

Deliverable 2.1: Report on the knowledge management platform specifications

Work Package 02

DOI: 10.5281/zenodo.24530896



Co-funded by the European Union under Grant Agreement n°101166718.

Document information

Project Acronym	**EURAD-2**
Project Title	**European Partnership on Radioactive Waste Management-2**
EC grant agreement No.	**101166718**
Work Package Title	**Knowledge Management**
Deliverable No.	**2.1**
Deliverable Title	**Report on the knowledge management platform specifications**
Lead Beneficiary	**ASNR**
Contractual Delivery Date	**01/10/2026**
Actual Delivery Date	**30/11/2026**
Dissemination level	**Public**
Authors	**Christophe Debayle (ASNR), Yevheniia Kudriashova (SSTC NRS).**

To be cited as:

Debayle C., Kudriashova Y. (2026): Report on the knowledge management platform specifications. Final version as of November 2026 of deliverable D2.1 of the European Partnership EURAD-2. EC Grant agreement n°:101166718

Disclaimer

All information in this document is provided "as is" and no guarantee or warranty is given that the information is fit for any particular purpose. The user, therefore, uses the information at its sole risk and liability. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Atomic Energy Community. Neither the European Union nor the granting authority or the individual Colleges of EURAD-2 can be held responsible for them.

Acknowledgement

This document is a deliverable of the European Partnership on Radioactive Waste Management 2 (EURAD-2). EURAD-2 is co-funded by the European Union under Grant Agreement N° 101166718.

Status of deliverable		
	By	Date
Delivered (Lead Beneficiary)	Christophe Debayle - ASNR (WP02 task 5 leader)	06/10/2026
Verified (WP Leader)	Alexandru Tatomir – BGE (WP02 leader)	16/10/2026
Reviewed (Reviewers)	Erika Holt – VTT (PMO representative)	21/09/2026
Approved (PMO)	Erika Holt – VTT (PMO representative)	24/11/2026
Submitted to EC (Coordinator)	Andra	25/11/2026

Executive Summary

The EURAD Knowledge Management Platform has been under specification development since 2020. Deliverable 11.1 of the EURAD(-1) programme gathered requirements, Deliverable 11.9 produced a full architectural and functional blueprint for a portal knowledge management system, Milestone 242 defined an implementation life cycle, and Milestones 245 and 333 reported on a prototype developed in the final weeks of EURAD(-1) in 2024. A SharePoint-based prototype was developed and tested but was not carried into service. EURAD-2 MS22 - Memo on KM platform established why no operational platform resulted, and the finding was unambiguous, the reason why it did not result in a sustainable operational platform was not technical. It followed from a scope that was too large to be funded in one step, from requirements that were captured from a small sample and then treated as settled, from a discontinuity between the people who defined the needs and the people who built the prototype, and from the absence of any entity able to guarantee that the result would still be running in ten years.

This deliverable therefore does not restate the blueprint. It specifies a first release that can be built, accepted and operated, and it defines the modular path by which the remaining capabilities are added.

The first operational release of the EURAD Knowledge Management Platform is now intended to be a public web section, hosted on European Commission web infrastructure, structured as a twin of the EU Nuclear Decommissioning Knowledge Management section operated by the Joint Research Centre (JRC). It should contain a landing page describing the EURAD knowledge management activity, a complete and continuously updated database of EURAD public deliverables and equivalent knowledge products, entry points towards EURAD training and events, and a named contact channel. Documents should not be organised into a fine-grained thematic classification, because the retrieval layer, not the folder tree, is the instrument of discovery.

On top of that corpus, the release specifies an artificial intelligence retrieval layer that answers natural language questions using only the EURAD corpus, and that returns, with every answer, the identity of the source documents and the location within them. The layer is specified against strict traceability, evaluation and human accountability requirements, because an unreliable answer about a safety relevant subject is worse than no answer at all.

Collaborative workspaces, forums and communities of practice tooling, granular access control for restricted material, project management functions, semantic ontology engineering, and the migration of existing EURAD tools are all excluded from the first release. They are not withdrawn from the vision. They are recorded as a prioritised backlog in Section 5.7 and are addressed by later modules once the foundation is stable and adopted.

1. Ownership of the knowledge must sit with the owner of the mission. No beneficiary organisation can commit to maintaining a knowledge asset beyond its own funded projects, and the first EURAD-2 annual event data management (DM) session feedback stated this explicitly. The European Commission funds the partnership, defines its policy purpose and outlasts each of its funding instruments, so hosting the platform on Commission infrastructure aligns the custody of the knowledge with the durability of the mandate.
2. The Commission already runs the equivalent activity for decommissioning, with a published web section, a knowledge products database, a training programme and active development on semantic tooling. The two communities share subject matter, stakeholders and preservation horizons. Building a second parallel infrastructure would duplicate public expenditure and would reproduce exactly the proliferation of unconnected portals that the community asked to stop.
3. A small system that exists outperforms a complete system that does not. Every capability added to the first release reduces the probability that any of it is delivered, and modularity is what allows ambition to be preserved without being front loaded.

Three hosting options are assessed in Section 6, hosting on European Commission infrastructure in association with the Joint Research Centre, hosting by an EURAD beneficiary, and hosting by a commercial supplier under a tender. The first is recommended, principally because it is the only one that satisfies the longevity requirement without a recurring financial commitment from the partnership.

Discussions with the Joint Research Centre are ongoing and have already established two working conclusions, that a transfer of existing EURAD websites as such is not the realistic route, and that a dedicated EURAD entry point within the Commission web presence together with interoperability between repositories is. Section 6.6 defines the portability guarantees that make the corpus independent of the hosting decision.

Four releases are specified. Release 0 consolidates the corpus and applies the minimum metadata set. Release 1 delivers the public web section and the deliverables database. Release 2 adds the retrieval layer, first as a restricted pilot and then publicly. Release 3 opens the modular extension programme, including legacy recovery from EURAD-1, PREDIS and earlier projects, and interoperability with the decommissioning knowledge products database. Acceptance criteria are defined for each release in Section 7.3, so that the transition between releases is a decision based on evidence rather than on optimism.

Keywords

EURAD-2 knowledge management open access, DataCite, persistent identifier, knowledge preservation, sustainability, long term hosting, European Commission, Joint Research Centre, FAIR principles, metadata, retrieval augmented generation, modular architecture, minimum viable product.

Table of content

Executive Summary 4
Keywords 5
Table of content 6
List of Tables 8
1. Introduction 9
1.1 Purpose and scope 9
1.2 Position within the knowledge management work package 9
1.3 Sources 9
1.4 How to read this document 10
2. Consolidated statement of needs 10
2.1 The EURAD-1 vision, as planned 10
2.2 The EURAD-2 reality check 10
2.2.1 Conserve 11
2.2.2 Cease 11
2.2.3 Create 11
2.3 The two chronical problems 11
2.3.1 The fragmented workflow. 12
2.3.2 The lost legacy. 12
2.4 Needs register 12
2.5 Answers to the open strategic questions 14
3. Rationale 14
3.1 Durability first, the KM-platform should outlast the project 14
3.2 Mutualisation for a single European nuclear knowledge management effort 15
3.3 Small, modular and scalable 16
3.4 Knowledge is the value, not its classification 16
3.5 Objections considered 17
4. Target concept 18
4.1 The reference model 18
4.2 A twin as first step 18
4.3 Structure of the first release 19
4.4 The knowledge extraction layer 19
4.5 Position within the existing EURAD landscape 20
5. Specifications 21
5.1 Architecture 21

5.2 Functional requirements 22
5.3 Minimum metadata set 23
5.4 Requirements on the knowledge extraction layer 23
5.5 Non functional requirements 25
5.6 Interoperability 26
5.7 Explicit exclusions and backlog 27
6. Hosting and governance 27
6.1 The hosting requirement 27
6.2 Options assessed 28
6.3 Recommendation and fallback 28
6.4 Roles and responsibilities 29
6.5 Continuity beyond EURAD-2 29
6.6 Portability guarantees 29
7. Implementation 30
7.1 Releases 30
7.2 Indicative schedule 31
7.3 Acceptance criteria 32
7.4 Resources 32
7.5 Risks 33
8. Conclusion 34
Appendix A. Traceability of requirements to needs 35
Appendix B. Abbreviations 36
References 37

List of Tables

Table 1. Consolidated needs register. 13
Table 2. Proposed answers to the open strategic questions of Milestone 2.2. 14
Table 3. Objections and responses. 17
Table 4. Objects of the first release. 19
Table 5. Position of the platform within the existing landscape. 20
Table 6. Functional requirements. 22
Table 7. Minimum metadata set. 23
Table 8. Requirements on the knowledge extraction layer. 24
Table 9. Non functional requirements. 25
Table 10. Interoperability requirements. 26
Table 11. Excluded capabilities and backlog. 27
Table 12. Hosting options. 28
Table 13. Roles and responsibilities. 29
Table 14. Release plan. 30
Table 15. Indicative schedule. Dates are indicative and are subject to the endorsement and hosting decisions identified above. 31
Table 16. Acceptance criteria. 32
Table 17. Risk register. 33

1. Introduction

1.1 Purpose and scope

This deliverable specifies the EURAD Knowledge Management Platform. Specification here means three things: a statement of the needs the platform must meet, a statement of the requirements that any acceptable implementation must satisfy, and a statement of the sequence in which those requirements are to be met.

