

fRamework for safE, opEn, collaboratiVe And incLUsive digitisAtion and managemenT of cultural heritagE

Grant Agreement No 101132389

DELIVERABLE 4.10:

Integrated REEVALUATE Platform.R2

Work Package 4

LEAD BENEFICIARY:

CERTH

Delivery Date: 30.05.2026

Document Sheet

Project acronym	REEVALUATE
Project full title	Framework for safe, open, collaborative and inclusive digitization and management of cultural heritage
Programme	Horizon Europe
Topic	HORIZON-CL2-2023-HERITAGE-01-03
Type of Action	HORIZON-Research and Innovation Actions
Grant Agreement	101132389
Start day	1 January 2024
Duration	36 months

LEGAL NOTICE
This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389. Neither the European Commission nor any person acting on behalf of the Commission is responsible for the use, which might be made, of the following information. The views expressed in this report are those of the authors and do not necessarily reflect those of the European Commission.

©REEVALUATE Consortium, 2024
Reproduction is authorised provided the source is acknowledged.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Document Information

Deliverable number	D4.10
Deliverable name	Integrated REEVALUATE Platform
Lead beneficiary	CERTH
WP	4
Dissemination Level	Public (PU)
Type	Report
Reviewers (Organization)	3DR, LINKS
Delivery date	30.05.2026
Main author(s)	CERTH
Contributor(s)	KULEUVEN, LINKS, FS, GVAM, HYP

Dissemination level

Level	Description
PU	Public

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Document history

Version	Date	Changes	Reviewer/Contributor
0.1	03.04.2026	TOC	Panagiotis Stalidis (CERTH)
0.2	08/04.2026	Human Driven contextualization enabler	Jaime Solano (GVAM)
0.3	29/04/2026	Collaboration enabler	Xuemin Duan (KUL)
0.4	01/05/2026	AI based contextualization enabler	Ruben Peeters (KUL)
0.5	11/05/2026	Demonstration	Nikolaos Kyriakidis (CERTH)
0.6	14/05/2026	Issues, Limitations and Improvements	Panagiotis Stalidis (CERTH)
0.7	18/05/2026	Context Validation enabler	Panagiotis Stalidis (CERTH)
0.8	18/05/2026	AI based contextualization enabler,	Iakovos Raptis (CERTH)
0.9	01/06/2026	Knowledge Base, AI based contextualization enabler	Giacomo Blanco (LINKS)
1.0	02/06/2026	Editorial Review	Panagiotis Stalidis (CERTH)
1.1	08/06/2026	Integrate suggestions from internal review	Panagiotis Stalidis (CERTH)
1.2	09/06/2026	Quality control	Christos Zonios (CERTH)
2.0	09/06/2026	Submitted	Panagiotis Stalidis (CERTH)

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Table of Contents

Table of Contents 5
List of Figures 6
List of Tables 6
Publishable summary 8
1 Introduction 9
2 Initial Integration, Testing and Validation 9
3 Deployment, Pilot Support and Feedback 10
3.1 Deployment 10
3.2 Pilot Support 11
3.2.1 Authentication and Authorization 12
3.2.2 Demonstration 12
3.2.3 Bulk Data Ingestion 14
4 Issues, Limitations and Improvements 14
4.1 Pilot Execution Feedback 14
4.2 Suggested Improvements 16
5 Updated MVP 18
6 Pre-Digitization Functionalities 20
6.1 Public Sensing Prioritization enabler 20
7 Data Ingestion and Contextualization 20
7.1 Knowledge Base 22
7.1.1 Knowledge Graph 22
7.1.2 Artefact Storage 22
7.1.3 Vector Database 22
7.2 Human Driven Contextualization enabler 23
7.3 AI Based Contextualization enabler 24
7.3.1 Visual modality 24
7.3.2 Aural modality 25
7.3.3 Textual modality 25
8 Marketplace Functionalities 27
8.1 Collaboration enabler 27
8.2 AI Driven Creative Reuse enabler 28

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

8.3 Distributed Ledger and IPR Management 28
8.4 Context Validation enabler 29
9 Conclusion & Future Work 29
References 31
Annex A 32

List of Figures

Figure 1 Reevaluate Deployment viewpoint 11
Figure 2 Screenshot of the first live demonstration session showing technical and pilot partners in discussion 13
Figure 3 Screenshot of the second live demonstration session showing technical and pilot partners in discussion 13
Figure 4 Updated MVP of the REEVALUATE Platform 19
Figure 5 Sequence diagram for data ingestion and contextualization workflow 21
Figure 6 Campaign Summary view of the HDCE 24
Figure 7 Handling a specific suggestion in the HDCE 24
Figure 8 Categories of Detected instruments 25

List of Tables

Table 1 Pilot Feedback and Integration Priorities Suggestion 16

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Abbreviations

Abbreviations	Full name
AI	Artificial Intelligence
REST	Representational State Transfer
API	Application Programming Interface
CHI	Cultural Heritage Institution
SSO	Single Sign On
JWT	JSON Web Token
RBAC	Role Based Access Control
SAML	Security Assertion Markup Language
LLM	Large Language Model
VLLM	Visual Large Language Model
RDF	Resource Description Framework

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Publishable summary

This deliverable presents the Integrated REEVALUATE Platform of the REEVALUATE project. The platform brings together a collection of interoperable enablers that collectively support the prioritization, digitization, contextualization, validation, and creative reuse of cultural heritage artefacts.

The Integrated Platform combines advanced artificial intelligence modules, semantic data modelling, blockchain-based intellectual property management, and participatory and flexible marketplace like interfaces into a unified ecosystem. It enables cultural heritage institutions and creative professionals to manage, enrich, and reuse digital artefacts while ensuring provenance, ethical data use, and interoperability. The platform's modular architecture facilitates distributed deployment, scalability, and interoperability among partners and external stakeholders.

This report describes the transition from the first MVP, presented in the previous iteration of the deliverable, to the second MVP of the REEVALUATE Platform. It includes the steps taken to support the deployment and testing from the pilots as well as the feedback that was provided by the users. Based on the feedback, the most important updates to the platform design and integration, allow the user to better navigate the platform and utilize the independent enablers, creating a more cohesive experience.

Overall, Deliverable D4.10 demonstrates how the REEVALUATE Platform realizes an updated, more interoperable environment, creating a future-proof digital ecosystem that bridges cultural heritage, technology, and society.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

1 Introduction

The REEVALUATE Platform represents the culmination of the project's efforts to develop an integrated, usable, open and collaborative ecosystem for the management and reuse of cultural heritage digital artefacts. REEVALUATE aims to empower cultural heritage institutions (CHIs) and the wider public to engage meaningfully in the digitization, contextualization, and creative reuse of cultural artefacts. Building upon a foundation of interoperable data structures, semantic technologies, and participatory tools, the platform brings together diverse components developed across different work packages into a cohesive operational environment.

At its core, the integrated REEVALUATE Platform serves as the technological backbone of the project, hosting and orchestrating a suite of enablers that address the end-to-end workflow of cultural artefact management. These enablers cover key functional areas such as artefact prioritization, contextualization, validation, intellectual property management, and creative reuse. Through modular architecture and microservice deployment, the platform ensures scalability, flexibility, and efficient collaboration among partners, institutions, and users, enabling seamless interaction between the technical and human dimensions of cultural heritage digitization.

The development of the REEVALUATE Platform has been driven by the requirements defined in earlier project phases, particularly those concerning user needs, interoperability, and ethical data handling. Integration has been achieved through a combination of API-based communication, shared ontologies, and federated data management within a distributed Knowledge Base. Each component—ranging from AI-based contextualization tools to blockchain-enabled intellectual property registries—has been designed to function autonomously while remaining interoperable within the broader platform ecosystem. This approach facilitates both decentralized innovation and coherent system behaviour across multiple platforms.

From an operational standpoint, the REEVALUATE Platform provides an accessible and user-friendly entry point through its Hub and Marketplace. These central interfaces allow curators, researchers, and creative professionals to access, manage, and reuse digital artefacts with confidence in their provenance and contextual integrity. The platform's layered architecture ensures secure authentication, traceable transactions, and compliance with European data governance standards, while encouraging public engagement and the co-creation of cultural knowledge through participatory tools and AI-assisted interfaces.

In section 2 we remind the reader how the REEVALUATE platform was designed, developed and integrated during the first half of the project and in section 3 we present the deployment scheme and operation during the first iteration of the pilot testing. Then in section 4 we briefly present the outcomes of the pilot testing, leading to the improvements of the platform and the redesign of the MVP in section 5. The details of the integrated REEVALUATE platform are presented in sections 6-8.

2 Initial Integration, Testing & Validation

During the first period, which was reported with deliverable *D4.3 – Integrated REEVALUATE Platform.R1,* significant progress was achieved in the integration of the REEVALUATE platform, covering preparation, enabler readiness, initial integration, and initial validation activities.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

First, the integration process was guided by the architectural framework defined in WP1. All platform components and enablers developed in WP2 and WP3 were aligned with this architecture, ensuring interoperability through standardized APIs, shared ontologies, and a microservices-based design approach.

All enablers reached a sufficient level of technical maturity prior to integration. Each component was developed, deployed, and validated independently by the responsible partners, including the definition of execution environments, API interfaces, and data exchange mechanisms. This ensured that all services were integration-ready and compliant with the overall system design.

The first Minimum Viable Product (MVP) of the REEVALUATE platform was successfully implemented, integrating the core system components, including the Hub (as the central access point), the Marketplace, and the Data Lake (comprising the Knowledge Graph, Artefact Storage, and supporting services), along with the full suite of enablers. Interoperability between components was achieved through API-based communication and unified authentication via a Single Sign-On mechanism.

