



D7.1 Service Management Tools – Technical Plan

Lead partner:	EGI Foundation
Version:	1
Status:	Under EC review
Dissemination Level:	Public
Keywords:	Service Management System, tool, accounting, monitoring, messaging, helpdesk
Document Link:	https://documents.eqi.eu/document/3814

Deliverable Abstract

The Service Management Tools support the management activities, the operations, and the other services of the EGI-ACE project. In this deliverable we present the technical roadmap of the Accounting Portal and Repository, ARGO Messaging and Monitoring Services, the Helpdesk service, the Configuration Database, the Operations Portal, and the Software Provisioning Infrastructure, describing the new features implemented during the first year of the project and the ones planned for the months to come.



EGI-ACE receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101017567.

go.eqi.eu/eqi-ace

COPYRIGHT NOTICE



This work by parties of the EGI-ACE consortium is licensed under a Creative Commons Attribution 4.0 International License. (<http://creativecommons.org/licenses/by/4.0/>).

EGI-ACE receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101017567.

DELIVERY SLIP

Date	Name	Partner/Activity
From:	Alessandro Paolini	EGI Foundation/WP7
Moderated by:	Malgorzata Krakowian	EGI Foundation / WP1
Reviewed by:	Malgorzata Krakowian Andrea Manzi	EGI Foundation / WP1 EGI Foundation / WP6
Approved by:	SDS	

DOCUMENT LOG

Issue	Date	Comment	Author
v.0.1	01/12/2021	first draft ready for external reviewers	A.Paolini, A.Coveney, G.Corbett, I.Diaz Alvarez, G.Grein, K.Koumantaros, T.Zamani, C.L'Orphelin, J.Pina
v.1.0	17/12/2021	addressed comments of external reviewers	A.Paolini, M.Krakowian, A.Manzi, G.Sipos

TERMINOLOGY

<https://confluence.egi.eu/display/EGIG>

Contents

Executive summary	5
1 Introduction	6
1.1 The Service Management Tools in the EGI-ACE architecture	6
1.2 Relationship with the EOSC Core Services	9
2 Accounting Portal	10
2.1 Functional overview	10
2.2 Roadmap.....	10
2.2.1 Implemented in 2021.....	11
2.2.2 Planned for 2022-2023.....	11
3 Accounting Repository	12
3.1 Functional overview	12
3.2 Roadmap.....	12
3.2.1 Implemented in 2021.....	13
3.2.2 Planned for 2022-2023.....	13
4 ARGO Messaging Service (AMS)	15
4.1 Functional overview	15
4.2 Roadmap.....	15
4.2.1 Implemented in 2021.....	16
4.2.2 Planned for 2022-2023.....	16
5 ARGO Monitoring Service.....	18
5.1 Functional overview	18
5.2 Roadmap.....	19
5.2.1 Implemented in 2021.....	19
5.2.2 Planned for 2022-2023.....	20
6 GGUS Helpdesk Service.....	22
6.1 Functional overview	22
6.2 Roadmap.....	23
6.2.1 Implemented in 2021.....	23
6.2.2 Planned for 2022-2023.....	23
7 GOCDDB Configuration database	24
7.1 Functional overview	24
7.2 Roadmap.....	26
7.2.1 Implemented in 2021.....	26

7.2.2	Planned for 2022-2023.....	26
8	Operations Portal.....	28
8.1	Functional overview	28
8.2	Roadmap.....	28
8.2.1	Implemented in 2021.....	29
8.2.2	Planned for 2022-2023.....	29
9	Software Provisioning Infrastructure.....	30
9.1	Functional overview	30
9.2	Roadmap.....	31
9.2.1	Implemented in 2021.....	31
9.2.2	Planned for 2022-2023.....	31

Executive summary

The Service Management Tools are one of the pillars composing the EOSC Compute Platform and they support the management and operational oversight of the EGI-ACE landscape. They cover areas such as configuration management, usage accounting, availability & reliability monitoring, infrastructure oversight, infrastructure security, and software release management. Their technical roadmap is directly driven by the needs of the EGI-ACE service providers, especially by the providers of the EOSC Compute Platform. Indirect influencers of the evolution of the Service Management Tools are the stakeholders of the broader EOSC landscape. A particularly important user is the EOSC-Future project, considering that some of the Service Management Tools that are in EGI-ACE are also part of the EOSC Core Services operated in EOSC Future (through different funding channels). Until now there are no particular requirements coming from the EOSC Future project for the Service Management Tools, therefore the current document captures the roadmaps of the tools based on the EGI-ACE requirements. We expect that over the lifetime of EGI-ACE new requirements formulated within EOSC Future will become relevant also for the next release of the roadmap in December 2022.