The scope of the deliverable is the platform itself. It does not cover the Data Management Plan, its associated training and guidance, or the resolution of partner access rights to the collaborative workspace. Those needs were raised with great force by the community and are genuine, but they are the object of other sub-tasks within Work Package 02 task 05 and they are solved by procedure and training rather than by software. Section 2.4 records them and identifies where they are handled, so that no need documented in Milestone 2.2 is silently dropped.

1.2 Position within the knowledge management work package

This deliverable is the direct successor of Milestone 2.2, issued in November 2025, which synthesised the EURAD-1 planning documents, the outcomes of the data management session held at the first EURAD-2 annual event DM session annual workshop in September 2025, and the strategic analysis conducted by Task 5. Milestone 2.2 concluded that the strategy had to shift from a single step implementation to a phased, user centred approach, and it announced that the resulting specification would be delivered here.

Milestone 2.2 also listed five open strategic questions requiring a decision from EURAD governance, concerning open access policy, long term hosting and budget, the structuring principle of the platform, the nature of the search capability, and the ownership of the system beyond the life of EURAD-2. This deliverable proposes an answer to each of them, in Section 2.5, and those answers are the foundation of the specification that follows. Where an answer requires a decision that Task 5 cannot take alone, the decision is identified as such and referred to governance.

1.3 Sources

The specification draws on the following material.

- EURAD-2 Grant agreement (task 5 activities description and D2.1 scope).
- Deliverable D7.7 of PREDIS.
- Deliverable 1.3 of EURAD-2, Data management plan.
- Deliverable 11.1 of EURAD-1, screening and review of existing knowledge management approaches and tools, which captured the initial expectations of the community through a survey of eleven partner organisations.
- Deliverable 11.9 of EURAD-1, specification of the EURAD KM platform, which produced the architectural, non-functional and functional blueprint of the proposed portal knowledge management system.
- Milestone 242 of EURAD-1, which defined the implementation stages of the prototype.
- Milestones 245 and 333 of EURAD-1, which reported the implementation status of the prototype and the lessons learned.
- Milestone 2.2 of EURAD-2, which established the current diagnosis and the strategic reorientation.
- The consolidated output of the data management session held during the EURAD-2 annual workshop in first EURAD-2 annual event DM session in September 2025, organised as “Conserve”, “Cease” and “Create”.
- The published description of the EU Nuclear Decommissioning Knowledge Management initiative operated by the Joint Research Centre, used in this document as the reference model.
- Exchanges held between Task 5 and the Joint Research Centre between February and June 2026 on sustainable hosting and on interoperability between repositories.

1.4 How to read this document

Section 2 states the needs. Section 3 states the reasoning, and it is the argumentative core of the deliverable, since the specification only makes sense once the reasons for its restraint are understood. Section 4 describes the target concept. Section 5 contains the requirements themselves, expressed as numbered and testable statements. Section 6 addresses hosting and governance. Section 7 gives the roadmap, the acceptance criteria and the risks. Appendix A maps every requirement back to the need it serves, so that the specification can be audited against the documented demand of the community.

2. Consolidated statement of needs

This section synthesises, in a form short enough to be usable, the needs established across the EURAD-1 specification work and the EURAD-2 community feedback. It does not repeat the detailed narrative of Milestone 2.2, to which the reader is referred for the full analysis.

2.1 The EURAD-1 vision, as planned

The initial survey conducted for Deliverable 11.1 gathered high level expectations from eleven partner organisations, six waste management organisations, two research entities and three technical safety organisations. It is important to record the limitation attached to that evidence base, the sample was small and the raw survey material is no longer available for re-analysis. The conclusions drawn from it are therefore a high level interpretation rather than a validated requirements baseline, a point that matters when those conclusions are used, as they have been, to justify an expensive system.

The survey nevertheless produced clear directives. Partners asked for a professional, integrated web based platform and explicitly rejected simple wiki based approaches and plain repositories as insufficient. They asked for a culture of collective contribution supported by interactive tools rather than a static library, for the integration of artificial intelligence tools to be planned from the outset, particularly for advanced search, for a shared organisational framework to prevent the system from degenerating into a data swamp, for robust and granular access control since not all material can be shared, and for harmonisation with the European Commission, the International Atomic Energy Agency and the OECD Nuclear Energy Agency in order not to reinvent the wheel. The same survey revealed that there was no consensus on how to build any of it. The structuring of information was described as unresolved, input and output formats were unsettled, and a central tension was left open between manual metadata curation and reliance on a powerful search engine.

Deliverable 11.9 then translated these expectations into a full blueprint, with seven architectural and non functional specifications, modularity and longevity, scalability, interoperability, security and compliance, performance and reliability, usability and accessibility, and long term preservation and portability, and with core functional specifications covering content and document management, search and retrieval, and collaboration and community features. The associated implementation strategy already recommended open source components, phased implementation, deferral of complex features such as full semantic search and deep artificial intelligence recommendation, and configuration rather than custom development. That advice was sound and it was not followed. The estimate associated with the full vision, of the order of several million euros for development and fifty years of support, is itself an explanation of why, since no such commitment was available.

2.2 The EURAD-2 reality check

The data management session held at the first EURAD-2 annual event DM session annual workshop in September 2025 produced feedback of a different character. It was operational, immediate and, on several points, unanimous. It is summarised here under the three headings used at the session. It is also important here to note that some topics considered to be conserved, are partially implemented and still need improvement.

2.2.1 Conserve

- The use of a shared collaborative workspace as the central place for work in progress, recognised as far superior to email exchange, together with its existing structure and labelling.
- The principle of a Data Management Plan, on the condition that it becomes a flexible framework that is properly trained and enforced rather than a document that is filed and forgotten.
- The commitment to open science, open access publication and the release of raw data alongside papers, seen as essential for transparency and for the evidence base needed in dialogue with regulators and civil society.
- Practices that force data standardisation in a natural way, in particular cross partner benchmark exercises.
- The use of established open infrastructures for archiving datasets with complete metadata, and the retention of large volumes in slow storage rather than their deletion.

2.2.2 Cease

- Local storage of raw data and results on personal computers, local servers or consumer file sharing services, identified as the root cause of the silo effect and as a permanent data loss risk.
- Exchange of data by email, condemned as insecure and as a primary source of fragmentation and version chaos.
- Development of data management procedures at work package level without central guidance, which produces parallel repositories.
- Creation of new, unconnected portals. The message was to improve and connect what exists rather than to add to it.
- Inefficient reporting, very large state of the art documents that are hard to find and read, and bureaucratic language.
- The debate on whether research data management is worthwhile. Its value is settled and the discussion must move to implementation.

2.2.3 Create

- A sustainable data repository with a guaranteed lifetime beyond EURAD-2, preferentially located on a central European Union platform.
- Universal access to information, resolution of the differential access rights that currently prevent some partners from reaching all levels of the collaborative workspace, and substantially improved search across deliverables and presentations.
- Artificial intelligence support, explicitly including a chatbot to simplify navigation and retrieval, assistance in classification and tagging, and support in synthesising large volumes of material.
- Persistent identifiers for datasets and deliverables so that they become citable first class research outputs, and the transfer of historical data to open archives.
- Reproducibility oriented preservation, in which input files, models and the measurement sources they derive from are stored together and explicitly linked.
- Practical guidance and training, in particular a short and usable version of the Data Management Plan and clear orientation material for newcomers.
- Open standards, open source components and non proprietary interoperable formats.
- Named responsibilities and allocated resources, since data management cannot be an unfunded additional task carried by technical experts.

The distance between Section 2.1 and Section 2.2 is the central fact of this deliverable. The EURAD-1 requirements describe a system. The EURAD-2 feedback describes a set of problems or needs. A specification that answers the first without answering the second will fail again.

2.3 The two chronical problems

Milestone 2.2 named the two failures that the platform exists to address, and they are restated here because every requirement in Section 5 is justified by one of them.

2.3.1 The fragmented workflow.

Knowledge produced by the partnership remains dispersed across personal storage, local servers, restricted workspaces and electronic mail. It is not that the knowledge is absent, it is that finding it requires knowing in advance who produced it and where they put it. This makes reuse dependent on personal networks, which is precisely the dependency that a knowledge management system exists to remove.

2.3.2 The lost legacy.

Knowledge does not reliably survive the project that produced it. Data supporting a deliverable becomes unavailable once the contract closes, the person who understood it moves on, and the storage that held it is decommissioned. On a mission whose horizon is measured in generations, and where the justification of a disposal concept may need to be re-examined decades after it was established, this is not an inconvenience. It is a direct loss of the safety case evidence base.