Initial integration testing and validation activities were carried out to verify the correct operation of the platform. This included validation of inter-enabler communication, execution of key workflows (e.g., artefact ingestion, contextualization, and reuse), and testing of the distributed deployment model. The results confirm that the platform operates as a coherent system while maintaining modularity of individual components. Overall, the work completed in the first phase resulted in a functional integrated platform that serves as the basis for pilot deployments.

3 Deployment, Pilot Support and Feedback

3.1 Deployment

The deployment of the REEVALUATE Platform for the first pilot iteration followed the high level architecture presented in Figure 1. The enablers are separated based on their functionality and the target user roles. Starting from the middle top, the Opinion Gathering app and the Prioritization enabler form a self-contained system that allows a curator to run campaigns based on physical artefacts for the selection of the order and necessity of their digitization.

The actual digitization process anticipated by the Digitization enabler is not in the primary focus of the Reevaluate project. It is included in the schematic to provide a complete overview of the process and denotes the source of digital artefacts for the platform. While in the long run the connection between physical artefacts used in the Prioritization enabler should be automatically linked to the digital artefact produced by the Digitization process, for this integration plan it remains a manual process performed by the curator/owner of the digital artefact using the Marketplace.

The main points of access to the platform for the users, based on their roles is through a single webpage, the Hub, that provides authentication for the users and easy access to the various Web applications of the platform, namely the Contextualization app, the Contextualization Suggestions app the Marketplace and the Opinion Gathering app.

As it is shown in the deployment diagram (Figure 1), the components are designed to be deployable as a distributed system. The core components will be deployed by the partners primarily responsible for the component development, running the components as microservices. This approach allows for fast updates and rapid development during the implementation phase, as partners are directly responsible for individual components and have direct access to and control over their deployment.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

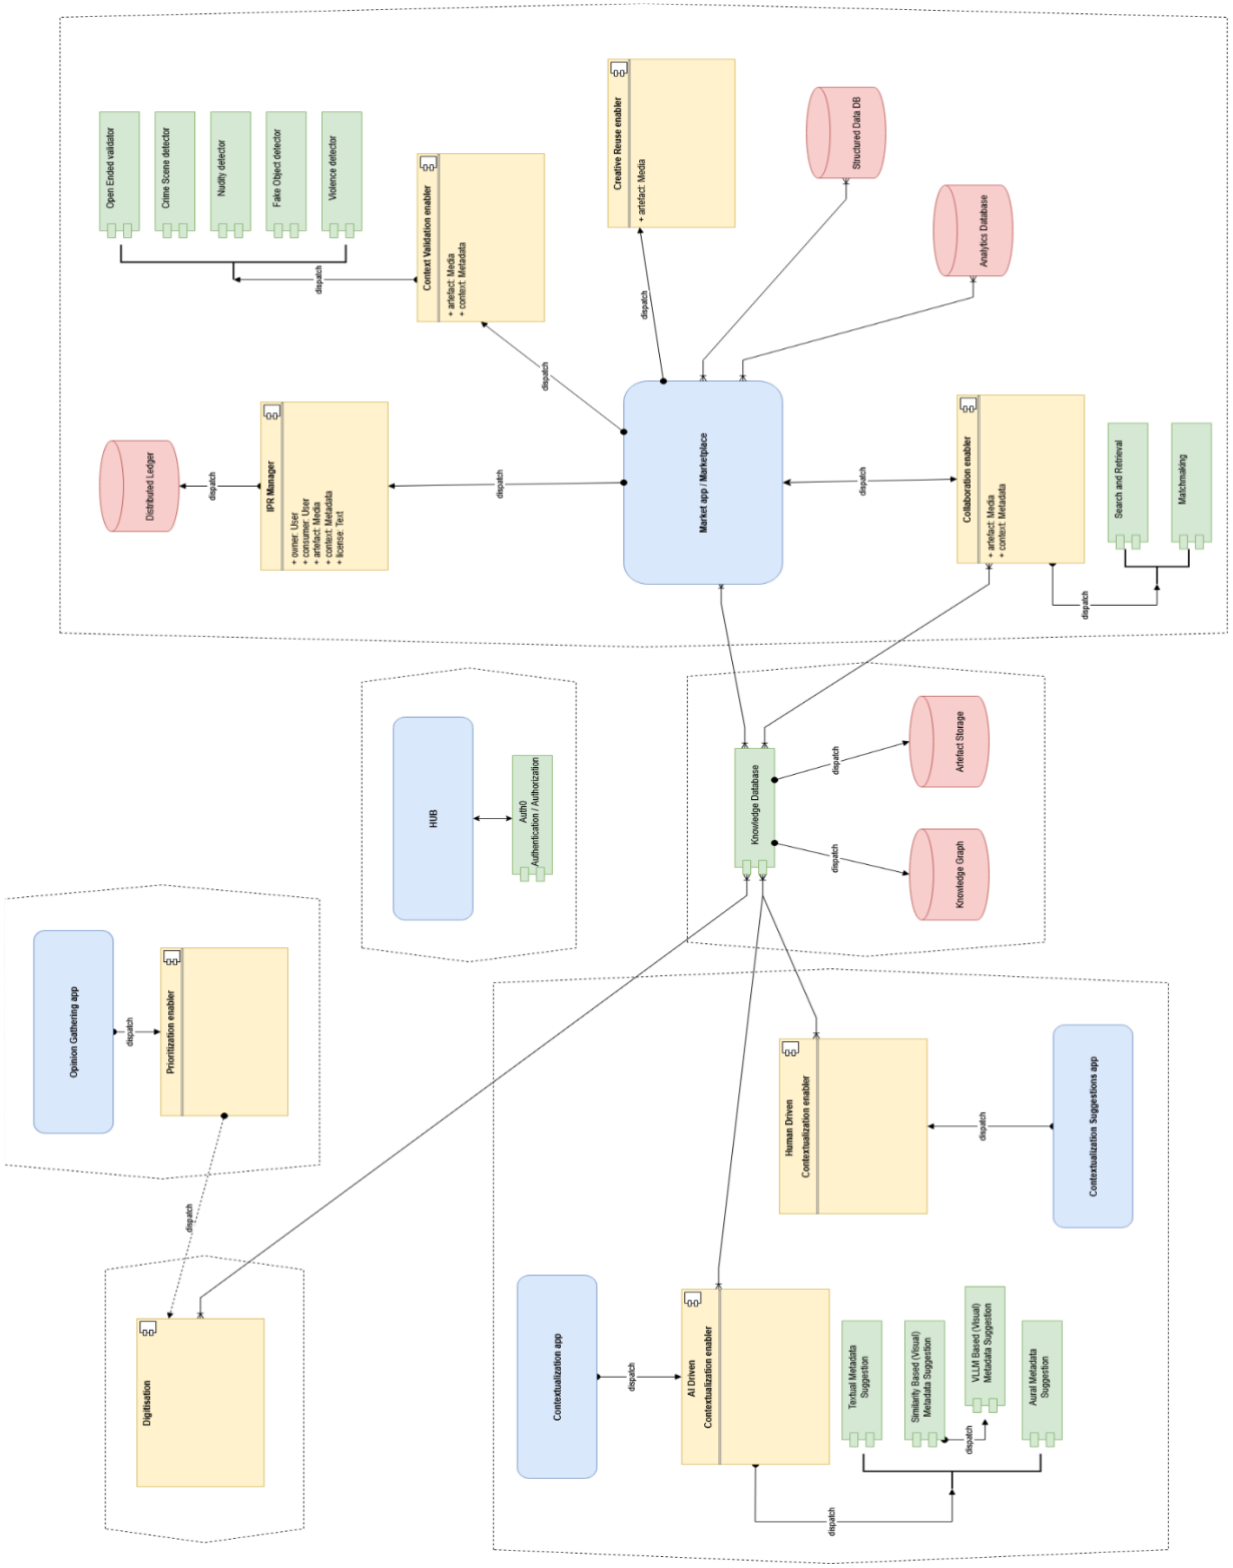


Figure 1 Reevaluate Deployment viewpoint

3.2 Pilot Support

In order for the users to be able to navigate and fully utilize the REEVALUATE platform during the pilot iterations, a number of steps needed to be fulfilled by the enabler providers.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

3.2.1 Authentication & Authorization

The REEVALUATE Platform employs Auth0 as its central identity and access management solution to provide a unified authentication and authorization framework across all platform components. Auth0 serves as the Single Sign-On (SSO) provider, enabling users—including cultural heritage professionals, project partners, and public contributors—to access multiple web applications and enablers within the REEVALUATE ecosystem using a single, secure set of credentials. Each platform service, such as the Hub, Marketplace, Contextualization apps, and AI-based enablers, is registered as a trusted Auth0 client application, ensuring consistent authentication flows and secure communication through OAuth 2.0 and OpenID Connect protocols.

In preparation for the pilot execution, 2 basic roles were created to support role-based access control:

1. The “institution” role, representing the cultural heritage institutions and in general the owners of the digital artefacts included on the platform, was assigned to representatives of the pilot partners responsible for providing the artefacts to be used as seed. Namely, 1 user from EFHA and 2 users from MoMu for pilot 1, 2 users from SMB and 4 users from OLY for pilot 2 and 3 users from AQUILEIA and 2 users from 3DR for pilot3 were assigned the role.
2. The “user” role was assigned automatically to all other accounts created in the platform, including 1 user from FFP and 1 user from CERTH for pilot1, 2 users from FFP for pilot2 and 2 users from 3DR for pilot3.