In this deliverable we present a brief functional description for each service, report the requirements that have been implemented over the first year of the project, and provide a technical development roadmap for the next period.. The document covers the following tools with the following upcoming main developments for 2022:

- Accounting Portal and Repository: the Storage and GPU Accounting will be fully supported as well as a new benchmark for normalising the HTC usage records.
- ARGO Messaging Service: besides several improvements to the service, a new instance is planned to be deployed in order to improve the High-Availability operation of the service.
- ARGO Monitoring Service: the addition of new probes and metrics is an ongoing activity, but it will be improved the mechanism for deploying them; a new view for displaying the outcomes of the tests and replacing the nagios webui will be added as well as a page summarising the status of a subset of services, and the notification system will be reviewed.
- GGUS Helpdesk service: regular releases improving the user experience, such as the optimisation of registering an incident or a service request, and of the searching of already recorded tickets.
- GOCDDB Configuration Database: migration to a new, open source database framework , enabling support for technologies based on OIDC authentication mechanisms, and introducing Identity Linking feature for the users.
- Operations Portal: simplification of the authenticated access to the service through EGI Check-in service, and merging multiple accounts that users may own into one profile.
- Software Provisioning Infrastructure: after handing over the delivery of the service to a new supplier, it has started the work to migrate the service on a more modern infrastructure to improve its provisioning and the user experience.

1 Introduction

1.1 The Service Management Tools in the EGI-ACE architecture

EGI-ACE delivers the EOSC Compute Platform as a fully integrated compute environment that federates distributed hybrid compute and storage facilities to support processing and analytics via a set of services for distributed data and compute access.

The EOSC Compute Platform architecture is organized in functional blocks as shown in Figure 1.

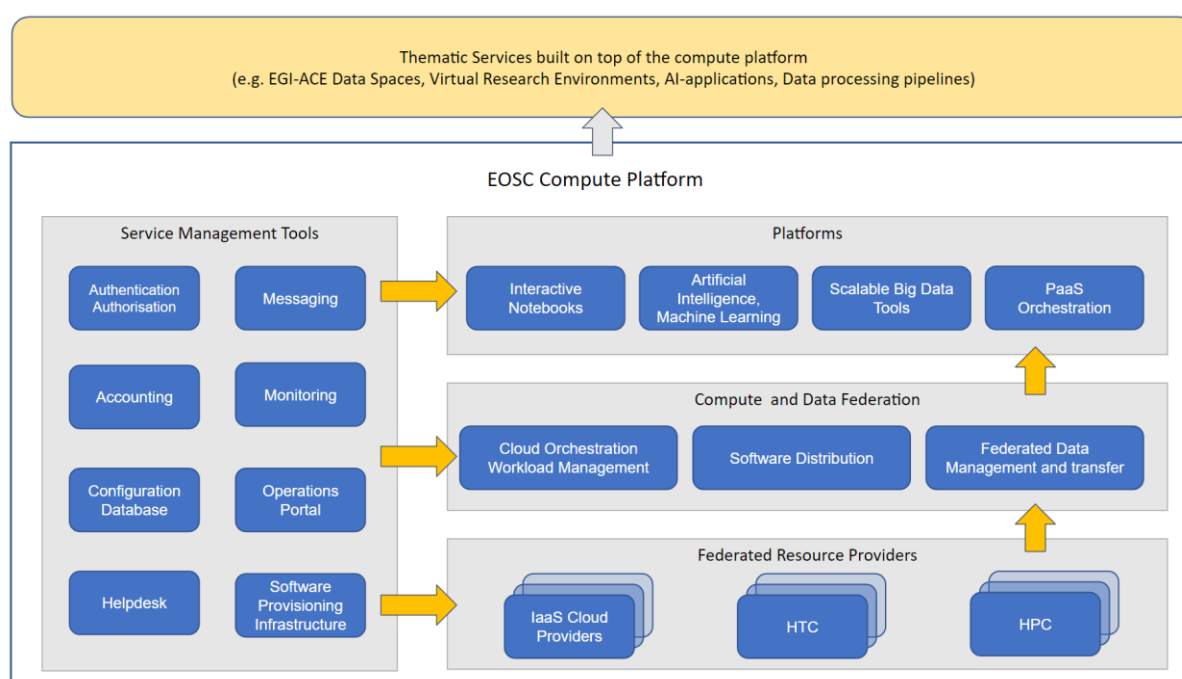


Figure 1. EOSC Compute Platform functional block diagram

At the bottom of the architecture, the Federated Resource Providers deliver a hybrid infrastructure for hosting research applications and data. Different types of providers are included in this layer:

