilitate knowledge transfer across domains.

The platform is conceived as a central meeting point where information, tools, and people come together in an orderly and reliable manner, combining a carefully maintained body of knowledge, web pages that adapt to the responsibilities of each type of user, openly documented connection points for other software, a set of small plug-in services that can be combined as needed, and guidance features that turn laboratory results and real-world context into practical recommendations for action.

In this model, the knowledge base acts like a **trusted library** that describes data, models, and procedures in a consistent way, while role-aware web screens ensure that decision-makers, operators, and developers each see exactly what is relevant to their tasks without unnecessary complexity or distraction.

The **open application interfaces** allow external systems to connect in a predictable fashion, which supports continual integration and testing, reduces duplication of effort, and makes it easier to swap or improve components over time as needs evolve.

The **modular services** are intentionally small and focused so they can be combined into workflows, allowing the platform to address both simple queries and complex chains of analysis without rebuilding the entire system for each new use case.

The **decision-support capability** brings these pieces together by summarising evidence, highlighting uncertainties and trade-offs, and presenting outcomes in a form that can be used to plan circular interventions in construction and manufacturing processes.

During the initial “Alpha” phase, the main priorities behind the scenes are to receive data safely, standardise it for consistent interpretation, and store both the raw information and its descriptive notes in a secure and orderly way, reusing high-quality open-source elements and operating a local connection point (a node) that can reach data and tools located within partner organisations.

On the user-facing side, early functions focus on creating and managing accounts, registering and describing assets, and preparing the path for later features that will provide developer workbenches and visual views for end users that report the results of the small plug-in services.

The overall plan, therefore, moves stepwise from trustworthy data handling toward coordinated, higher-value services that can support large-scale, evidence-based circular decisions across projects and regions.

Error! Reference source not found. shows the components of the DWMP.




Figure 4 Components of the CIRCULess waste management platform.

A decentralised network requires shared services that verify identities, protect communications, describe and publish resources, and formalise agreements between independent parties, so that collaboration can proceed without placing all authority in a single organisation. The following components deliver these functions in a coherent and reusable manner across partners and use cases.

Identity and access management. This service establishes how people and systems are registered, verified when they sign in, and granted permissions that match their responsibilities, ensuring that only authorised actions can be performed and that activities can be traced to accountable parties. It supports the full life cycle of an account, including creation, role assignment, credential renewal, suspension, and removal, with records suitable for audit and compliance needs.

Trust service. This service functions as a certification authority for the network, issuing and managing digital certificates that allow machines and services to recognise each other and to exchange information over encrypted channels with strong assurance of identity. Certificates are created, renewed, and revoked using established open-source cryptographic tools, and a programming interface is provided so partner systems can request and validate credentials automatically during operations. Prior research outcomes from the Horizon

Europe project SPADE are reused where suitable to reduce risk and accelerate delivery while adapting to the specific needs of this project.

Catalogue. The catalogue coordinates how partner connection points and their resources are described and published so others can discover what exists, under which conditions it may be accessed, and who is responsible for it. Building upon a collaboration catalogue developed in the SPADE project, CIRCULess will redesign the user interface, enrich descriptions for both data assets and computational services, and integrate the catalogue with the agreement management service to support end-to-end discovery-to-access workflows

Contract manager. This service enables organisations to publish offers of data or services, request access to such offers, and conclude formal agreements that define permitted uses, obligations, and time limits in a manner readable by people and enforceable by software. It extends prior work from the SPADE project by strengthening usage-policy management and enforcement at network entry points, ensuring that only requests compliant with agreed terms can reach protected resources and that access can be renewed, suspended, or withdrawn as conditions change.

Integrated outcome. Together, these components provide a consistent foundation for identity, trust, description, discovery, and agreement, enabling secure and transparent collaboration across a decentralised environment without compromising control or accountability.

4.2 Information model

The information model establishes shared descriptive rules to preserve meaning across partners, enable a reliable combination of data, and align with evolving European Union data sharing practices. Its design integrates a simple text-based representation for operational efficiency with computer-interpretable meaning to support precise exchange within and beyond the platform. CIRCULess will adopt recognised, established standards and conceptual vocabularies rather than introducing new ones, thereby reducing implementation risk and strengthening compatibility.

4.2.1 Waste Management Platform Data Models

Within the platform, a lightweight text format is employed to represent data because it is efficient to process and broadly supported, particularly for communication among modular software components that each perform a specific function. When data is shared across organisations or must carry unambiguous meaning, a linked data form of that text format is used to attach explicit context and authoritative references, ensuring consistent interpretation by computers and people. Discoverability is achieved by applying a widely used web catalogue standard so that datasets, services, and related resources are

described in a uniform, searchable manner within the platform's catalogue, including published assets and partner connection points.

4.2.2 Digital product passport data model

The digital product passport functions as the principal record, assigning each product a persistent digital identity across its life cycle, encompassing composition, production, use, reuse, and end-of-life management. It records information on materials, environmental performance, and responsible end-of-life handling, using clear data structures for creation, access, and update so that actions are auditable and consistent across partners.

4.2.3 Advanced use of semantic capabilities

To advance discovery and analytical capabilities, CIRCULess will prototype the use of graph databases for storing data with semantic meaning. Specifically, a domain-specific knowledge graph centred on circularity concepts will be developed, with an emphasis on assessing its usability and value in real-world scenarios. The platform will incorporate semantic tags to classify data assets at multiple levels: Materials (e.g., organic, mineral, hybrid), Process (e.g., logistics, separation) and Product. This approach supports precise classification and efficient access to relevant assets and services, thereby enhancing the decision-making process for circularity-oriented use cases and paving the way for a more robust and efficient decision support system in subsequent stages.

4.3 Concept model

The concept model specifies the principal entities (as shown in **Error! Reference source not found.**) within the platform and their relationships, namely organisations, human users assigned to defined roles, connection points termed nodes that link local infrastructures to the shared environment, and assets comprising data or applications subject to governance and controlled sharing.

Each organisation exercises authority over its users, nodes, and assets, and determines which assets remain private, which are made discoverable, and which may be shared under specified conditions.

Every asset is accompanied by descriptive metadata to support search and comparison, a defined privacy classification, and a policy that stipulates permitted actions, authorised parties, and temporal limits. When an asset is published in the catalogue, third parties may locate it and initiate an access request through a managed workflow, rather than retrieving or invoking it directly without authorisation.

This arrangement reconciles the openness required for collaboration with robust governance that safeguards ownership and ensures accountability




Figure 5 Entity Relationships within the CIRCULess DWMP.

4.4 Logical view

The logical architecture describes the functionality the system should provide to its users. Based on CIRCULess technical concept and the requirements (See Section 3), the logical view shows what functions the system performs and how they fit together, from bringing data in, to describing and storing it, to helping users find, trust, and use it safely. The part that users see provides simple screens for registration, configuration, asset management, and discovery, while the part they do not see handles data pipelines, catalogues, policy decisions, and secure gateways. The shared cloud services provide identity, trust, and contracting functions that everyone in the network can rely on, while each organisation's node registers local assets and enforces policies close to where the data or service lives. This design allows partners to offer data and tools without relocating them, reducing duplication and security exposure. The product passport capability plugs into this structure so that lifecycle data can be created and consulted alongside other assets when making circular decisions.

4.5 Process view

This chapter translates routine operational needs into precise, stepwise procedures so that responsibilities, prerequisites, and outcomes are explicit and

verifiable, thereby removing ambiguity from both execution and oversight. The onboarding of organisations and their users comprises the creation of digital identities, assignment of clearly bounded roles, and establishment of administrator privileges, supported by validation procedures and tamper evident activity logs suitable for audit and compliance. The processes will be grouped in the following functionality domains:

- Joining CIRCULess, users and organisations
- Connecting a partner infrastructure or service
- Contracts and access management
- Data discovery and consumption
- Data ingestion harmonisation

Together with Chapter 4.6, which specifies the interfaces through which these procedures are invoked and integrated, this Chapter 4.5 constitutes the practical core of the software construction process.

This chapter not only documents how the platform is intended to operate, but also provides the actionable specifications that developers implement, integrators connect to, and operators govern.

4.6 Interfaces view

This chapter presents the Interface View of the CIRCULess architecture, detailing the communication patterns and data exchanges between system components. The focus is to define the specific interactions that are essential for the system's functionality.

Interfaces are grouped by responsibility so that integrators know where to connect and what to expect. The platform interfaces cover bringing data in, managing assets, searching the catalogue, and handling administrative functions in a consistent web-friendly format. The shared cloud interfaces handle identity verification, trust services, catalogue operations, and digital contracting so that these sensitive capabilities are centralised and uniform. The node interfaces handle local asset registration, policy application, and secure routing to the actual data or service endpoints inside an organisation. Consistency of request and response structures across all interfaces lowers integration effort and reduces errors.

4.7 Information view

The information view explains where different kinds of information travel and where they are kept so that the system is both efficient and transparent. In

general, three data flows will be considered: the flow of real data, the metadata and the configuration data.

Descriptive metadata flows to catalogues to make discovery possible without exposing the underlying raw data or internal services. Policies and decisions flow to gateways and nodes so that access control is enforced as close as possible to the protected asset. Raw data and service endpoints remain under the control of the owning organisation, with only the minimum necessary references shared through the catalogue. This separation of description, policy, and content improves both performance and security in a distributed setting.

4.8 Deployment view

The system is deployed as code containers across three main parts: **Cloud**, **Platform**, and **Nodes**, as shown in **Error! Reference source not found..** Deployment is typically managed with a container orchestrator on servers or virtual machines, while Nodes use single-host container setups. Deployment separates shared coordination from local control so the network can scale and remain trustworthy.

The shared cloud layer runs identity, trust, catalogue, and contracting services that every participant needs and that benefit from being centrally managed. Each organisation runs its own node inside its environment to connect local data and services to the network while keeping ownership and compliance responsibilities in place. The platform runs the storage, processing, and user-facing components needed to deliver day-to-day value, and it can be scaled up independently as usage grows. This hybrid model supports different information technology policies across partners while maintaining a coherent experience for collaboration.




Figure 6 CIRCULess Cloud, Platform, and Node Deployment Architecture.

4.9 Development view

This section covers the implementation work done in T1.4 and T1.5.

The Alpha version concentrates on the essentials: bringing data in reliably, describing it well, managing identities, and exposing open interfaces so that partners can start integrating. Early web screens focus on the tasks that matter first, such as registration, configuration, and basic asset management, to gather feedback and improve usability. The shared cloud services are developed with secure storage of credentials and clear boundaries between identity, trust, and catalogue functions to reduce coupling and risk. Organisational nodes are kept lightweight so they can be deployed and updated easily while still enforcing strong policies at the edge. The Beta version in WP2 will build on this foundation by adding advanced services for environmental and social assessment, logistics decision-making, quality control, and materials performance, all orchestrated through the same governance and interface patterns.

The plan is to have a first version in Q4 2025.

4.10 Security and quality considerations

This section deals with security and quality considerations for CIRCULess DWMP. The goal is to indicate the principal aspects that CIRCULess DWMP should address. Specifically, the evaluation focuses on the following issues:

- Security and Privacy
- Availability and Reliability
- Scalability and Performance
- Maintainability and Observability
- Deployment and Node Distribution

- Data Integrity and Accountability

4.10.1 Security Considerations

Security is a fundamental aspect of the CIRCULess DWMP, ensuring that data, services, and interactions across partners remain protected and trustworthy. The key mechanisms implemented to guarantee secure operation include identity management, access control, data protection, and accountability.

Authentication and Authorisation. Each user and service must be uniquely identified and authenticated before accessing the platform. For users, CIRCULess DWMP relies on Identity Management systems and standardised mechanisms, ensuring that human users are properly verified. Short-lived tokens are employed to reduce risks and enforce regular credential rotation.

When it comes to services (machine-to-machine interactions), mTLS and X.509 certificates are used to secure service-to-service communication, strengthening trust between distributed components of the system. Also, each service (An also data assets) is uniquely identified in CIRCULess.

