



# FENIX

RESEARCH INFRASTRUCTURE

## D4.8 R&D Results

<b>Work package:</b>	<b>WP4 - Procurement, deployment and operation</b>	
<b>Author(s):</b>	Thomas Leibovici	CEA
	Björn Hagemeier	JUELICH-JSC
	Dorian Krause	JUELICH-JSC
	Stefano Gorini	ETHZ-CSCS
	Michele Carpené	CINECA
<b>Reviewer #1</b>	Dirk Pleiter	JUELICH-JSC
<b>Reviewer #2</b>	Sadaf Alam	ETHZ-CSCS
<b>Dissemination Level</b>	Public	
<b>Nature</b>	Other	

Date	Author	Comments	Version	Status
30.11.2020	Thomas Leibovici	Document structure	V0.1	Draft
19.02.2021	Thomas Leibovici	All contents collected	V0.2	Draft
22.03.2021	Thomas Leibovici	Ready for review	V0.5	Review
31.03.2021	Thomas Leibovici	Taking review comments into account	V0.6	Review
06.04.2021	Thomas Leibovici	Taking review comments into account	V0.7	Reviewed
13.04.2021	Valentina Armuzza	Final editorial updates	V1.0	Final

The ICEI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 800858.

© 2021 ICEI Consortium Partners. All rights reserved.



## Executive Summary

Task T4.7 of ICEI is responsible for procuring R&D services to develop solutions for implementing different Fenix services. Depending on the progress of the various calls for tenders and the complexity of the developments, some of these R&D services are now in operation, while others will not produce results until the first half of 2021. Hence, this document (deliverable D4.8) summarises the current progress of the tender procedures, the development of the Fenix services, the results obtained to date, the next implementation steps, and how the services are used within the Fenix e-infrastructure. In order to comply with confidentiality obligations, the main part of the document presents only an overview. Details are provided in a confidential annex.

The report shows that the Authorisation and authentication infrastructure (AAI) as well as the User and Resource Management Service (FURMS) are at an advanced stage. AAI, for example, is planned to be in production by the end of the first quarter 2021, while for FURMS, a first prototype with limited functionality has been deployed for testing and evaluation by Fenix sites.

The Data mover service, the OpenStack Swift enhancements, as well as the service for Swift over open-source file systems, are progressing. These are being implemented, delivered, and tested on the Fenix site managing the respective R&D procurement. Status updates will be periodically checkpointed in the coming quarters.

The tendering process for the service allowing for interactive computing will close in April 2021.

## Contents

### Table of Contents

Executive Summary .....	2
Contents .....	3
List of Acronyms .....	4
1. Introduction.....	5
2. Authentication and Authorization Infrastructure (AAI) .....	6
2.1. Tender procedure and status .....	7
2.2. Development progress and results .....	8
3. User and Resource Management service (FURMS) .....	8
3.1. Tender procedure and status .....	8
3.2. Development progress and results .....	9
4. Data mover service .....	9
4.1. Tender procedure and status .....	9
4.2. Development progress and results .....	10
5. OpenStack Swift enhancements.....	10
5.1. Tender procedure and status .....	10
5.2. Development progress and results .....	10
6. Swift over open-source parallel file system.....	11
6.1. Tender procedure and status .....	12
6.2. Development progress and results .....	12
7. Interactive computing service .....	12
7.1. Tender procedure and status .....	13
7.2. Development progress and results .....	13
8. Conclusion and next steps.....	13
9. References.....	14