- IaaS Cloud Providers provide access to Virtual Machine-based computing with associated Object and Block storage. These deliver a very flexible and customisable platform where users have complete control over the software and the supporting compute capacity. This flexibility of the computing platform enables the support of a variety of workloads: user gateways or portals, interactive computing platforms and almost any kind of data- and/or compute-intensive workloads.
- HTC (High Throughput Compute) provides access to large, shared computing systems for running computational jobs at scale. These allow researchers to analyse large datasets in an 'embarrassingly parallel' fashion, i.e. by splitting the data into small pieces, and executing thousands, or even more independent computing tasks

simultaneously, each processing one piece of data. HTC means the execution and management of many independent tasks over longer times.

- HPC (High Performance Compute) (to be available in 2022) supports very optimised application of machines that have a lot of interconnected processing units. with many dependent tasks that need large amounts of parallel computing along with a low latency and high bandwidth interconnection network. The integration of HPC services is currently ongoing, and based on the findings they do not seem to require new capabilities from the Service Management Tools. However, should the ongoing tests reveal lack of features to support HPC in the services, the development roadmaps will be updated.

The Compute Federation services orchestrate the execution of user workloads in the Federated Resource Providers. They exploit data locality by moving computing near data and facilitate application portability with the support of a diverse range of computing platforms (Cloud IaaS, HTC, HPC) and the interaction with software distribution tools (as VM images, container images or binaries directly). There are three services in this layer of the architecture:

- Hybrid cloud orchestration for the deployment of custom virtual infrastructure over multiple IaaS cloud backends;
- Workload Manager for the scheduling and execution of jobs in the federated resource providers (both cloud and HTC/HPC); and a
- Software distribution, for making software available at the Federated Resource Providers (e.g., as VM images).

The Federated Data services support exposing discoverable datasets and staging data into/out of the EOSC Compute Platform Cloud. The Federated Data Management services control the raw storage capacity offered by the Federated Resource Providers to deliver data products that can be transferred among the EGI-ACE providers, and between EGI-ACE providers and external data repositories. The Federated Data Management function uses the Data Transfer service to perform the transfers.

A Platforms service area provides generic added-value services for scientific communities to build thematic services for end-users (typically for researchers). The platforms rely on the existing Compute Federation and Data Federation services to access the Federated Resource Providers and deliver Interactive Notebooks, PaaS Orchestration to facilitate the deployment of complex applications, and Artificial Intelligence and Machine Learning and Scalable Big Data Tools that can be reused in several research disciplines.

Thematic Services are built by combining services of all these areas to provide capabilities for simulation, machine learning and data analytics that are tailored to the needs of a specific research domain. Thematic Services build on the overall EGI-ACE service stack by external communities and projects. Each thematic service has a custom combination of applications/tools and data specific for the given disciplines or scientific questions they address. EGI-ACE also includes such thematic services in WP5 (Called Data Spaces). Thematic Services (incl. EGI-ACE Data spaces) share the EOSC Compute Platform as a common architecture, but the rest of their setup is custom therefore they are covered in this deliverable.

The Service Management Tools pillar delivers the functionality for services of all other areas to be integrated in the Federation. They support the management activities and the operation of the whole EOSC Compute platform and integrate and interoperate with the EOSC Core that is run and further developed in the EOSC Future project. This building block comprises the following technical services:

- Authentication and Authorisation: enables the federated access to all the EGI services and manages memberships and groups in communities and Virtual Organisations (VOs)
- Accounting: to collect, store, and report the usage of services like HTC, cloud, and storage
- Configuration database: a central registry to record topology information about all the services provided by the EGI infrastructure.
- Helpdesk: the central access point for providing technical support to both the users and the operators of resource centres.
- Messaging: to support the delivery of messages between the other EGI services
- Monitoring: to monitor the status of the EGI Infrastructure and produce performance reports based on predefined metrics.
- Operations Portal: it provides tools supporting the daily running of operations of the entire infrastructure, such as Infrastructure oversight, security operations, VO management, broadcast, availability reporting.
- Software provisioning infrastructure: it provides the technical tools and support for the release process of the middleware deployed in the resource centres.