A third problem, more specific to this activity, must be added. Successive attempts to solve the two chronic problems have taken the form of comprehensive platform designs that were never implemented. Each attempt has consumed effort, produced documentation, and left the underlying problems untouched, while adding to the community the impression that the subject is discussed rather than acted upon. The specification proposed here treats that history as a design constraint.

An important point to note is that this deliverable should first address the document and knowledge component, while the data component is managed separately through dedicated tools such as GitHub and Zenodo.

2.4 Needs register

The following register consolidates the needs into a form against which requirements can be traced. The priority column states the treatment applied in this specification. Deferred does not mean rejected, it means scheduled after the foundation is operational, and the corresponding entries are developed in Section 5.7.

ID	Need	Principal origin	Treatment
N-01	Guaranteed long term hosting and preservation of EURAD knowledge assets by a stable European entity, independent of any single beneficiary funding cycle	First EURAD-2 annual event DM session «Create», MS2.2 section 6.2	Release 1, foundational
N-02	A single, publicly findable entry point giving access to all EURAD public outputs	D11.1, First EURAD-2 annual event DM session «Create»	Release 1, foundational
N-03	Substantially improved discovery of deliverables, presentations and knowledge products	First EURAD-2 annual event DM session «Create»	Release 1
N-04	Reliable extraction of knowledge from the document corpus, including natural language and assistant style access	D11.1, First EURAD-2 annual event DM session «Create», MS2.2 open question 4	Release 2
N-05	Persistent, citable identification of deliverables and datasets	First EURAD-2 annual event DM session «Create»	Release 1 for deliverables, Release 3 for datasets
N-06	Preservation and recovery of legacy knowledge from EURAD-1, PREDIS and earlier European projects	MS2.2 section 6.3, First EURAD-2 annual event DM session «Create»	Release 3
N-07	Open, non proprietary formats and demonstrable portability of all content	D11.9 non functional specification 7, First EURAD-2 annual event DM session «Create»	Release 0, foundational
N-08	Granular access control and defined user roles for restricted material	D11.1, D11.9 non-functional specification 4	Deferred, Release 1 is public only
N-09	Collaboration features, forums and communities of practice spaces	D11.1, D11.9 functional specification 3	Deferred, see section 5.7
N-10	Resolution of differential partner access to the collaborative workspace	First EURAD-2 annual event DM session «Create»	Out of scope, handled by WP2 and the programme management office
N-11	Practical Data Management Plan guidance, training and a short usable guideline	First EURAD-2 annual event DM session «Create»	Out of scope, handled by other WP2 tasks
N-12	Harmonisation with European Commission, IAEA and OECD/NEA activity, and avoidance of duplication	D11.1	Design constraint, all releases
N-13	No proliferation of new unconnected portals	First EURAD-2 annual event DM session cease	Design constraint, all releases
N-14	Named responsibilities, curation and allocated resources for maintenance	D11.1 lessons learned, First EURAD-2 annual event DM session «Create»	Release 1, governance
N-15	Directory of experts with expertise keywords, and partner contact points	First EURAD-2 annual event DM session «Create»	Deferred, see section 5.7
N-16	Reproducibility oriented preservation linking data, models and measurement sources	First EURAD-2 annual event DM session «Create»	Deferred, dependent on work package pilot

Table 1. Consolidated needs register.

2.5 Answers to the open strategic questions

Milestone 2.2 referred five questions to EURAD governance. Leaving them open is itself one of the causes of the current situation, since no specification can be completed while they remain unanswered. The following answers are proposed by Task 5 and are the basis of this specification. Those marked as requiring GA endorsement are submitted to governance for formal confirmation.

Question	Proposed answer
Open access. What must be fully open and what is restricted to partners.	The platform, in its first release, publishes only material that is already public or that has been cleared for publication. Restricted material stays in the collaborative workspace and is not in scope. This removes the access control requirement from the critical path without changing the confidentiality regime of any document. Requires endorsement.
Long term hosting and budget. Whether to commit to a large development and fifty-year support cost.	No. The partnership commits instead to a minimal, open standards-based system hosted on European Commission infrastructure, with content held in open formats that can be migrated at any time. The fifty-year commitment is placed on the format and on the custodian, not on the software. Requires endorsement.
Structuration. Whether the platform is optimised for specialists or for a broader audience.	For specialists and informed professionals, with a public presentation layer that is legible to a broader audience. The corpus is not restructured for outreach purposes, since outreach is served by the programme communication site.
Search engine. Whether a semantic engine is required or keyword search suffices, and whether a chatbot interface should be considered.	Both, in sequence. Release 1 provides faceted metadata search, which is the reliable baseline. Release 2 adds a retrieval layer with a natural language interface, specified in Section 5.4 with mandatory source citation. The chatbot never replaces the ability to browse and download the underlying documents.
Governance and maintenance. Who owns and maintains the platform beyond EURAD-2.	The custodian of the infrastructure is the European Commission. The editorial responsibility during EURAD-2 should rely on WP02 KM but this is not absolutely clear at the moment, and the transfer arrangement to be applied at the end of the partnership is defined in Section 6.5. Requires endorsement.

Table 2. Proposed answers to the open strategic questions of Milestone 2.2.

3. Rationale

This section sets out why the platform is specified in the way that it is. It is placed before the requirements because the requirements are unremarkable in isolation and only become defensible in the light of the reasoning. Four arguments are developed, followed by an explicit treatment of the objections that can reasonably be raised against them.

3.1 Durability first, the KM-platform should outlast the project

Radioactive waste management operates on timescales that have no equivalent in ordinary information technology planning. A disposal concept developed today may be revisited, challenged and re-justified over a century. The knowledge that supports it must remain available, intelligible and attributable across that period, and across the disappearance of the individuals, the software and in some cases the organisations that produced it.

Against that horizon, the arrangement under which EURAD knowledge is currently held is inadequate in a specific and remediable way. The tools that already form part of the knowledge management landscape, the programme website, the training platform, the wiki and several dispersed initiatives, are operated under the responsibility of beneficiaries whose commitment is bounded by the projects that

fund them. Milestone 2.2 stated the consequence plainly, no partner funded by the European Commission can guarantee the sustained financing and maintenance of these tools over the long term. This is not a criticism of any organisation, it is a structural property of project funding and spread over numerous projects as also considered by Spichtinger D. et al., 2022.

The community reached the same conclusion independently during the first EURAD-2 annual event data management session, where the requirement was expressed as a data repository with a guaranteed lifetime beyond EURAD-2, preferentially located on a central European Union platform. The convergence of a top-down analysis and a bottom-up demand on the same point is the strongest signal available in the whole evidence base, and this specification treats it as the primary constraint.

The reasoning that follows is straightforward. The European Commission funds the partnership, defines the policy purpose that the partnership serves, and persists across successive funding instruments in a way that no individual beneficiary does. Placing the custody of the knowledge with the entity that owns the mission aligns the durability of the asset with the durability of the mandate. It also removes a recurring negotiation, since the alternative arrangement requires that some organisation volunteer, at the end of every programme, to carry a cost that brings it no direct benefit.

A secondary argument reinforces the first. Hosting within the Commission web presence brings the platform inside an infrastructure that already applies corporate policies on accessibility, multilingual presentation, security, personal data protection and web archiving. These are obligations that the partnership would otherwise have to specify, procure and audit for itself, at a cost that is invisible in project planning and substantial in practice.

3.2 Mutualisation for a single European nuclear knowledge management effort

The European Commission already operates, through the Joint Research Centre, a knowledge management initiative for nuclear decommissioning (EU Nuclear Decommissioning Knowledge Management). That initiative has a published web section, a knowledge products database, a training programme, a summer school, established relations with the International Atomic Energy Agency and the OECD Nuclear Energy Agency, and current development work on semantic tools and on a platform prototype.

Waste management and decommissioning are distinct activities, but they are adjacent in every dimension that matters for knowledge management. They share subject matter, since decommissioning produces the waste that waste management characterises, conditions and disposes of. They share stakeholders, since the same operators, technical safety organisations, regulators and research entities appear in both communities. They share the preservation horizon, and they share the same international reference organisations. They also share, at this moment, the same technical problem, how to make a large heterogeneous document corpus reliably searchable.

A first consequence is financial because building a second European infrastructure for the same function, from the same public source of funding, is duplication, and it would be difficult to justify to the funder at the moment when the partnership is asking that funder to host it. The second is substantive. The EURAD community asked, in the clearest terms recorded at first EURAD-2 annual event DM session, that the creation of new unconnected portals stop. A EURAD platform built independently of the Commission activity would be exactly such a portal, however well built, and it would carry the additional defect of being invisible to the neighbouring community.