## List of Acronyms

AAI	Authentication and Authorization Infrastructure
ACD	Active Data Repositories
API	Application Programming Interface
ARD	Archival Data Repositories
BSC	Barcelona Supercomputing Center
CEA	Commissariat à l'énergie atomique et aux énergies alternatives
CINECA	Consorzio Interuniversitario
CSCS	Centro Svizzero di Calcolo Scientifico
DM	Data Mover Service
FURMS	Fenix User and Resource Management Services
GRDP	General Data Protection Regulation
HBP	Human Brain Project
HPC	High Performance Computing
IAC	Interactive Computing Services
ICCP	Interactive Computing Cloud Platform
ICEI	Interactive Computing E-Infrastructure for the Human Brain Project
ICN	Interactive Computing Node
IdP	Trusted Identity Provider
JSC	Jülich Supercomputing Centre
MS	Monitoring Services
PRACE	Partnership for Advanced Computing in Europe
R&D	Research & Development
SCC	Scalable Computing Services
TGCC	Très Grand Centre de calcul du CEA
US	User Support Services
VM	Virtual Machine Services

## 1. Introduction

Five High Performance Computing (HPC) centres, namely BSC<sup>1</sup> in Spain, CEA-TGCC<sup>2</sup> in France, CINECA<sup>3</sup> in Italy, ETHZ-CSCS<sup>4</sup> in Switzerland and JUELICH-JSC<sup>5</sup> in Germany, committed to set up and to operate a new set of federated e-infrastructure services within the ICEI project. This is a long-term effort organised under the umbrella of Fenix, an organisation that currently involves the aforementioned HPC centres plus the computing center CSC in Finland. The procurement of R&D services conducted within the ICEI project is important for realising Fenix.

The Fenix e-infrastructure comprises the following services defined in [D3.1]:

- Scalable Compute Services (SCC);
- Interactive Compute Services (IAC);
- VM Services (VM);
- Active Data Repositories based on fast memory and active storage tiers (ACD);
- Archival Data Repositories (ARD);
- Data mover services.

To allow for these distributed computing and storage resources to be perceived by the users as a coordinated, single infrastructure, various additional federated services need to be deployed:

- Authentication and Authorisation Infrastructure (AAI) services;
- User and resource management services (FURMS).

To realise these services, partners in the ICEI project have awarded R&D and other service contracts to enhance and operate the AAI services as well as to develop the FURMS service. Further R&D services are needed to create or adapt and enhance available solutions for:

- Enabling data movement between ACDs and ARDs (Data Mover service);
- Realising Swift<sup>6</sup>-based Archival Data Repositories in a cost-efficient manner;
- Realising Interactive Computing Services and improving the utilisation of clusters, which are partially used for such services.

The subcontractors who will develop these services have been chosen through open competitive tenders. The tender packages for R&D services have been previously described as part of the deliverable D4.15 [D4.15].

Each R&D tender and the resulting service is led by one of the project beneficiaries, which runs the tender and manages the awarded contract with the selected service provider. However, all partners are involved in the execution phase of the R&D services, e.g. through joint meetings with the contractors

---

<sup>1</sup> <https://www.bsc.es/>

<sup>2</sup> <http://www-hpc.cea.fr/en/complex/tgcc.htm>

<sup>3</sup> <https://www.cineca.it/>

<sup>4</sup> <https://www.cscs.ch/>

<sup>5</sup> [https://www.fz-juelich.de/ias/jsc/EN/Home/home\\_node.html](https://www.fz-juelich.de/ias/jsc/EN/Home/home_node.html)

<sup>6</sup> <https://wiki.openstack.org/wiki/Swift>

and involvement in the process of approving contractor deliverables as well as periodic testing, evaluation and training workshops. This ensures coherence of the deployed solutions within the overall infrastructure, and the compliance with site-specific architectures. The following table lists the R&D tenders and the responsible partner:

<b>R&amp;D tender</b>	<b>Responsible partner</b>
Authorization and authentication infrastructure (AAI)	CINECA
User and resource management service (FURMS)	JUELICH-JSC
Data mover service	JUELICH-JSC
OpenStack Swift enhancements	ETHZ
Swift over open-source parallel file system	CEA
Interactive computing service	CINECA

*Table 1: Overview of ongoing tenders*

The development of the services is conducted in several steps:

- Preparation of tender documents
- Execution of the tender process
- Awarding of contracts
- Development of services
- Deployment and site-local integration as well as the operation of the services.

