Organisation: CERN

Project Charter

CERN Node

revised 31 March 2025

1. Project Summary

CERN's participation in the EOSC Federation is aimed at strengthening the EU Node with key high-TRL (Technology Readiness Level) services that are essential for advancing high-energy physics (HEP) research and promoting open science. These services include Zenodo for data storage, REANA for running reproducible analysis workflows, Indico for managing research events, as well as robust Data and Literature repositories. The integration of these services, along with the deployment of a Virtual Research Environment (VRE), will enhance the accessibility and usability of high-value data and research tools, not only for HEP but for a wide range of scientific disciplines.

CERN's involvement in the EOSC Federation is structured in two distinct phases. In the first phase, we will focus on deploying and interconnecting CERN-provided services on the EU Node, gathering critical insights and operational experience while testing the virtual credit model. This model will allow the management of service usage based on virtual credits allocated to users, ensuring fair access to resources. The goal is to establish a functioning, interoperable node that serves as a foundation for broader integration into EOSC.

In the second phase, we will extend the virtual credit model to CERN's digital infrastructure, allowing us to control the allocation of resources within CERN's services. This will include the creation of a CERN Node Portal where users can access a subset of services with defined quotas, enabling us to manage the capacity that can be allocated concurrently to all users. The introduction of quota limits for users will ensure resource fairness and optimize the management of the overall infrastructure within the EOSC ecosystem.

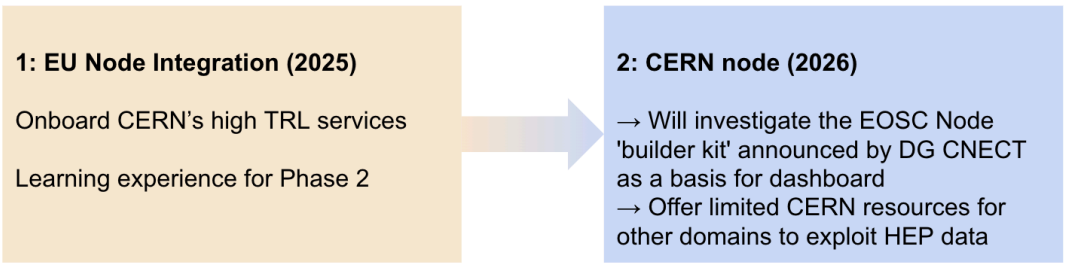

This phased approach is crucial for building the EOSC infrastructure, as it will help refine the integration of services, demonstrate the viability of the virtual credit system, and ultimately support the scalable, open, and FAIR data access that is central to the EOSC's mission. By aligning CERN's services with EOSC's objectives, we will contribute to the seamless integration of research workflows across institutions, foster collaboration, and enhance the accessibility of cutting-edge scientific data and tools for the global research community.

2. Value Proposition

- **Main Goal:**
  - CERN's participation in the EOSC Federation is structured in two phases.
  - In the first phase, we will enhance the EU Node by integrating key CERN services (such as Zenodo, REANA, Indico, ESCAPE VRE, and our data