3.2.2 Demonstration

The demonstration of the project’s enablers represented a critical bridge between technical development and practical institutional application, helping the pilot partners understand how to use the enablers, while also giving them the chance to make comments and propose useful changes. The demonstration phase was designed to showcase them as integrated solutions within the broader ecosystem of the tool. The demonstration utilized two separate means to exemplify the enablers to the pilot partners, a series of live demonstrative sessions and the distribution of comprehensive analytical documentation.

The first phase of this validation took the form of two live demonstration sessions conducted in as online meetings between the pilot partners and the technical partners. During these meetings, each technical partner presented a complete, end-to-end example of their specific enabler, moving beyond simple feature lists to demonstrate how a user would interact with the enabler from the initial data entry to the final output. This live format was essential for transparency, as it allowed the pilot partners to witness the real-time performance of the tools within simulated workflows that aimed to mirror their own institutional needs or utilization cases.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

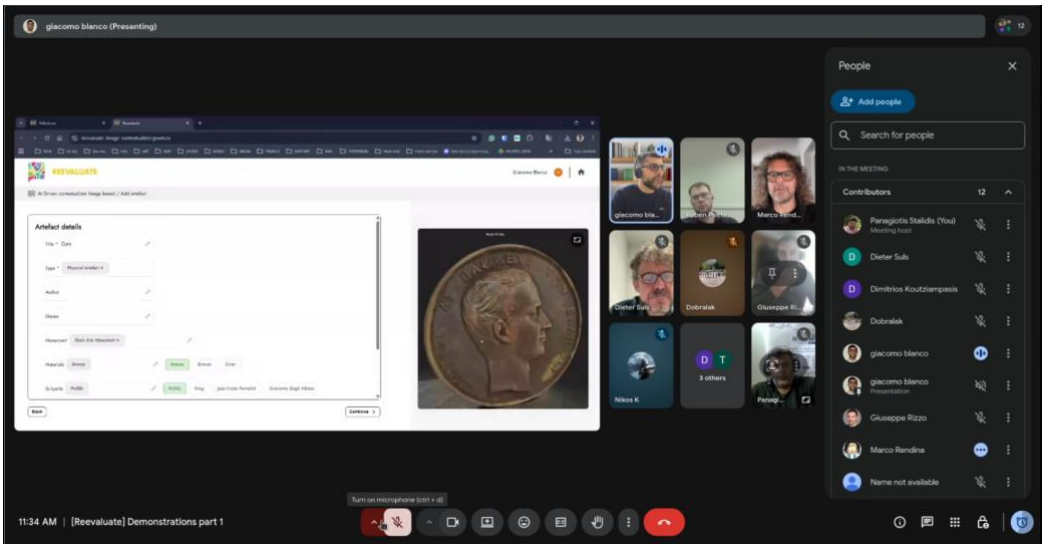


Figure 2 Screenshot of the first live demonstration session showing technical and pilot partners in discussion

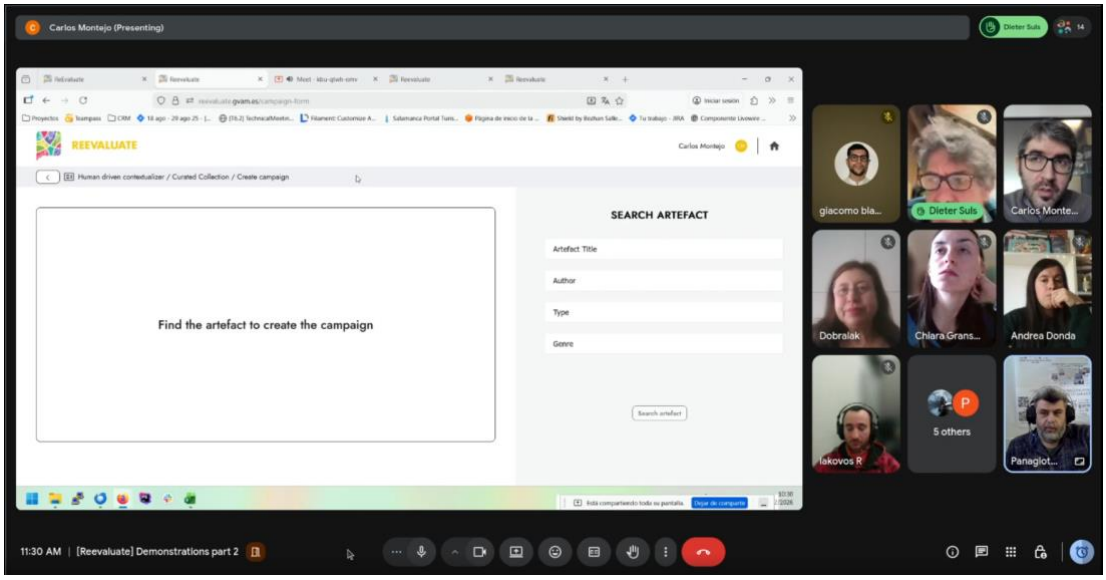


Figure 3 Screenshot of the second live demonstration session showing technical and pilot partners in discussion

The partners had the chance to provide immediate commentary and feedback following each demonstration. This dialogue was crucial, as it allowed the developers to receive direct insights into how the enablers might be received in a real-world setting, while the pilots gained a deeper understanding of the technical logic driving the tools. These sessions better captured the nuances of user experience, often lost in other types of reports, so that the final refinements to the enablers were grounded in the practical realities of cultural heritage management. These meetings were recorded and stored, so that every partner could be able to attend to the video back and find any details they might have missed or want to refresh.

As part of the demonstration process, analytical documents by the technical partners were created. These documents were designed to serve as a comprehensive instructional framework, helping guide the pilot partners through using the enablers. The documentation was summarized to focus on the essential aspects of each enabler's role and scope.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Each document provided a clear roadmap for usage, featuring step-by-step instructions that were supported by interface screenshots to ensure a low barrier to entry for non-technical staff. Furthermore, these materials included real-life use cases that illustrated the enablers in action, providing a narrative context that helped pilot partners visualize how the tool could solve specific challenges within their organizations. The documentation also included the contact information for the responsible technical partners, so that pilot partners could know with whom to speak when the need arises. The aim was that, through this combination of live engagement and instructional material, the demonstration phase could showcase the enablers to the pilot partners in a concise and easy to follow way, that would allow them to use, comment and help in maturing them further.

3.2.3 Bulk Data Ingestion

In order to populate the Knowledge Base (KB) with the artefacts provided by the pilot partners, a structured bulk data ingestion process was established. Each pilot partner provided their artefacts in the form of digital images, accompanied by descriptive metadata properties typically compiled in Excel spreadsheets. These spreadsheets contained fields such as title, creator, material, date of production, dimensions, and hosting institution, aligned with the information model supported by the platform.

The collected metadata were then processed and converted into RDF triples compliant with the CACAO ontology (the domain-specific ontology developed within the REEVALUATE project), ensuring that each artefact was semantically represented in a structured, interoperable format. Alongside the semantic data, the corresponding artefact images were ingested into the object storage and vector database components of the KB, where each image was assigned a persistent identifier linking it to its corresponding entity in the knowledge graph. This identifier is also used to store the visual embedding of each image in the vector database, enabling AI-based retrieval services. Once ingested, the artefacts become immediately available across all services exposed by the KB, providing a complete and operational data foundation for the pilot use cases.

4 Issues, Limitations and Improvements

The first iteration of the pilots marks an important milestone in the validation of the overall platform architecture and its core functionalities. Through the execution of these initial pilot activities, both the integrated REEVALUATE platform and the individual enablers were deployed and tested under realistic conditions, enabling the consortium to assess their performance, interoperability, and technical readiness. The outcomes of this phase provide valuable insights for refining the solutions, identifying optimization opportunities, and ensuring a solid foundation for the subsequent pilot iterations and large-scale deployment activities.

4.1 Pilot Execution Feedback

The following Table 1 summarizes the feedback of the pilot execution provided in deliverable D4.5. The main conclusion is that the REEVALUATE platform has reached a level of maturity sufficient for controlled pilot execution. The Hub, Marketplace, authentication and account management processes enabled access to key services, although pilot partners required support to understand account roles and workflow logic. The Knowledge Graph and artefact storage components provided the basis for connecting digital assets to platform services, but the value of this infrastructure depends on the quality, completeness and domain fit of the metadata entered into the system. While the enablers were operational and could be tested through pilot workflows, several components require adaptations.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

The first iteration also confirmed that the platform should be assessed as an ecosystem rather than as a set of isolated components. Even though individual enablers can be operated on their own, the true value of the REEVALUATE platform depends on how smoothly they are interconnected with one another.