Depending on the progress of the various calls for tenders and the complexity of the developments, some of these services are currently in the integration or operational stage, while others will not produce results until later in 2021. This document describes the status of these various developments.

## 2. Authentication and Authorization Infrastructure (AAI)

In order to provide access to Fenix infrastructure services, the federation needs to rely on a robust and reliable Authentication and Authorization Infrastructure (AAI), a trustworthy environment where users can be managed and granted access to resources in a secure and seamless way. For seamless access it is intended for a user with an identity from a trusted Identity Provider (IdP) and granted to consume a certain budget of resources, to access federated services using his/her credentials (i.e., username/password, X.509, etc.) without going through a new registration or authentication process.

The main functionalities supported by the two services are:

- **SERVICE 1: The Fenix Central Proxy IdP**
  - Federation of multiple IdPs by proxying authentication requests
  - Users identification and authentication
  - Validation of user profiles
  - Policy registry and management of principles of engagement and access terms and conditions
  - User profile attributes enrichment by interacting with external attributes providers
- **SERVICE 2: FURMS**
  - Group/budgets membership management
  - Authorization attributes provider

- SSH public keys management
- Managing site specific Usage Agreements
- Reporting and metering

The reason behind the separation of functionalities across two distinct services is twofold:

- Keep the Fenix Central Proxy IdP as lean as possible in order to provide high operation performance and flexibility while federating new IdPs;
- Improve infrastructure security by decoupling highly critical functions, e.g. authentication processes, from less critical ones, e.g. groups management.

## 2.1. Tender procedure and status

A contract was awarded to GÉANT<sup>7</sup>. Based on the technical requirements put forward by the Fenix consortium and detailed in [D3.2], the contractor proposed eduTEAMS as the solution to meet the needs of the Fenix community, and in order to implement the Fenix Central Proxy.

The Fenix AAI will be operational for 5 years as per contract by Cineca. A service level agreement (SLA) will be in place for the whole period starting from the moment in which the Fenix Central Proxy IdP goes into production. GÉANT will ensure that maintenance and support for the same duration according to what has been defined in the awarded contract.

One of the aspects that needed to be addressed for the operation of the AAI services was the definition of roles and responsibilities related to processing of personal data. Following the GDPR, CINECA has been identified as a Data Controller, while GÉANT will be considered as a Data Processor for the data records associated to Fenix users. The contract with CINECA has been signed in December 2020. It does foresee the first version of the production service to be delivered in March 2021.

---

<sup>7</sup> <https://www.geant.org/>

## 2.2. Development progress and results

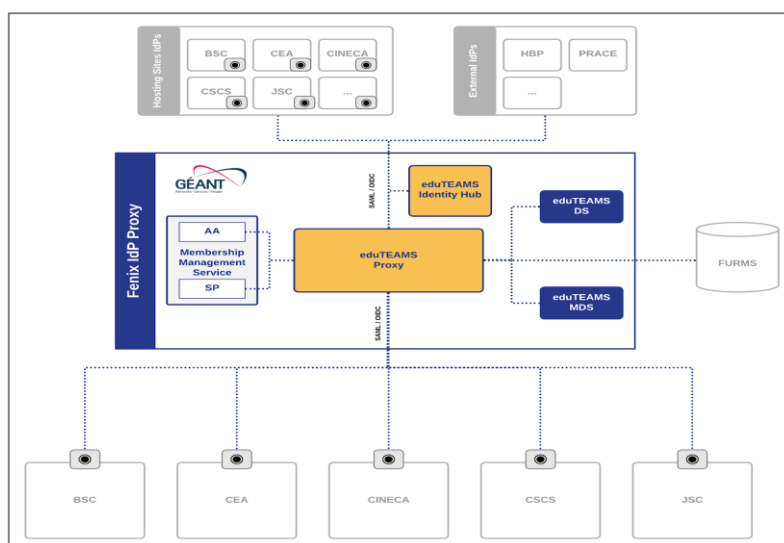


Figure 1: The eduTEAMS Fenix Central Proxy at Géant

A test version of the Fenix Central Proxy is available since June 2020 in an acceptance environment that was setup by GÉANT. April 2021 is the official date for the final production version of the eduTEAMS Proxy, which is currently based on the SATOSA<sup>8</sup> technology developed by GÉANT. At the same time, the integration with the authorization service (FURMS) will be tested.