Besides them, there are also a number of non-technical services which are fundamental to sustain the daily running of the infrastructure covering management, oversight, and coordination areas, like Operations and Security Coordination and Support, Strategy and Policy Development, Training, Project Management and Planning, Community and Communication Coordination.

For what concerns this deliverable, we will focus on the roadmap of the service management tools included in the WP7.6: Accounting, Configuration Database, Helpdesk, Messaging, Monitoring, Operations Portal, and Software Provisioning Infrastructure. In the following sections in addition to a functional overview, we will describe for each one of these services the requirements implemented over the first year of the project (2021) and the ones that at the moment have been planned for the coming months and years.

All the changes to the Service Management Tools are evaluated by a team of experts prior to their deployment in production in order to assess their risks and minimize any potential negative impact to services: this is part of the Change Management Process¹ implemented by EGI Foundation which will be described in more detail in the deliverable “D7.2 Status of the SMS Processes” due by M15 of EGI-ACE project.

¹ <https://confluence.egi.eu/display/EGIPP/EGI+Change+Management+CHM>

1.2 Relationship with the EOSC Core Services

Some of the EGI-ACE Service Management Tools are also part of the EOSC Core services delivered through the EOSC-Future project²: Accounting, Configuration Database, Messaging, Monitoring, and Operations Portal.

Both the projects collect and discuss the service requirements from and with their stakeholders, and the implementation of new features within one project can be beneficial also for the other. This document covers the requirements gathered within the EGI-ACE project, anyway we would like to quickly mention, as an example, some new features that might be relevant also for the EOSC-Future project:

- the support for Storage and GPU accounting to collect the usage of new resources
- the enabling of Federated Identities and the identity linking to make it easier the access to the services
- detection of failures and related analysis of service performance trends to support the investigation on incidents and their solution

Viceversa, the primary development direction of the EOSC Core services within the EOSC-Future project is given by the High Level Technical Roadmap, defined during the proposal preparation and taking into account the requirements collected in all the major past EOSC projects and initiatives (EOSC-hub, EOSC Enhance, OpenAIRE-Advance, etc.) and the wide experience on EOSC of the 4 e-Infrastructures (EGI, EUDAT, GEANT, OpenAIRE) and the 5 science clusters (ENVRI-Fair, EOSC-Life, ESCAPE, PANOSC, SSHOC) that are the core partners of the EOSC Future project.

At the moment of writing this deliverable, we don't have particular requirements and new features coming from EOSC-Future to add to the technical roadmap of the EGI-ACE Service Management tools (mostly because they already fulfill the EOSC-Future High Level Technical Roadmap), however we expect that over the lifetime of the projects new requirements developed within EOSC-Future will become relevant also for satisfying the EGI-ACE needs.

² <https://eoscfuture.eu/>

2 Accounting Portal

2.1 Functional overview

The EGI Accounting Portal³ is a data aggregation and visualization tool for the data gathered and processed by the Accounting Repository. Its function is to present this data in a summarized form easily usable for end users in the form of web pages with tables and graphs, and also simplified interfaces for external tools that may use all or a portion of the data.

The Portal has separate views for HTC and Cloud data (Fig. 2), and allows the visualization of the data based on geographic or organizational topologies. VOs can be filtered by categories or with fine-grained controls. Each view allows the user to select a metric to display and two freely selectable variables that are used to group results, including date, country, by node, by number of processors, and much more depending on the view.