Access Control. CIRCULess will ensure that resources are only accessible to authorised entities, following the principle of least privilege. To achieve this, the platform applies ODRL-driven rules that define permissions, prohibitions, and constraints. These policies are evaluated at the gateway or orchestrator level and consistently enforced at the node level, ensuring that data and services remain protected even in a distributed environment. CIRCULess Nodes translates the ODRL policies into REGO rules that are then enforced at the gateway level.

Privilege Management. CIRCULess DWMP enforces the principle of least privilege to minimise security risks and limit access to only what is strictly necessary. Each tenant is assigned dedicated credentials, ensuring unauthorised access and preventing cross-tenant exposure. For service-to-service interactions, mTLS strengthens trust and confidentiality.

Encryption & Data Confidentiality. CIRCULess DWMP protects sensitive data through encryption mechanisms that secure both communication and storage.

Data Integrity. CIRCULess DWMP will ensure that data and policies are protected against tampering or unauthorised changes. Using TLS encryption for all the communication prevents message modifications and man-in-the-middle type of attacks, while it would be possible to use digital signatures of the payloads for enhanced security. For DPP records, additional integrity checks are applied to guarantee trust in exchanged information.

Audit & Accountability. CIRCULess DWMP maintains append-only audit logs to provide evidence of system operations and policy decisions. These records cover access requests, contract approvals, and other relevant transactions,

ensuring transparency and traceability. Audit mechanisms also contribute to accountability by enabling verification of how data and services have been used across the platform.

Secure Communication Channel. CIRCULess DWMP requires secure communication across all nodes and control plane components. All external interactions are mediated through gateways that manage ingress traffic, applying rate limiting and other techniques to mitigate misuse or attacks. These measures strengthen confidentiality and resilience in the distributed environment.

Privacy by Design. CIRCULess DWMP follows a privacy-by-design approach, ensuring that personal and sensitive data are protected by default. Data sharing is minimised, and wherever possible, computation is moved to the Node so that raw information does not leave its original storage location. Users have the option to keep their data private by default and only share what they feel necessary, with the right to revoke the permission at any time.

4.10.2 Quality Considerations

Quality considerations ensure that the CIRCULess DWMP can deliver its services reliably and efficiently under varying conditions. They cover aspects such as availability, scalability, robustness, monitoring, and deployment distribution. The following subsections describe how these qualities are addressed to support stable and dependable operation across all partners.

Availability. CIRCULess DWMP aims to provide continuous service with clearly defined availability targets. The platform is expected to operate 24/7, ensuring that users and partners can always access its services. Acceptable downtime is limited, with thresholds defined to keep unavailability within agreed margins (e.g., not exceeding 2%). To support this, the system incorporates mechanisms for notification and recovery after out-of-service periods, allowing disrupted operations to be restored quickly and minimizing negative impact on ongoing processes.

Scalability. CIRCULess DWMP is designed to handle growth in both workload and participants through scalable deployment models. Horizontal scalability allows the platform to expand by adding new nodes across partners, distributing processing and storage needs across the network. Vertical scalability, on the other hand, allows existing nodes to be upgraded with additional computing resources such as CPU, memory, or storage capacity.

Reliability & Robustness. CIRCULess DWMP targets resilient operation across distributed nodes. Services should tolerate component failures, degrade gracefully, and restore quickly after outages, while peer-to-peer and independent node operation limits the blast radius of local faults. In line with

prior practice, robustness also means avoiding data corruption, handling high loads without severe slowdowns, and supporting safe upgrades during operation.

Performance Monitoring. CIRCULess DWMP integrates monitoring capabilities to ensure the platform's performance can be continuously observed and maintained. Logs and metrics are collected across system components. These mechanisms allow early detection of anomalies and support the reliable operation of the platform.

5 Conclusions

This deliverable sets out the “Digital Platform Alpha Version” for the CIRCULess initiative, delineating the foundational architecture, the initial implementation artefacts, and the prescribed pathway for progression to the Beta release. It frames the Digital Waste Management Platform as a federated, standards-conformant substrate for the upcycling of mineral and timber-based waste in construction and manufacturing, assigning responsibilities unambiguously across partners and maintaining a sustained focus on interoperability, security, and lifecycle transparency.

At its core, the deliverable provides the architectural definition and Alpha instantiation of three interlocking elements: the value-creating Digital Waste Management Platform, the decentralised environment that connects partner nodes through shared identity, trust, catalogue, and contracting services, and the Digital Product Passport capability for product lifecycle records. Alpha priorities include secure data ingestion and harmonisation, rigorous metadata and catalogue description, identity and access management, certificate-based trust, agreement workflows, and user experiences for registration, configuration, and asset management.

As outlined in Chapter 4, the system architecture is presented in Annex II through a set of key views, articulated and supported by UML models and diagrams. These are further complemented by quality considerations that are critical to ensuring the robustness, scalability, and effective operation of the system.

The Alpha version of the DWMP development initiated in WP1 continues under WP2, which will concentrate on the integration of microservices and decision-making workflows. Insights derived from Deliverable 1.2, “Construction Value Chain Review”, capturing stakeholder expectations and functional requirements, will inform the evolving business requirements guiding WP2 activities. Updates to the architecture resulting from the work in WP2 and, particularly, to the development of the Waste Management Platform, will be provided alongside the Beta version of the platform in Deliverable 2.1 in the form of additional documentation or part of the related GitHub repositories.

The CIRCULess consortium is fully committed to delivering outcomes that not only fulfil the established project objectives but also effectively address stakeholder requirements wherever feasible. To support this ambition, WP2 will organise and coordinate structured cross-work package meetings involving the leaders of all technical WPs, encompassing both the physical and digital domains. These “cyber-physical” meetings will serve as a central mechanism to ensure that the development of the digital components is consistently aligned with the practical needs of end users, while being informed and validated by the scientific

knowledge and professional expertise of researchers and practitioners directly engaged with CDW materials. This approach is intended to reinforce user-driven innovation, foster cross-disciplinary integration, and maximise the impact and relevance of project results in line with Horizon Europe objectives.

Therefore, this deliverable constitutes a foundational reference for WP2 and the Beta platform release, serving as both a guideline and technical baseline for the team responsible for building value-added functionalities atop the CIRCULess DWMP.

6 References

- [1] CIRCULess (2025). Deliverable 1.2 "Construction value chain review"
- [2] CIRCULess (2025). Deliverable 1.3 "Product passport tool"
- [3] Wuyts, Wendy & Liu, Yongping & Huang, X. & Huang, Lizhen. (2023). Evaluating existing digital platforms enabling the reuse of reclaimed building materials and components for circularity. 10.1201/9781003354222-8.

ANNEX I: Review of DPP tools in circular construction

Name	BIM/CAD Support	ID Tech	ID Creation	Data Storage	Data Transfer	Access Control	Data Access
BlockMaterials	Flexible	Digital link	Centralized	Centralized	API	Access policy	Data port
BlueRoom	N/A	QR; RFID	Decentralized	Decentralized	Download	N/A	N/A
Circularise	Yes	QR; ID	Centralized	Decentralized	API; Download	RBAC	Cloud
CoBuilder	Yes	Digital link	Centralized	Centralized	API	RBAC	Cloud; Data port
DigiPass	N/A	QR; RFID; Digital link; Bluetooth	Decentralized	Decentralized	API	RBAC	Data port
Toxnot Exchange	No	QR; Bluetooth	Centralized	Centralized	API	RBAC	N/A
FinBoot	Flexible	RFID	Decentralized	Decentralized	API	N/A	Cloud; Data port
LyonDellBasell	N/A	QR	N/A	Decentralized	N/A	N/A	Cloud; Data port
Madaster	Yes	Digital link	Centralized	Centralized	API; Download	RBAC	Cloud
OK Trade	No	QR; RFID; Digital link	Centralized	Decentralized	API	RBAC	Data port
Oliot	No	QR; RFID; Bluetooth	Decentralized	Decentralized	API	RBAC	Wallet
Origin Trail	N/A	QR; RFID	Decentralized	Decentralized	API	Access policy	Wallet
Pico Next	N/A	QR	Centralized	Decentralized	API; Download	RBAC	Cloud; Data port
PixelCities	Yes	QR; RFID; Digital link; Bluetooth	N/A	N/A	Download	Access policy	Cloud
Protokol	N/A	QR	Decentralized	Decentralized	Download	RBAC	Cloud
Saifty	No	QR	Decentralized	Decentralized	API	Undisclosed	Wallet
WebIQ	No	QR; RFID; Digital link; Bluetooth	Centralized	Decentralized	API	RBAC	Wallet
Whatt.io	Flexible	QR; RFID	Centralized	Decentralized	API	RBAC	Cloud

Table 3 DPP tools

ANNEX II: Technical annex covering the architecture of CIRCULess DWMP

This annexe provides an overview of the overall architecture with greater detail. Developers and integrators will find extended documentation in the various repositories listed in Section AII.9.

AII.1 Overall design

AII.1.1 Concept and main characteristics

The technical concept of the CIRCULess DWMP is designed to facilitate the seamless integration of data and services across multiple organisations, aiming to collaboratively develop advanced workflows, visualisations, and, more broadly, added value for platform users. The architecture shown in Figure 7, distinguishes three main components: (I) the CIRCULess DWMP, responsible for value creation; (II) infrastructural components that enable partner connectivity and data sharing; and (III) the Digital Product Passport, which offers access to secure, tamper-proof digital records documenting the lifecycle of selected CIRCULess assets and processes. This architectural section focuses primarily on the first two components, while the Digital Product Passport is addressed in dedicated deliverables. Only the interfaces of this architecture with the Waste Management Platform relevant to this architecture are documented herein.




Figure 7 Architecture of the CIRCULess DWMP.

AII.1.2 Targeted users

CIRCULess has identified two main groups of users:

- End users, primarily professionals in the circularity domain, will be able to utilise the functionalities provided by CIRCULess. Examples among these are:
 - Material producers from the construction industry interested in learning about new recycled/upcycled materials, or how to improve sustainability in existing production processes.
 - Structural engineers or architects who can learn new solutions in eco-design approaches.
 - Recyclers seeking innovative measures for both micro- and macro-level logistics.
 - Construction industry professionals or employees.
- IT developers, who are the technology experts or IT firms building the microservices and workflows that bring value to the CIRCULess digital platform. Some examples are SW providers, system integrators, SMEs with specific digital solutions or developers from academia and RTOs.

Within WP2, a more refined selection of target users will be indicated based on the final list of microservices and functionalities available. Furthermore, in D1.3 “Construction Value Chain Review” (and through the work of partners closer to the final user, such as OMTRE and HOLCIM), CIRCULess is investing efforts in understanding who the final user is and what it interested in.

AII.1.3 Digital Waste Management Platform

CIRCULess will develop a Digital Waste Management Platform (DWMP) that supports and enhances physical operations. This platform will serve as a central integration hub, hosting a (I) **database** combining laboratory data, open datasets, and computational simulations, thereby ensuring that all experimental and contextual knowledge is consistently stored and accessible. To ensure broad usability, the digital platform will provide (II) dedicated **frontends** tailored to different user types, enabling researchers, industry stakeholders, and policy makers to interact with the data in ways meaningful to their roles. The platform will also expose (III) **open APIs** to allow continuous integration of external data streams, tools, and services. At the same time, its modular (IV) **microservices** architecture will permit flexible and scalable expansions over time. In addition to these foundational elements, the digital platform will incorporate (V) **decision-support intelligence**, leveraging experimental results, simulations, and integrated datasets to generate insights, guide future experimentation, and facilitate knowledge transfer across domains. This synergy of physical and digital

capabilities will create a dynamic ecosystem where real-world experiments and computational intelligence continuously reinforce each other.

From a development management perspective, the platform architecture is structured into two principal parts: backend and frontend. The implementation will proceed through two phases, corresponding to WP1 and WP2. The initial phase centres on data ingestion as well as the establishment of search and management functionalities. The subsequent phase is dedicated to generating user value and integrating microservices to enhance overall system capabilities.