End of 2021 is the final date for completing the delivery of the Fenix AAI, including integration of all the sites IdPs and Service Providers (SPs) with the Fenix Central Proxy.

## 3. User and Resource Management service (FURMS)

FURMS is intended to be the central service for managing resource allocations within the Fenix infrastructure. Sites indicate available resources within the service, which are subsequently distributed to communities and projects. By joining a project, users can consume site-specific resources that have been allocated to the project previously. The service also includes relevant actions related to accessing services at sites, such as credential management and policy approval. Additionally, FURMS covers functionality related to AAI like managing SSH keys as described in the previous section.

### 3.1. Tender procedure and status

The first call for tender has been published on December 11, 2019. Only a single bidder participated in the procedure, the offer was not satisfactory.

The second call for tender was published on February 17, 2020. Three bidders participated in the competitive procedure with negotiation. The contract was finally awarded to Bixbit, a Polish company with experience in AAI and working with public funded research. The contractual work commenced in September 2020.

<sup>8</sup> <https://github.com/IdentityPython/SATOSA>



### 3.2. Development progress and results

First phase deliverables, including a refined work plan and detailed documentation about the envisioned developments, have been submitted to partner JUELICH-JSC on December 3, 2020. The deliverable for phase 2 has been delivered on February 15, 2021. It consists of the first version of the software with an initial set of features: FURMS 1.0.0. This version was installed and passed the acceptance criteria on March 3, 2021. The expected delivery dates and contents of the phases are summarized in table 2. The main FURMS design and concepts documents are attached in annex.

Phase	Delivery date	Contents
2	February 15, 2021 (delivered)	Integration with central IdP; authentication and authorization; API for central IdP, FURMS user interface (UI) with user management features; basic UI views; first version of admin and user manuals.
3	April 2021	Implementation of resource allocation; broker setup; communication with local sites; remaining UI views to support aforementioned features; site local agent example; training for site admins on API; update of admin and user manuals
4	August 2021	Policy documents support; advanced FURMS API; branding; command line interface (CLI); audit log; user invitations and applications to projects; roles management; generic groups support; update of admin and user manuals
5	January 2022	Accounting features; performance tests; backup/restore procedure; hooks; trainings; update of admin and user manuals

Table 2: FURMS: Expected delivery dates

## 4. Data mover service

The data mover infrastructure will be a programmable, secure high-performance service for transferring data between the archival and active data repositories at the individual Fenix sites. The service will transform user data from objects (archival data repository) to files (active data repository) and vice versa, ensuring AAI compatibility across two domains. It hence bridges the gap between site-local parallel file systems, that are close to the HPC resources but inaccessible from outside for the federation, and the federated object-storage which provides important data sharing capabilities but is not suitable for data processing directly from computing resources.

### 4.1. Tender procedure and status

The development service was procured using a competitive procedure with negotiations followed by an open selection of candidates. The initial request for concepts was published on December 12, 2019 with a deadline January 22, 2020 for submitting applications to participate in the tendering procedure. Two applications had been submitted such that no down-selection was performed. Both candidates were

invited to submit proposals. The first full proposals were due by March 02, 2020. Two negotiation rounds were performed with meetings in April and June 2020. Final proposals from both candidates were received by July 07, 2020.

Following the evaluation of both final offers, the contract was awarded to Akquinet AG, a German IT-consulting and development company, which submitted the offer together with Belgium company MT-C. MT-C develops the Nodeum software, that is the basis for the planned solution for the data mover service.