repositories), which will also allow us to gather insights into the virtual credit system.
    - In the second phase, we will establish the CERN Node, offering controlled access to a subset of services and HEP data, enabling researchers from all scientific fields to expand their analyses in a managed environment, controlling usage through quotas.
- **Needs addressed:**
    - Ensure accessibility of HEP data: Make High Energy Physics (HEP) data easily discoverable and accessible for all EOSC users, supporting open science principles and collaboration across disciplines.
    - Provide access to high-TRL tools and services: Offer wider access to high-TRL general-purpose tools and services originally developed for HEP, such as REANA, InvenioRDM (through Zenodo), and Indico, ensuring they are reusable by other research communities while maintaining interoperability with the EU Node.
    - Support non-trivial workflows: Enable the execution of complex workflows that process large datasets, allowing researchers to run scalable and reproducible analysis directly in the EOSC ecosystem.
    - Manage scientific events and shared information: Provide solutions to manage scientific events (e.g., conferences, workshops) and ensure that all associated information is readily accessible to the research community.
    - Process curated HEP datasets without downloads: Enable users from any discipline to process and analyze curated HEP datasets within the EOSC framework, reducing the need for large data downloads and optimizing resource usage.
- **Key Benefits:**
    - Extended access to curated HEP datasets: Expands the availability of highly curated HEP datasets to a broader range of research communities, promoting interdisciplinary collaboration.
    - Facilitates analysis for all users: Provides an accessible means for a wider range of researchers, including those in the "long tail" of science and citizen scientists, to analyze and process HEP datasets.
    - Enhanced services beyond HEP: By integrating CERN services into the EOSC EU Node, and potentially other nodes like ICSC, we will offer researchers reliable access to FAIR data, collaborative tools, analysis workflows, and reproducible research environments.
    - Support for scalable, reproducible research: The deployment of high-TRL tools such as REANA and Zenodo will empower researchers to execute large-scale, reproducible workflows and data-sharing practices across disciplines.
    - Streamlined resource management: The virtual credit model will ensure efficient and equitable allocation of computational resources, improving the overall user experience and promoting sustainable usage of shared services within the EOSC ecosystem.
- **Who Benefits:**
    - Researchers and educational institutions in HEP: Those currently utilizing HEP data but lacking the necessary tools to process and analyze large datasets will benefit from the availability of CERN's highly curated datasets and associated services.
        - The CERN Open Data Portal provides access to data from multiple HEP experiments (e.g., ALICE, ATLAS, CMS, DELPHI, LHCb, OPERA, PHENIX, TOTEM), offering an invaluable resource for HEP research communities. More information at: https://opendata.cern.ch/docs/ten-years-of-cern-open-data-portal
        - INSPIRE-HEP platform users: INSPIRE-HEP, as a one-stop information platform for the HEP community, comprises 8 interlinked databases on

literature, conferences, institutions, journals, researchers, experiments, jobs, and data. This integration provides a unified access point for research publications, collaboration opportunities, and scientific knowledge. See more: https://inspirehep.net/

        - HEP-Data users:
          Researchers and institutions utilizing HEP-Data, which provides records for non-covered literature (e.g., preprints, theses) and high-level tabular data, will benefit from easier access to over 130,000 data tables and more than 10,000 HEP publications. See more: https://www.hepdata.net/
    - Research communities beyond Physics:
        - Other research fields beyond high-energy physics will benefit from the tools and services CERN developed for HEP, such as REANA, Zenodo, and Indico, which can be repurposed for data processing and collaborative workflows in other disciplines.
    - ESCAPE Virtual Research Environment (VRE) users:
        - CERN's contributions to the ESCAPE VRE will benefit all members of the ESCAPE cluster, including infrastructures like CTAO, KM3NeT, EGO-Virgo, ESO, EST, FAIR, JIV-ERIC, SKAO, and ET. This VRE will enable a diverse set of research communities to leverage the tools and infrastructure for managing large-scale, distributed scientific data and workflows. See more: https://www.escapescience.eu
- **Access policy:**
    - For services like Zenodo, CERN Open Data, INSPIREHEP, and HEPData, access policies are clearly defined by their already established terms of use. However, for services that will be onboarded to the EU Node, such as REANA, Indico, and VRE, access policies are not yet defined and will depend on the ongoing design phase with the EU node responsible parties. In Phase 2, the development of access policies will be a key task for the upcoming year, with no clear policy currently outlined.

3. Use Case(s) in descending order of priority

Use Case ID	Use Case Description	Federation Contributions & Value to Users
a: REANA	Reproducible research data analyses	REANA is a free and open source reproducible data analysis platform that allows researchers to run declarative computational workflows (using Common Workflow Language, Snakemake, or Yadage) on containerised compute clouds (using Kubernetes, HTCondor, or Slurm). The researchers can structure their analyses, submit many parameterized jobs, keep track of the results, and organise their analysis process in an actionable manner to facilitate future reuse.
b: ESCAPE VRE	The ESCAPE Virtual Research Environment (VRE) is a modular, scalable analysis platform that integrates CERN's EU onboarded services, enabling users to access data, execute workflows, and interact with research tools.	Leveraging the composability and interoperability of the VRE, use-cases requiring access to distributed data management systems, reproducibility pipelines, and public repositories will benefit from deployment in the EOSC Federation. This includes Dark Matter and Gravitational Wave searches, EOSC Science Projects, Open Data analyses and detector related R&D

The following scoping aspects apply to all the listed use cases:

In scope:

- CERN is actively involved in the development of the ScienceMesh and the Open Cloud Mesh protocol for federating storage systems, which has been adopted by the EOSC Federation. CERN will explore with other CS3 community partners how these developments can be leveraged in the context of the EOSC Federation.
- CERN will share its deployment experiences with the WLCG community and discuss how the full research lifecycle, from data acquisition through to publishing of results, can be made more fluid by potential connections between WLCG and the EOSC Federation.