Figure 8 shows the components of the DWMP. Those are organised into three main groups: **Group 1.** the core components that should be developed or adapted from open-source tools (in green). **Group 2.** components that integrate with the Waste Management Platform (in blue). While not exclusive to it, they are also developed or improved within the CIRCULess project. **Group 3.** the microservices or tools that will be integrated with the platform to build the decision-making or services for the end-users (in yellow).




Figure 8 Components of the CIRCULess waste management platform.

AII.1.3.1 CIRCULess DWMP backend

In the initial phase of the project (WP1 – Alpha version), the platform's backend will encompass APIs for data ingestion, transformation modules for data processing, and storage components for both raw data and contextual metadata,

supporting search and filtering capabilities. In the subsequent phase (WP2), microservices will be integrated alongside the available datasets, enabling the generation of derived data and expanded functionality for end-users.

Whenever possible, open-source components will be reused to develop the backend. Among these, at least a proxy, a tool to process and distribute data, and storage for data and metadata will be needed. The backend of the CIRCULess waste management platform will also run one of the nodes described in Section AII.1.3.2, providing access to the distributed network and thus enabling access to data and services stored remotely.

This part corresponds to the work done in T1.4 “Back-end alpha version development”, which will be extended in chapter 4.8 “Development view” and continued in T2.1 “Back-end beta-version development”.

AII.1.3.2 CIRCULess DWMP frontends

When it comes to the frontend, three areas will be differentiated: (I) management of users and assets, which will be the focus in WP1; (II) UIs for developers to interact with the data and generate new datasets, through the integration of popular open software such as Jupyter notebook²³ or n8n²⁴; and finally, (III) end-user visualization or interactive interfaces to see the outputs of the microservices. The last two points are expected to start development during WP2.

This part corresponds to work developed in T1.5 “Front-end alpha version development”, which will be extended in Section AII.8 “Development view” and continued in T2.4 “Front-end beta-version development”.

AII.1.4 Decentralised environment

Specific datasets and microservices will remain external to the core platform, necessitating a decentralised architectural scenario in which the owners of these assets can interconnect their respective remote infrastructures. In this federated setting, participating entities will operate independent nodes that are integrated through a common control plane. This control plane is designed to be data-agnostic, ensuring neutrality with respect to data formats and sources, while simultaneously providing the essential functions of trust management, identity provisioning, policy enforcement, and interaction with federated catalogues. These catalogues enable the discovery and interoperability of distributed resources across the network.

Figure 9 outlines the main components of the Cloud and the Node. All components are within the scope of CIRCULess, except for the Data Plane, which

²³ <https://jupyter.org>

²⁴ <https://n8n.io>

represents a generic data infrastructure. Partners are responsible for integrating their respective data infrastructures or services with the CIRCULess network through the Node.




Figure 9 Decentralised Environment and CIRCULess Node Architecture.

To create the decentralised environment, CIRCULess will rely largely on open-source components and the results of previous projects AURORAL H2020²⁵ and SPADE Horizon Europe²⁶. The developments needed to create the CIRCULess Node version will be included in the T1.4 “Back-end alpha version development”, which will be continued by T2.1 “Back-end beta-version development”. The main components are the following:

AII.1.4.1 CIRCULess DWMP Cloud: Security and governance services

The decentralised network will require several tools to coordinate and provide trust and security. To provide that functionality, CIRCULess will deploy:

- **Identity and access management:** A tool to provide support for user registration and authentication. Keycloak will be used for this purpose.

²⁵ <https://www.auroral.eu/#/>

²⁶ <https://spade-horizon.eu/>

- **Trust module:** This serves as a certificate authority for CIRCULess. We will generate and manage certificates using OpenSSL and feature an API to access the resources. It will reuse results from HORIZON SPADE project.
- **Catalogue:** Coordination and publication of nodes and their assets. This component will be developed during the CIRCULess project but reuses and extends the component from the SPADE HORIZON Europe project. In the SPADE HE project, the Collaboration Catalogue already offers similar functionality; however, in the CIRCULess project, it is planned to change the UI, upgrade the catalogue of assets and services functionality, and ultimately provide integration with the Contract Manager (See below).
- **Contract Manager:** This module will be used to publish data or service offerings, request access to them, and make agreements between entities. This component will be developed during CIRCULess and will extend the work done in the SPADE HE project, focusing on ODRL management and policy enforcement at the gateway level.

Both the catalogue and the contract manager could also be part of the Node (in a more decentralised architecture). However, for the sake of simplicity, in this version, their nodes will interact with a centralised version of them over APIs.

AII.1.4.2 CIRCULess DWMP Node

The node refers to the client application that will be distributed among partners who want to connect with the CIRCULess DWMP. The node will have three core elements:

- **Control Plane:** It will be used to register assets, allow interaction with the cloud services, and configure the data plane to route access to published data resources. This component will be developed during CIRCULess and will also be part of T1.4, later continued by T2.1.
- **Gateway:** The entry point of the node, it routes incoming requests and supports TLS termination and authentication, among other typical tasks of a proxy gateway.
- **Policy engine:** Component to enforce policies.

AII.2 Information model

Establishing an information model for the platform constitutes an important step towards ensuring semantic coherence, interoperability, and compliance with emerging European data-sharing frameworks. Its design strategically combines lightweight syntactic formats with machine-readable semantics, thereby enabling both efficient processing within the platform and seamless integration with external systems. In CIRCULess, there will be no ambition to extend any ontologies but only to adopt existing standards.

AII.2.1 Waste Management Platform Data Models

For data representation, the CIRCULess DWMP utilises JSON as the baseline serialisation format due to its efficiency and broad adoption in software development ecosystems. To enrich this representation with explicit semantics, JSON-LD is employed wherever data is expected to be exchanged, or it is needed to add meaning and context. In this way, JSON supports internal microservice communication and performance-oriented tasks, while JSON-LD ensures unambiguous semantic interoperability by linking data objects to established ontologies.

To manage metadata and ensure the discoverability of datasets, services, and associated resources, the platform incorporates the DCAT. DCAT provides a standardised description layer for datasets, their distributions, and access services. Within CIRCULess, DCAT will be used in the DWMP catalogue to describe assets and nodes.

```
{
  "@context": "https://www.w3.org/nN/Acat.jsonld",
  "@type": "Dataset",
  "title": "Waste Management Records 2025",
  "description": "Dataset documenting lifecycle records for
waste management assets in the CIRCULess platform.",
  "identifier": "http://example.com/dataset/waste-2025",
  "publisher": {
    "name": "CIRCULess Project Consortium"
  },
  "issued": "2025-08-22",
  "distribution": [
    {
      "type": "DataDistributionService",
      "accessURL": "http://example.com/download/waste-
2025.json"
    }
  ]
}
```

To ensure that data usage complies with contractual and legal requirements, the platform integrates the ODRL. ODRL enables the formal encoding of permissions, prohibitions, and obligations in a machine-readable form, allowing for precise and automated enforcement of data-sharing policies. In CIRCULess, basic policies will be used—such as policies to allow READ/WRITE/USE actions—which will be predefined for the users to select and assign to their assets or endpoints. This step simplifies complexity for the user while allowing for future extensions to the usage policies feature.

```
{
  "@context": "http://www.w3.org/ns/odrl.jsonld",
  "uid": "http://example.com/policy/123",
  "type": "Set",

```

```

"permission": [
  {
    "target": "http://example.com/dataset/my-data",
    "action": "odrl:read",
    "assignee":
"http://example.com/organization/EUProjectPartner"
  }
]
}

```

Finally, interaction with the platform and its components is described with standardised API description mechanisms. The OAS3 is applied to document and manage REST-based services, offering both human-readable developer documentation and machine-processable service descriptions. For environments involving connected devices and IoT-based functionality, the W3C Web of Things Thing Description (WoT TD) is employed, thus enabling discovery, semantic annotation, and interoperability of heterogeneous physical and digital entities. In CIRCULess project, we will follow a dual approach to API description where OAS3 will be used in app development to document all APIs, while WoT TD will be used for data exchange and configuring the Node's gateways, as it allows for integrating semantics if needed. WoT TD will be key for orchestrating assets and services over APIs, as it will serve as the standard for interacting with the data plane. While WoT TD offers three interaction patterns: properties (pull), events (push), and actions (asynchronous), in CIRCULess DWMP we will focus on the properties and give support to other patterns only if needed (WP2 will define the requirements of the microservices).

```

{
  "@context": ["https://www.w3.org/2019/wot/td/v1"],
  "title": "TemperatureSensor",
  "id": "urn:dev:ops:32473-TempSensor-001",
  "properties": {
    "temperature": {
      "type": "number",
      "readOnly": true,
      "forms": [
        {
          "href": "https://example.com/sensors/temp",
          "contentType": "application/json",
          "op": "readproperty"
        }
      ]
    }
  }
}

```

AII.2.2 Digital Product Passport data model

The DPP represents the central data structure of the platform, designed to ensure traceability, transparency, and interoperability across product life cycles.

Each DPP instance provides a unique digital identity for a product, enriched with information on materials, environmental impact, and end-of-life handling. The following subsections illustrate the core JSON representations that govern the creation, retrieval, and updating of DPPs, as well as their more detailed internal structure within the platform.

The most fundamental operation is retrieving a DPP by its UUID. This response captures the essential attributes of a product, including manufacturer details, material composition, and sustainability/circularity indicators. The structure provides a concise but complete snapshot of a product's passport.

```
{
  "uuid": "123e4567-e89b-12d3-a456-426614174000",
  "productName": "EcoSmart Hair Dryer",
  "manufacturer": "GreenTech GmbH",
  "materials": [
    { "type": "Plastic", "origin": "Recycled", "certified": true }
  ],
  "carbonFootprint": "1.2 kg CO2",
  "repairabilityScore": 7.8,
  "endOfLifeInstructions": "Separate plastic and metal components before disposal"
}
```

This baseline representation provides a lightweight yet structured method for exchanging passport data across systems. It captures both product-level descriptors (e.g., name, manufacturer) and sustainability-related attributes (e.g., carbon footprint, repairability score). While simple in form, the schema establishes the minimum common ground for interoperability and can be extended in future iterations without breaking compatibility.

When creating a new passport, the schema allows a client system to submit product information to the platform. The UUID may be pre-defined or automatically generated, while lifecycle attributes situate the product in its current state of production or use.

```
{
  "productName": "Circular Laptop Stand",
  "manufacturer": "Zerofect SA",
  "materials": [
    { "type": "Aluminum", "origin": "Recycled", "certified": true }
  ],
  "lifeCycleStage": "Manufactured",
  "uuid": "optional-if-predefined"
}
```

This creation request illustrates how essential product data is initialized in the platform. By combining identifiers, material composition, and lifecycle stage, the model ensures that each passport is anchored to both a product identity and its current contextual state. The optional UUID mechanism offers flexibility while safeguarding uniqueness across systems.

Beyond creation and retrieval, the platform also allows updating or appending information to an existing DPP. This is accomplished through a patch-based approach, where only the new or modified fields need to be provided rather than resubmitting the complete record. In this way, the model enables gradual enrichment of the passport as products advance through different stages of the supply chain.

```
{
  "carbonFootprint": "0.8 kg CO2",
  "transportInfo": {
    "from": "Athens, Greece",
    "to": "Berlin, Germany",
    "mode": "rail"
  }
}
```

This example shows how sustainability and logistics information can be progressively added to an existing record. The patch model enables direct linking of production and transport data to the product's passport, thereby reinforcing its role as a living document that evolves throughout the lifecycle.

