SURF

Project Charter

EOSC pilot node SURF

Version 0.9 – 31 March 2025

1. Project Summary

SURF is the national ICT cooperative of Dutch education and research institutions. Members (100+) are all universities, universities of applied sciences, UMCs, and a lot of the research institutions. SURF provides IT services like identity & access management, procurement & contracting, security, network connectivity, storage and data management, publishing and compute. SURF innovates, among other things, in the area of quantum compute and AI fabrics.

As EOSC-A MO, SURF brings key stakeholders together in a national working group to explore building a national node offering a wide range of services to EOSC Federation, from compute, data and networking services to services in support, consultancy, training and contracting.

As a national organisation, collaborating with Dutch Large Scale Research Infrastructures (LSRI), some of those being members of European Research Infrastructures and/or Science Clusters, SURF will explore how services and resources can be offered through Thematic EOSC Nodes.

As a member of e-Infrastructures like GEANT, EGI and EUDAT, SURF will explore how services and resources can be offered through e-Infrastructure EOSC Nodes.

Depending on how the discussions on the EOSC Core evolve, some of the SURF services could be considered part of the core of the EOSC Federation.

As a national pilot node, SURF will offer computing, data and storage resources. For example, through SURF Research Cloud, SURF will offer virtual machines, containers and jupyter notebook resources and through ResearchDrive, File Sync and Share capabilities. SURF also operates a FAIR data repository service to support the publication of large-scale datasets supporting multiple communities. Metadata from the SURF data repository is made discoverable through EUDAT, B2FIND, and the OpenAIRE Research Discovery Graph, contributing to the Web of FAIR data. SURF is developing through FAIRCORE4EOSC an European-level RAiD service for managing project-related information.

2. Value Proposition

SURF is an IT cooperative serving the needs of its members. In this capacity, SURF fulfils 3 roles.

1. An **association** in which members work together across the boundaries of their sector/campus and, together with the SURF organisation, develop, combine and share knowledge about the optimal use of IT in education and research.
2. A **service provider**, in which the SURF organisation provides a reliable, state-of-the-art range of services that has been created in consultation with the members.
3. An **innovation workspace**, where SURF creates an environment in which members can collaborate on complex innovation issues with each other and with the SURF organisation in an optimal way. We choose to tackle these issues with an **ecosystem approach** that brings together various parties, agreements, and technologies to achieve a solution.

In contributing to the EOSC Federation and participating as a pilot node, SURF sees added value in all these roles:

As an **association** and as mandated organisation within the EOSC-Association SURF has a coordinating role, bridging EOSC contributions from the Netherlands (at institutional (e.g., RPOs) and national level) with the European level. In the coming years, SURF is seen, by its Scientific Advisory Board, as a neutral ICT organisation to take the lead in realising a new data-oriented infrastructure together with its members and partners.

SURF has taken the initiative in the Netherlands to set up a EOSC National Working Group that investigates how the technical and social-human elements can be integrated into a junction. The Thematic DCCs, DANS, OSNL, eScience Center are represented in this working group.

As a **service** provider, SURF offers a rich portfolio of services to its members to support research and education in the Netherlands. The services range from a variety of computing, data and networking services to human services in support, consultancy, training and contracting. Through the EOSC Federation and through the pilot nodes, SURF will assess offering these services cross-border to the wider EOSC community.

As an **innovator**, SURF will bring together key stakeholders to build a national node and contribute to the development of the EOSC federation as an EOSC pilot node. These activities will bring together resources and capacity with the aim of making them available to the EOSC Federation and will assess how resources from other nodes can be made available to the research community in the Netherlands.

One of the main challenges for EOSC and for public-funded organisation is cross-border service provisioning in a sustainable way. This is due to how public and semi-public organisations as SURF are nationally funded with a national scope. By participating as a pilot node in the development of the EOSC Federation, SURF wants to co-create the procedures and conditions under which nationally funded services can be offered cross-border to the wider EOSC community.

Another challenge is how to integrate existing services to support end-to-end (multiple nodes) workflows. What rules of participation & access, standards and interoperability, governance, human capacity, engagement are needed to establish such workflows.

3. Use Case(s)

In this table, SURF lists three potential use cases. Contributing to the general use cases which were discussed during the kick-off meeting of the EOSC Federation build-up phase:
(A) tools to the data, (B) Collecting data and (C) Access to sensitive data.