Out of Scope:

- It is not the intention of CERN to establish a node with generic services that are readily available on the market. The services offered via the CERN node will support the exploitation of CERN's scientific results by the wider research and educational communities.

Each listed use-case is now described in more detail.

a. REANA: Reproducible research data analyses

- **Use Case Description:**
  REANA is a free and open source reproducible data analysis platform that allows researchers to run declarative computational workflows (using Common Workflow Language, Snakemake, or Yadage) on containerised compute clouds (using Kubernetes, HTCondor, or Slurm). The researchers can structure their analyses, submit many parameterized jobs, keep track of the results, and organise their analysis process in an actionable manner to facilitate future reuse.

  The first use case concerns the deployment of REANA at the EOSC EU Node, enabling individual researchers to deploy a personal REANA instance. The goal could be represented by offering an "Analysis Workflows" menu alongside the existing "Interactive Notebooks" menu at the EOSC EU Node, specifically for researchers who prefer declarative workflows over interactive notebooks for running their analyses. Activating this option would launch the REANA computational instance with backend storage, making it available to the researcher for running containerized workflows, transferring data in and out as needed, and shutting down the system when no longer in use. Additionally, if the EOSC EU Node Groups feature allows the creation of a shared REANA instance, where the original creator can invite other researchers to join, this feature could be valuable for small research teams or subject-specific communities working on the same analysis to share their workflows.

  The second use case involves a generic EOSC Node interested in deploying REANA to enable the execution of computational data analysis workflows close to the data they disseminate. For example, if a researcher publishes data on Zenodo along with an example analysis workflow demonstrating how these data should be analyzed, this workflow could be ingested and exposed through the EOSC EU Node's Resource Hub, which could include a new "Run on CERN Node" button alongside the "Cite" button. Researchers interested in these data would be able to click the "Run on CERN Node" link, authenticate successfully, and have the analysis launched on the given CERN Node

in the EOSC federation. This would make data and software resources easily actionable by bridging the data with automatically spawned analysis environments for users interested in further data explorations.

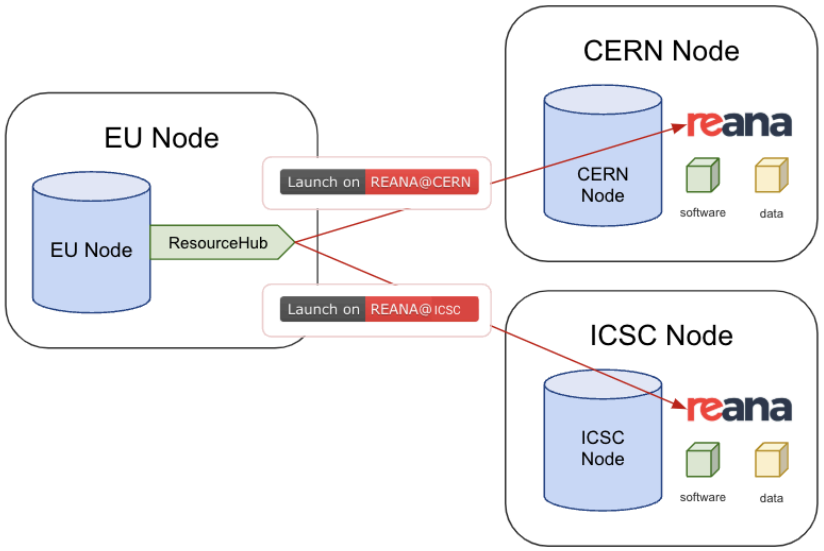

A more complex example of multiple EOSC Node workflow execution would involve hybrid distributed workloads, where one part of the workflow runs on one EOSC Node, while another part runs on a different EOSC Node. This is especially useful for communities with multiple data centers hosting datasets, in which case it would be most efficient for the researcher to send the analysis code to run close to the location of the data in order to shorten computation times. The use case aims to establish collaborative workflows between two or more EOSC Nodes, each running both data storage and compute elements. The technical options could be assessed in collaboration with interested partners (e.g., EU Node, CERN, ICSC/INFN, PUNCH/NFDI, and perhaps SURF), although full production-ready implementation would likely require additional resources.