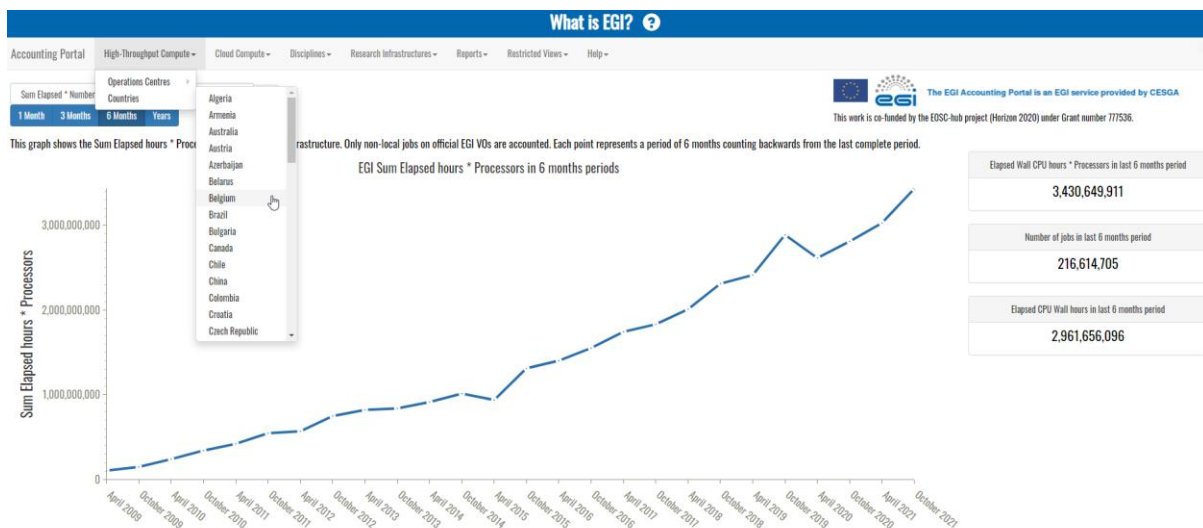


Figure 2. From the home page of the Accounting Portal users can access different views.

2.2 Roadmap

Summary of what was achieved and what's planned for the remaining time in the project.

Over the last year, some changes were applied to the underlying software in order to improve the overall service performance and to increase the security of the service. For the coming months new features will be released to introduce the Storage Accounting view in the Portal and the metrics associated to the usage of GPU resources, and to support the new benchmark type that will be used to normalise the HTC accounting records. Besides, there is a plan to work on a mechanism to exclude from the publication the inconsistent data until they are fixed.

³ <https://accounting.egi.eu/>

2.2.1 Implemented in 2021

Summary	<i>Python 3 code migration</i>
Details and impact	<p>With the release of Python 3, some parts of the code and the libraries used by the Accounting Portal, which was depending on Python 2, needed to be updated.</p> <p>This change increased the security aspects and introduced some performance improvements.</p>
Summary	<i>PHP 8 code migration</i>
Details and impact	PHP code in the backend was updated from version 5 to version 8, resulting in significant performance improvements for all users.
Summary	<i>Improve performance and responsiveness of the Portal (ongoing task)</i>
Details and impact	Further improvements in performance beyond those cited above.

2.2.2 Planned for 2022-2023

Summary	<i>Accounting data anomaly detection</i>	Expected date of completion	June 2022
Motivation and details	Filters in the Portal that warn of incoherent and clearly wrong data. It should reduce the impact of erroneous published data when collecting the periodic accounting figures.		
Summary	<i>New benchmark for HTC services</i>	Expected date of completion	May 2022
Motivation and details	A new benchmark proposed by the WLCG community will be introduced to replace the old HEPSPC06 one. The Portal will be able to present the accounting data normalised with the new benchmark.		
Summary	<i>Storage Accounting view</i>	Expected date of completion	March 2022
Motivation and details	The Storage Accounting records will be displayed on the Accounting Portal. This will make it easier to collect usage figures of the EGI storage resources.		
Summary	<i>GPU Accounting</i>	Expected date of completion	June 2022
Motivation and details	Implementation of a view displaying the GPUs usage information.		

3 Accounting Repository

3.1 Functional overview

The Accounting Repository aggregates and stores compute (serial and parallel jobs), storage, and cloud resource accounting data collected from Resource Centres of the EGI Federation (Fig. 3). Accounting information is gathered from a variety of distributed sensors, some of which are developed by the APEL team, into a central Accounting Repository where it is processed to generate summaries that are available through the EGI Accounting Portal. Depending on the use case the data may go via intermediate repositories that collate accounting data for particular regions, infrastructures or communities.

The Accounting Repository is based on the APEL software and provides interfaces for the exchange of accounting data in a standard format. The Accounting Repository has a database backend, and needs to ensure the exchange of accounting information with peer e-Infrastructures. The Accounting Portal receives and stores the resource centre, user, and user groups (e.g. Virtual Organisation/VO) level aggregated summaries generated by the Accounting Repository and provides views via a web portal.