Enabler or platform component	Readiness and performance observations	Priorities for iteration 2
Public Sensing Prioritisation Enabler	Operational and able to collect preference data. However, interactions outside the formal workflow were not captured. Campaign editing, result interpretation, artefact-level analytics and ranking require improvement. Support for audio and video is still limited. Campaign navigation should be restricted to relevant artefacts.	Improve analytics, ranking, campaign management, social media integration, multimedia support and campaign boundaries.
Human-Driven Contextualisation Enabler	Operational and useful for collecting structured contextual input. Issues included workflow rigidity, search and retrieval transparency, category management, uppercase email handling and temporary mobile input problems.	Improve usability, mobile stability, invitation management, contribution workflows, category handling and integration with Marketplace records.
AI-driven text-based Contextualisation Enabler	Operational and promising, but outputs require further reliability checks. Users noted that the results can appear solid at first glance but are not yet ready for production use without validation. Batch processing and clearer UI support are needed.	Improve accuracy, explainability, batch processing, user guidance and integration with pilot-specific metadata and the Knowledge Graph.
AI-driven image-based Contextualisation Enabler	Easy to use, but functionality remained limited. Suggested metadata were often incorrect or not aligned with the pilot content. The metadata model did not sufficiently match fashion heritage needs.	Improve domain adaptation, metadata alignment, batch upload, feedback loops and integration with CHI-provided image collections.
Marketplace	Operational as a central access point for artefact management and reuse workflows. Most user scenarios could be executed, but interface behaviour was not always transparent. Some bugs, screen-size issues and metadata limitations were identified.	Improve usability, account role clarity, collection and artefact workflows, metadata flexibility, batch upload and cross-enabler integration.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Collaboration Enabler	Search and recommendation functions are conceptually central, but full validation across all pilots is still partial. Its effectiveness depends on richer metadata, larger artefact sets and more complete user interaction data.	Test with larger and more diverse collections, validate search relevance, improve recommendation logic and connect results to concrete reuse scenarios.
AI-Driven Creative Reuse Enabler	Operational, but outputs sometimes violated conditions of use or did not align sufficiently with expectations. More controlled prompts, clearer restrictions and stronger context validation are needed.	Improve prompt guidance, output quality, restriction handling, integration with context validation and use in complete creative reuse workflows.
Context Validation Enabler	Operational, but outputs were sometimes confusing or inconsistent. Some problematic contexts were not always recognised, while warnings could appear too broadly. Stand-alone use was less effective than integrated use would be.	Improve clarity of outputs, detection reliability, threshold handling, user guidance and integration with Marketplace and creative reuse workflows.
DLT-based Tokenisation and IPR Management Enabler	Operational at platform level, but not yet fully validated through complete reuse agreements or real production workflows.	Test complete rights-management workflows, including permissions, reuse duration, agreements, artefact access and production-related reuse.

Table 1 Pilot Feedback and Integration Priorities Suggestion

4.2 Suggested Improvements

The second iteration depends on a set of technical and platform improvements that should be prioritised according to their impact on pilot objectives. The most important priority is not only to make individual enablers more stable, but to improve the integration between them. The first iteration showed that cultural heritage users experience the REEVALUATE framework through workflows, not through technical components. For this reason, improvements should focus on user journeys that connect several enablers in a coherent sequence.

The Marketplace should remain the central access point for artefact management and reuse workflows, but its usability should be improved before broader testing. Users need clearer guidance on account roles, collection creation, artefact upload, metadata entry, publication status, licensing, reuse permissions and the relationship between Marketplace actions and connected enablers. System behaviour should be made more transparent through clearer messages, confirmations and error explanations. The second iteration should also test batch upload or at least a more efficient ingestion process, since cultural heritage institutions rarely manage digital assets one by one in real operational settings.

Metadata flexibility should be a priority across the platform. Pilot 1 requires better support for fashion-specific information, including garment type, construction, pattern, materials, techniques, sizes, historical

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

period and reuse context. Pilot 2 requires better support for audio, sound recordings, musical instruments, audio-visual material and intangible cultural heritage contexts. Pilot 3 requires stronger links between artefact-level metadata, site-level interpretation, spatial context and virtual tour points of interest. These domain-specific needs should be addressed through configurable metadata fields or pilot-specific extensions that remain mapped to the common semantic structure of the platform.

The Public Sensing Prioritisation Enabler requires several improvements before the second iteration. Campaign setup and editing should become more flexible. Campaign analytics should be easier to interpret, with clearer artefact-level results, automatic ranking based on votes or interactions, exportable data and transparent separation between quantitative and qualitative outputs. Campaign navigation should be restricted to the artefacts included in the specific campaign, avoiding user confusion caused by access to unrelated content. The enabler should also improve the handling of multimedia assets, especially audio and video, which are essential for Pilot 2. Finally, the relationship between social media outreach and measurable campaign participation should be clarified technically and operationally.

The Human-Driven Contextualisation Enabler should be improved in terms of workflow clarity, mobile stability and contributor management. The first iteration identified issues related to email handling, search and retrieval transparency, category systems and temporary mobile input problems. For the second iteration, invitation management should be smoother, text input should be stable across devices, contribution categories should be easier to understand, and curators should have clearer moderation and validation options. The link between contextualisation contributions, Marketplace records and the Knowledge Graph should also be made more visible to users, so that they understand how their input becomes part of the wider platform.

The AI-driven text-based Contextualisation Enabler should be improved in reliability, explainability and user guidance. Pilot 1 feedback indicated that outputs can appear useful at first glance but still require careful validation. This confirms the need for a curator-in-the-loop approach. The enabler should support clearer review steps, indicate uncertainty where relevant, and make it easier for users to correct, approve or reject structured outputs before they are stored. Batch processing should also be considered, as cultural heritage institutions often work with larger sets of descriptions and metadata records.

The AI-driven image-based Contextualisation Enabler should be further adapted to the domain-specific needs of the pilots. In Pilot 1, suggested metadata were often not aligned with fashion heritage content. This indicates a need for better training data, stronger use of CHI-provided reference collections, clearer selection of transferable properties and more transparent feedback loops. The tool should also clarify how user validation improves future outputs, since testers need to understand the value of correcting or confirming suggestions.

The Collaboration Enabler should be tested with larger and more diverse collections during the second iteration. Its information retrieval and matchmaking functions depend on the richness of the Knowledge Graph, the quality of metadata, the availability of embeddings and the presence of meaningful user profiles or interaction histories. The second iteration should therefore include structured tests of search relevance, image-based retrieval, combined text and image search, recommendation quality and usefulness for creative reuse. These tests should be connected to real pilot tasks, such as finding suitable fashion assets for design inspiration, selecting music-related artefacts for the video production, or retrieving archaeological objects for educational storytelling.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

The AI-Driven Creative Reuse Enabler should be tested within more controlled creative workflows. The first iteration showed that outputs may violate conditions of use or fail to align with user expectations. The second iteration should therefore include stronger prompt guidance, clearer restrictions, improved integration with context validation and a review mechanism for generated outputs. This is particularly important for the Fashion Time Machine and for FFP's video production, where creative reuse must remain aligned with the cultural meaning and rights conditions of the source assets.

The Context Validation Enabler should be improved in terms of output clarity and detection reliability. Testers need to understand why a context is considered problematic, whether the result is a warning or a blocking issue, and how confident the system is. The tool should include clearer examples, more transparent categories, improved threshold handling and better integration with the Marketplace and Creative Reuse Enabler. Stand-alone use may be useful for testing, but the second iteration should prioritise integrated use within concrete reuse workflows.

The DLT-based Tokenisation and IPR Management Enabler should be tested through complete rights-management scenarios. The first iteration provided only partial evidence of this workflow. In the second phase, pilots should test the sequence from artefact upload to rights definition, permission request, agreement or licensing, reuse duration, access to files, proof of ownership or permission, and documentation of the reuse event. This is essential for validating the REEVALUATE promise of safe and legally supported creative reuse.

Finally, the consortium should introduce a more structured issue-tracking and release-management process for the second iteration. Shared documents were sufficient for the first testing round, but a more systematic approach is needed for broader validation. Each issue should include the pilot, enabler, user role, steps followed, expected behaviour, actual behaviour, severity, screenshot or evidence where possible, and status of resolution. Technical partners should also provide short release notes before each testing phase, indicating what has changed, what is known to be limited and what should be tested.

5 Updated MVP

One of the major outcomes of the first pilot iteration is the conclusion that the REEVALUATE platform can increase its value by strengthening the integration between enablers and prioritising end-to-end workflows rather than isolated technical improvements. The REEVALUATE platform allows four major workflows to be accomplished through the use of its individual enablers: the prioritization of physical artefacts that should be digitized and uploaded to the platform, the upload of a digital artefact to the platform and the enhancement of contextual metadata, the discovery of digital artefacts in the marketplace and the contract signing between the owner and the producer, and finally the verification that the production does not violate the signed contract.

Aiming to enhance the usability and coherence of the workflows, in Figure 4 we present the updated minimum viable product (MVP) of the REEVALUATE Platform. The major change from the first version is the direct connection between the marketplace and the contextualization enablers, bypassing the necessity to navigate through the hub. The gain from this change is that the user keeps perspective of the digital artefact for which the metadata are suggested and there is a single place from which a new artefact can be uploaded to the platform, removing confusion for the user. Another update is the connection of the context validation enabler to the analytics database through a REST API from the marketplace, allowing the user validating an

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

artefact to link the validation of the production to an existing digital artefact and smart contract relevant to the user (i.e. owned or purchased by the user).