- **Value proposition:**
  - Individual researchers interested in running computational workflows (rather than notebooks) could deploy their own REANA instances on the EOSC EU Node without needing extensive knowledge of Kubernetes technology.
  - Self-organized research teams interested in reproducible data analysis practices could use REANA on the EOSC EU Node as a team, sharing their data and results.
  - EOSC Nodes interested in promoting reproducible and reusable data analysis practices could offer the REANA platform to their users for running analyses at scale.
  - Collaborative REANA deployments across two or more EOSC Nodes would enable researchers to run distributed workflows by sending code to the data centers where the data is located.
  - Demonstrators built for particle physics could easily be replicated for other scientific communities, as REANA itself and the supported workflow languages are general and not specific to the particle physics discipline.

- **In-scope:**
  - Potential deployment of REANA on several EOSC Nodes to support multi-node use cases.
- **Out of scope:**

- Fine-grained accounting of data ingress/egress costs and of the user computations across EOSC Nodes in the case of distributed workflows.
- Full implementation of the more complex example of multi-node workflow execution that would involve hybrid distributed workloads. This could move to in-scope if the right funding and capacity is identified.

b. ESCAPE VRE (Virtual Research Environment)

- **Use Case Description:**

  The ESCAPE Virtual Research Environment (VRE) is a modular, scalable analysis platform that integrates CERN's EU onboarded services, enabling users to access data, execute workflows, and interact with research tools. It connects distributed storage systems (via Rucio), manages analysis workflows (via REANA), and provides access to interactive notebooks through the JupyterLab interface. This setup allows for seamless execution of reproducible workflows and large-scale data analysis across federated services.

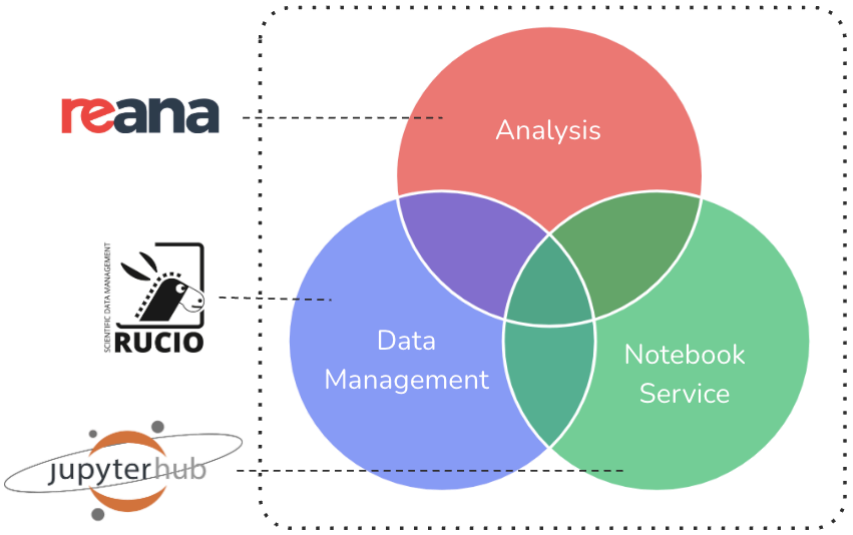

- **Value proposition:**
  - The VRE will showcase the connection and integration of heterogeneous services and resources under a single platform, using only the federated EOSC AAI. It will seamlessly allow access and interaction with Rucio for data access and REANA for reproducible analysis workflows.
    - Rucio is a federated data management system (a Data Lake), that connects distributed storage elements under the same framework. Via the VRE interface, users will be able to explore and access data assets available in any connected storage. Rucio can add and remove storage elements to the framework, allowing the possibility of including EOSC storage resources as part of the system.
  - Leveraging the composability and interoperability of the VRE, use-cases requiring access to distributed data management systems, reproducibility pipelines, and public repositories will benefit from deployment in the EOSC Federation. This includes Dark Matter and Gravitational Wave searches, EOSC Science Projects, Open Data analyses and detector related R&D. Test deployments of the VRE on the EU Node show the cloud container platform limits the use-cases that can be executed so extending the deployment to CERN, ESCAPE, and National Nodes would enable more advanced workflows and attract more users.
- **In-scope:**
  - Deployment and integration of the ESCAPE VRE with key EOSC onboarded services, including:

- EOSC AAI for federated authentication and authorization.
- REANA for executing analysis workflows in a reproducible, containerized environment.
- Integration of the VRE extensions (REANA and Rucio) into the EU Node Jupyter framework to provide access to the aforementioned services.
- Explore the integration of an EOSC Storage element into the Rucio framework to make it part of the ESCAPE Data Lake.

- **Out of scope:**
  - There will be no Rucio deployment on EOSC resources. Any VRE will be connected to a distributed storage element, but the Data Lake instance will be running on CERN resources.

4. External Dependencies & Key Risks

Risks	Actions	Deadline
Delays in AAI integration or interoperability with EU Node	Collaborate with DG CNECT and the EOSC AAI working groups to ensure timely resolution	Q2 2025
Limitations on configuration of the EU node services inhibit the deployment of the CERN tools/services	Reach agreement with DG CNECT how CERN tools/services could be best onboarded into the EOSC EU Node 'Other Services' for automatised single-user deployments; if unsuccessful, then document the procedure the researchers would take manually	Q2 2025
Potential mismatch between metadata models of INSPIRE-HEP/OpenData and OpenAIRE graph, making integration into the resource hub difficult	Work with OpenAIRE to overcome the issue, potentially reducing the scope of exposed metadata if needed	Q2 2025
Deployment of REANA on EU Node requires more OKD access rights than is currently provided for Cloud Container Platform deployments	Discuss early with EU Node team about the technical details and any “hard” OKD limits and see about possible special deployment rights for REANA applications	Q2 2025
EU Node offers too limited computational resources (virtual cores) to be useful to researchers’ complex workflows	Build early demonstrators and explore gradually expanding OKD credit offering towards using many computing nodes behind the cluster	Q3 2025
Limited scalability of existing infrastructure for multi-node workflows	Gradual expansion based on initial deployment learnings and scalability testing	Q3 2025
EOSC Node builder kit is not available or not suitable for usage by CERN	Use REANA/VRE as the portal or find alternative open source software as the basis for EOSC node (i.e. we understand SURF has something similar)	2026
Access policy does not permit EOSC federation users (i.e. those without a CERN account) to access rivalrous services on CERN infrastructure	Reduce the scope of services offered by the CERN node	2026

Experience gained during phase 1 indicates the resources available for phase 2 deployment are insufficient	Find alternative funding sources, extend the time available for the deployment or reduce the scope	2026
virtual credit model not available or not suitable for use by CERN	Define home-grown means of limiting and controlling access to CERN resources	2026
Legal and privacy concerns regarding data sharing between nodes	Align with legal and privacy policies across participating nodes	2026

- **External Dependencies:**
  - AAI challenges
    - CERN operates multiple AAI components for the laboratory and experiments including:
      - An Identity Provider for verified CERN users to access eduGAIN Service Providers. This is already connected to MyAccess ID through eduGAIN. [Software = Keycloak]
      - A Service Provider proxy for sirtfi-compliant eduGAIN Identity Providers to access CERN Service Providers. This is equivalent to a community proxy. [Software = Keycloak]
      - OAuth 2 Token Issuers for grid computing authorisation (for CERN based experiments and those based at selected partner organisations). These are equivalent to community proxies. [Software = INDIGO IAM]
      - Additionally, a partner organisation operates AAI for the ESCAPE project. We are able to modify these services to release or accept attributes with the appropriate format for EOSC.
    - Several of the services to be onboarded during phase 1 have an inbuilt notion of authorisation, for example mapping users to appropriate storage components in the Rucio extension of Jupyter Notebooks. It is unclear at this stage where this information should come from when these services are connected to the EOSC infrastructure proxy directly, and the lack of such information may significantly decrease the potential impact of the service. The only system with such authoritative information for experiment members is the relevant community proxy. There will likely be additional complications with verification of a user's home institute. For many CERN researchers it will be unclear whether they should authenticate to MyAccess ID with CERN or with their home institute, each individual feels a unique degree of association with their university and with CERN.
    - Implementing the remote token introspection endpoint (required for certain integration configurations) will be a non-trivial extension of either Keycloak, INDIGO IAM or both. It is understood that there is already some ongoing effort by the wider community, to which we would participate.
  - Dependencies during phase 1:
    - Integration with EU Node Infrastructure and Services: Successful integration with the EU Node infrastructure and services is essential to ensure smooth interoperability between CERN's services (e.g., REANA, Zenodo) and the EOSC ecosystem.
    - Access to Technical Documentation and Resources from the EU Node: Access to comprehensive technical documentation and resources from