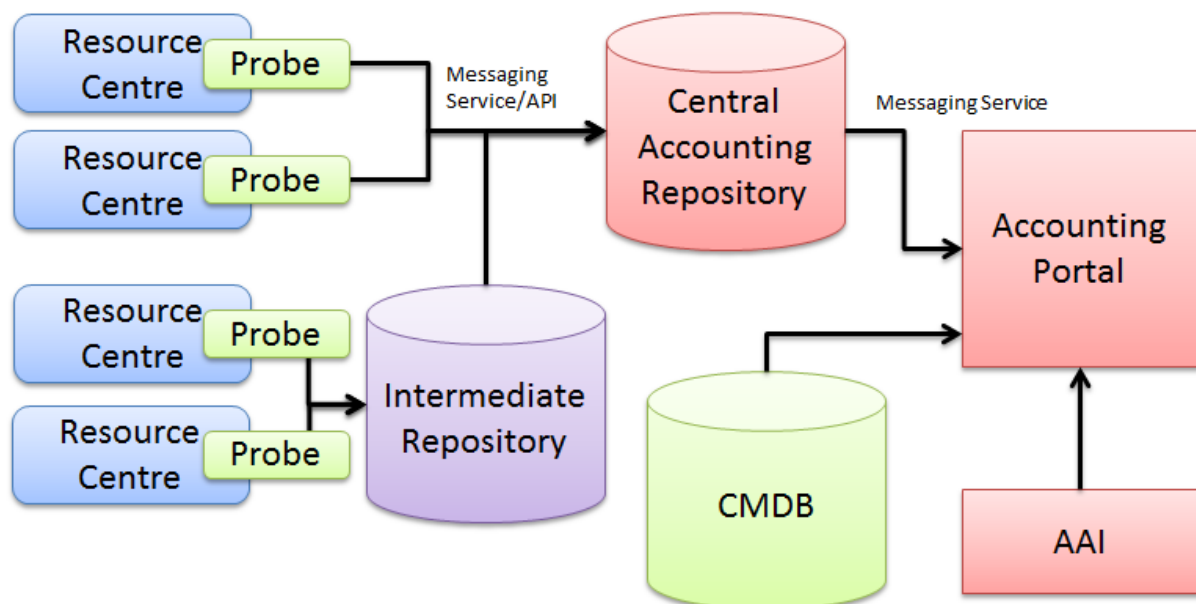


Figure 3. A diagram showing the collection of the accounting data in the EGI Infrastructure.

3.2 Roadmap

In the first year of the project it was introduced in the APEL software the support to the new messaging infrastructure delivered by ARGO Message Service for sending and collecting the accounting records. During the coming months the support to GPU accounting and to a new HTC benchmark is going to be implemented, as well as some improvements to the system that monitor the accounting data publication. Moreover, a prototype of datasets accounting will be set up.

3.2.1 Implemented in 2021

Summary	<i>APEL software regular release</i>
Details and impact	The old Message Brokers network based on ActiveMQ was dismissed in favour of new protocols and technologies provided by ARGO Message Service (AMS), so it was necessary to add the support to AMS

3.2.2 Planned for 2022-2023

Summary	<i>Sanity checks in the APEL loader to the repository</i>	Expected date of completion	February 2022
Motivation and details	Addition of some sanity checks when the accounting records are loaded to the central repository to notify about malformed data and to support the Accounting Portal in preventing the publication of incorrect figures.		
Summary	<i>Monitoring of the accounting data</i>	Expected date of completion	February 2022
Motivation and details	Introduction of a new system supporting the monitoring of the accounting records publication: new pages (last publication, accounting data synchronisation) will be deployed to support the monitoring probes.		
Summary	<i>Summarisation of the storage accounting data</i>	Expected date of completion	February 2022
Motivation and details	The process of summarising the storage accounting records needs to be performed on the Accounting Repository: this will allow the Accounting Portal to publish them.		
Summary	<i>New benchmark for HTC services</i>	Expected date of completion	May 2022
Motivation and details	A new benchmark proposed by WLCG community will be introduced to replace the old HEPSPC06 one. The Repository will be able to receive and store the accounting data normalised with the new benchmark.		
Summary	<i>GPU accounting</i>	Expected date of completion	June 2022
Motivation and details	The support of GPU accounting will allow us to collect information about the usage of GPU resources.		
Summary	<i>Full Python 3 compatibility</i>	Expected date of completion	June 2022
Motivation and details	APEL SSM and Client software will fully support Python 3 version after the decommission of Python 2		

Summary	<i>Datasets accounting</i>	Expected date of completion	December 2022
Motivation and details	Setting-up a prototype of datasets accounting working with EGI DataHub, and definition of the metrics to use.		