Mutualisation is therefore specified as a design constraint rather than as an aspiration. The EURAD knowledge management section is specified as a twin of the decommissioning section, which means that it adopts the same structural pattern, the same publication infrastructure and, where possible, the same retrieval tooling. Interoperability between the two corpora is a requirement, not an option, and it is specified in Section 5.6. The objective is that a professional working on decommissioning waste should be able to reach the relevant EURAD material without knowing that EURAD exists.

3.3 Small, modular and scalable

The third argument concerns sequencing, and it is the one on which this deliverable departs most sharply from its predecessors.

EURAD-1 produced a comprehensive specification, an implementation life cycle, and a prototype developed in the final weeks of the programme. The result was not an operational platform. Milestone 2.2 identified the reasons, and none of them was a failure of technical design. The scope could not be funded in one step. The requirements had been captured from a small sample and were then treated as settled. Continuity between the people who understood the needs and the people who built the prototype was lost. The technology choice for the prototype was not clearly justified and overlapped functionally with a tool the community already used. And underneath all of it, there was no answer to the question of who would still be running the system in ten years.

The lesson recorded in Deliverable 11.1 is worth repeating exactly because it predicted this outcome, knowledge management systems fail for organisational reasons, not technical ones, and the reasons cited for abandoning previous systems were lack of sponsorship, poorly organised content provision, poor promotion and training and failure to keep the system updated and moderated. Each case is a failure of sustainability and the bigger the project, the more likely these failures become.

The specification therefore adopts an explicit rule. The first release contains only what is necessary to solve the most urgent problem, and each additional capability must earn its place by demonstrating that the previous release is stable and used. This is not caution for its own sake. It reflects three properties of small first releases that matter here.

1. They are fundable. A bounded scope can be costed, approved and delivered within a single programme, which removes the dependency on a decision that nobody can take.
2. They generate real requirements. The requirements gathered so far are declarations of preference collected in surveys and workshops. Requirements derived from observing what people actually search for, and fail to find, in a live system are of a different quality, and they are the only reliable basis for deciding what the second release should contain.
3. They are recoverable. If a small release is wrong, it is corrected. If a comprehensive release is wrong, the correction is unaffordable and the system is abandoned, which is the history this specification is trying not to repeat.

Modularity is what makes this restraint compatible with ambition. The architecture specified in Section 5.1 separates the corpus and its metadata from the presentation layer and from the retrieval layer. Because the corpus is held in open formats with a documented metadata schema, any later module, a collaboration space, an expert directory, a training module, a reproducibility package, can be added without disturbing what already works, and the corpus survives the replacement of any tool built on top of it. The full EURAD-1 vision is not abandoned. It is placed on a path where each step can actually be taken.

3.4 Knowledge is the value, not its classification

The fourth argument settles the tension that Deliverable 11.1 identified and left open, whether documents must be categorised if the search capability is good enough.

The position taken here is that the completeness of the corpus and the quality of its descriptive metadata are foundational, while fine grained thematic classification is not. There are three reasons. First, thematic classification is expensive, it is performed by the experts whose time is the scarcest resource in the partnership, and the first EURAD-2 annual event DM session feedback was explicit that data management cannot be an unfunded additional task. Second, classification schemes designed by one generation are routinely unusable by the next, whereas a complete corpus with accurate provenance metadata remains interpretable indefinitely. Third, the retrieval technology that was speculative when Deliverable 11.1 was written is now available, and it operates on document content rather than on an imposed hierarchy.

This does not license the absence of metadata. Provenance metadata, who produced the document, within which work package and project, when, at what version, under what licence and dissemination level, is mandatory, because it is what makes a document citable, auditable and preservable. Section 5.3 defines the minimum set, and it is deliberately short enough to be applied without specialist effort. The distinction is between describing a document, which is cheap and durable, and classifying it into a subject taxonomy, which is expensive and perishable.

The practical consequence is that the first release presents a flat, exhaustive and well described corpus rather than a structured knowledge tree, and that discovery is delivered by faceted search over the metadata in Release 1 and by the retrieval layer in Release 2.

3.5 Objections considered

The approach specified here is open to reasonable objection, and the objections are addressed directly rather than left implicit.

Objection	Response
The survey conducted for Deliverable 11.1 recorded an explicit rejection of simple repository approaches. This specification proposes a repository.	What is proposed is a curated public knowledge section with a complete deliverable corpus, persistent identification, faceted search and an artificial intelligence retrieval layer. That is materially different from the ad hoc wiki that the survey rejected. It is also worth recalling the limitation of that evidence, eleven organisations, a survey whose raw material is no longer available, and a preference expressed before any implementation cost was known. The 2025 feedback, gathered from a larger and more operational audience, asked first for things to be findable and preserved.
Reducing the first release below the EURAD-1 blueprint is a loss of ambition.	The blueprint has produced no operational capability in six years. Ambition that is not delivered is not ambition. Every element of the blueprint is retained in the backlog in Section 5.7 with its origin recorded, and the modular architecture is designed to accept them. The change is to the order of construction, not to the destination.
Hosting on Commission infrastructure reduces the freedom of the partnership to shape the platform.	It does impose constraints, on the content management system, on publication workflow and on corporate design. Those constraints are accepted knowingly, and they are the price of durability, of corporate compliance obligations being met by default, and of the removal of a recurring maintenance liability. Section 5.5 and Section 6.6 specify the portability guarantees that ensure the corpus is never captive to the hosting arrangement.
An artificial intelligence layer over an unclassified corpus will produce unreliable answers, which is dangerous in a safety related domain.	This is the most serious objection and it drives the requirements in Section 5.4. The retrieval layer is restricted to the EURAD corpus, is required to cite its sources at document and location level for every assertion, is required to decline rather than to speculate when the corpus does not support an answer, is evaluated against an expert validated question set before public release and never removes the ability to reach the underlying document. It is specified as an instrument for finding and attributing knowledge, not as an authority on it.
Making the platform dependent on a hosting discussion with the Joint Research Centre creates a blocking dependency.	The dependency is real and is carried in the risk register as R-01. It is mitigated structurally. Release 0, the consolidation of the corpus and the application of the metadata set, is where most of the durable value lies and it is entirely independent of the hosting decision. The presentation layer is specified so that it can be implemented on an alternative infrastructure without any change to the corpus, and Section 6.3 defines the fallback.
Decommissioning and waste management are different communities with different needs, so a shared infrastructure will fit neither.	Shared infrastructure is not shared editorial control. The specification calls for a twin section with its own identity, its own corpus and its own curator, hosted on the same platform and interoperable at the data level. The elements being shared, publication infrastructure, retrieval tooling, metadata conventions and persistent identification, are precisely the elements that are domain independent.

Table 3. Objections and responses.

4. Target concept

4.1 The reference model

The EU Nuclear Decommissioning Knowledge Management section, operated by the Joint Research Centre within the Commission web presence, is used in this deliverable as the reference model. It is not cited as an inspiration but as a template, because it demonstrates, in production and on the infrastructure being requested, that the pattern works.

The section has a simple and stable structure. A landing page explains the purpose of the initiative and its two objectives, knowledge preservation and exchange, and the development of ties between the actors. A small number of activity blocks lead to the substantive content, a knowledge products database holding the outputs produced by the operators involved, a training programme, and a recurring school. Events are listed. A named functional mailbox provides a single contact point. Related links connect the section outwards, to the Euratom research and training programme, to the Commission publication repository and data catalogue, and to the corresponding activities of the International Atomic Energy Agency and the OECD Nuclear Energy Agency. The page carries a visible last update date.

Three properties of this model are worth naming explicitly, because they are the ones this specification adopts. It is small, and therefore maintainable by a curator rather than by a development team. It leads with the corpus, since the knowledge products database is the reason the section exists and everything else is context. And it is embedded, so it inherits the corporate infrastructure, its accessibility and language handling, its archiving and its permanence, without operating any of it.

4.2 A twin as first step

The EURAD knowledge management section is specified as a twin of that model. Twin has a precise meaning in this document.

- The same host infrastructure and the same corporate content management environment, so that the EURAD section is a first-class part of the Commission web presence rather than an external site linked from it.
- The same structural pattern, a landing page, a corpus database as the principal object, entry points to training and events, and a named contact channel.
- The same descriptive and persistent identification conventions, so that a record in one corpus and a record in the other can be processed by the same tools.
- A distinct identity, a distinct corpus and a distinct editorial owner. The two sections are siblings, not a merger. EURAD retains full editorial responsibility for its own content.
- Reciprocal links at section level from the outset, and interoperability at record level as specified in Section 5.6.

4.3 Structure of the first release

The first release consists of the following pages and objects.