## 4.2. Development progress and results

A series of design workshop meetings were held in autumn 2020 to clarify remaining open questions concerning the service semantics and integration into the federation, in particular concerning the interaction with the Fenix AAI. The full technical specifications, which constitute the first deliverable of the contract, have been finalized in January 2021. This specification is attached in annex. The second deliverable will be due in Q2 2021. This first versions of the data mover service may lack full integration with Fenix AAI, therefore, it will unlikely be offered as an official Fenix service.

## 5. OpenStack Swift enhancements

This R&D effort intends to develop additional features for the OpenStack Object Storage Software, SWIFT, that have been identified by the ICEI project as critical for Fenix. The areas where the R&D will focus for improving usage of the tape library are:

1. Enabling SWIFT user to save long term storage on inexpensive media such as tape, hierarchical storage management (HSM) functionality.
2. Developing a SWIFT backup-like functionality for objects. The software must enable object data duplication (pre-migration) on separate media (in particular tape storage).

### 5.1. Tender procedure and status

The tender has been issued following the Swiss public procurement rules. A first call for tender was published on November 21, 2019. Only a single bidder participated in the procedure and the offer did not satisfy the mandatory criteria.

A second call for tender was published on June 20, 2020. A single bidder participated, and the offer was compliant with the mandatory criteria. The tender was awarded to IBM Switzerland in September 2020. The final contract was signed in November 2020.

### 5.2. Development progress and results

After the signature of the contract, IBM and ETHZ-CSCS had a preliminary meeting early December 2020 preparing the team setup and format to properly engage all the Fenix partners. Late December 2020, CSCS organised a meeting with all the partners and IBM to present the outcome of the tender and to explain how the collaboration would work during the period of the contract (until March 2023). The following terms are agreed:

- ETHZ-CSCS will act as the main proxy to IBM when a major coordination will be needed.
- Each Fenix site appoints a local contact for quick interactions.
- To facilitate the collaboration during the development ETHZ-CSCS started a mailing list.
- The project will follow an Agile project management approach with periodic and regular updates.
- IBM started a Kanban board for project management with all the tasks foreseen to achieve a satisfactory development.

- IBM is setting up “sprints” and there will be a monthly meeting at a technical level (it could become more frequent if needed).
- In the first half of 2021 IBM will provide a first prototype to be tested by the partners.

## 6. Swift over open-source parallel file system

Lustre is an open-source parallel file system with a strong community of storage vendors, HPC sites, and public institutions, which makes Lustre a sustainable solution without a vendor lock-in. This R&D service addresses the needs of the HBP community for accessing shared data repositories, high-performance computing and data analysis facilities. The infrastructure can take advantage of Lustre for the following usages:

- Performance optimised Active Data Repositories (ACD): this component can take advantage of Lustre scalability to handle large-scale and data-intensive computations;
- Capacity optimised Archival Data Repositories (ARD): this can take advantage of Lustre tiering features to store large amounts of data in a cost-effective way;
- Flexible storage for virtual machines (VMs): Lustre can be used as a large and high performance data pool where storage resources can be allocated on-demand for virtual machines.

OpenStack Swift is the protocol that enables the federation of ARDs across Fenix sites. To enable the use of Lustre for the federated ARD Fenix service, a goal of this procurement is implementation and optimization of a Swift interface, which is currently lacking in Lustre. The R&D service consists in implementing 3 main features:

- Enable the use of Lustre as a storage backend for VM services;
- Enable the use of Lustre as a storage backend for Archive data repository (ARD);
- Provide the user with a unified view between a Lustre Active data repository (ACD), and the same storage accessed through Swift as an Archive data repository (ARD).