Finally, the DPP model accommodates more complex process-oriented data structures used internally by the platform. These structures capture multi-step supply chain activities, including sequence indices, verification states, and diverse input types. The schema provides the necessary detail to represent operational evidence and link processes to predefined templates.

```
{
  "dpp": {
    "id": null,
    "supplyChainTemplate": "1f0315cb-6e71-6d2e-a0ea-e10786ad6b3d",
    "node": "1f0315d9-dcc8-644a-afc8-e5507ad95df7",
    "company": null,
    "logistics": [],
    "fromDpps": [],
    "steps": [
      {
        "id": "1f031568-5574-6dfc-ba6c-15967e7279c8",
        "name": "Plastic sorting",
        "sequenceIndex": 1,
        "quantityIndex": 1,
        "productTemplate": "1f03155c-1e6d-62a2-ac9f-a3b31898331f",

```

```

"verified": false,
"confirm": null,
"inputs": [
  { "name": "Cycle time (Hours)", "type": "numerical",
"numericalValue": null },
  { "name": "Pictures from Sorting", "type": "images" },
  { "name": "Final weight of sorted plastic (Kg)", "type":
"numerical", "numericalValue": null }
]
},
{
  "id": "1f031569-a76c-6424-8a7d-39545711e850",
  "name": "Plastic Recycling",
  "sequenceIndex": 2,
  "quantityIndex": 1,
  "productTemplate": "1f03155c-1e6d-62a2-ac9f-a3b31898331f",
  "verified": false,
  "confirm": null,
  "inputs": [
    { "name": "Energy consumed (KWh)", "type": "numerical",
"numericalValue": null },
    { "name": "Waste in the process (Kg)", "type": "numerical",
"numericalValue": null },
    { "name": "Colour", "type": "text", "textValue": null }
  ]
}
]
}
}
}

```

AII.2.3 Advanced use of semantic capabilities

To advance discovery and analytical capabilities, CIRCULess DWMP will prototype the use of graph databases for storing data with semantic meaning. Specifically, a domain-specific knowledge graph centred on circularity concepts will be developed, with an emphasis on assessing its usability and value in real-world scenarios for the circularity domain.

The platform will incorporate semantic tags to classify data assets at multiple levels:

- Materials (e.g., organic, mineral, hybrid)
- Process (e.g., logistics, separation)
- Product

This approach supports precise classification and efficient access to relevant assets and services, thereby enhancing the decision-making process for circularity-oriented use cases and paving the way for a more robust and efficient decision support system in subsequent stages.

AII.3 Concept model

This section presents an abstraction of the entities interacting within the CIRCULess DWMP, employing a crow's foot notation to illustrate their relationships, shown in Figure 10. The hierarchy of entities and their associations is clearly visible. The Organisation operates as the top-level entity. In CIRCULess, multiple organisations act as umbrellas for users (the virtual representation of human actors) and nodes (representing infrastructure connectors or data brokers).

The asset constitutes the virtual representation of a data source, service, or application. Its definition inherits key properties from the Thing entity described in the WoT Thing Description specification, for example, properties related to interactions. Nevertheless, additional properties are included to fulfil the specific functional requirements of CIRCULess. These include components such as policies, privacy information, and other metadata relevant to secure, compliant, and interoperable operations across the platform. Furthermore, assets can be published to catalogues.




Figure 10 Entity Relationships within the CIRCULess DWMP.

AII.3.1 Actors

Users are managed within a hierarchical organisational structure. This structure is integral to the system's permission model, as access rights for data and services are granted at organisation level. User permissions are further defined by roles, with an initial distinction between a "Regular User" and an "Admin User" who holds elevated privileges.

Assets are the digital entities managed by the system. They represent either data assets (e.g., datasets, IoT connections) or application assets (e.g., API-based services). Access to all Assets is strictly controlled through attached privacy levels and usage policies.

AII.4 Logical view

The logical architecture should describe the functionality the system should provide to its users. Based on CIRCULess DWMP technical concept and the requirements (See Section 3), the logical architecture is divided in the following sections:

- CIRCULess DWMP
 - Backend
 - Frontend
- Decentralized environment
 - Cloud
 - Node

AII.4.1 CIRCULess DWMP

The CIRCULess Waste Management Platform is the key technical result that should bring innovations to the final users.

AII.4.1.1 Backend

The main goal of the Waste Management Platform backend is to establish a robust data infrastructure leveraging the support of a decentralised network. All data should be protected and respect the sovereignty of its owners, but at the same time, be findable and easy to retrieve once access is granted. Therefore, the main functionalities are the following:

- Data ingestion
- Data harmonization
- Data storage & maintaining CIRCULess Knowledge Base
- Enhanced discovery functionalities
- Integration with Digital Product Passport
- Authentication and access control

During WP2, the microservices and support tools necessary to build workflows or work with the data will be deployed. The Alpha version (WP1) of the backend will be assessed and extended as appropriate based on the microservices' needs. For instance, based on the data collected and the utility for the consortium, it may be possible to integrate a knowledge graph to enhance the discovery within the platform.

AII.4.1.2 Frontend

The frontend apps will not bring functionality on their own, however, they will simplify the interaction with the platform.

AII.4.2 Decentralised environment

The main goal of a decentralised environment is to provide a way to share resources between participants of CIRCULess DWMP. The most used capability will be to offer data and services to the Waste Management Platform, without the need to store or run them physically on the CIRCULess DWMP servers.

AII.4.2.1 CIRCULess DWMP Cloud

The CIRCULess DWMP cloud is a service shared by all nodes in the decentralised ecosystem, providing governance and security features. Concretely, they will provide the following main functionality:

- User registration
- Management of users' organisations
- Management of user roles
- Discovery of assets and services registered in the platform
- Registration and discovery of services
- Management of services
- Management of contracts and data sharing rules and agreements
- Generation of credentials for CIRCULess DWMP Nodes
- Management of CIRCULess DWMP Nodes
- User notifications for relevant events and synchronisation of CIRCULess DWMP nodes
- User activity auditing
- Generation of client certificates and credentials

AII.4.2.2 CIRCULess DWMP Node

The CIRCULess Node is the client application that allows the connection of user infrastructures decentralised network. These are its main functionalities:

- Registration of assets
- Management of asset lifecycle
- Publication of assets and assets' policies
- Support the generation of semantic descriptors and context metadata

- Expose API for asset discovery
- Configuring (API) routes in the data plane
- Manage client certificates and establish secure communications
- API gateway classical functionality, such as authentication (OAuth2/OIDC, JWT validation, mTLS), rate limiting, telemetry, request/response transformation, authorization and enforcement of privacy policies
- ODRL policy store and logging of access decisions

AII.4.3 How authority and access to data works

The CIRCULess platform employs two principal modules, hosted within its cloud infrastructure, to ensure robust identity management and trust establishment:

- Identity and Access Management (IAM) module: The IAM module supports user registration and authentication processes, issuing JWT token authorised entities.
- Trust module: This module manages certificates as signed by the CIRCULess authority, facilitating secure and verified communication between users and services.

Both modules constitute the technical framework that enables the verification of user identities and services attempting to access platform resources.

As outlined in Section AII.3.1, users and digital assets are logically assigned organisations. At this stage of implementation, access control is enforced at organisational level, with future releases envisaged to provide increasingly granular access rights. The role model currently distinguishes between administrative and standard user access; additional roles will be defined and incorporated according to the requirements identified within WP2.

The CIRCULess architecture distinguishes three types of access to platform resources:

- Human user access: Human users authenticate via web interfaces using their credentials and obtain access tokens. Resource visibility is determined through the token claims, which specify the user identity's organisational affiliation.
- Core component interactions: Machine-to-machine (M2M) interactions typically involve client applications (e.g., microservices or core components) authenticating with other platform services by presenting client certificates. These interactions primarily support configuration activities or the exchange of metadata.
- Platform and Node data plane interactions: For the exchange of operational data, client applications initiate requests to Nodes or the platform Gateway. The credentials of the requester (JWT or client certificate) are validated at the Gateway in a two-step process: (I) digital

certificate signature validation and (II) evaluation of the requester's attributes against policy rules (REGO²⁷ policies) maintained on the Node. Upon successful validation, the Gateway authorises access to the data plane, enabling data retrieval. Then, interactions with external components that do not use the CIRCULess Node or Gateway will be out of scope.

AII.4.3.1 Privacy levels

The digital assets in CIRCULess DWMP will have three levels:

- Private: Their metadata will remain in the local node and data access is only possible for organisation members.
- Restricted: Their metadata is published to the catalogue and it is possible for organisation members to request access.
- Public: Their metadata is published in the catalogue and access is free (Subject to existing licenses).

AII.4.3.2 Contracts and data sharing agreements

A centralised contract manager component manages access requests within CIRCULess DWMP. While decentralised contract management, with independent managers operating at individual nodes, is feasible, a centralised approach has been adopted in CIRCULess to streamline operations.

Users can manage data sharing requests via the Management User Interface (UI). Within the UI, users can review available datasets and services, identify asset owners, and inspect the terms and conditions of relevant contracts. Contracts are specified according to the ODRL standard, enabling the precise definition of permissions, obligations, and restrictions associated with each resource. To facilitate contract creation and management, CIRCULess provides a suite of contract templates that support common actions such as READ, WRITE, and service USE.

ODRL rules are assigned at the Node level and propagated to the cloud infrastructure. When a contract request is initiated, asset owners receive notifications if the requested asset is not publicly accessible. Owners may then approve or decline the access request; approved access may also be revoked subsequently. Upon termination or completion of an agreement, the cloud system notifies the relevant Node of contractual changes, resulting in an update of the corresponding REGO policy at the Gateway layer.

Finally, when a requester seeks to access data governed by an active contract, the request is directed to the Node Gateway at any of the interaction URLs specified in the WoT Thing Description. The Gateway is responsible for validating

²⁷ <https://www.openpolicyagent.org/docs/policy-language>
CIRCULess

the requester's credentials and, upon successful validation, authorising access to the requested data.

AII.5 Process view

In this section, the main processes driving interactions in CIRCULess DWMP are presented with their diagram flows. In the diagrams, the components involved in each process will be included. The processes will be grouped in the following functionality domains:

- Joining CIRCULess, users and organisations
- Connecting a partner infrastructure or service
- Contracts and access management
- Data discovery and consumption
- Data ingestion

AII.5.1 Joining CIRCULess

Here the interaction processes that allow joining the CIRCULess community are described.

AII.5.1.1 Register User & Organisation

An overview of the Organisation registration process is shown in Figure 11Figure 11 User and Organisation Registration Process..

Superuser initiates organisation creation:

- The Superuser (S) decides to create an organisation in the system.
- The superuser sends a request to the Catalogue service (CAT), providing the organisation's name.
- The Catalogue (CAT) first validates that the requester has superuser rights to per organisation creation.
- After validation, the Catalogue (CAT) forwards the request to the Identity Management service (IDM) to create organisation entity.

Identity Manager (IDM) creates organisation:

- The IDM creates the organisation entity in its internal system, establishing an organisation context.
- Once successful, the IDM confirms the organisation creation back to the Catalogue (CAT).
- The Catalogue (CAT) then records the organisation in its internal database.
- Finally, the Catalogue (CAT) notifies the Superuser (S) that the organisation has been successfully created.

Superuser creates an admin user:

- After the organisation is set up, the Superuser (S) requests to create a new admin user, providing the user's email.
- The Catalogue (CAT) validates that the requester has superuser rights to create admin users.
- The request is then forwarded to the IDM for registration of the new user with admin rights within its organisation.

Identity Manager processes admin registration:

- The IDM registers the new admin user and associates the account with the organisation.
- After successful registration, the IDM confirms the user creation back to the Catalogue (CAT).
- The Catalogue (CAT) then stores the new user in its internal database.
- Finally, the Catalogue (CAT) confirms to the Superuser (S) that the admin user creation was successful.

Completion:

- A new organisation is now fully set up in the system.
- A new admin user has been created and associated with the organisation.
- The superuser receives confirmation of successful organisation and user creation.




Figure 11 User and Organisation Registration Process.

AII.5.1.2 Invite user

Figure 12 illustrates the overview of the user invitation process.

Admin initiates user invitation:

- The Company Admin User (A) decides to invite a new user to the system.
- The admin sends a request to the Catalogue service (CAT), providing the email address of the user to be invited.
- The Catalogue service (CAT) first checks and validates that the requester has sufficient admin rights to invite users on behalf of the company.