Object	Content and function
Landing web page, EURAD knowledge management / website	Purpose of the EURAD knowledge management activity, its objectives, its position within the European partnership on radioactive waste management, the funding statement, and navigation to the activity blocks below. Carries a visible last update date and a named contact.
EURAD knowledge products database	The principal object of the section. A complete, browsable and searchable listing of EURAD public deliverables and equivalent knowledge products, each record carrying the minimum metadata set of Section 5.3, a persistent identifier and a direct download of the document in an open format. Filterable by project generation, work package, year, document type and dissemination level. This is the corpus.
Training and capacity building	Entry point describing EURAD training activity and linking to the EURAD School and to the relevant Commission and international training initiatives, including the decommissioning learning initiatives. No training content is migrated in the first release.
Events	Listing of EURAD annual events, workshops and schools, using the standard event object of the host environment.
Contact	A single functional mailbox for the EURAD knowledge management activity, with a named responsible role behind it, together with the postal reference of the hosting service.
Related links	Outward links to the EURAD programme site, to the Euratom research and training programme, to the open archive holding EURAD datasets, to the Commission publication repository and data catalogue, to the decommissioning knowledge management section, and to the corresponding IAEA and OECD/NEA activities.

Table 4. Objects of the first release.

The corpus is presented flat. There is no thematic tree, no goals breakdown navigation and no curated knowledge hierarchy in the first release, for the reasons developed in Section 3.4. Discovery is provided by metadata facets in Release 1 and by the retrieval layer in Release 2.

4.4 The knowledge extraction layer

The second release adds an artificial intelligence layer whose function is to make the corpus answer questions. The intended user is a professional who knows the domain but does not know the corpus, and whose question is of the form what does EURAD say about a given subject, rather than which deliverable should I open.

The layer is specified as retrieval over the EURAD corpus, in which the system identifies the passages of the corpus relevant to a question and produces an answer grounded in them, accompanied by the identity and location of every source used. It is not specified as a general purpose assistant, it does not answer from outside the corpus, and it does not replace the documents. The full requirements are in Section 5.4 and they are deliberately demanding, because a retrieval tool that is confidently wrong in a safety related domain destroys the credibility of the whole platform and of the partnership that operates it.

Two implementation routes are open and the choice is deferred to the point where the corpus exists, which is the only point at which it can be made on evidence. The first is to adopt or share the semantic and retrieval tooling being developed within the Commission for the decommissioning knowledge management activity, which is the route consistent with Section 3.2 and is the preferred one. The second

is a self-contained component procured or assembled from open-source parts and operated over the same corpus. Section 5.4 is written to be neutral between them, and Section 5.6 ensures that the corpus can be consumed by either.

4.5 Position within the existing EURAD landscape

The specification adds one object to the EURAD digital landscape and is intended to reduce the total number over time. The table below states the intended role of each element, which also answers the concern recorded at first EURAD-2 annual event DM session about the proliferation of portals.

Element	Role	Effect of this specification
EURAD programme website	Programme communication, project description, partner information, news and public outreach	Retained and unchanged during EURAD-2. Progressively relieved of the archival function, which moves to the knowledge management section. Its long term hosting is addressed by the same route.
Collaborative workspace (ProjectPlace)	Restricted work in progress collaboration within work packages	Unchanged and out of scope. It remains the place where work is done, and the knowledge management section is where validated output is published.
EURAD School	Training delivery	Linked from the training block in Release 1. Candidate for consolidation in a later module, subject to a separate assessment.
Wiki and dispersed initiatives	Assorted thematic content	Assessed case by case during Release 0. Each is either absorbed into the corpus, linked, or retired with its content preserved. No dispersed initiative is left unaddressed.
Open archive for datasets	Long term open archiving of datasets and publications, and minting of persistent identifiers	Used as the archival back end for data. The knowledge management section links to it and does not duplicate it.
Decommissioning knowledge management section	Commission knowledge management activity for nuclear decommissioning	Sibling section. Reciprocal links in Release 1, record level interoperability in Release 3.
EURAD knowledge management section	Permanent public reference point for validated EURAD outputs, plus knowledge extraction	New. The only element added by this specification.

Table 5. Position of the platform within the existing landscape.

5. Specifications

Requirements are expressed as numbered statements. The obligation column uses three levels, essential means that the release is not acceptable without it, expected means that it is required unless a documented justification is accepted by the work package, and optional means that it is desirable and may be traded against effort. The release column states the release in which the requirement applies.

5.1 Architecture

The platform is specified in three separable layers. The separation is the mechanism that delivers the modularity argued for in Section 3.3, and it is itself a requirement.

1. Corpus layer: the documents, their files in open formats, their metadata records and their persistent identifiers. This layer is the asset. It is specified so that it can be exported in full, understood without the software that presents it, and migrated to any future system.
2. Presentation layer: the public web section that renders the corpus, provides browsing and faceted search, and carries the contextual content. This layer is provided by the host content management environment and is expected to be replaced at least once during the life of the corpus.
3. Retrieval layer: the knowledge extraction service that consumes the corpus and answers natural language questions with sourced results. This layer is expected to be replaced most frequently of the three, since the underlying technology is evolving rapidly, and it is therefore specified as strictly detachable.

The governing rule is that no layer above the corpus may hold information that is not derivable from the corpus. If the presentation layer and the retrieval layer are both removed, nothing of value is lost. This rule is what makes the technology choices of 2026 reversible in 2036.

5.2 Functional requirements

ID	Requirement	Obligation	Release
F-01	The platform provides a single public landing page presenting the EURAD knowledge management activity, its purpose and its contact point, within the host web presence.	Essential	R1
F-02	The platform holds a database of EURAD knowledge products in which every public deliverable of EURAD-1 and EURAD-2 is represented by one record.	Essential	R1
F-03	Each record carries the minimum metadata set defined in Section 5.3, complete and validated before publication.	Essential	R1
F-04	Each record provides direct download of the document without registration, authentication or acceptance of terms.	Essential	R1
F-05	Each record carries a persistent identifier that resolves to the record and does not change if the presentation layer changes.	Essential	R1
F-06	The corpus is browsable as a complete list and filterable by project generation, work package, publication year, document type and dissemination level.	Essential	R1
F-07	The corpus supports full text search across document content and metadata, with results ordered by relevance and filterable by the facets of F-06.	Essential	R1
F-08	Where a deliverable has successive versions, all versions remain retrievable and the current version is unambiguously identified.	Expected	R1
F-09	The platform provides an entry point to EURAD training activity and links to related European and international training initiatives. (optional because these activities stop when EURAD stops).	Optional	R1
F-10	The platform lists EURAD events using the standard event object of the host environment (optional because these activities stop when EURAD stops).	Optional	R1
F-11	The platform links to the corresponding decommissioning knowledge management section and to the relevant IAEA and OECD/NEA activities.	Essential	R1
F-12	Every published page carries a visible date of last update and identifies the author.	Essential	R1
F-13	The corpus is exportable in full, metadata and files together, in a documented open format, by an operation that requires no proprietary tool.	Essential	R1
F-14	The platform provides a natural language query interface returning answers grounded in the corpus, subject to Section 5.4.	Essential	R2
F-15	The platform accepts a submission route by which a work package can propose a knowledge product for inclusion, with curator validation before publication.	Expected	R2
F-16	The corpus is extended to hold selected legacy material from EURAD-1, PREDIS and earlier European projects on radioactive waste management.	Expected	R3
F-17	The corpus exposes a machine-readable interface allowing an external system to harvest records and metadata.	Expected	R3

Table 6. Functional requirements.

5.3 Minimum metadata set

The metadata set below is mandatory for every record. It is deliberately short. Every field is either necessary to cite the document, necessary to preserve it, or necessary to filter the corpus, and no field requires subject matter judgement to complete, which is what makes it applicable without expert effort. Fields marked as controlled take their values from a fixed list maintained by the curator.

Field	Description	Controlled
Title	Full title of the document as published	No
Identifier	Persistent identifier of the record	No
Document reference	EURAD reference, for example the deliverable or milestone number	No
Authors	Named authors, with organisational affiliation	No
Producing organisations	Organisations responsible for the content	Yes
Programme	EURAD-1, EURAD-2, PREDIS or the named predecessor project (Grant Agreements identifiers)	Yes
Work package	Work package reference and title	Yes
Document type	Deliverable, milestone, state of the art, guidance, dataset description, presentation, other	Yes
Date of issue	Date of the published version, ISO 8601	No
Version	Version designation, and status as current or superseded	No
Dissemination level	Public, or the level under which it was cleared for publication	Yes
Licence	Reuse licence applying to the document	Yes
Language	Language of the document	Yes
Abstract	Abstract as published, or the executive summary if none exists	No
Keywords	Free keywords as supplied by the authors	No
Related resources	Identifiers of associated datasets, publications, superseded versions or successor documents	No
Format	File format and, for preservation, the open format in which the content is held	Yes

Table 7. Minimum metadata set.

Two points of principle apply. Keywords are captured as supplied by the authors and are not normalised into a taxonomy, consistent with Section 3.4. Related resources is the field that carries the reproducibility ambition recorded as N-16, since it is the mechanism by which a deliverable is linked to the dataset that supports it once datasets enter the corpus in Release 3.