the EU Node is crucial for understanding integration procedures and ensuring alignment with EOSC standards.
        - Collaboration with Other EOSC Nodes (REANA, Indico, etc.): Collaboration with other EOSC nodes, particularly for shared services such as REANA and Indico, will be necessary to maintain consistent service provision across the EOSC ecosystem and enable the execution of cross-node workflows.
        - Zenodo Integration (EOSC AAI): EOSC AAI is needed to authenticate and authorize EOSC users who upload materials to Zenodo. This is particularly important for ensuring that users can interact with the platform and upload their data seamlessly. Potential integration with EU project member quotas may require an understanding of how EOSC AAI can provide sufficient data to determine the user's associated community or funding agency (e.g., EU).
        - EU Node Compute & Storage Services: The ability to store data and run computational tasks is fundamental for the successful deployment of REANA and the ESCAPE VRE use cases. Access to EU Node compute and storage services is necessary to support these workflows.
        - VRE Dependencies: Integration of the EOSC AAI, storage, and compute resources is essential to facilitate the operation of the ESCAPE VRE use case.
        - Support from other EOSC nodes, such as EU Node, ICSC, and PUNCH4NFDI (if participating in the NFDI node) will be crucial for the cross-node REANA workflows.
    - Dependencies during phase 2:
        - EOSC Node Builder Kit (DG CNECT): The EOSC Node Builder Kit from DG CNECT will be required to set up and explore it for the CERN Node, providing the necessary infrastructure and tools to deploy and manage services for researchers beyond HEP.
        - Virtual Credit Model: The implementation of the virtual credit model will be critical for managing resource usage and allocation within the CERN Node. This will require collaboration with the EU Node and other participating nodes to ensure consistent and equitable resource distribution.
        - Other Nodes Supporting REANA Use Case: as in Phase 1, continued collaboration with other nodes supporting the REANA use case, including ICSC and PUNCH4NFDI, will be vital to ensure that distributed workflows and computational tasks can be executed efficiently across multiple EOSC nodes.

5. Contributions

- **Role in Project:**
    - CERN will serve as a *data provider* and *infrastructure* facilitator within the EOSC Federation. We will provide access to CERN's HEP data, supporting its use in cross-disciplinary research. Additionally, CERN will contribute limited compute resources through the Virtual Research Environment (VRE) and REANA, enabling researchers to execute reproducible workflows and analyze large datasets directly where the data is hosted.
- **Main Contributions:**
    - CERN:
        - Existing production quality repositories & services operated by CERN (Zenodo, Indico, REANA, CERN open data portal, HEPData, INSPIREHEP, CERN implementation of ESCAPE VRE)

        - Open source software (InvenioRDM, INDICO, REANA, RUCIO, CVMFS, etc.)
    - Durham university (UK): originator of HEPData
    - INSPIRE-HEP collaboration members: DESY (Germany), Fermilab (USA), IHEP (China), CNRS IN2P3 (France), SLAC (USA)