Catalogue triggers the identity management:

- Once the admin's permissions are confirmed, the Catalogue (CAT) forwards the request to the IDM.
- The request instructs the IDM to generate and send an invitation email to the specified address.

Invitation email delivery:

- The IDM composes an invitation email that contains a unique registration link.
- The email is sent to the New User (U).

New user starts registration:

- The New User (U) receives the invitation email and clicks the registration link.
- The link redirects the user to the IDM.
- The IDM responds by displaying a registration form to the user.

Registration form submission:

- The New User (U) fills in the required details (e.g., name, password, additional info) and submits the form to the IDM.

Identity Management processes registration:

- The IDM processes the submitted data.
- It creates a new user account and associates that account with the correct company based on the invitation context.
- After successful registration, the IDM sends confirmation back to the New User (U).

Completion:

- The New User (U) receives confirmation of successful registration;
- The IDM then redirects the user to the application, allowing them to log in and begin using the system.




Figure 12 User Invitation Process.

AII.5.1.3 Remove user

Figure 13 shows the process for User Removal.

Admin initiates user removal:

- The Company Admin User (A) decides to remove a user from the system.
- The admin sends a request to the Catalogue service (CAT), specifying the user ID or email of the target user.
- The Catalogue (CAT) first validates that the requester has sufficient admin rights to remove users.

Catalogue requests removal from Identity Management:

- Once admin rights are confirmed, the Catalogue (CAT) forwards the request to the IDM to remove the user from the system.

Identity Management processes removal:

- The IDM removes the user from its internal database, effectively revoking all access rights.
- After successful removal, the IDM confirms to the Catalogue (CAT) that the user has been removed.

Catalogue finalises removal:

- The Catalogue (CAT) removes any references to the user from its own internal database, ensuring the user is fully removed from the system.
- The Catalogue (CAT) sends a confirmation to the Admin (A) indicating that the user removal was successful.
- Optionally, the Catalogue (CAT) can also notify the Target User (U) that their account has been removed.

Completion:

- The user is now fully removed from both the Identity Management system and the Catalogue system.
- The admin receives confirmation, and the system state is updated to reflect the removal.




Figure 13 User Removal Process

AII.5.2 Connecting a partner infrastructure or service

Here, the interaction processes that allow registering assets (Data or applications/services) in CIRCULess are described.

AII.5.2.1 Registering a node

Figure 14 illustrates the Node Registration Process.

Admin requests creation of a new node:

- The Company Admin User (A) initiates a request to create a new node in the system.
- The request is sent to the Catalogue service (CAT) and includes the node name and domain information.
- The Catalogue (CAT) first validates that the requester has admin rights, ensuring only authorised users can create new nodes.

Catalogue requests a client certificate:

- Once admin rights are confirmed, the Catalogue (CAT) sends a request to the Trust Module (TRUST) to issue a client certificate for the new node.
- The Trust Module (TRUST) generates a new certificate for the node and signs it with the system's root certificate, ensuring trust and authenticity.
- The signed client certificate is returned from the Trust Module (TRUST) to the Catalogue (CAT).

Catalogue registers the new node:

- The Catalogue (CAT) stores information about the new node, including its name, domain, and associated certificate, in the system database.
- After successful registration, the Catalogue (CAT) returns the signed certificate (CA) to the Company Admin User (A). Now the user can run the node with the client certificate.

Completion:

- The Admin now receives the signed certificate and can configure the new node to securely communicate with the system.
- The new node is officially registered and trusted within the system.
- The node (Client app) can be installed by the user and initialised with the client certificate.




Figure 14 Node Registration Process.

AII.5.2.2 Removing a node

Figure 15 shows the Node Removal Process.

Admin initiates node removal:

- The Company Admin (A) decides to remove a node from the system.
- The admin sends a request to the Catalogue service (CAT), specifying the node ID or name.
- The Catalogue (CAT) first validates that the requester has sufficient Company Admin rights to remove nodes.

Catalogue requests certificate revocation:

- Once admin rights are confirmed, the Catalogue (CAT) forwards a request to the Trust Module (TRUST) to revoke the client certificate associated with the node.
- The Trust Module (TRUST) revokes the certificate and updates the Certificate Revocation List (CRL) to prevent any future use of the certificate.
- The Trust Module (TRUST) confirms back to the Catalogue (CAT) that the certificate has been successfully revoked.

Catalogue removes node from database:

- After receiving confirmation from the Trust Module, the Catalogue (CAT) removes the node record from its internal database, ensuring it is no longer recognised by the system.

- The Catalogue (CAT) then notifies the Company Admin (A) that the node removal has been successfully completed.

Completion:

- The node is fully removed from both the Trust Module (certificate revoked) and the Catalogue system (record deleted).
- The admin receives confirmation, and the system state is updated to reflect that the node no longer exists.




Figure 15 Node Removal Process.

AII.5.2.3 Registering an asset

Figure 16 shows the Asset Registration Process.

User authenticates with Identity Management:

- The User (U) initiates authentication with the IDM.
- The request is handled via Single Sign-On (SSO).
- The IDM validates the credentials internally.
- If successful, the IDM issues an access token to the User (U), which will be required for further actions.

User registers an asset on the Node:

- The User (U) sends an asset registration request to the Node (NODE).
- The request includes the Thing Description (Following Web of Things) metadata describing the asset, along with the access token.

- The Node (NODE) validates the token by contacting the IDM.
- The IDM confirms that the token is valid, and the user is authorised.

Node processes asset registration:

- Once authentication is confirmed, the Node (NODE) processes the registration request:
 - It stores the provided local Thing Description metadata.
 - It creates a new Thing Description entry representing the asset within its internal registry.
- After successful registration, the Node (NODE) returns a confirmation to the User (U), including the newly assigned asset ID.

Completion:

- The asset is now officially registered in the Node (NODE);
- The User (U) receives the asset ID and can use it for future interactions (updates, removals, or queries).




Figure 16 Asset Registration Process.

AII.5.2.4 Removing an asset

Figure 17 illustrates the process for removing an asset.

User authenticates with Identity Management:

- The User (U) initiates authentication with the IDM.
- The request is handled via Single Sign-On (SSO).
- The IDM validates the credentials internally.
- Upon success, the IDM issues an access token to the User (U) for use in subsequent requests.

User requests asset removal:

- The User (U) sends a request to the Node (NODE) to remove a specific asset, providing the asset ID along with their access token.
- The Node (NODE) forwards the token to the IDM for validation.
- The IDM verifies that the token is valid and confirms the user's authorisation.

Node processes removal:

- Once the token is validated, the Node (NODE) proceeds to delete the asset.
- This involves removing the Thing Description and any associated metadata from its internal records.
- After successful removal, the Node (NODE) sends a confirmation back to the User (U).

Completion:

- The asset is fully removed from the system.
- The User (U) receives confirmation that the operation was successful.




Figure 17 Asset Removal Process.

AII.5.2.5 Obtaining service credentials

Figure 18 illustrates the process for obtaining service credentials.

Admin authenticates with Identity Management:

- The Admin User (A) initiates authentication with the IDM.
- The request is handled via Single Sign-On (SSO).
- The IDM validates the credentials internally.
- Upon success, the IDM issues an access token to the Admin (A) for use in subsequent requests.

Admin requests service credentials from Catalogue:

- The Admin (A) sends a request to the Catalogue (CAT) to obtain service credentials (e.g., for a new service or integration).
- The request includes the access token received from IDM.
- The Catalogue (CAT) forwards the token to the IDM for validation.
- The IDM verifies the token's authenticity and responds with "Token valid".

Catalogue verifies admin rights:

- The Catalogue (CAT) checks that the authenticated user has the required administrator rights to request service credentials.
- If the user lacks permissions, the process would be rejected at this step.

Catalogue requests certificate from Trust Module:

- Once rights are verified, the Catalogue (CAT) sends a request to the Trust Module (TRUST) to issue a client certificate.
- The request includes relevant attributes such as:
 - Company name
 - Email address
 - Service identifier (if applicable)

Trust Module generates client certificate:

- The Trust Module (TRUST) generates a new client certificate.
- It signs the certificate with its root CA to ensure authenticity and trust.
- The signed client certificate is then returned to the Catalogue (CAT).

Catalogue delivers credentials to Admin:

- The Catalogue (CAT) delivers the signed client certificate back to the Admin User (A).
- This certificate serves as the service credentials, enabling the Admin to authenticate services securely in the platform.

Completion:

- The Admin User (A) successfully receives the client certificate from the Catalogue (CAT).
- This certificate now acts as the official service credential, enabling secure mutual authentication between services and the platform.
- The Admin can now use the certificate to configure services, establish trusted connections, and interact securely with platform components.




Figure 18 Service Credential Acquisition Process

AII.5.3 Contracts and access management

Here, processes related to access management and privacy are described.

AII.5.3.1 Asset privacy management

Figure 19 illustrates the asset privacy management process.

Admin authenticates with Identity Management:

- The Admin User (A) initiates authentication with the IDM.
- The IDM validates the user credentials against its internal authentication mechanisms.
- If validation is successful, the IDM issues an access to the Admin User (A), which will be used to authorize subsequent requests.

Admin calls Node API:

- The Admin User (A) uses the access token to send a request to the Node (NODE).
- The request specifically targets the Node's API to change the privacy setting of an item.
- The token is included in the request headers to prove identity and authorisation.

Node processes request:

- The Node (NODE) verifies the token to ensure it is valid and that the Admin User (A) has the necessary permissions to update item privacy;
- Once validated, the Node updates the item's privacy setting to the new value (e.g., private, restricted, or public);
- After completing the update, the Node sends a success response back to the Admin User (A), confirming that the privacy change was applied successfully.

Completion:

- The Admin User (A) has successfully updated the item's privacy.
- The IDM ensured secure authentication, and the Node enforced proper access control before applying the change.




Figure 19 Asset Privacy Management Process.

AII.5.3.2 Assign policy to asset

Figure 20 shows the asset policy assignment process.

Admin authenticates with Identity Management:

- The Admin User (A) initiates authentication with the IDM.
- The request uses Single Sign-On (SSO).
- The IDM validates the credentials internally.
- Upon successful validation, the IDM issues an access token to the Admin User (A) for use in further requests.

Admin assigns a policy to an asset:

- The Admin User (A) sends a request to the Node (NODE) to assign or update a policy for a specific asset.
- The request includes:
 - The asset identifier (ID)
 - The policy identifier (policy_ID)
 - The access token for authentication and authorisation.

Node validates the request:

- The Node (NODE) forwards the access token to the IDM to verify its validity and the admin's rights.
- The IDM validates the token and confirms it as valid.

Node processes policy assignment:

- With authentication confirmed, the Node (NODE) attaches or updates the policy associated with the asset.
- This ensures the asset's usage, visibility, or access rights are controlled according to the defined policy.

Node responds to the Admin:

- After successful processing, the Node (NODE) sends a success response back to the Admin User (A).
- The response confirms that the policy has been successfully assigned or updated.

Completion:

- The asset now has an associated policy governing its access and usage;
- The Admin User (A) receives confirmation that the operation completed successfully.




Figure 20 Asset policy Assignment Process.

AII.5.3.3 Publishing offering (DCAT)

An overview of the publishing offering process is illustrated in Figure 21.

User authenticates with Identity Management:

- The Normal User (U) initiates authentication with the IDM.
- The request is handled via Single Sign-On (SSO).

- The IDM validates the credentials internally.
- Upon successful validation, the IDM issues an access token to the user for authenticated API calls.

User publishes an offering:

- The User (U) calls the Node (NODE) API to publish an item as an offering, providing:
 - The item ID
 - The access token for authentication
- The Node (NODE) forwards the token to the IDM to validate it.
- The IDM confirms that the token is valid, and the user is authorised.

Node converts and forwards the offering:

- After authentication, the Node (NODE):
 - Converts the item data into DCAT format (Data Catalogue Vocabulary) for standardisation and interoperability.
 - Sends the DCAT-formatted item to the Catalogue (CAT) service as an offering.
- The Catalogue (CAT) acknowledges receipt of the offering.