5.4 Requirements on the knowledge extraction layer

These requirements govern the artificial intelligence component. They are more demanding than those applied to the rest of the platform, and deliberately so. The community asked for this capability with some insistence, and it is the element most likely to determine whether the platform is trusted or discredited.

ID	Requirement	Obligation	Release
AI-01	Answers are generated exclusively from the EURAD corpus and from material explicitly declared as part of the retrieval scope. The service does not answer from general knowledge held in an underlying model.	Essential	R2
AI-02	Every substantive assertion in an answer is accompanied by a reference to its source at document level and at a location within the document sufficient to verify it, such as a page or section. This requires to have an unique source of information.	Essential	R2
AI-03	Every cited source is directly reachable from the answer, so that the user can open the document and read the passage in context.	Essential	R2
AI-04	Where the corpus does not support an answer, the service states that the corpus does not contain the information and does not offer an unsourced answer.	Essential	R2
AI-05	The service presents an explicit and visible statement that answers are generated automatically, that they are not validated EURAD positions, and that the source documents prevail.	Essential	R2
AI-06	Before public release, the service is evaluated against a set of questions with expert validated reference answers, spanning the work packages, and the evaluation report including the observed rates of unsupported and incomplete answers is published with the release.	Essential	R2
AI-07	The evaluation is repeated at every material change to the model, the retrieval configuration or the corpus, and the results are retained.	Essential	R2
AI-08	The version of the model, of the retrieval configuration and of the corpus index used to produce an answer is recorded and can be reported, so that an answer can be reproduced or explained after the fact.	Essential	R2
AI-09	The corpus is not transferred to any third party for the training of general models, and any contractual arrangement with a supplier states this explicitly.	Essential	R2
AI-10	The service is detachable. Removing it leaves the corpus, the browsing and the faceted search of Release 1 fully operational. The retrieval layer holds no content, no metadata and no canonical text that is not present in the corpus layer; everything it holds is derived and rebuildable from a corpus export.	Essential	R2
AI-11	Queries and answers are logged in a form that supports quality improvement, subject to the personal data protection rules of the host, and the log of queries is used as the primary evidence base for prioritising later modules.	Expected	R2
AI-12	The service accepts queries in the official languages of the European Union and answers in the language of the query, while citing sources in their original language.	Optional	R2
AI-13	A user feedback mechanism allows an answer to be reported as incorrect, and reported answers are reviewed by the curator.	Expected	R2
AI-14	The retrieval scope can be extended to the decommissioning knowledge products corpus, subject to agreement with its owner.	Optional	R3

Table 8. Requirements on the knowledge extraction layer.

Requirement AI-06 deserves comment, since it is the one that carries a real cost. Constructing an expert validated question set requires time from work package experts. That cost is accepted because it is the only means of knowing whether the service works, and because the question set becomes a permanent asset, it is reused at every subsequent evaluation and it constitutes, in itself, a structured record of what the community considers the important questions to be.

5.5 Non functional requirements

ID	Requirement	Obligation	Release
NF-01	The platform is hosted on infrastructure whose operator provides a maintenance and preservation commitment that is not bounded by the duration of EURAD-2.	Essential	R1
NF-02	All corpus content is held in open, non proprietary, documented formats. Where a source document exists only in a proprietary format, an open format rendition is produced and both are retained.	Essential	R0
NF-03	The corpus and its metadata can be exported in full at any time, and the export is sufficient to reconstitute the corpus on a different platform without loss of information.	Essential	R1
NF-04	The public presentation complies with the accessibility standard applicable to the host environment.	Essential	R1
NF-05	Published pages and records are archived by the web archiving arrangements of the host, so that superseded states remain consultable.	Expected	R1
NF-06	The platform is operable and maintainable by a curator without development skills for all routine editorial tasks, including the publication of a new record.	Essential	R1
NF-07	The platform introduces no licence cost per user and no dependency on a proprietary component that would prevent migration.	Essential	R1
NF-08	Availability and performance are those of the host environment, and no separate service level is contracted by the partnership.	Expected	R1
NF-09	Personal data appearing in records, principally author names and affiliations, is processed under the data protection regime of the host, and the applicable privacy statement is linked from the section.	Essential	R1
NF-10	Every component developed or configured for the platform, together with its configuration and its metadata schema, is documented to a level sufficient for a third party to reproduce it.	Essential	R1
NF-11	The metadata schema is published and versioned, and changes to it are backward compatible or accompanied by a documented migration.	Essential	R1

Table 9. Non functional requirements.

5.6 Interoperability

ID	Requirement	Obligation	Release
IN-01	The metadata schema is expressed as a mapping to a widely used general purpose metadata standard, so that records can be understood outside EURAD.	Essential	R1
IN-02	Records referring to datasets held in an external open archive link to them by persistent identifier rather than duplicating them.	Essential	R1
IN-03	The section is reciprocally linked with the decommissioning knowledge management section.	Essential	R1
IN-04	The corpus is exposed through a documented harvesting interface allowing an external service to retrieve records and metadata systematically.	Expected	R3
IN-05	Alignment of metadata conventions with the decommissioning knowledge products database is pursued so that both corpora can be queried by a common retrieval service.	Expected	R3
IN-06	Where a EURAD output is also held in the Commission publication repository or data catalogue, the record links to it rather than creating a competing copy.	Expected	R1

Table 10. Interoperability requirements.

5.7 Explicit exclusions and backlog

The following capabilities were specified during EURAD-1 or requested at first EURAD-2 annual event DM session and are not part of the first release. They are recorded here with their origin and the condition under which they are reconsidered. This register is the mechanism by which restraint remains compatible with the retention of ambition, and it is maintained by the work package as a live document.

Capability	Origin	Condition for reconsideration
Collaborative workspaces, forums, communities of practice tooling	D11.1, D11.9 functional specification 3	After Release 2 is adopted, and only if the existing collaborative workspace is demonstrated to be insufficient. Duplication of an existing tool is not acceptable.
Granular access control, user roles, restricted areas	D11.1, D11.9 non-functional specification 4	When a defined and material body of restricted content requires publication outside the collaborative workspace.
Semantic ontology, controlled subject taxonomy, goals breakdown navigation	D11.1, D11.9	When query log evidence from Release 2 demonstrates a retrieval failure that metadata facets and the retrieval layer cannot address.
Directory of experts with expertise keywords	First EURAD-2 annual event DM session «Create»	Release 3, subject to consent and data protection assessment, which is the principal obstacle rather than the technology.
Reproducibility packages linking data, models and measurement sources	First EURAD-2 annual event DM session «Create»	Release 3, driven by a work package pilot as recommended in Milestone 2.2.
Migration of the EURAD School and of the wiki into the section	MS2.2	After Release 1 is stable, subject to a separate assessment per tool.
Project management and reporting functions	D11.9 boundary issue	Not reconsidered. Milestone 2.2 identified the blurring of the boundary between a knowledge management platform and a project management tool as a defect to be corrected.
Automated classification and tagging by artificial intelligence	First EURAD-2 annual event DM session «Create»	Release 3, and only as a curator assistance function with human validation, never as an unsupervised process.

Table 11. Excluded capabilities and backlog.

6. Hosting and governance

6.1 The hosting requirement

Hosting is treated in this specification as a requirement of the same order as the functional requirements, and not as an implementation detail to be settled afterwards. Requirement NF-01 states it, the operator of the infrastructure must provide a maintenance and preservation commitment that is not bounded by the duration of EURAD-2. Everything else in the specification is conditional on it, since a corpus that disappears when a grant ends does not solve the lost legacy problem, whatever its functional richness.

6.2 Options assessed

Three options were assessed against five criteria, durability beyond the partnership, absence of a recurring cost falling on the partnership, compliance with corporate obligations by default, avoidance of duplication with existing European Commission activity, and time to a first operational release.

Option	Assessment
A. Hosting within the European Commission web presence, in association with the Joint Research Centre knowledge management activity	Durability derives from the institution that funds and owns the mission rather than from a project. No recurring hosting cost falls on the partnership. Corporate obligations on accessibility, languages, security, data protection and archiving are met by the environment. Duplication with the existing decommissioning activity is avoided and mutualisation becomes possible. The constraints are real, the content management environment, the publication workflow and the corporate design are imposed, and the schedule depends on an agreement that is not yet concluded.
B. Hosting by an EURAD beneficiary	Fast to arrange and fully under the control of the partnership. It fails the primary criterion, since no beneficiary can commit beyond its own funded projects, a point established in Milestone 2.2 and confirmed by the community at first EURAD-2 annual event DM session. It reproduces the current situation of the existing tools, which is the situation this deliverable exists to correct.
C. Commercial hosting under a tender	Capable of delivering any specified functionality and of meeting a contractual service level. It converts durability into a recurring financial commitment that must be renewed at every programme boundary, which is the commitment the partnership has already declined to make. It carries a supplier dependency and, on the evidence of the EURAD-1 experience, the highest risk of producing a system that is delivered and then not maintained.