4 ARGO Messaging Service (AMS)

4.1 Functional overview

The ARGO Messaging Service (AMS) is a Publish/Subscribe Service, which implements the Google PubSub protocol. Instead of focusing on a single Messaging API specification for handling the logic of publishing/subscribing to the broker network the API focuses on creating nodes of Publishers and Subscribers as a Service. It provides an HTTP API that enables Users/Systems to implement message oriented service using the Publish/Subscribe Model over plain HTTP.

The Main features of the service are the following:

- **Ease of use:** It supports an HTTP API and a python library so as to easily integrate with the AMS.
- **Push Delivery:** AMS instantly pushes asynchronous event notifications when messages are published to the message topic. Subscribers are notified when a message is available.
- **Replay messages:** replay messages that have been acknowledged by seeking a timestamp.
- **Schema Support:** on demand mechanism that enables a) the definition of the expected payload schema, b) the definition of the expected set of attributes and values and c) the validation for each message if the requirements are met and immediately notify client
- **Replicate messages on multiple topics:** Republisher script that consumes and publishes messages for specific topics (ex. SITES)

Apart from all these the Messaging Service supports:

- **Argo-ams-library:** A simple library to interact with the ARGO Messaging Service.
- **Argo-AuthN:** Argo-authn is a new Authentication Service. This service provides the ability to different services to use alternative authentication mechanisms without having to store additional user info or implement new functionalities. The AUTH service holds various information about a service's users, hosts, API urls, etc, and leverages them to provide its functionality.
- **AMS Metrics:** Metrics about the service and the usage.

At the moment, AMS is consumed by accounting, the AAI Federation Registry, monitoring, and cloud services.

4.2 Roadmap

Besides the regular and ongoing work to maintain and update the AMS Library, over the past year the documentation has been improved and new metrics have been added to better control the usage and the load of the service. For the coming months, most of the focus will be on deploying a second instance of AMS to be used as backup.

4.2.1 Implemented in 2021

Summary	<i>Documentation</i>
Details and impact	The documentation ⁴ is updated and matches the user needs. This new documentation area helps the users understand how the service is working.
Summary	AMS Metrics
Details and impact	One of the main functionalities of AMS is the metrics for every resource. During this period we updated the metrics and added some metrics for project resources.
Summary	AMS library
Details and impact	This is an ongoing task. AMS Library is following the functionalities of AMS. During this period the AMS Library was updated to support the AMS authorization header (x-api-key in the header) instead of using the key in the url of the requests. It is also updated upgraded to python3

4.2.2 Planned for 2022-2023

Summary	AMS library	Expected date of completion	on going
Motivation and details	AMS Library is a simple library written in python to interact with the ARGO Messaging Service. This is an ongoing task to maintain and further develop it.		
Summary	AMS Metrics	Expected date of completion	March 2022
Motivation and details	The new metric will be to add Messages Per Hour. This metric will be counted for all different levels and resources (ex. project, topic)		
Summary	Update the documentation	Expected date of completion	February 2022
Motivation and details	New content will be added to match the different types of clients and / or users.		
Summary	1st version of a second instance of the message service	Expected date of completion	March 2022
Motivation and details	A new instance will be deployed as backup to the one provided by SRCE in order to improve the HA configuration of the service, This will be the proof of concept		

⁴ <https://argoeu.github.io/argo-messaging/>

Summary	Update the documentation	Expected date of completion	November 2022
Motivation and details	A final updated version of the documentation based on the users requirements.		
Summary	2nd version of a second instance of the message service	Expected date of completion	September 2022
Motivation and details	The updated deployment of the instance as backup to the one provided by SRCE.		