- **Deliverables:**

ID	Deliverable Name	Deliverable Description	Deliverable Owner
D1.1	Indico Service onboarded to the EU Node	Indico: Integration with the EU Node as a tier 2 service for managing research events	IT-CA
D1.2	REANA Service onboarded to the EU Node	REANA: Deployment and integration into the EU Node as a tier 2 service, enabling reproducible analysis workflows	IT-CA
D1.3	VRE Service onboarded to the EU Node	VRE (Virtual Research Environment): Deployment of the VRE as a tier 2 service, enabling the execution of distributed workflows and integration with CERN services	IT-TD
D1.4	Catch-all Repository for Research Output Publication	Zenodo: Enable the publication of research outputs, ensuring integration with EOSC AAI for seamless access to upload and share data	IT-CA
D1.5	CERN Open Data Portal	Provide access to HEP data and related resources, ensuring visibility and discoverability via EOSC	IT-CA
D1.6	HEP-Data	Integration of high-level data related to HEP publications, making it available for the wider EOSC community	RCS-SIS
D1.7	INSPIRE-HEP	Integration of HEP literature and related resources, enhancing the discoverability and accessibility of research data	RCS-SIS
D1.8	Presentation on Phase 1	A comprehensive presentation documenting the experience gained during the first phase deployment	Coordinator

- **Resources Provided (if any):**
    - CERN digital infrastructure hosting the repositories and services. CERN's digital infrastructure, hosted in 2 dedicated data centres, is dimensioned to support its scientific programme. We are confident it can handle the additional load of EOSC Federation users once an acceptable access policy has been agreed.
    - Integration expertise & effort via CERN's EOSC federation task force:
        - IT department
            - Collaborative Applications group (Jose B.G.L.) addressing
                - Zenodo (Alex I.),
                - Indico (Adrian M.),
                - REANA (Tibor S.),
                - CERN open data portal (Pablo S.),

- Platforms & Workflows group addressing AAI integration (Hannah S. & Berk B.)
- Technical Coordination and Governance groups addressing
  - the ESCAPE Virtual Research Environment deployment (Xavier E., Giovanni G. & Enrique G.G.)
  - and cybersecurity (Stefan L.)
- Storage and Data Management group (Giuseppe L.)

- Scientific Information Service
  addressing HEPData and INSPIRE (Micha M.) as well as alignment with CERN's open science policy (Kamran N.)

6. Timing and Milestones

ID	Milestone Description	Target Delivery Date
M1.1	CERN open data portal exposes all necessary metadata for OpenAIRE harvesting and is registered as a provider	Q2 2025
M1.2	HEPData exposes all relevant metadata to OpenAIRE by improving its DataCite metadata export	Q2 2025
M1.3	INSPIRE-HEP exposes all necessary metadata for OpenAIRE harvesting and is registered as a provider	Q3 2025
M1.4	Indico integrated with EOSC AAI	Q2 2025
M1.5	Indico onboarded to EU Node	Q3 2025
M1.6	REANA onboarded to EU Node as a single-user deployment	Q2 2205
M1.7	REANA integrated with EOSC AAI	Q3 2025
M1.8	REANA Assessment of federated workflow execution with several EOSC Nodes	Q3 2025
M1.9	ESCAPE VRE integrated with EOSC AAI	Q2 2025
M1.10	ESCAPE VRE onboarded to EU node	Q3 2025
M1.11	Zenodo integrated with EOSC AA	Q2 2025
M1.12	Zenodo link from EU Node	Q3 2025
M1.13	Access policy for phase 2 agreed with CERN management	Q3 2025
M1.14	Assessment of EOSC Node builder kit	Q3 2025
M1.15	Assessment of EU node virtual credit model for use by CERN node	Q3 2025
M1.16	Completion and assessment of phase 1 deployment	Q3 2025
M1.17	Plan for phase 2 deployment approved by CERN management	Q4 2025

ID	Milestone Description	Target Delivery Date
M1.18	Start of Phase 2	Q1 2026

- Phase 2 milestone & deliverables will be defined as part of M1.17
- **Start Date:**
  - Phase 1: 31 March 2025
  - Phase 2: Q1 2026
- **Expected Duration:**
  - Phase 1: 8 months
  - Phase 2: TBD