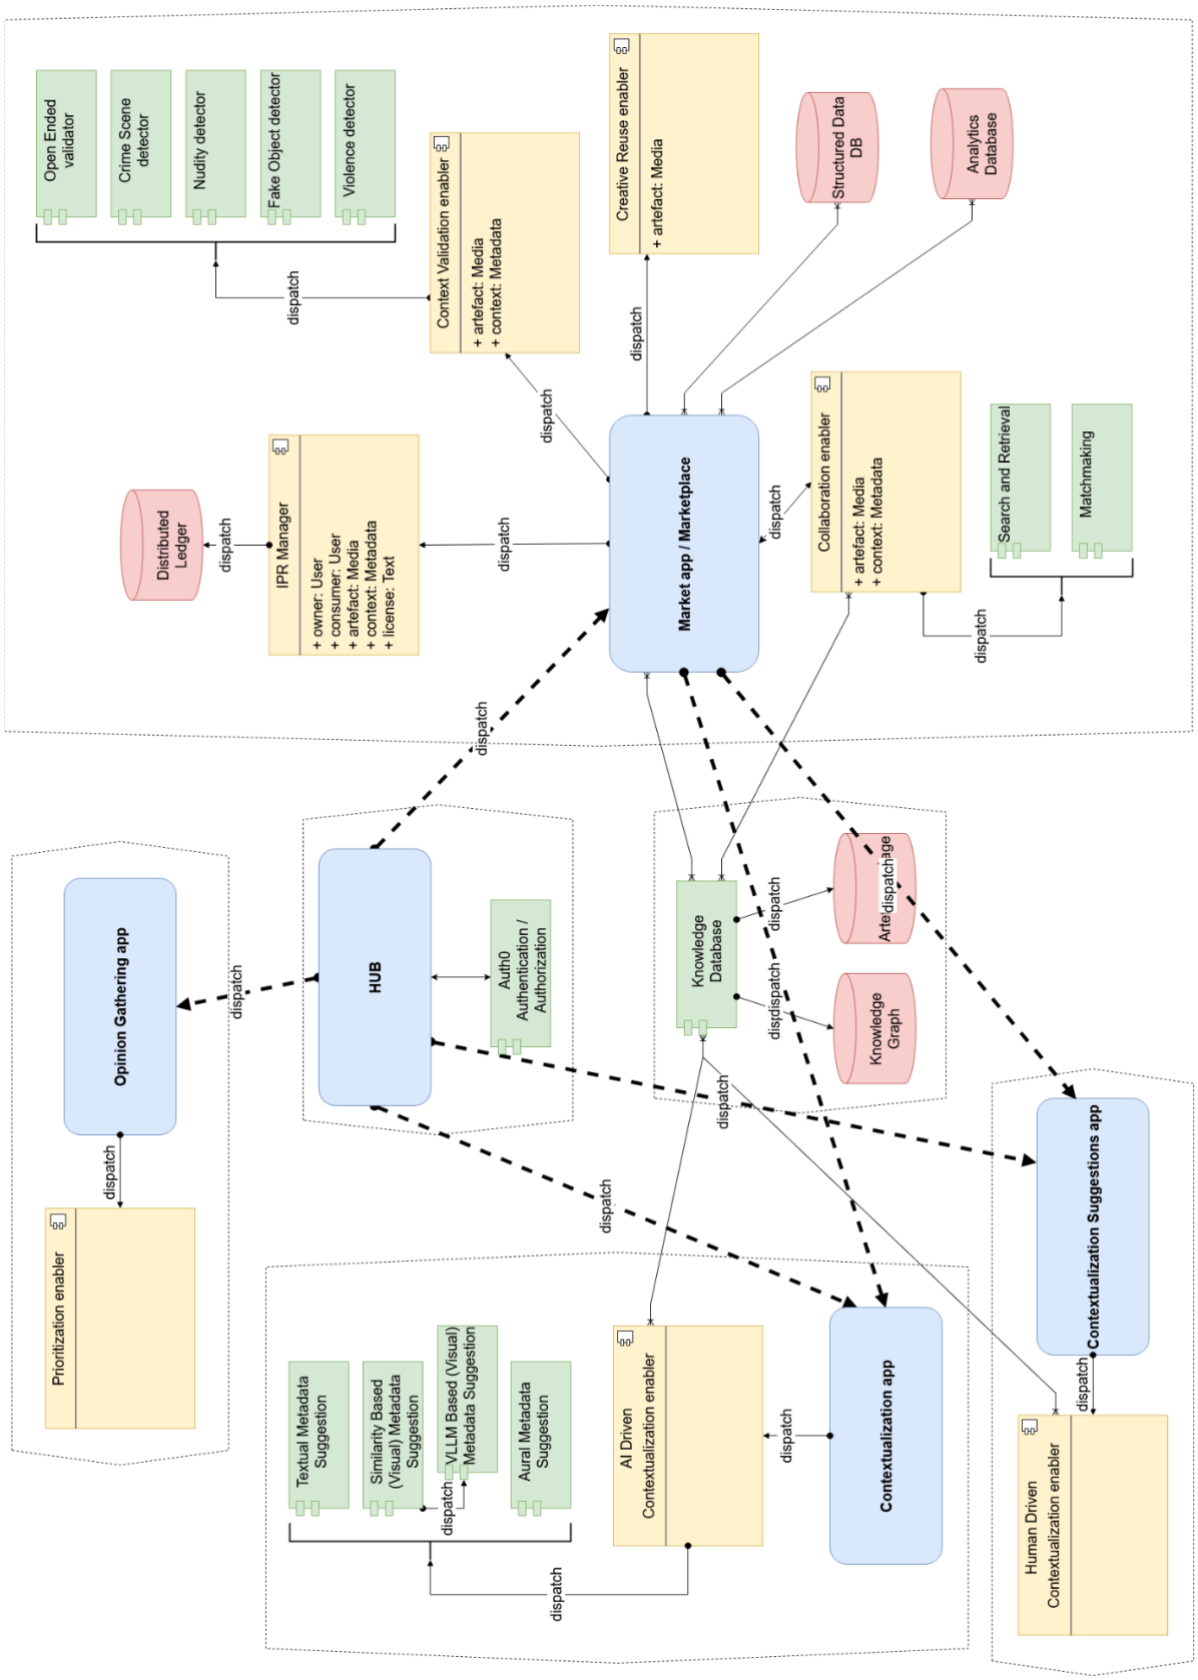


Figure 4 Updated MVP of the REEVALUATE Platform

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

The following sections offer details in the implementation and integration for each of the four workflows and the enablers that are used in them.

6 Pre-Digitization Functionalities

The selection of which physical artefacts should be digitized by a CHI is a multi-factor process that depends on the cost of digitization, the public interest for the digital form of the artefact and the marketability to mention only a few. The REEVALUATE platform, through the Public Sensing Prioritization enabler, allows a CHI to gather information from the public, by running campaigns in social media platforms.

6.1 Public Sensing Prioritization enabler

The first pilot phase provided useful feedback for the Public Sensing Prioritization enabler. Pilot partners stress tested the enabler in real life conditions in order to find technical bugs or ways for a more streamlined experience. Both of them were addressed in a development effort in order to have a better enabler for the second pilot phase.

Minor bugs that were identified were fixed. These included the erroneous appearance of a link in a campaign post, that is relevant to the gamification system. Now, this link only appears when the gamification system is selected. Another technical limitation that was addressed was the failure of posts that included emojis in the text body.

Lastly, the user interface was updated. In the older version metrics such as emotion and sentiment analysis were presented holistically for the whole campaign. However, this proved to not be as useful as envisioned and pilot partners wished to be presented with the same information per artefact. This provides them with more specific feedback. The user interface now is able to provide this information.

7 Data Ingestion & Contextualization

This section presents the integrated workflow for the ingestion and contextualization of digital artefacts within the REEVALUATE platform. Based on the architecture, the first MVP distributes the storage of information to the individual enablers that gather it. In effect, the Artefact Storage module gathers the binary data related to an artefact, the Knowledge Graph gathers the context information and user-artefact interaction relational data are gathered inside the Marketplace.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

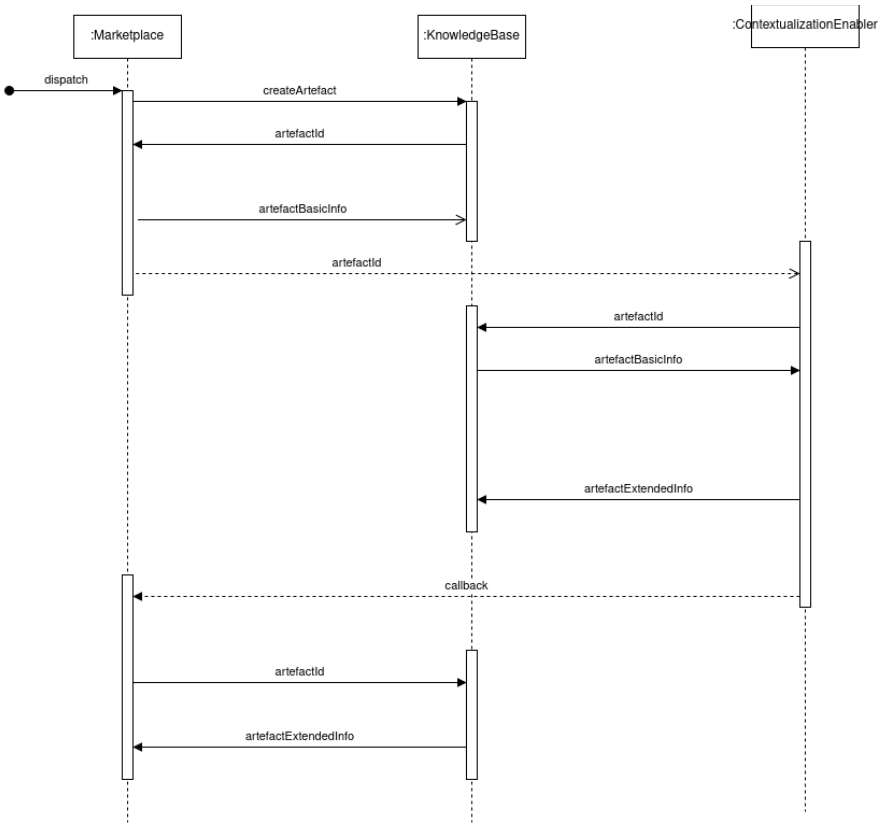


Figure 5 Sequence diagram for data ingestion and contextualization workflow

While this schema allows more flexibility and minimal communication overhead for storing and retrieving data, in practice, during the first pilot iteration deployment it became apparent that multiple entry points for data ingestion decreased the user experience. The users were confused as to which tool they should use to add new artefacts, how they could enhance the context information of existing artefacts and which enabler should be used for which artefact.