The final selection of the use cases depends on the use cases/ available capabilities of the other nodes during the Federation build-up phase, internal available resources and alignment in planning with ongoing activities.

Use Case ID	***Use Case Description***	***Federation Contributions & Value to Users***
1 **(A)** **Tools to the data**	**Shared software stack across EOSC nodes** In this use case, we will promote the adoption and usage of EESSI within the EOSC federation. This use case would consist of three activities: Deploy and maintain a server that hosts the software stack for use within the EOSC federation. More concretely, we	This use case enables the following tools-to-data workflow in the EOSC Federation: a user develops a computational analysis in one EOSC node using the scientific software provided by EESSI. When this user, wants to run this analysis in a different EOSC node because a large dataset is available there, the software dependencies of the analysis are guaranteed to be available because EESSI provides the same, consistent software stack at each EOSC node. This is regardless of underlying

	would provide a CVMFS Stratum 1 server for the EESSI project, and would like to collaborate with other nodes, so they could also host a Stratum 1 server (see "Suggestions for multinode collaborations" below). 2. Support other nodes in making use of the software stack. For an EOSC node to make the EESSI software stack available to their users, system administrators need to install and configure software on their infrastructure. Through this activity, we would support them in this process. Furthermore, we have experience with integrating EESSI into SLURM, Kubernetes, Jupyter, RStudio, and other Virtual Research Environments, which improves the user experience of using the shared software stack in their workflow tooling. Note: in the pilot phase, enabling EESSI would consist of an installation and configuration process. In later phases, we could provide guides for creating a more performant infrastructure solutions per node, such as deploying additional caching servers or a local Stratum 1 server. 3. Currently, we are collaborating with the life sciences, hydraulics engineering, and radioastronomy community to assist in the use and adoption of EESSI. As EESSI is based on EasyBuild which contains a large collection of scientific software build recipes, there are many opportunities and existing resources to promote the adoption of EESSI in several research domains. In this activity, we would provide support to researchers who would like to contribute software to EESSI or make use of EESSI software in their scientific workflows.	computer architectures, operating systems, and virtualization layers. Moreover, the software will be optimized for the hardware it is running on, which is ideal in a heterogeneous federation. To summarize, this use case provides a uniform software stack which is available at each EOSC node, helps users develop and execute federated workflows more easily, as they can assume each node provides software in the same way. The software provided by EESSI is optimized for a large selection of (micro)architectures, so this use case will also benefit the energy consumption within the EOSC federation, as well as the time it takes to run analyses. **Resources:** Dröge B, Holanda Rusu V, Hoste K, van Leeuwen C, O'Cais A, Röblitz T. EESSI: A cross-platform ready-to-use optimised scientific software stack. *Softw Pract Exper*. 2023; 53(1): 176–210. doi:10.1002/spe.3075 **Access policy:** In the pilot phase, the infrastructure (i.e, the Stratum 1 server) will be made freely accessible to all EOSC nodes based on an informal fair use policy. When scaling up this use case in subsequent phases, the fair use policy should be formalized. In addition, there should be multiple nodes in the EOSC Federation providing Stratum-1 servers to ensure a highly available service that can distribute the load across the EOSC Federation. **Suggestions for multinode collaborations:** 1. We would like to collaborate with CERN and CSC in the first activity of this use case. CERN because of their leading expertise in this area, and CSC because the use case aligns well with their involvement in the EuroHPC Federation Platform. Concretely, we propose that each of us provides a CVMFS Stratum 1 server of the EESSI project that can be used within the EOSC Federation.