Node confirms to the user:

- After receiving acknowledgement from the Catalogue (CAT), the Node (NODE) sends a success response back to the User (U).
- This confirms that the offering has been successfully published and is now registered in the catalogue.

Completion:

- The item is now officially published as an offering in the Catalogue.
- The User (U) receives confirmation that the operation completed successfully.




Figure 21 Publishing Offering Process (DCAT).

AII.5.3.4 Removing offering

An overview of the offering removal process is shown in Figure 22.

User authenticates with Identity Management:

- The Normal User (U) initiates authentication with the IDM.
- The request is handled via Single Sign-On (SSO).
- The IDM validates the credentials internally.
- Upon successful validation, the IDM issues an access token to the user for authenticated API calls.

User requests to remove an offering:

- The User (U) calls the Node (NODE) API to remove a published offering, providing:
 - The item ID of the offering
 - The access token for authentication
- The Node (NODE) forwards the token to the IDM to validate it.
- The IDM confirms that the token is valid, and the user is authorised to perform the removal.

Node forwards removal request to Catalogue:

- After token validation, the Node (NODE) sends a request to the Catalogue (CAT) to remove the offering in DCAT format.
- The Catalogue (CAT) confirms that the offering has been successfully removed.

Node confirms removal to the user:

- The Node (NODE) sends a success response back to the User (U), confirming that the offering has been removed.

Completion:

- The item is no longer listed as a published offering in the Catalogue.
- The User (U) receives confirmation that the removal operation completed successfully.




Figure 22 Offering Removal Process

AII.5.3.5 Requesting access

Figure 23 illustrates the requesting access process.

User authenticates with Identity Management:

- The Normal User (U) initiates authentication with the IDM.
- The request is handled via Single Sign-On (SSO).
- The IDM validates the credentials internally.
- Upon successful validation, the IDM issues an access token to the user for authenticated API calls.

User requests access to a DCAT item:

- The Requester User (U) sends a request to the Catalogue (CAT) to access a specific DCAT item.
- The request includes the item ID and the access token for authentication.
- The Catalogue (CAT) first checks whether the item is publicly accessible.

Access approval process (if item is not public):

- If the item is not public, the Catalogue (CAT) notifies the item owner (O) about the access request.
- The Owner (O) reviews the request and decides whether to approve the access.
- Upon approval, the Catalogue (CAT) proceeds to update the item's access policies.

Node updates access policy:

- The Catalogue (CAT) notifies the corresponding Node (NODE) to update the access policy for the item.
- The Node (NODE) applies the new policy, granting access to the requester.

Node confirms access to requester:

- After the policy is updated, the Node (NODE) sends an access granted confirmation to the Requester User (U).

Completion:

- The Requester User (U) now has authorised access to the DCAT item;
- The system ensures proper owner approval and policy enforcement, maintaining secure and controlled access.




Figure 23 Access Request Process.

AII.5.3.6 Revoking access

Figure 24 shows an overview of the revoking access process.

DCAT Owner authenticates with Identity Management:

- The DCAT Owner (O) initiates authentication with the IDM.
- The request is handled via Single Sign-On (SSO).
- The IDM validates the credentials internally.
- Upon successful validation, the IDM issues an access token to the user for authenticated API calls.

Owner requests access revocation:

- The DCAT Owner (O) sends a request to the Catalogue (CAT) to revoke access for a specific DCAT item.
- The request includes the item ID and the access token for authentication.
- The Catalogue (CAT) validates that the request is legitimate and checks the current access policy for the item.

Node updates item access policy:

- After validation, the Catalogue (CAT) notifies the corresponding Node (NODE) to update the access policy for the item.
- The Node (NODE) modifies the policy to remove access for the specified user(s).
- The Node (NODE) confirms back to the Catalogue (CAT) that the access has been revoked.

Catalogue confirms revocation to the owner:

- The Catalogue (CAT) informs the DCAT Owner (O) that the access revocation was successful.
- Optionally, the Catalogue (CAT) can notify the user (U) whose access was revoked.

Completion:

- The target user no longer has access to the DCAT item.
- The system ensures that access is controlled, traceable, and revocable by the owner, maintaining secure and compliant item sharing.




Figure 24 Access Revocation Process.

AII.5.4 Data Discovery and Consumption

Here, processes that describe how metadata and data are accessed are described. The data discovery sections will be covered by sequence flows, while the data consumption will be explained as different mechanisms are possible.

AII.5.4.1 Platform data discovery (User)

Figure 25 shows the process for the platform data discovery from the user's perspective.

User authenticates with Identity Management:

- The Normal User (U) initiates authentication with the IDM.
- The request is handled via Single Sign-On (SSO).
- The IDM validates the credentials internally.
- Upon successful validation, the IDM issues an access token to the user for authenticated API calls.

User requests published datasets:

- The User (U) sends a request to the Catalogue (CAT) asking for the list of published datasets.
- The request includes the access token to prove the user's authenticated session.

Catalogue validates the request:

- The Catalogue (CAT) forwards the provided token to the IDM for validation.
- The IDM verifies that the token is valid and not expired, then responds with "Token valid".

Catalogue retrieves dataset list:

- The Catalogue (CAT) queries its internal database to retrieve all datasets that are currently published and available for access.

Catalogue returns datasets to the user:

- The Catalogue (CAT) sends back the list of published datasets to the User (U);
- The response may include details such as:
 - Dataset identifiers
 - Titles and descriptions
 - Metadata (owner, creation date, etc.)
 - Access conditions (if applicable)

Completion:

- The User (U) successfully receives the list of published datasets.
- The process ensures that only authenticated users with a valid token can access this information, maintaining controlled and secure dataset discovery.




Figure 25 Platform Data Discovery Process (user).

AII.5.4.2 Federated data discovery (Node)

Figure 26 shows the federated data discovery process.

Local Node requests assets from other nodes:

- The Local Node (LN) sends a request to Other Nodes (RN) to retrieve available assets;
- The request can target both public assets and contracted assets (i.e., assets shared under contract agreements).

Other Nodes retrieve Thing Descriptions (TDs):

- The Other Nodes (RN) process the request by checking the contract and retrieving the WoT Thing Descriptions (TDs) of the assets they are willing to share;
- These TDs describe the assets in detail, including metadata, interaction affordances, and access conditions.

Other Nodes return asset descriptions:

- The Other Nodes (RN) send back the retrieved WoT TDs to the Local Node (LN);
- The response includes all relevant assets that match the request (public or contracted).

Local Node caches received Thing Descriptions:

- The Local Node (LN) stores the received Thing Descriptions locally;
- These cached TDs are treated as share objects, ensuring that the Local Node can access and serve them without repeatedly querying remote nodes.

Local availability of shared assets:

- The Local Node (LN) makes the shared objects available within its local environment;
- This allows local users and services to discover and interact with assets originally hosted on other nodes as if they were available locally.

Completion:

- The Local Node has successfully retrieved and cached the shared assets from other nodes;
- This mechanism enables federated asset discovery and reuse, ensuring that nodes in the ecosystem can collaborate while respecting public visibility and contractual sharing agreements.




Figure 26 Federated discovery request diagram.

AII.5.4.3 Platform data consumption (API)

This process will be defined in WP2, as it is based on the final microservices deployed and user needs; the different data access methods will be prepared. At least, RESTful APIs will be available.

The minimum flow will support:

- User or service authentication (Using credentials or a certificate)
- Interaction with the platform gateway
- Gateway authorises access to selected resources
- Resource is returned

AII.5.4.4 Remote data consumption

Remote data access refers to the interaction of a user or service directly with one of the CIRCULess nodes. In this case, two scenarios are possible:

- Access over standard APIs defined and managed by CIRCULess, in which case authentication will also be managed in-band.
- Access over different protocols or custom endpoints not protected by CIRCULess Node Gateway, in which case authentication and credential exchange are managed out-of-band.

The minimum flow will support:

- User or service authentication (Using credentials or a certificate)
- Interaction with the remote Node gateway
- Gateway authorises access to selected resources

- Resource is returned

AII.5.5 Data Ingestion and Harmonisation

Here, the process for ingesting and transforming data in the Waste Management Platform is described.

AII.5.5.1 Pushing data to the platform

Figure 27 **Error! Reference source not found.** shows the data ingestion process.

Service authenticates with Identity Management:

- The Service (S) initiates authentication by sending its client certificate to the IDM;
- The IDM validates the certificate against its trusted CA and revocation lists;
- If the certificate is valid, the IDM confirms successful authentication. Token will be issued for secure communication and used to authenticate against the platform gateway.

Service uploads data to Node:

- After authentication, the Service (S) sends the dataset or file to the ETL tool (ETL).
- The ETL Component (ETL) processes the document and extracts relevant metadata, such as:
 - MIME type
 - Origin
 - Owner
 - Storage path
 - Other relevant descriptive attributes
- The ETL Component (ETL) sends the processed metadata to the Node (NODE).

Node registers the dataset:

- The Node (NODE):
 - Assigns a unique ID to the dataset;
 - Registers a Thing Description entry representing the dataset and its associated metadata.
 - Node (NODE) returns the unique ID to the ETL (ETL).

ETL stores dataset in Platform Storage:

- The ETL (ETL) stores the actual dataset/file in the Platform Storage (PS), using the unique ID to match with the metadata stored by the Node;

- The Platform Storage (PS) confirms successful storage back to the ETL (ETL).

Node responds to the Service:

- After the dataset is processed, registered, and stored, the ETL (ETL) returns a response to the Service (S);
- The response includes:
 - Confirmation of successful upload
 - The unique dataset ID
 - The extracted metadata

Completion:

- The dataset/file is now fully authenticated, processed, registered, and stored in the platform.
- The Service (S) has a persistent ID and metadata for reference in future interactions.




Figure 27 Data Ingestion Process (pushing data to the platform).

AII.6 Interfaces view

This chapter presents the interface view of the CIRCULess architecture, detailing the communication patterns and data exchanges between system components. The focus is to define the specific interactions that are essential for the system's functionality. The accompanying architectural diagram provides a visualisation of the overall system structure.

All inter-component communication is secured, utilising the HTTPS protocol for encrypted data transit. Authentication is enforced universally, but is handled based on the nature of the interaction:

- Human-to-System Interaction: Authenticated via the OIDC/OAuth2 framework, ensuring secure user access.
- M2M Interaction: Authenticated using mTLS, providing strong, certificate-based identity verification for automated services.

Note: In the diagram shown in Figure 28, components rendered in grey represent ancillary or partner-provided services. While they may integrate with the platform, they are not considered part of the core CIRCULess system.




Figure 28 Interfaces View of the CIRCULess Architecture.

Below sections AII.6.1-AII.6.3 present the external interfaces of the main components. A more detailed version of the interfaces documented using OAS3 will be part of the code repositories of each component (See section AII.9).

AII.6.1 CIRCULess DWMP interfaces

Table 4 presents the main external interfaces offered by the CIRCULess DWMP, which will be extended by the Node interfaces detailed in section AII.6.3 (As a Node is also part of the DWMP).

Interface name	Used by	Integration pattern	Data exchanged
Data upload	User client app or UI	Request/Response	Files or structured data
Data consumption	User client app or value-added services	Request/Response	Files or structured data

Data discovery	User client app or value-added services	Request/Response	Structured data (JSON/JSON-LD)
----------------	---	------------------	--------------------------------

Table 4 CIRCULess waste management platform interfaces.

AII.6.2 CIRCULess DWMP Cloud interfaces

Table 5 presents the external interfaces exposed by the Cloud component.

Interface name	Used by	Integration pattern	Data exchanged
Registration of users	User client app or UI	Request/Response	Files or structured data
Authentication of users	User client app or value-added services	Request/Response	Files or structured data
Contract management	User client app or value-added services	Request/Response	Structured data (JSON/JSON-LD)
Discovery of nodes and assets	User client app or UI	Request/Response	Structured data (JSON/JSON-LD)
Certificate management	User client app or UI	Request/Response	Structured data (JSON/JSON-LD)
Web resources	User client app or UI	Request/Response	API to read and update data from the web UI

Table 5 CIRCULess Cloud Interfaces.