Therefore, the REEVALUATE platform was redesigned in order to provide the user with a single point of ingestion for new artefacts. The new integration design recognizes that the central piece of the architecture is the Marketplace. On the one hand, the Marketplace is the most logical module to handle the single ingestion task, as it already has a data ingestion workflow and the existing data lake schema places the most constraints on the relational data that are generated from the Marketplace. On the other hand, fully incorporating the contextualization enablers inside the Marketplace would create a major redesign in the UI and an unnecessary extra burden in task T4.2.

The proposed solution is a compromise that still uses the Marketplace as the focal point in the workflow but instead of ingesting the REST APIs of the contextualization enablers inside the UI of the Marketplace, it passes control to the UI of each enabler through buttons. Each of the enablers will receive this redirection in a new or updated landing page, taking as parameters the logged in user and the digital artefact. After the enabler gathers the information to enrich the artefact and stores it in the KB, the control is passed back to the Marketplace using a callback address. The sequence diagram for this workflow is presented in Figure 5.

The modified workflow constitutes a smooth transition from a disjointed collection of independent tools to a unified system that can properly play the role of being the core component of data ingestion and management.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

7.1 Knowledge Base

The Knowledge Base (KB) of the REEVALUATE platform a modular end-to-end infrastructure that unifies semantic modelling, multimodal representation, and operational access to cultural heritage artefacts. The architecture follows a component-agnostic design principle and integrates four main components: a graph database storing the knowledge graph, an object storage preserving the digital representations of artefacts, a vector database storing AI-generated visual embeddings, and an API gateway providing a unified and controlled access layer to all underlying services. Together, these components establish an integrated infrastructure that bridges semantic knowledge with visual and data-driven representations of cultural artefacts, enabling hybrid retrieval by meaning or by visual similarity and providing a practical example of neuro-symbolic integration.

7.1.1 Knowledge Graph

The knowledge graph forms the semantic core of the knowledge base, modelling entities such as artefacts, artists, materials, artistic movements, and their contextual relations according to the CACAO ontology, a domain-specific extension of CIDOC-CRM developed within the REEVALUATE project. The graph is stored and managed through a GraphDB instance, which provides robust support for RDF data and exposes both a SPARQL endpoint for expressive semantic querying and a GraphQL endpoint for API-oriented consumption. Through SPARQL, it is possible to retrieve multi-step relationships, such as chains linking an artefact to the artist who created it, the materials employed, or the broader artistic movement it belongs to. The knowledge graph was constructed by extracting data from Wikidata and project pilot partners and aligning it with the CACAO ontology, and currently comprises more than 8 million RDF triples representing 1.4 million unique entities across 82 classes from CACAO, CIDOC-CRM and ODRL, including approximately 76K artists, 307K physical artefacts, 2.5K museums, and over 4.8K artistic movements, genres and materials.

7.1.2 Artefact Storage

The object storage component is responsible for storing and serving the binary content associated with each artefact, including high-resolution images, videos, 3D scans, and other digital assets. It is implemented using a MinIO-based solution, which guarantees scalability, data durability, and efficient access management for large volumes of heterogeneous files. The choice of a general-purpose object storage is deliberate: the KB is designed to manage any digitised representation of a cultural artefact, not only images, and MinIO natively supports this heterogeneity of media types within a single unified infrastructure. Each stored object is assigned a persistent identifier that links it to its corresponding entity in the knowledge graph, enabling direct bidirectional navigation between semantic and multimedia representations.

7.1.3 Vector Database

The vector database introduces a complementary representation layer that encodes the visual features of artefacts, enabling similarity-based retrieval and linking beyond explicitly modelled semantic relations. It is implemented as an instance of Qdrant, which allows the creation of multiple collections of vector representations derived from different AI models or feature extraction strategies. Each vector is associated with the persistent identifier of the corresponding entity in the knowledge graph, allowing seamless navigation from perceptual representations back to their semantic counterparts. Cosine similarity is employed as the retrieval metric. The vector database serves as the backbone for AI-based services such as image-to-image retrieval, text-to-image retrieval, and the neural-aided metadata suggestion service.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

7.2 Human Driven Contextualization enabler

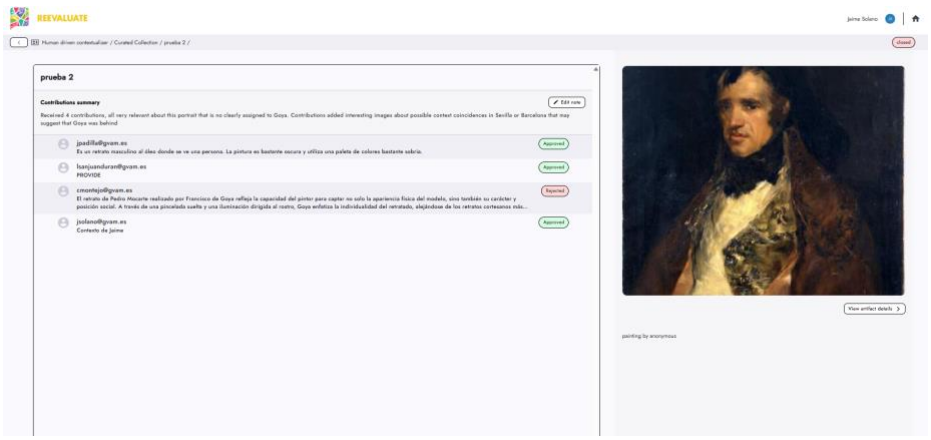
The Human Driven Contextualization Enabler (HDCE) has been fully developed and connected to the Knowledge Base to assist curators in collecting context (images, texts, URLs, references and other supporting material) about any artefact through contributions provided by the general public, domain experts or other cultural heritage professionals. Contributions can be gathered through mailing campaigns, web forms or social network actions, allowing institutions to engage different stakeholder groups depending on the contextualisation objectives.

UI has been adapted to the new visual and experience directives defined in T1.5 to make the use case more fluent and accessible. UX has also been oriented to the "job to be done" agile usability methodology so that call-to-action steps, system responses and user expectations are better aligned throughout the contribution process. Several improvements requested during the first pilot iteration have also been incorporated, including the final summary form, which can be prepared manually by curators or assisted by AI tools before being attached to the artefact as validated contextual information.

The redesigned HDCE has evolved into a multi-stakeholder contextualisation environment capable of supporting contributions from curators, researchers, domain experts and members of the public. Through configurable campaigns linked to specific artefacts stored in the Knowledge Base, cultural heritage institutions can gather complementary perspectives, local knowledge and thematic interpretations that enrich the contextual representation of artefacts beyond what can be achieved through automated approaches alone.

Particular emphasis has been placed on improving campaign scalability and participation workflows. Curators can define campaign objectives, select target audiences, distribute invitations through different channels and moderate incoming submissions through a unified interface. Contributions may include textual descriptions, references, images, URLs and other contextual resources, all of which can be reviewed, validated and consolidated before becoming part of the official artefact record.

The integration of the HDCE with the Knowledge Graph has also been reinforced to ensure that validated contributions are captured as structured contextual knowledge. By leveraging the ontology-driven interoperability framework adopted by REEVALUATE, contextual information collected through campaigns can be linked to existing artefacts, reused by other platform services and made available to future contextualisation, validation and creative reuse processes. This creates a continuous knowledge enrichment cycle where human expertise complements AI-generated suggestions and contributes to the long-term growth of the platform ecosystem.



Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Figure 6 Campaign Summary view of the HDCE

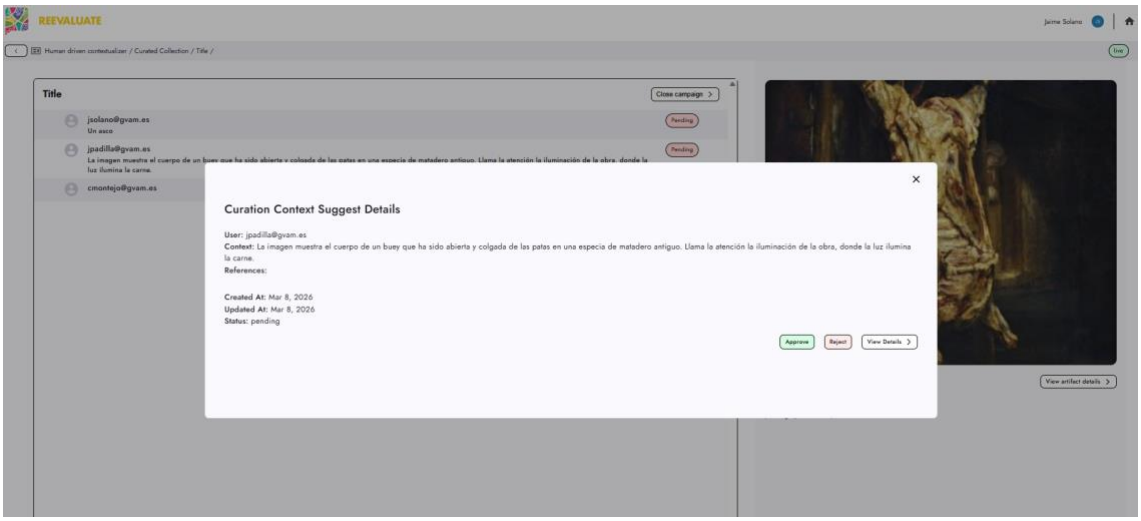


Figure 7 Handling a specific suggestion in the HDCE

Next iterations will focus on other enablers' interconnections, to demonstrate and facilitate how REEVALUATE Enablers represent a relevant toolset for context curation.