		2. For the second activity, we would also like to collaborate with other EOSC nodes by offering support to those nodes that would like to make the EESSI software stack available to users of their infrastructure. Based on the software currently distributed via EESSI, we believe the Life Science Research Node and CNR (Blue-Cloud 2026) could be great collaborators, but we are open to supporting any node interested in adopting EESSI.
2 **(B)** **Collecting data**	**Federating Enterprise File Sync and Share (EFSS) services** Multiple EOSC nodes are offering an EFSS service for their users to collaborate with fellow researchers and to share data and research outputs. In this use case we are federating EFSS services within the EOSC Federation. We focus on services supporting the OpenCloudMesh (OCM) API. This use case will be demonstrated by researchers working on EFSS services from different EOSC nodes sharing and collaborating research data and outputs. A potential scientific cross-domain use-case can be based on the use case being developed in Data Terra and ELIXIR.	This use case enables researchers to easily share research data and outputs across research communities and within the EOSC Federation. It avoids the need for researchers to access the different EFSS services where the data resides or to use a shared Big Tech EFSS service to share the data. Sharing data through EFSS services provides a researcher with easy access to shared data, analysis tools, and services to which he or she has access, whether they are located on a local node or through services offered through another node. For example, running an analysis workflow on a shared data set in a VRE environment and sharing the analysis outputs again in a similar way as the data set has been shared. This method could even be part of a TRE workflow. To build up the EOSC federation of EFSS services the EFSS services need to support, and have the OCM API enabled, and to allow sharing of each other's account directories. **Resources:** ResearchDrive (https://www.surf.nl/en/services/research-drive) **Access policy:** No additional access policy is required, researchers have access to local EFSS service and data is shared across EFSS services. **Suggestions for multinode collaborations:** The following nodes offer EFSS services: EUDAT, EOSC EU Node, CSC, EOSC-PL, SURF
3 **(C)** **Access to**	**Processing and Analysis of Sensitive data within EOSC** SURF Research Cloud (SRC) is a service through which cloud computing	Sensitive data is, by default, not open and accessible and is therefore limited available/accessible. Access can only be provided in safe and controlled environments. For creating the safe and controlled

Sensitive data	resources from different providers are made available to run VM, Containers, Notebooks, including the SANE Trusted Research Environment for processing and analysing sensitive data. This use case is focussing on integrating the cloud computing resources from the EOSC EU Node within SRC and to run the SANE Trusted Research Environment (TRE) on cloud resources provided through the EOSC EU node. In this use case we are planning to collaborate with the ELIXIR node to allow processing and analysis of sensitive data made available through ELIXIR node on the SANE TRE. The focus will be on providing the researcher with access to the sensitive data while the data provider remains in control of who has access to the data and of how the data is being used. This will be done by using secure data upload methods through which the data provider prepares the data for the researcher (user). In this use case, we leverage the guidelines provided through the EOSC-ENTRUST project.	environments it would also be possible to use the EESSI software. In this case we can combine some of the work from use case 1 into this use case. Trusted Research Environments (TREs) offer such safe and controlled environments. By offering TREs within the EOSC Federation it opens the federation to data providers and offers value to researchers working with sensitive data. **Resources:** SURF Research Cloud: (https://www.surf.nl/en/services/compute/surf-research-cloud) **Access policy:** Available for SURF members. Available for other users through a different cost model, and subject to review against SURF target group policy. For existing users no additional access policy is required. Researchers will have access to SURF Research Cloud to enable them to deploy a TRE on EOSC EU node. Providing EOSC EU node can provide sufficient resources for the TRE. **Suggestions for multinode collaborations:** The following nodes make use of sensitive data within their infrastructures. Galaxy.eu, BBMRI, Elixir, CSC. Although the EOSC EU node does not allow for processing of sensitive data because of the absence of a processing agreement, the goals is to demonstrate this TRE functionality with synthetic data on the EOSC EU node. This way we can still demonstrate this functionality without the need for the processing agreements.

4. External Dependencies & Key Risks

External Dependencies & Risks	Actions	Deadline
The basic components for building up the node platform are available as presented during the 2nd dialogue meeting. With these components we are able to implement the use cases. How these components are being offered to the federation depends on the requirements of the policies and procedures of the Federation.	The creation of the node platform and the integration of the separate components within the node platform is one of the activities planned during build up phase. An assessment of the requirements will be made to assess the efforts for required developments. Internally a strategic discussion is required to transition from individual components to a node platform.	An initial version of the node platform will be ready M6 of the build phase.
Technical alignment of the node capabilities with the federating capabilities. The EOSC Federation Federating Capabilities are not yet defined at sufficient details, therefore it is difficult to estimate and plan the required development efforts.	An assessment of the requirements will be made to assess the efforts for required developments.	An initial version of the node capabilities aligned with the EOSC Federation Federating Capabilities will be ready at the time of the EOSC Symposium
Aligning SURF policies and processes (e.g. contracting, access provisioning, accounting, billing, reporting) with the EOSC Federation policies and processes is a risk because the EOSC Federation policies and processes are not yet defined in sufficient details. Planning strongly depends on the outcome of the discussions during build up phase, the established EOSC Federation policies and processes, discussions within the National working group and the required changes within the SURF organisation.	SURF contributes to defining and detailing the EOSC Federation policies and processes. This activity is part of the build up phase of the EOSC Federation. We expect significant time to be spend on aligning the Life Sciences AAI, MyAcces ID and the SURF AAI solutions in the early stages of the project in order to be able to provide the TRE environment. The main focus should be on the node enrolment option since this option provides the most flexibility in using the already existing collaboration functions of the nodes when deploying the TRE.	No deadline is set.