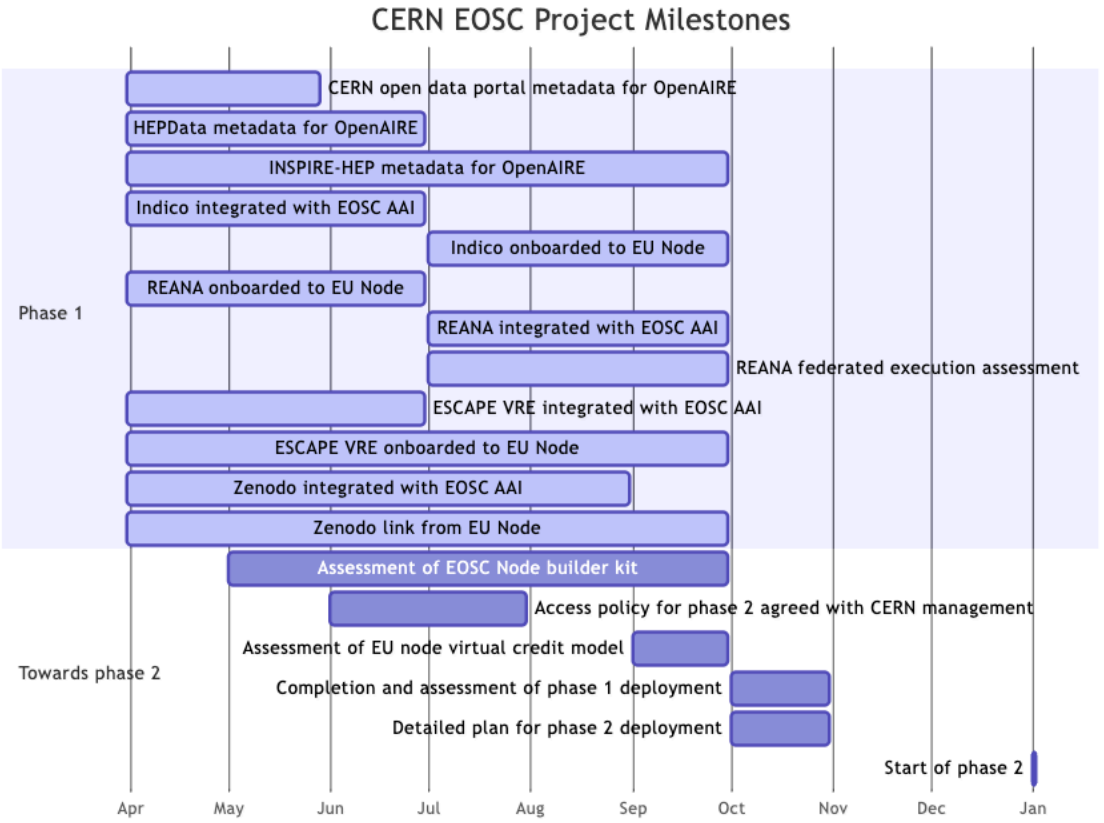

7. Contact & Submission

Role	Name	Email
Coordinator	Jose Benito Gonzalez Lopez	Jose.Benito.Gonzalez@cern.ch
Operation Manager	Jose Benito Gonzalez Lopez	EOSC-task-force@cern.ch
Security Officer	Stefan Lueders	Stefan.Lueders@cern.ch
Scientific Officer	Micha Moskovic & Xavi Espinal	EOSC-task-force@cern.ch
Legal/privacy Officer	Jose Benito Gonzalez Lopez in consultation with CERN's legal service	Jose.Benito.Gonzalez@cern.ch
Communication Officer	Antonella Del Rosso	Antonella.Del.Rosso@cern.ch

Beyond the official roles, this is how to contact each service provider:

- Zenodo: https://zenodo.org/support (or via info@zenodo.org)
    - EC-related requests can be labelled
- CERN Open Data portal: https://opendata-forum.cern.ch/
- INSPIRE-HEP: https://help.inspirehep.net/knowledge-base/contact-us
- HEPData: https://hepdata-forum.cern.ch/
- REANA:
    - User forum: https://forum.reana.io/
    - Third line support for onboarded service in EU Node: To be agreed with EOSC Federation.
- VRE:
    - Email: cern-vre-support@cern.ch
    - User forum: https://escape-talk.web.cern.ch
    - Third line support for onboarded service in EU Node: To be agreed with EOSC Federation.
- Indico:
    - User forum: https://talk.getindico.io/
    - Third line support for onboarded service in EU Node: To be agreed with EOSC Federation