7.3 AI Based Contextualization enabler

The AI-Based Contextualisation Enabler enriches digitised cultural heritage artefacts with contextual metadata through three specialised submodules operating on different input modalities: image, text, and audio. The image enabler retrieves visually similar artefacts from the Knowledge Base and aggregates their metadata through a similarity-weighted voting scheme to produce ranked suggestions for categorical fields. The text enabler transforms free-text descriptions into CACAO-compliant RDF triples via a sequential multi-agent LLM pipeline with SHACL and DbC validation constraints. The audio enabler identifies contextual metadata from audio recordings. All submodules operate as decision-support tools, presenting ranked candidate values to the curator for validation before any metadata is committed to the Knowledge Base.

7.3.1 Visual modality

The image-based contextualisation submodule provides AI-assisted metadata suggestion for newly digitised artefacts uploaded to the REEVALUATE platform. Its core idea is to streamline the curatorial effort required to annotate a new artefact by leveraging the artefacts already present in the Knowledge Base as a reference database.

The pipeline operates in two stages. In the first stage, an image search module processes the query image through a SwinV2 transformer backbone to compute its feature vector, which is then compared via cosine similarity against the embeddings of all images stored in the vector database. The top N most visually similar images are retrieved together with their associated metadata. In the second stage, a similarity-weighted majority voting scheme aggregates the metadata values of the retrieved images into a ranked set of suggestions for each categorical attribute, such as artwork type, material, artistic movement, genre, and depicted subject, where values shared by more similar images receive higher scores.

To address cold-start scenarios, where the Knowledge Base may contain few or no artefacts sufficiently similar to the query image, a complementary VLM-based submodule based on the Ovis model is available

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

as a fallback. Ovis generates candidate metadata values directly from the visual content of the artefact through zero-shot inference, without relying on retrieved examples, and is triggered automatically whenever the confidence scores of the retrieval-based suggestions fall below a defined threshold.

Evaluation on the MET dataset showed that the retrieval-based component achieves Accuracy@3 above 50% for most metadata attributes, with particularly strong performance on classification (59.9%) and medium (68.9%). The VLM-based component complements these results in cold-start conditions, achieving Accuracy@1 of 76.6% for classification and 58.7% for medium in fully zero-shot settings. Both submodules were additionally tested on a small internal dataset of pilot artefacts, confirming their practical utility in the real-world scenarios of the REEVALUATE project.

7.3.2 Aural modality

The audio-based contextualization enabler was overhauled after the first round of testing. Both the back end and the front end had extensive additions and redesigns. The first extension in the addition of a second contextualization pipeline. This pipeline instead of a traditional classification model, uses a more robust and easier extended method; similarity search. This system creates latent representations of all the available audio tracks, and matches a newly uploaded track with the closest audio in the system. Then it provides the metadata. Lastly, the user interface was updated. Instruments are now grouped in categories based on the "Systematic of Instruments from Hornbostel and Sachs". These categories include: Idiophones, Membranophones, Chordophones, Aerophones and Electrophones as seen in Figure 8.



Figure 8 Categories of Detected instruments

7.3.3 Textual modality

Since R1, the internal architecture of the text contextualization tool has been formalised into a six-agent sequential pipeline orchestrated through Google ADK[1]. The stages are: ontology-aware Named Entity Recognition; Entity Linking against Wikidata[2], the Getty Art & Architecture Thesaurus (AAT) [3]and an internal index built over the project's Knowledge Graph; Coreference Resolution; ontology-constrained Relation

[1] https://adk.dev/, last accessed 29/04/2026
[2] https://www.wikidata.org/, last accessed 29/04/2026
[3] https://www.getty.edu/research/tools/vocabularies/aat/, last accessed 29/04/2026

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Extraction; Ontology Mapping into RDF/Turtle aligned to the CACAO ontology defined in D1.6; and a final RDF Validator.

While some minor changes were performed in the endpoint, targeting the structural quality of the produced graph, they do not alter the integration interface with the other components of the platform. The runtime requirements remain as documented in D4.3. The endpoint is hosted at KU Leuven facilities.

In addition to the backend service, a dedicated user interface for the text enabler[4] has been developed and integrated with the platform's central authentication: curators reach the tool through REEVALUATE's Single Sign-On. The interface exposes the knowledge graph produced by the pipeline as an editable view, in which the curator can accept, modify or discard individual triples. Pushing the curated output into the project Knowledge Graph is deliberately not yet enabled: pilot and benchmark results indicate that the current quality is not high enough to merge directly into the shared KG without risking pollution, so this final step awaits further improvement of the underlying pipeline.

A validation layer has been added on top of the baseline pipeline. The RDF produced by the Ontology Mapping agent is checked against the CACAO SHACL [1] shapes delivered in D1.6, and graphs that violate the shapes are returned to a repair agent that revises the output before re-validation, up to a configurable iteration cap. Gemini 2.5 Flash[5] is used as the default extraction model; it was selected over higher-F1 alternatives because of its stronger structural adherence and lower hallucination rate, which are the properties that matter when integrating into the rigid CIDOC-CRM patterns required by CACAO.

Between March and April 2026, the baseline plus SHACL configuration was tested by a panel of cultural heritage professionals from REEVALUATE pilot partners operating on institutional texts from their own collections. Their structured feedback has been integrated into D2.5 (R2 of the Contextualisation Enabler) and will be used to inform the future version for the second pilot run.

In parallel, a domain-specific Cultural Heritage benchmark has been constructed to evaluate the tool against CACAO-aligned gold standards rather than only general-purpose triple extraction. The benchmark contains thirty texts spanning four pilot domains, each with a gold graph produced independently by three knowledge engineers, mandatory event-mediated CIDOC-CRM [2] extraction, and required links to Wikidata, AAT and ICONCLASS[6]. Evaluation against this benchmark exposes a substantial structural gap, referred to in D2.5 as the distance problem, between performance on general extraction benchmarks [3] (Triple F1 in the 0.58–0.63 range) and CH-aligned performance (Triple F1 below 0.10 for the best baseline variant). The gap is driven by CIDOC-CRM's multi-hop event-centric patterns, such as E65_Creation and E15_Identifier_Assignment, which have no direct analogue in the single-hop relation patterns of general benchmarks.

Work building on the D2.5 results is already in progress. A series of prompt and architectural revisions has been carried out targeting the failure modes surfaced by the benchmark and the pilot feedback: an event-centric prompt rewrite addressing namespace consistency and predicate selection; a recall-oriented two-pass NER paired with a deterministic entity-linking confidence filter aimed at the completeness weakness reported by pilot users; deterministic post-processing of namespace prefixes; and an LLM-as-Judge [4] layer. A multi-phase variant of the pipeline has also been developed for long documents, introducing

[4] https://reevaluate.cs.kuleuven.be/playground, last accessed 29/04/2026
[5] https://deepmind.google/models/gemini/, last accessed 29/04/2026
[6] https://iconclass.org/, last accessed 29/04/2026

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

dedicated Temporal and Provenance agents to produce CIDOC-CRM event nodes and PROV-O-compatible attribution [5] alongside the core extraction.

8 Marketplace Functionalities

The remaining 2 functional workflows are integrated in the platform behind the Marketplace. The REEVALUATE Marketplace brings together the Cultural Heritage Institutions (CHIs) who own the digital artefacts with the Producers, who want to acquire the license to use or reuse the artefacts. The platform, using the advanced technologies of the Collaboration enabler, enhances the discoverability of digital artefacts through search and matchmaking functionalities and through the use of the AI based creative reuse enabler allows a producer to preview a digital artefact in the desired scenario. The end goal is the signing of a contract clearly stating the permissions and restrictions transferred from the CHI to the Producer. The final workflow allows both sides of the signed contract to verify if the production based on a digital artefact adheres to the signed contract.

8.1 Collaboration enabler

The Collaboration Enabler comprises two core systems: an Information Retrieval System and a Matchmaking System. Both systems are deployed as backend APIs for integration with the REEVALUATE marketplace, operating on the Knowledge Graph that stores the artefacts uploaded by pilot museums. To ensure data consistency, the system automatically synchronizes with the Knowledge Graph every two minutes, with the synchronization interval being configurable. Additionally, manual synchronization can be triggered via API (POST /sync/artifacts) to accommodate immediate updates when required.

The Information Retrieval System is built on a knowledge-enhanced multimodal retrieval architecture that integrates three complementary modules. The CLIP module, domain-adaptively fine-tuned on the project's ArtKB KG, supports both text-to-image and text-to-text semantic matching over artefact visual content and metadata. The Text2SPARQL module employs Mistral-small LLM to translate natural language queries into executable SPARQL queries, enabling precise retrieval based on structured metadata constraints. The MagicLens module supports Composed Image Retrieval, allowing users to search using a reference image combined with modification text; this module has been fine-tuned on ArtCIR, a newly constructed dataset of approximately 70,000 samples built from ArtKB image pairs sharing at least two metadata fields, with instructions generated via LLM. The retrieval results from CLIP and Text2SPARQL are combined through a weighted linear fusion strategy, with weights empirically optimized on a validation set. The system exposes a unified endpoint (POST /search/artifact) supporting three retrieval modes: text search, knowledge-enhanced text search, and composed image retrieval.