Implementing the cross node end-to-end use cases. This strongly depends on the use cases and available capabilities of the other nodes during the Federation build-up phase. The initial use cases will be set up between 2 or a few nodes. Converting a use case into a federating capability requires additional discussions and agreements at the federation level. This is an external dependency and is out-of-scope for implementing the use cases.	Develop a joint work plan with the nodes involved in the use-cases.	At the EOSC Symposium 1 or more use cases can be demonstrated.

5. Contributions

ID	Deliverable Name	Deliverable Description	Deliverable Owner (the partner within the organisation responsible for producing the deliverable, if applicable)
1	**IT Governance documentation** A. Project/Programme Charter B. Architecture Design document describing the node platform and required changes. C: IT Security Plan: security model, security architecture and IT security impact assessment D: Evaluation of the IT Security Plan: provide input to the IT Security Risk Report	A. Project_Charter EOSC pilot node SURF (this document) B. Node architecture design document C: Node security plan D: SURF is ISO27k certified and has established procedures to regular evaluate of security plans.	A. Magchiel Bijsterbosch, innovation manager research B. Mark van de Sanden, Domain Architect Research C: Ivar Janmaat, Teamlead SURF Research Cloud D: Ivar Janmaat, Teamlead SURF Research Cloud
2	***Data protection policies and procedures***	To be defined during the buildup phase.	
3	***Deployment strategies, integration plans, charters, progress, risks and timelines***	To be defined during the buildup phase.	
4	***Operational quality plans, service reviews, verification***	To be defined during the buildup phase.	

	and test session results and defects status and resolution		
5	***Stakeholder and community engagement strategy*** A. Communications Plan	A: EOSC-A MO Workplan 2025.	A. Laurents Sesink, programma manager Open Science
6	***Production roll-out of service components including web-service API's and associated documentation*** A. Initial node platform	A. Initial node running production	A: Ivar Janmaat, Teamlead SURF Research Cloud
7	***Software, tools, services, datasets, prototypes.*** A. Prototypes Use case 2: ResearchDrive	A. 1 or more of the proposed use cases demonstrated	A. Depends on the to be realised use cases

6. Timing and Milestones

ID	Milestone Description	Target Delivery Date
1	A. Project_Charter EOSC pilot node SURF (this document)	April 2025
2	Node architecture design document	April 2025
3	Node security plan	Sept 2025
4	Initial node	Sept 2025
5	1 or more of the proposed use cases demonstrated	Oct 2025

ID	Use case 1: Shared software stack across EOSC nodes	Target Delivery Date
1	The preconditions for working on this use case are already met, and we expect that within a 3–6-month period we could have the first deliverable ready.	Oct 2025

ID	Use case 2: Federating Enterprise File Sync n Share services	Target Delivery Date
1	Design scientific use case	May 2025
2	Demonstrate PoC Federating EFSS	June 2025
3	PoC scientific use case	Sept 2025
4	Demonstrate scientific use case	Oct 2025
5	Demonstrate scientific use case at EOSC Symposium	Nov 2025

ID	Use case 3: Processing and Analysis of Sensitive data within EOSC	Target Delivery Date
1	Design scientific use case	May 2025
2	Demonstrate PoC TRE	Sept 2025
3	Demonstrate scientific use case	Oct 2025
4	Demonstrate scientific use case at EOSC Symposium	Nov 2025

7. Contact & Submission

During the initiation phase of constructing the project charter the coordinator will be the link to the security officer and legal advisor.

Role	Name	Email
Coordinator	Laurents Sesink	Laurents.sesink@surf.nl
Operation Manager	Ivar Janmaat	Ivar.janmaat@surf.nl
Node architect	Mark van de Sanden	Mark.vandesanden@surf.nl
Program manager	Raymond Oonk	Raymond.oonk@surf.nl