AII.6.3 CIRCULess DWMP Node interfaces

Table 6 presents the external interfaces of the Node.

Interface name	Used by	Integration pattern	Data exchanged
Registry of assets	User client app or UI	Request/Response	Structured data (JSON/JSON-LD)
Publishing assets	User client app or UI	Request/Response	Structured data (JSON/JSON-LD)
Discovery	User client app or UI	Request/Response	Structured data (JSON/JSON-LD)

Data consumption	User client app or UI	Request/Response	Files or structured data
------------------	-----------------------	------------------	--------------------------

Table 6 CIRCULess Node Interfaces.

AII.7 Information view

This section describes the flow of information between components. In general, three data flows will be considered. The flow of real data, the metadata and the configuration data. In the diagram shown in Figure 29 **Error! Reference source not found.**, we can see which components exchange data.




Figure 29 Information Flows.

AII.7.1 CIRCULess Waste Management Platform information flows

Table 7 shows the DWMP's main information flows.

Information flow	Destination	Description	Requirement
Store data in CIRCULess DWMP	Node (metadata) and storage in the platform (data)	User pushes data into CIRCULess managed platform	FR11 Data ingestion, harmonisation and storage

Use data from CIRCULess DWMP	Value-added services	Applications or users access data through unified interfaces	FR8 Interoperability and open APIs
Query metadata stores	User or Value-added services	Services can use metadata to prepare search engines or advanced discovery	FR7 Search & Discovery
Access Digital Product Passport	User or Value-added services	Obtain information from DPP to be used by microservices	FR10 Integration of the Digital Product Passport

Table 7 CIRCULess waste management platform information flows.

AII.7.2 CIRCULess DWMP Cloud information flows

Table 8 shows the main information flows with the cloud.

Information flow	Destination	Description	Requirement
Obtaining client certificates	User or value-added services	Generate certificates to register nodes and do M2M authentication.	FR6 Control plane and node fabric
Sync contracts with access management in nodes	Node policy engine & Contract Manager	After a contract is signed, policy administration point receives new rules	FR5 Asset access management and privacy
User registration	Catalogue in cloud	Fill a form with user info and generate credentials and setup the organisation	FR1 User Accounts
Manage organisations	Catalogue in cloud	Updates to information of the organisation	FR2 Organisations setup
Node registration	Catalogue in cloud	After node installation the cloud is notified, and the info needed for coordination is stored	FR6 Control plane and node fabric
Microservice registration	Catalogue in cloud	Microservices metadata can be registered in the catalogue	FR12 Microservices orchestration

Table 8 CIRCULess Cloud information flows.

AII.7.3 CIRCULess DWMP Node information flows

Table 9 shows the main information flows with the Node.

Information flow	Destination	Description	Requirement
Registering assets Thing Description	Node	Generate certificates to register nodes and do M2M authentication.	FR4 Asset Lifecycle & FR3 Assets management
Publishing DCAT document with assets catalogue	Catalogue in cloud	After a contract is signed, policy administration point receives new rules	FR4 Asset Lifecycle & FR3 Assets management
Providing ODRL policies to published assets	Contract manager in cloud	Assign policies to published assets	FR9 Policies, consent and collaboration environment
Update routes to data plane in Gateway	Node - Gateway	Update configuration with routes based on WoT TD of the registered assets or microservices, thus offering resources over API	FR12 Microservices orchestration

Table 9 CIRCULess Node information flows.

AII.8 Deployment view

In this section are described the main aspects of the deployment of the CIRCULess DWMP components. The whole system is deployed using containerized software (Docker²⁸) orchestrated in a single or multiple host environments, depending on the component. This approach ensures a consistent runtime environment across all components. The architecture is container-based by design and can be migrated to Kubernetes²⁹ clusters if advanced scheduling, autoscaling, or multi-host orchestration is required. In the Figure 30 there is an illustration of the scalability options, vertical and horizontal.

²⁸[Docker](#)

²⁹[Kubernetes](#)




Figure 30 CIRCULess Cloud, Platform, and Node Deployment Architecture.

AII.8.1 CIRCULess DWMP Cloud

In the Cloud, containers run the trust module, identity management, collaboration catalogue, and contract management. These can be deployed together using orchestrators of containers on dedicated servers. Stateless containers scale horizontally, while databases or other stateful parts scale vertically. For production environments, cloud services are deployed redundantly across multiple hosts to increase availability and fault tolerance. Persistent data services require regular backups, and monitoring is centralised for observability. Migration to Kubernetes would allow more fine-grained scaling and automated recovery for these services if needed.

AII.8.2 CIRCULess DWMP Node

The Nodes are deployed at the edge as single-host Docker instances. Each Node runs a control plane, data plane, and a gateway with a policy engine. Nodes scale horizontally by adding more of them across sites, and vertically by assigning more resources to a single host. Although Nodes typically run standalone, they can also be integrated into lightweight Kubernetes distributions (e.g., k3s³⁰) if orchestration at the edge becomes necessary.

AII.8.3 CIRCULess DWMP

The CIRCULess DWMP runs ETL, platform storage, discovery services, a gateway with a policy engine, as well as microservices, dashboards, and workflows. These are packaged as Docker containers and deployed on servers or cloud instances. Processing components such as ETL and gateways scale horizontally, while storage and discovery scale vertically or through partitioning. To support resilience, Platform deployments are typically distributed across multiple

³⁰ [K3S](#)

servers, with logging, monitoring, and backup services ensuring reliable operation and recovery. Migration to Kubernetes would provide additional elasticity and make it easier to manage workloads across multiple hosts.

AII.9 Development view

This section covers the implementation work done in T1.4 and T1.5.

For more details on technology, technical documentation or status, GitHub³¹ repositories are provided.

AII.9.1 CIRCULess Digital Waste Management Platform

AII.9.1.1 Backend




Figure 31 CIRCULess Waste Management Detail View.

Purpose

The CIRCULess DWMP backend represents the infrastructure supporting the data ingestion, harmonisation and valorisation. It should be modular and provide the means to extend with the integration of new functionality over APIs.

The diagram in Figure 31 we can see the first version.

Functions

As introduced in Section AII.4.1.1, the main functionalities are the following:

- Data ingestion
 - The system is able to receive diverse types of data over APIs and batch process that data to extract its main characteristics. The

³¹ <https://github.com/>

proposed architecture employs APACHE APISIX³² as an API gateway to validate and extract token claims and route traffic. APACHE NiFi³³, serves as the processing tool, providing a flexible environment for developing data ingestion pipelines.

- Data harmonization
 - An application will be developed to map the metadata extracted during ingestion with the data models used in the CIRCULess platform. Then the ingestion pipeline will be able to register assets in the Node.
- Data storage & maintaining CIRCULess Knowledge Base
 - Raw data will be stored directly in an S3-compatible storage, while the metadata will be part of the Node storage. In WP2, the extracted metadata will be semantically enriched (detailed in Section AII.2.3), enabling the use of advanced storage mechanisms such as knowledge graphs.
- Enhanced discovery functionalities
 - As detailed in Section AII.2.3, the project will assess the integration of advanced query functionality through SPARQL, providing powerful tools for developers and data scientists working with the CIRCULess platform. SPARQL's capability to query graph-based and linked datasets will enable sophisticated data analysis and knowledge extraction for expert users. In parallel, core Node functions will offer resource discovery based on JSON-Path, ensuring that general users retain straightforward and efficient access to platform data.
- Integration with Digital Product Passport
 - During WP2, the platform will develop an interface to interact with the DPP, allowing data to be accessed and utilized by the microservices.
- Authentication and access control
 - Using the resources from the CIRCULess Cloud (IAM and Trust module), the platform will be able to authenticate and match identities, as well as verify and enforce access permissions.

Dependencies

The CIRCULess DWMP depends on the decentralised network components, concretely, it runs one node and therefore interacts with the Cloud as well. Authentication and trust for all components are provided by the Cloud as well.

Data

³² <https://apisix.apache.org>

³³ <https://nifi.apache.org>

In addition to the information flows described in Section AII.7, this section will cover the storage deployed within the platform.

The platform features one central storage for the data uploaded by the users. It stores the raw data in an object storage. The uploaded data is processed to extract metadata, such as owner, MIME type, and other useful properties, which will enable categorization and enforce access rules. The Node component will store the metadata.

In WP2, the storage capabilities may be extended with specific databases for given data types (e.g., timeseries) or to enhance search functionalities (e.g., triplestores).

Repository

<https://github.com/bAvenir/circuleless-platform>

Status

The first step has been selecting the technology for the ingestion pipeline, from the gateway that authenticates and exposes the APIs, to the tool for processing the pipelines, to the storage of raw data.

The second step has been developing the basic metadata extraction.

Between September and December 2025, the project will integrate with the Node and implement automatic metadata registration following the selected standards (WoT TD and DCAT). In 2026, work will commence on advanced metadata extraction and discovery, followed by the integration of microservices and decision-making capabilities.

AII.9.1.2 Frontend

Purpose

The frontend development during the WP1 (Alpha version of the platform) will be dedicated to the platform management part, so that will be the focus of this section. As explained in previous sections, other frontends will be integrated during WP2, which may involve original development or the reuse of open-source tools.

Functions

As described in Section AII.4.1.2, the frontend provides a graphical interface to the functionality provided by the backend tools. In this section, we will focus on describing which of those functionalities will be covered by the management frontend.

- Authentication Views

- Authentication (Login and Registration) is not implemented directly in the frontend application. Instead, these views are provided by the Identity Management Module (IdM), which handles user identity and access management.




Figure 32 Identity Management Service - Login Screen.




Figure 33 Identity Management Service - Register Screen.

- List Nodes View
 - The List Nodes View is implemented in the frontend application and is responsible for displaying all Nodes that belong to the currently authenticated user's organisation.




Figure 34 Management UI - List nodes.

- List Items View
 - The List Items View is implemented in the frontend application and is responsible for displaying all local Items connected to the user's selected Node.





Figure 35 Management UI - List items.

- Discovery View
 - The Discovery View is implemented in the frontend application and is responsible for displaying Items connected to other Nodes that are either publicly visible or contracted with the user's Node via an offering.

- **Credentials View**
 - The Credentials View is implemented in the frontend application and is responsible for managing client certificates associated with the user's Node. These credentials are used by external services to communicate with the Node securely.
- **Company Detail View**
 - The Company Detail View is implemented in the frontend application and is responsible for managing the identity and information of the user's organisation and its members.



The screenshot shows the 'Organisation settings' page in the CircuLess dashboard. The top right corner displays 'Circuless dashboard' and a user profile 'Matej K.' with a dropdown arrow. The left sidebar is divided into 'General' and 'Management' sections. Under 'General', there is a link to 'bAvenir' and a 'Settings' button. Under 'Management', there are links for 'Users', 'Invitations', and 'Partnerships'. The main content area is titled 'Organisation settings' and contains a 'Name' field with the value 'bAvenir' and an 'alias: bavenir' label. Below this is an 'Identity' section showing 'Realm: circuleless' and 'Keycloak ID: 5bee4b41-84aa-4c26-8121-98ba713ff7a7'. At the bottom of the main area are 'Save' and 'Cancel' buttons.

Figure 36 Management UI - Configure organisation.




Figure 37 Management UI - Manage organisation users.

- List Contracts View
 - The Contracts View is implemented in the frontend application and is responsible for displaying all contracts established between the user's Node and Items belonging to other Nodes.
- List Offerings View
 - The List Offerings View is implemented in the frontend application and is responsible for displaying all available Items that can be contracted by the user's Node.

Dependencies

The following backend components and services are dependencies for the UI:

- Identity Management
- Catalogue
- Node
- Trust Module

Data

Information flows described in section AII.7.

Repository

<https://github.com/bAvenir/circuleless-management-ui>

Status

The development of the frontend application is progressing in phases. At this stage, the core functionality for onboarding users and managing organisations

is in place, along with basic support for viewing Nodes and their Items. This establishes the foundation for user and organisation lifecycle management within the platform.