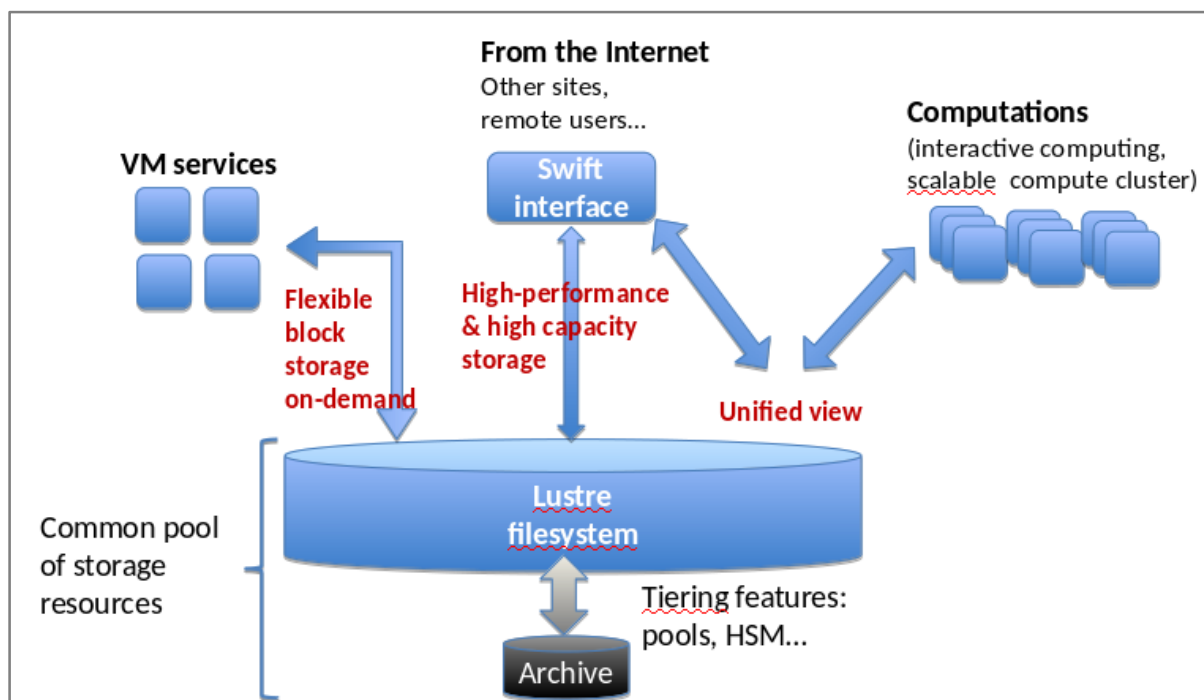


Figure 2: Overall schema of Swift over Lustre features

## 6.1. Tender procedure and status

The tender documents for this procurement were submitted as part of deliverable D4.15 (version 1), in September 2019.

The procurement has been performed as an open procedure. The Call for Tender was published in January 2020, and final bids were due in August 2020. Unfortunately, due to a technical issue on the platform that manages public tenders, some candidates did not receive answers to their questions before the due date for submitting the bid. To ensure fairness between candidates, and to avoid later recourse, the decision was taken to cancel the first call for proposal and issue a new call following an accelerated schedule. The second Call for Tender was published in September 2020. Two bidders submitted valid offers. In February 2021, the tender was awarded to LINAGORA, a company specialised in open-source software development. The contract is expected to be signed soon.

## 6.2. Development progress and results

An initial technical study has been led by LINAGORA as part of the procurement procedure. The outcome of this study is part of the confidential annex.

The first deliverable is expected during the first half of 2021. It will enable a basic use of Lustre as a storage backend for Swift. This will then be improved during the next months to provide additional features and performance improvements. The complete R&D service is expected to take 9 to 12 months, followed by a period to upstream the features to open-source projects, thus ensuring the long-term support of the implemented software. To facilitate regular exchange between the service provider and the other Fenix sites, periodic communication is planned with the Fenix sites.