The Matchmaking System provides personalized artefact recommendations through dual-modality user profiling. The system combines semantic embeddings capturing conceptual relationships with visual embeddings encoding appearance-based characteristics. Since the previous deliverable, the semantic embedding approach has been updated from RDF2Vec to Sentence Transformer to address out-of-vocabulary issues encountered with unseen entities. Visual embeddings continue to be generated using TinyNeXt, a lightweight model trained via knowledge distillation from DINOv2. User profiles are constructed using the Frequent Directions (FD) algorithm, which maintains a compact representation of user preferences that evolves with interaction history. The system supports cold-start scenarios by initializing semantic profiles from user-declared preferences at registration, while different interaction types (e.g., purchases versus clicks) are assigned different weights to reflect varying preference signals. The

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

matchmaking functionality is exposed through three endpoints: GET /recommend/artifacts for personalized artefact recommendations, GET /recommend/users for discovering users with similar interests, and POST /sync/users for synchronizing user interaction data from the Analytics API.

8.2 AI Driven Creative Reuse enabler

The AI-Driven Creative Reuse Enabler aims to promote the reuse of digitized cultural artefacts by allowing users to generate new versions of historical items. By combining descriptive text prompts with optional image references from digital archives, users can creatively reimagine heritage materials while maintaining ties to their historical roots. For example, for the concept of historical fashion users can provide a prompt like “an emerald green formal dress from the 1700s” and receive a newly generated image reflecting that clothing style. The enabler is powered by Stable Diffusion 3 model which is fine-tuned to address the diverse visual and conceptual needs of the REEVALUATE project.

The enabler operates through a cloud-based, high-performance infrastructure supporting the computational demands of AI-driven image generation. It is accessible via a secure web API, allowing external services to integrate it easily. At the heart of the image generation process is the RunPod Instance, an on-demand GPU environment that dynamically activates when needed and powers down during inactivity to optimize resource usage. Specifically, the system utilizes a single 16GB VRAM GPU such as the RTX A4000, RTX A4500, RTX 4000 Ada, and RTX 2000 Ada and requires 20GB of storage.

The Backend API acts as a secure intermediary between external services and the RunPod Instance. Hosted separately on a virtual machine with 4 CPU cores, 8GB of RAM and 20 GB of storage, it manages authentication through Auth0 SSO and ensures structured request handling. Together, the RunPod Instance and the Backend API form the core components of the AI-Driven Creative Reuse Enabler.

8.3 Distributed Ledger and IPR Management

The REEVALUATE IPR Management Enabler uses Non-Fungible Tokens (NFTs) as the core mechanism for managing provenance, ownership, authenticity, and licensing of cultural-heritage (CH) assets. Each digitised artefact ingested into the platform is represented by a unique NFT that is cryptographically linked to the asset’s metadata and digital files. In this way, the NFT acts simultaneously as proof of authenticity and origin, a container for licensing and rights information, and a persistent identifier that enables controlled sharing and reuse across REEVALUATE services.

When a museum or cultural institution digitises an artefact and registers it in the Marketplace, the Marketplace backend invokes the IPR Enabler to mint an NFT on the Hyperledger Fabric network. The token embeds the artefact identifier together with the hash of the stored digital record, ensuring long-term integrity and verifiability of the asset. Derivative content produced through the Creative Reuse Enabler is also registered as linked NFTs, establishing explicit relationships between new creations and the rights obligations of the original material.

All NFT-related actions—including minting, ownership transfer, licensing, and proof-of-ownership verification—are initiated through the Marketplace frontend and executed by the Marketplace backend as signed blockchain transactions. Integration with the distributed ledger is achieved through a REST Smart Contract Gateway that connects the Marketplace backend to the Hyperledger Fabric chaincode, while the Community Management Service (CMS) enforces authentication and authorisation policies for NFT operations.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

This architecture enables REEVALUATE to provide a seamless and user-friendly interface without requiring end users to possess blockchain expertise. At the same time, the distributed ledger guarantees a tamper-evident and auditable lifecycle for all cultural-heritage assets and their derivatives. RESTful APIs further allow platform services and external applications to invoke NFT functions and query token state or transaction history for rights verification and provenance tracking. Full technical details of the ledger configuration, smart contracts, and integration workflows are provided in Deliverable D3.4.

8.4 Context Validation enabler

The Context Validation Enabler aims to reliably validate multimodal content before final upload to online platforms and marketplaces. Within the REEVALUATE platform, the Context Validation enabler analyses uploaded content before publication or reuse in the Marketplace. It supports automated platform workflows triggered by the AI driven creative reuse enabler as well as user-driven validation scenarios by providing contextual analysis, confidence-based outputs, and feedback mechanisms that help identify potential violations or inconsistencies early in the content lifecycle. The enabler analyses multiple modalities and can process images, videos, audio, and text in order to detect contextual violations, such as crime instances, policy conflicts, propaganda and fake content.

The need for such a service arises from the increasing volume of multimodal digital content and the growing use of creative reuse workflows across online platforms. Existing validation and moderation solutions are often fragmented, modality-specific, or dependent on manual review processes. The Context Validation enabler addresses these challenges by combining computer vision, audio analysis and text processing techniques into a unified validation framework. Additionally, checking and validating the content becomes much faster, especially for long videos and audio files, since the enabler is designed to process large multimodal inputs efficiently.

The updated Context Validation enabler addresses several limitations identified by the end users during the pilots (see Table 1), improving the clarity of outputs, the detection reliability, providing clearer explanations of what each output is, and proving user guidance during the validation process. The updated version also strengthens integration with Marketplace infrastructures and creative reuse workflows, enabling more reliable validation and verification of reused or generated content within the REEVALUATE platform. Additionally, the system supports contextual feedback mechanisms that help users better understand and resolve potential issues before publication.

The core components of the system include multimodal AI services responsible for content analysis and contextual validation, as well as a RESTful API for integration with the Marketplace and potentially third-party systems and a web-based user interface allowing autonomous content analysis. All of the components are dockerized and served from servers in CERTH's premises.

9 Conclusion & Future Work

The integration of the REEVALUATE Platform marks a significant milestone in the project's progression. By consolidating the technological outputs of multiple work packages into a single, interoperable environment, the platform demonstrates the feasibility of an ecosystem that bridges artificial intelligence, semantic data representation, participatory engagement, and blockchain-based intellectual property management. It provides a robust foundation upon which cultural heritage institutions can digitize, enrich, and share their collections responsibly and innovatively.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

The modular microservice-based architecture ensures that each enabler—ranging from artefact prioritization to contextualization, validation, and creative reuse—operates as an independent yet interconnected service. This design choice enhances flexibility, scalability, and maintainability while allowing for the seamless incorporation of new tools and functionalities. The Marketplace and Hub serve as central access points, ensuring that end-users, including curators, researchers, and creative practitioners, can interact with the system intuitively and securely.

The transition of the platform from the first to the second viable version focused on the improvement of interoperability between the enablers, enhancing performance, and extending the platform's usability taking into account the outcomes of the project pilots. The pilot implementations and stakeholder feedback guided the development of advanced user interfaces and integration pathways of internal and external repositories and platforms.

In conclusion, the Integrated REEVALUATE Platform provides a forward-looking model for how Europe's cultural heritage can be digitized, managed, and creatively reused in an ethical and collaborative digital ecosystem. Its architecture and philosophy embody the project's vision of an inclusive, sustainable, and technologically empowered future for cultural heritage management.

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

References

[1] Knublauch, H., Kontokostas, D. (eds.) Shapes Constraint Language (SHACL). W3C Recommendation, 20 July 2017. https://www.w3.org/TR/shacl/

[2] Doerr, M. The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata. AI Magazine 24(3), pp. 75–92, 2003.

[3] Mihindukulasooriya, N., Tiwari, S., Enguix, C. F., Lata, K. Text2KGBench: A Benchmark for Ontology-Driven Knowledge Graph Generation from Text. International Semantic Web Conference (ISWC), 2023.

[4] Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E., Stoica, I. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. NeurIPS, 2023.

[5] Lebo, T., Sahoo, S., McGuinness, D., Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., Zhao, J. (2013). PROV-O: The PROV Ontology. W3C Recommendation. World Wide Web Consortium. http://www.w3.org/TR/2013/REC-prov-o-20130430/

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

Annex A

PARTNER		SHORT NAME
CERTH CENTRE FOR RESEARCH & TECHNOLOGY HELLAS	CENTRE FOR RESEARCH & TECHNOLOGY HELLAS	CERTH
Fraunhofer FOKUS	FRAUNHOFER INSTITUTE FOR OPEN COMMUNICATION SYSTEMS FOKUS	FOKUS
.fs	FROMSCRATCH DESIGN STUDIO BV	FS
FONDAZIONE links PASSION FOR INNOVATION	FONDAZIONE LINKS - LEADING INNOVATION & KNOWLEDGE FOR SOCIETY	LINKS
FFP PRODUCTIONS	FFP FILM- & FERNSEHPRODUKTION GMBH	FFP
GVAM MOBILE GUIDES FOR ALL	GVAM GUIAS INTERACTIVAS SL	GVAM
Stiftung Preußischer Kulturbesitz	STIFTUNG PREUSSISCHER KULTURBESITZ	SPK
Hypertech	HYPERTECH	HYP
Thessaloniki Olympic Museum	OLYMPIAKO MOUSEIO	OLYMPIC
ARTHUR LEGAL, STRATEGIES & SYSTEMS	ARTHUR'S LEGAL BV	ARTHUR
KU LEUVEN	KATHOLIEKE UNIVERSITEIT LEUVEN	KU Leuven
MODEMUSEUM ANTWERPEN	AG CULTURELE INSTELLINGEN ANTWERPEN/ ERFGOED	MoMU

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389

	EUROPEAN FASHION HERITAGE ASSOCIATION	EFHA
	3D RESEARCH SRL	3DR
	FONDAZIONE AQUILEIA	AQUILEIA

Funded by the European Union

This project has received funding from the European Union Horizon Research and Innovation programme under grant agreement No 101132389