Already Implemented

- User onboarding flows;
- Organisation management, including the creation of organisations and inviting users;
- Views to display Nodes and their associated Items.

Further Development

- Integration with external Nodes and the ability to add them into an organisation;
- Item visibility controls to manage whether Items are private, restricted, or public;
- A contracting system to formalise relationships between Nodes and external Items;
- An offering catalogue for browsing assets available for contracting from other Nodes.

AII.9.2 Decentralised environment

AII.9.2.1 Cloud




Figure 38 CIRCULess Cloud detail view.

Purpose

CIRCULess Cloud provides federated services that coordinate and establish trust to CIRCULess. The system comprises four main components: the Trust module, inherited from the SPADE EU HE project; the IAM component, implemented using the open-source tool Keycloak; and the Catalogue and Contract Manager, for which a modified version of SPADE HE project will be prepared.

Functions

As introduced in Section AII.4.2.1, the main functionalities are the following

- User registration; Management of users and organisations; Management of user roles
 - Leveraging Keycloak capabilities, the catalogue will build a wrapper to facilitate easy interaction with all the necessary identity management functionality.
- Discovery of assets registered in the platform
 - The Nodes will publish DCAT documents, which will be managed by the catalogue and made available for search through the Management UI.
- Registration and discovery of services
 - Similar to the discovery of assets, microservices and applications will be discoverable in the catalogue, with the difference that they do not require registration through the Node. However, it will be possible to register and obtain credentials directly in the catalogue for simplicity.
- Management of services
 - Service descriptions and access rules will be modifiable through the Management UI by their owners.
- Management of contracts and data sharing rules, and agreements
 - Functionality to establish or revoke contracts.
- Generation of credentials for CIRCULess Nodes & Management and synchronisation of CIRCULess nodes
 - Process to generate Node credentials and support the installation of the client on the user's machine. Later, the nodes will be mapped to their respective DNSs/IPs and synchronised with the cloud.
- User notifications for relevant events
 - The relevant notifications will be generated by the Cloud and either displayed in the Management UI or forwarded as events to the affected Nodes.
- User activity auditing
 - Audit trails regarding important changes (Especially with contracts) will be stored in the Cloud for simplicity. The Node will also keep access logs.
- Generation of client certificates and credentials
 - Interface to the Trust Module for generating client certificates signed by the CA.

Dependencies

The CIRCULess Cloud has no dependency on other CIRCULess components.

Data

In addition to the information flows described in section AII.7, this section will cover the storage deployed within the cloud.

The Identity and Access Management (IAM) service utilises a dedicated PostgreSQL³⁴ database instance to securely manage identity data. A separate PostgreSQL database instance underpins both the Catalogue and Contract Manager services, ensuring data isolation and operational efficiency. For the secure storage of certificates and secrets, the deployment of a vault solution (e.g., HashiCorp Vault³⁵) is planned, providing enhanced protection for sensitive assets.

Repository

<https://github.com/bAvenir/circuleless-cloud>

Status

Most functionalities of the cloud are ready and deployed as they will be needed to implement and test other components.

Outstanding features are those related to the Contract Manager and some Catalogue functionality, especially that related to the microservices.

AII.9.2.2 Node




Figure 39 CIRCULess Node detail view.

³⁴ <https://www.postgresql.org>

³⁵ <https://www.hashicorp.com/en/products/vault>
CIRCULess

Purpose

CIRCULess Node enables the connection of data infrastructures or services to the CIRCULess network. It inherits some of the work from the SPADE HORIZON project; however, in CIRCULess, a more lightweight version will be developed with a clear separation between the control plane and data plane.

As illustrated in the diagram above, the CIRCULess architecture implements a clear separation between South-Bound and North-Bound APIs. This logical division distinguishes between functionalities accessible by the cloud or other nodes (South-Bound APIs) and those available to human users who own or operate the Node (North-Bound APIs). Such an architecture supports security, modularity, and maintainability in access and control flows.

Functions

As introduced in Section AII.4.1.1, the main functionalities are:

- Registration of assets & Management of asset lifecycle
 - The Node will allow the registration of WoT TD descriptors to define how to access resources. This metadata information will be possible to be updated, enriched or removed. It will be stored only in the nodes unless the asset is shared.
- Publication of assets and assets' policies
 - Each node publishes information about itself and its public resources to the cloud in DCAT format.
- Support for the generation of semantic descriptors and context metadata
- Generating metadata according to schemas such as WoT TD or NGSI-LD is typically problematic and time-consuming. To address this, the Node deployed in the platform will include plugins to facilitate this task. Expose API for asset discovery
 - Each registered asset's metadata will be stored in the Node and made discoverable through an endpoint that supports the JSON-Path query language.
- Configuring (API) routes in the data plane
 - Each registered asset will have all its available interactions registered, routed and protected by the Node Gateway.
- Manage client certificates and establish secure communications
 - The Node will use the certificates generated by the Cloud to encrypt, sign or validate communications.
- API gateway classical functionality, such as could be authentication (OAuth2/OIDC, JWT validation, mTLS), rate limiting, telemetry, request/response transformation
- Authorisation enforcement of privacy policies

- The Gateway and the Policy Engine will ensure access requests are evaluated and resources protected. The Gateway can interact with the IAM and verify signatures for authentication, while the Policy Engine will validate that the identity trying to access a resource has the right to do so.
- ODRL policy store and logging of access decisions
 - ODRL policies will be stored and assigned to assets. Also, the access attempts will be logged by the Gateway and the Policy Engine.

Dependencies

The CIRCULess Waste Management Platform depends on the Cloud to offer its full functionality.

Data

In addition to the information flows described in Section AII.7, this section will cover the storage deployed within the platform.

The Node features a PostgreSQL database to keep relevant configuration and asset descriptors. The Gateway, depending on the deployment mode chosen, will use a file or etcd³⁶.

Repository

<https://github.com/bAvenir/circuleless-node>

Status

The node components have been chosen, but the implementation of the original developments and integration is still at an early stage.

The plan is to have a first version in Q4 2025.

AII.10 Security and quality considerations

This section deals with security and quality considerations for CIRCULess DWMP. The goal is to indicate the principal aspects that CIRCULess DWMP should address. Specifically, the evaluation focuses on the following issues:

- Security and Privacy
- Availability and Reliability
- Scalability and Performance
- Maintainability and Observability
- Deployment and Node Distribution
- Data Integrity and Accountability

³⁶ <https://etcd.io>

This section describes the measures that will be taken by CIRCULess developers to secure DWMP and respond to the non-functional requirements NFR1-6 and NFR8-12. A review of the security features of the digital ecosystem will be conducted in the task “Functional and security validation of requirements of the digital platform” during WP2.

AII.10.1 Security Considerations

Security is a fundamental aspect of the CIRCULess DWMP, ensuring that data, services, and interactions across partners remain protected and trustworthy. The following subsections describe the key mechanisms implemented to guarantee secure operation, covering identity management, access control, data protection, and accountability.

AII.10.1.1 Authentication and Authorisation

Each user and service must be uniquely identified and authenticated before accessing the platform. For users, CIRCULess DWMP relies on Identity Management systems and standardised mechanisms such as OAuth2 and OpenID Connect, ensuring that human users are properly verified. Short-lived tokens, such as JWT, are employed to reduce risks and enforce regular credential rotation.

When it comes to services (machine-to-machine interactions), mTLS and X.509 certificates are used to secure service-to-service communication, strengthening trust between distributed components of the system. Also, each service (An also data assets) is uniquely identified in CIRCULess.

AII.10.1.2 Access Control

CIRCULess must ensure that resources are only accessible to authorised entities, following the principle of least privilege. To achieve this, the platform applies ODRL-driven rules that define permissions, prohibitions, and constraints. These policies are evaluated at the gateway or orchestrator level and consistently enforced at the node level, ensuring that data and services remain protected even in a distributed environment. CIRCULess Nodes translates the ODRL policies into REGO rules that are then enforced at the gateway level.

AII.10.1.3 Privilege Management

CIRCULess DWMP enforces the principle of least privilege to minimise security risks and limit access to only what is strictly necessary. Each tenant is assigned dedicated credentials, ensuring unauthorised access and preventing cross-tenant exposure. For service-to-service interactions, mTLS strengthens trust and confidentiality. Additionally, each dedicated service developed in the CIRCULess project might extend the permission layer with dedicated roles or permissions.

AII.10.1.4 Encryption & Data Confidentiality

CIRCULess DWMP protects sensitive data through encryption mechanisms that secure both communication and storage. All data exchanged across the platform is encrypted in transit, using modern standards such as TLS 1.3 to prevent interception or tampering. For stored information, encryption at rest might be applied based on the use case requirements or the level of confidentiality of the data. In addition, selective disclosure mechanisms can be applied in scenarios such as the DPP, allowing only the necessary parts of data to be shared.

AII.10.1.5 Data Integrity

CIRCULess DWMP must ensure that data and policies are protected against tampering or unauthorised changes. Using TLS encryption for all the communication prevents message modifications and man-in-the-middle type of attacks, while it would be possible to use digital signatures of the payloads for enhanced security. For DPP records, additional integrity checks are applied to guarantee trust in exchanged information.

AII.10.1.6 Audit (& Accountability)

CIRCULess DWMP maintains append-only audit logs to provide evidence of system operations and policy decisions. These records cover access requests, contract approvals, and other relevant transactions, ensuring transparency and traceability. Audit mechanisms also contribute to accountability by enabling verification of how data and services have been used across the platform.

AII.10.1.7 Secure Communication Channel

CIRCULess DWMP requires secure communication across all nodes and control plane components. Mechanisms such as mTLS ensure that service-to-service connections are both authenticated and encrypted, preventing unauthorised access. All external interactions are mediated through gateways that manage ingress traffic, applying rate limiting and other techniques to mitigate misuse or attacks. These measures strengthen confidentiality and resilience in the distributed environment.

AII.10.1.8 Privacy by Design

CIRCULess DWMP follows a privacy-by-design approach, ensuring that personal and sensitive data are protected by default. Data sharing is minimised, and wherever possible, computation is moved to the Node so that raw information does not leave its original storage location. Users have the option to keep their data private by default and only share what they feel necessary, with the right to revoke the permission at any time.

AII.10.2 Quality Considerations

Quality considerations ensure that the CIRCULess DWMP can deliver its services reliably and efficiently under varying conditions. They cover aspects such as

availability, scalability, robustness, monitoring, and deployment distribution. The following subsections describe how these qualities are addressed to support stable and dependable operation across all partners.

AII.10.2.1 Availability

CIRCULess DWMP aims to provide continuous service with clearly defined availability targets. The platform is expected to operate 24/7, ensuring that users and partners can access its services at all times. Acceptable downtime is limited, with thresholds defined to keep unavailability within agreed margins (e.g., not exceeding 2%). To support this, the system incorporates mechanisms for notification and recovery after out-of-service periods, allowing disrupted operations to be restored quickly and minimizing negative impact on ongoing processes.

AII.10.2.2 Scalability

CIRCULess DWMP is designed to handle growth in both workload and participants through scalable deployment models. Horizontal scalability allows the platform to expand by adding new nodes across partners, distributing processing and storage needs across the network. Vertical scalability, on the other hand, allows existing nodes to be upgraded with additional computing resources such as CPU, memory, or storage capacity.

AII.10.2.3 Reliability & Robustness

CIRCULess DWMP targets resilient operation across distributed nodes. Services should tolerate component failures, degrade gracefully, and restore quickly after outages, while peer-to-peer and independent node operation limits the blast radius of local faults. In line with prior practice, robustness also means avoiding data corruption, handling high loads without severe slowdowns, and supporting safe upgrades during operation.

AII.10.2.4 Performance Monitoring

CIRCULess DWMP integrates monitoring capabilities to ensure the platform's performance can be continuously observed and maintained. Logs and metrics are collected across system components. These mechanisms allow early detection of anomalies and support the reliable operation of the platform.