## 7. Interactive computing service

Interactive Computing in ICEI refers to a software solution capable to provide interactive access to Fenix resources. Traditionally, users interact with high performance computing systems through a batch job system. Jobs submitted to the system are allocated on the system via a workload manager application, that has the primary goal to maximize the throughput of the HPC system at a cost of typically long response times. This may not be adequate for modern workflows - composed of different simulation and processing steps - and may reduce the capability of users to react quickly and steer the planned simulations. Interactive Computing in HPC tries to cover the needs of these modern workflows with a different usage model where the system is designed to respond promptly to user requests. This approach could be summarized in the following steps: i) users ask for interactive sessions, ii) the system provides state-of-the-art computing and storage resources in reasonable time (minutes) iii) users have access to the interactive sessions and all necessary data to be processed.

A typical usage scenario for interactive computing involves the visualization, processing, and reduction of large amounts of data, especially when interacting with data cannot be automatized or implemented in batch workflows. The time spent in this activity is a non-negligible part of the “time to science”. In neuroscience and other science fields the computational scientists show growing interest in using interactive frameworks such as R, Stata, Matlab/Octave or Jupyter Notebook.

To be effective in this context, the Interactive Computing Service would need to facilitate users’ interaction with the infrastructure, support efficient handling of interactive sessions, support staging of data across multiple memory and storage tiers and improve the overall usage experience of an HPC system. To this goal, this R&D activity should aim at a solution able to provide users with a single interaction mechanism, e.g. Web, CLI, API capable to handle different utilization requests, from the submission of batch jobs, the creation of interactive visualization sessions, to the transparent set up of notebook environments, e.g. Jupyter Hub web server. The tender focuses on R&D activities for the realisation of supported solutions that address several of the following needs:

1. Fast and easy access to Interactive Computing Services through servers that are part of an HPC cluster;
2. Optimised use of available compute and storage resources, possibly including support for efficient suspend-resume mechanisms, for HPC jobs running on Scalable Compute Services;
3. Easy-to-use tools for efficient remote visualisation through the Interactive Computing Services.

## 7.1. Tender procedure and status

The technical requisites were agreed upon by the project partners and submitted as part of deliverable D4.15. The selected procedure for the procurement process is a competitive dialogue. The procurement procedure was published on January 13, 2021, and the request for participation was open until February 15. The selection criteria to participate were clearly stated in the contract notice and the tender documents of the procedure.

## 7.2. Development progress and results

The procurement procedure is still in progress. After organising competitive dialogue meetings, the final technical requirements will be released so that a best and final offer for this service can be submitted.

## 8. Conclusion and next steps

The following table summarises the status of current developments detailed in the previous sections:

<b>Service</b>	<b>Status as of March 2021</b>	<b>Next expected results</b>
Authorization and authentication infrastructure (AAI)	Acceptance environment operational	In production April 2021
User and resource management service (FURMS)	Phase 2 features delivered	Phase 3 features to be delivered in April 2021
Data mover service	Development	First test version in Q2 2021
OpenStack Swift enhancements	Development	First prototype available in the first half of 2021, followed by regular updates
Swift over open-source parallel file system	Contract signing	First delivered features: first half of 2021
Interactive computing service	Tendering process	Awarding in May 2021

*Table 3: Overview on status of current developments*

As highlighted in table 3, the procurements and the development of services are at various stages of

progress. Most of the contracts have been successfully awarded. First services, which are implemented by the contractors, are close to the production stage (e.g. the AAI services). Within all awarded contracts more features are being implemented, delivered, and tested across Fenix sites.

## 9. References

[D3.1] ICEI Deliverable 3.1: Common Technical Specifications

<https://drive.ebrains.eu/smart-link/ab3abc1d-7edc-4383-abef-e0c150385358/>

[D3.2] ICEI Deliverable 3.2: Initial federated AAI infrastructure (M1-M13)

<https://drive.ebrains.eu/smart-link/5af9a020-f187-447d-98d7-8e968971b6a6/>

[D4.15] ICEI Deliverable 4.15: Tender Documents (Part 2)

<https://drive.ebrains.eu/smart-link/88442a94-6fbf-41ed-8e90-4d771588f089/>