Table 12. Hosting options.

6.3 Recommendation and fallback

Option A is definitely recommended because it is the only option that satisfies the durability requirement without a recurring financial commitment from the partnership, and it is the only one consistent with the mutualisation argument of Section 3.2. The constraints it imposes on presentation and workflow are accepted knowingly, and the portability guarantees of Section 6.6 ensure that they are not irreversible.

The dependency created by this recommendation is managed rather than ignored. Discussions between Task 5 and the Joint Research Centre have been held between February and June 2026 and are continuing. They have established two working conclusions that shape this specification. The first is that a transfer of existing EURAD websites as such into the Commission web space is not the realistic route, and that the practical route is a dedicated EURAD entry point within the Commission web presence. The specification in Section 4 is written to that conclusion. The second is that the objective on the Commission side is to connect existing repositories and make them interoperable rather than to duplicate them, which is the same objective as the one stated here. Development work on a semantic tooling prototype is under way on the Commission side, and EURAD material has been shared to support its testing. No commitment has been concluded at the date of this deliverable, and the arrangements, including any financial aspects, remain to be settled.

The fallback is structural rather than contractual. Release 0, the consolidation of the corpus, the production of open format renditions and the application of the minimum metadata set, carries most of the durable value of the whole programme of work and is entirely independent of the hosting decision. If Option A cannot be concluded within the schedule of Section 7.2, Release 1 is implemented on an interim infrastructure using the same corpus, the same metadata schema and the same persistent identifiers, and the transfer to Option A is executed later as a migration of the corpus rather than as a

redevelopment. The cost of that scenario is a delay and a duplicated presentation effort, not the loss of the asset.

6.4 Roles and responsibilities

Role	Responsibility
Infrastructure custodian	Operates the hosting environment, applies corporate policies on accessibility, security, data protection and archiving, and guarantees continuity beyond the partnership.
Content curator	A named role within Work Package 2. Responsible for the completeness and accuracy of the corpus, for the validation of metadata before publication, for the currency of the contextual content, and for responding through the contact channel. This role must be resourced, which is the single most frequently repeated condition in the lessons learned of Deliverable 11.1 and in the first EURAD-2 annual event DM session feedback.
Work Package 2, Task 5	Owns the specification, prioritises the backlog, prepares release decisions and reports to the partnership governance.
IT Experts Group	Provides technical guidance, validates the metadata schema and the interface conventions, and ensures that developments in individual work packages remain compatible with the general interface. Its activation was recommended in EURAD-1 and is a prerequisite of Release 0.
KM Ambassadors	Act as the channel between work packages and the curator, identify the knowledge products of their work package, and supply the question set required by AI-06.
Programme governance	Endorses the strategic answers of Section 2.5, approves release transitions against the acceptance criteria of Section 7.3, and secures the resources for the curator role.

Table 13. Roles and responsibilities.

6.5 Continuity beyond EURAD-2

The arrangement to be applied at the end of the partnership is specified now rather than at the end, because the failure documented in Section 2.3 occurs precisely at project boundaries and because an arrangement negotiated under time pressure is the one that fails.

- The corpus remains published and reachable at its persistent identifiers after the end of EURAD-2, under the responsibility of the infrastructure custodian.
- Editorial responsibility transfers to the successor knowledge management activity of the European partnership, or, in the absence of one, to the custodian, which continues to publish the corpus as a closed and dated collection.
- A closing operation is carried out in the final year of the partnership, in which the corpus is completed, verified against the metadata schema, exported in full under NF-03 and deposited as a dated snapshot in an independent open archive (Zenodo).
- The documentation required by NF-10 and the published metadata schema are deposited with the same snapshot, so that the corpus remains reconstructible independently of any surviving software.

6.6 Portability guarantees

The following guarantees are the answer to the objection that hosting on external infrastructure creates a dependency, and they are binding requirements rather than intentions.

1. The corpus is held in open documented formats, NF-02, and can be exported in full at any time in a form sufficient to reconstitute it elsewhere, NF-03.
2. The metadata schema is published, versioned and mapped to a general purpose standard, NF-11 and IN-01, so that it is intelligible outside its original context.
3. No layer above the corpus holds information that is not derivable from the corpus, Section 5.1, so that the loss of the presentation or retrieval layer is not a loss of knowledge.

4. Persistent identifiers resolve independently of the presentation layer, F-05, so that citations of EURAD material survive a change of platform.
5. A complete export is produced and verified at each release transition, not only at the end, so that the export path is known to work rather than assumed to.

7. Implementation

7.1 Releases

Release	Content	Purpose
R0, foundation	Inventory of all EURAD public outputs across EURAD-1 and EURAD-2. Case by case assessment of the wiki and dispersed initiatives. Production of open format renditions. Application of the minimum metadata set. Publication and versioning of the metadata schema. Activation of the IT Experts Group. Appointment and resourcing of the curator.	Creates the asset. This release delivers most of the durable value and is independent of the hosting decision, which is why it is executed first.
R1, public section	The public knowledge management section as specified in Section 4.3, with the knowledge products database, faceted search, persistent identifiers, training and events entry points, contact channel and reciprocal links.	Makes the corpus public, findable and citable. Solves the fragmented workflow problem for validated outputs and, together with NF-01, the lost legacy problem.
R2, knowledge extraction	The retrieval layer as specified in Section 5.4, released first to a restricted pilot group drawn from the KM Ambassadors and possible end-users, then publicly once AI-06 is satisfied. Submission route for new knowledge products.	Delivers the capability the community asked for and, through the query log, produces the first evidence based statement of what users actually need.
R3, modular extension	Legacy recovery from EURAD-1, PREDIS and earlier European projects. Dataset records and links to the open archive. Harvesting interface. Interoperability with the decommissioning corpus. Backlog items promoted on the evidence of R2.	Begins the progressive realisation of the wider vision, on a stable foundation and with real usage evidence.

Table 14. Release plan.

7.2 Indicative schedule

Period	Activity	Responsible
Q4 2026	Endorsement by governance of the strategic answers of Section 2.5 and of the hosting recommendation. Appointment and resourcing of the curator. Activation of the IT Experts Group.	Programme governance, WP2 Task 5
Q4 2026 to Q1 2027	R0. Corpus inventory, open format renditions, metadata schema published, metadata applied. Assessment of the dispersed initiatives.	WP2 Task 5, curator, KM Ambassadors
Q1 2027	Conclusion of the hosting arrangement, or activation of the interim fallback of Section 6.3.	WP2 Task 5, programme governance
Q2 to Q4 2027	R1. Implementation and publication of the public section. Acceptance against Section 7.3.	WP2 Task 5, curator, infrastructure custodian
Q4 2027 to Q1 2028	Construction of the expert validated question set required by AI-06.	KM Ambassadors, work package experts
2028	R2. Retrieval layer, restricted pilot, evaluation, then public release. Communication campaign.	WP2 Task 5, curator
From 2029	R3. Modular extension, legacy recovery, interoperability. Closing operation of Section 6.5 in the final year.	WP2 Task 5, successor activity

Table 15. Indicative schedule. Dates are indicative and are subject to the endorsement and hosting decisions identified above.

7.3 Acceptance criteria

A release is not considered delivered until its criteria are met, and a subsequent release is not started until the preceding one is accepted. This rule is the operational expression of Section 3.3 and it is the principal safeguard against the accumulation of unfinished work that characterised the previous attempts.

Release	Acceptance criteria
R0	The inventory is complete and reconciled against the deliverable and milestone lists of both programme generations. Every entry has an open format rendition and a complete metadata record validated against the published schema. The schema is published and versioned. Every dispersed initiative has a documented decision. The curator is appointed and resourced.
R1	Every public deliverable of both programme generations is published as a record with a resolving persistent identifier and a direct download. Faceted browsing and full text search operate across the whole corpus. Reciprocal links with the decommissioning section are live. A full export under NF-03 has been produced and verified by reconstituting the corpus in a test environment. The section is reachable from the EURAD programme site.
R2	The evaluation required by AI-06 has been completed, published and reviewed by the work package. No answer is produced without a resolving source citation. Behaviour under AI-04 is verified on a set of questions known to be outside the corpus. Detachability under AI-10 has been demonstrated by disabling the service and confirming that Release 1 functionality is unaffected. The feedback mechanism of AI-13 is operational. Accessibility conformance and machin readability checks should be also performed.
R3	Criteria are defined at the point of scoping, on the basis of the query log evidence from Release 2 and of the backlog priorities of Section 5.7.

Table 16. Acceptance criteria.

7.4 Resources

No costed budget is proposed in this deliverable, since the hosting arrangement that determines the largest part of the cost is not concluded. The shape of the resource requirement is nevertheless clear and is stated here so that it can be planned.

- The dominant and recurring cost is human, the curator role. It is a fraction of a full time position, it is continuous rather than one off, and the historical record is unambiguous that a knowledge management system without a resourced curator is abandoned. This is the one line of the resource plan that cannot be traded away.
- Release 0 is an effort cost concentrated in a single period, distributed between the curator and the KM Ambassadors, and it is the largest one off item.
- Release 1 is expected to be a configuration effort rather than a development effort, consistent with the recommendation already made in EURAD-1 to favour configuration over custom development, and its cost depends on the hosting arrangement.
- Release 2 carries a development or licensing cost for the retrieval layer, an operating cost that scales with usage, and the expert time required by AI-06. If the tooling being developed within the Commission for the decommissioning activity can be shared, the first two are substantially reduced, which is a further practical argument for Section 3.2.
- The expenditure declined by this specification is equally relevant. The full EURAD-1 vision was associated with an order of magnitude of six million euro for development and support. The approach specified here does not require that commitment and does not ask the partnership to make it.

7.5 Risks

ID	Risk	Mitigation
R-01	The hosting arrangement with the European Commission is not concluded, or is concluded later than the schedule assumes.	Release 0 is independent of it and proceeds regardless. The interim fallback of Section 6.3 allows Release 1 to be implemented on another infrastructure with the same corpus, schema and identifiers, so that the later transfer is a migration and not a redevelopment.
R-02	The curator role is not resourced and the corpus becomes stale within a year of publication.	The appointment is an acceptance criterion of Release 0 and a prerequisite of Release 1. Governance endorsement in Q4 2026 is explicitly framed as a resourcing decision, not only a technical one.
R-03	The retrieval layer produces confident but unsupported answers and damages the credibility of the platform.	Requirements AI-01 to AI-08, the restricted pilot before public release, the published evaluation, and the detachability requirement AI-10 which allows the service to be withdrawn without loss.
R-04	Scope creep during Release 1 reproduces the pattern of the previous attempts.	The exclusion register of Section 5.7 is the formal instrument for refusing additions. Every proposed addition is recorded there with its condition for reconsideration rather than being debated on its merits at the moment it is raised.
R-05	Work packages do not supply their knowledge products and the corpus remains incomplete.	The inventory of Release 0 is built from the formal deliverable and milestone lists rather than from voluntary submission, so completeness does not depend on solicitation. The KM Ambassadors channel is used for anything outside those lists.
R-06	Clearance, licensing or intellectual property issues prevent the publication of some material.	Release 1 is restricted to material already public or explicitly cleared, which removes the issue from the critical path. Material that cannot be cleared is recorded in the inventory with its status, so that its existence is known even when its content is not published.
R-07	The constraints of the corporate content management environment prevent some specified functionality.	The reference model demonstrates that the specified pattern is achievable in that environment. Where a constraint bites, the functional requirement is relaxed rather than the corpus requirements, since the corpus is the asset and the presentation is replaceable.
R-08	The platform is delivered but not used, repeating the shortcomings recorded in the EURAD-1 lessons learned.	The section is linked from the EURAD programme site and from the decommissioning section, a communication campaign accompanies Release 2, and the query log provides an early and objective measure of adoption rather than an assumed one.

Table 17. Risk register.

8. Conclusion

EURAD holds a large and valuable body of knowledge, and it holds it in a way that does not guarantee its survival, but more particularly long-term accessibility. Two programme generations of specification work have described, with increasing precision, a platform that would solve the problem, and no such platform exists. Milestone 2.2 established that the obstacle was never technical, and this deliverable is written on that basis.

What is specified here is deliberately modest in its first release, a public knowledge management section hosted on European Commission infrastructure and built as a twin of the decommissioning knowledge management activity, holding the complete corpus of EURAD public outputs as a flat and well described collection, and equipped in a second step with a retrieval layer that answers questions from that corpus with traceable sources. It does not attempt collaboration, access control, taxonomies or project management, and it records each of those as a backlog item with the condition under which it returns.

The reasoning that supports this restraint is the substance of the deliverable rather than a preamble to it. Durability requires that the custodian of the knowledge outlast the projects that create it, which points to the institution that funds the partnership and not to any of its beneficiaries. Mutualisation requires that the European Union operate one knowledge management effort for its nuclear legacy rather than two, which points to the infrastructure and the tooling already in place for decommissioning. Incrementality requires that the first step be small enough to be taken, because the alternative has been tried and has produced documentation instead of a platform. Modularity is what reconciles that restraint with the ambition expressed in EURAD-1, since a corpus held in open formats with a published schema can carry any number of later modules and can survive all of them.

Some new aspects will have to be implemented in the potential future governance of possible future partnerships, endorsement of the strategic answers proposed in Section 2.5, endorsement of the hosting recommendation of Section 6.3, and the allocation of resources to the curator role. The third is the least visible and the most important. The record of comparable initiatives, in EURAD and elsewhere, is consistent on the point that knowledge management systems are not abandoned because they were badly built, they are abandoned because nobody was paid to keep them alive.

The work described in Release 0 can begin as soon as those decisions are taken, and its value does not depend on any of the choices that remain open.

Appendix A. Traceability of requirements to needs

The table maps every need of Table 1 to the requirements that satisfy it, so that the specification can be audited against the documented demand of the community, and so that no requirement exists without a recorded justification.

Need	Requirements	Release
N-01, long term hosting and preservation	NF-01, NF-02, NF-03, NF-05, Section 6.5	R0, R1
N-02, single public entry point	F-01, F-02, F-11, IN-03	R1
N-03, improved discovery	F-06, F-07, F-03, Table 7	R1
N-04, reliable knowledge extraction	F-14, AI-01 to AI-14	R2
N-05, persistent citable identification	F-05, IN-02, IN-06	R1, R3
N-06, legacy preservation and recovery	F-16, IN-02	R3
N-07, open formats and portability	NF-02, NF-03, NF-07, NF-11, F-13, Section 6.6	R0, R1
N-08, granular access control	Deferred, Table 11	Backlog
N-09, collaboration features	Deferred, Table 11	Backlog
N-10, partner access to the collaborative workspace	Out of scope, handled by WP2 and the programme management office	Not applicable
N-11, Data Management Plan guidance and training	Out of scope, handled by other WP2 tasks	Not applicable
N-12, harmonisation and non duplication	F-11, IN-01, IN-03, IN-05, IN-06, Section 3.2	All
N-13, no new unconnected portals	Section 3.2, Table 5, IN-03	All
N-14, named responsibilities and resourced curation	NF-06, Table 13, acceptance criteria of R0	R0, R1
N-15, expert directory	Deferred, Table 11	Backlog
N-16, reproducibility oriented preservation	Related resources field of Table 7, F-16, deferred implementation	R3

Table 18. Traceability matrix.

Appendix B. Abbreviations

Abbreviation	Meaning
DMP	Data Management Plan
EC	European Commission
EURAD	European Joint Programme, then European Partnership, on Radioactive Waste Management
FAIR	Findable, Accessible, Interoperable, Reusable
IAEA	International Atomic Energy Agency
IPR	Intellectual Property Rights
JRC	Joint Research Centre of the European Commission
KM	Knowledge Management
NDKM	Nuclear Decommissioning Knowledge Management
OECD/NEA	Nuclear Energy Agency of the Organisation for Economic Co-operation and Development
p-KMS	Portal Knowledge Management System, the platform concept specified in EURAD-1
RWM	Radioactive Waste Management
WP	Work Package

References

1. Abbasova D., Arnold T., Brendler V., Franzen C. (2021). Screening and review of existing and available knowledge management approaches and tools. Deliverable D11.1 of the Horizon 2020 project EURAD, EC Grant Agreement n°847593.
2. Abbasova D., Arnold T., Brendler V., Franzen C. (2023). Specification of the EURAD KM platform (p-KMS). Deliverable D11.9 of the Horizon 2020 project EURAD, dissemination level public, issued 9 October 2023.
3. EURAD (2023). Specification of the EURAD p-KMS prototype. Milestone MS242, Work Package 11, issued 27 January 2023.
4. EURAD (2024). Implementation status, results, lessons learned and improvements on Domain Insight and State of Knowledge document production. Milestones MS245 and MS333, Work Package 11, issued 29 May 2024.
5. EURAD-2 (2025). Memo about specifications of the Knowledge Management platform. Milestone MS22, Work Package 02, issued 28 November 2025.
6. EURAD-2 (2025). Consolidated output of the data management session, annual workshop, First EURAD-2 annual event DM session, September 2025.
7. European Commission, Joint Research Centre. EU Nuclear Decommissioning Knowledge Management. Web section, joint-research-centre.ec.europa.eu, projects and activities, consulted August 2026.
8. European Commission. Nuclear Decommissioning Assistance Programme. energy.ec.europa.eu, consulted August 2026.
9. Spichtinger D. Data Management Plans in Horizon 2020: what beneficiaries think and what we can learn from their experience [version 2; peer review: 2 approved, 1 approved with reservations] Open Research Europe 2022, 1:42 https://doi.org/10.12688/openreseurope.13342.2