5 ARGO Monitoring Service

5.1 Functional overview

Monitoring is a key service needed to gain insights into an infrastructure. It needs to be continuous and on-demand to quickly detect, correlate, and analyse data for a fast reaction to anomalous behaviour. The challenge of this type of monitoring is how to quickly identify and correlate problems before they affect end-users and, ultimately, the productivity of their organizations. The features of a monitoring system are:

- monitoring of services,
- reporting availability and reliability,
- visualization of the services status,
- providing dashboard interfaces,
- and sending real-time alerts.

Management teams, administrators, service owners can monitor the availability and reliability of the services from a high-level view down to individual system metrics and monitor the conformance of multiple SLAs.

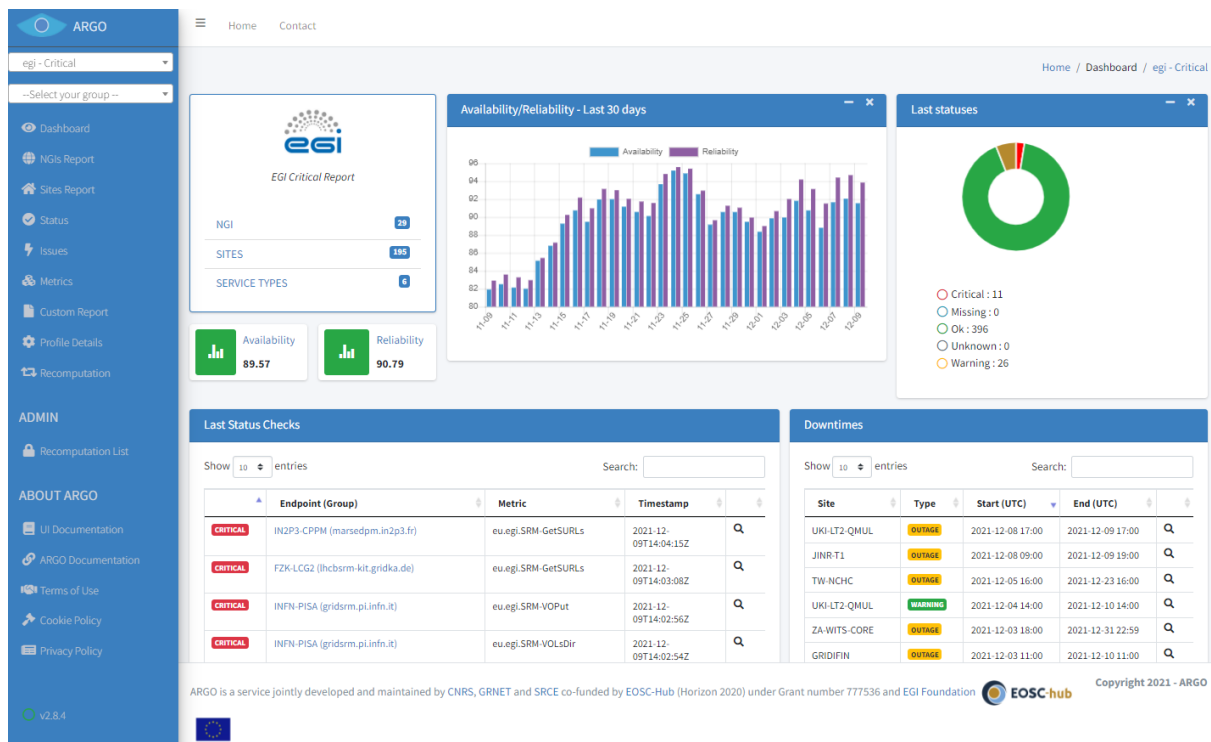


Figure 4. ARGO web dashboard

EGI provides a monitoring service based on the ARGO⁵ system. This ARGO Service collects status results from one or more monitoring engine(es) and delivers status results and/or

⁵ <https://argo.egi.eu/>, see also EGI webinar <https://indico.egi.eu/event/5496/>

monthly availability (A) and reliability (R) results of distributed services. Both status results and A/R metrics are presented through a Web UI (Fig. 4), with the ability for a user to drill-down from the availability of a site to individual services, to individual test results that contributed to the computed figure. ARGO is capable also to send notifications to the service admins in case of a failure/warning on one of the services monitored.

5.2 Roadmap

Over the past year the ongoing task of deploying new monitoring probes and updating the existing ones was carried out, according to the Infrastructure needs. In addition, the detection and the related report of services that change status too frequently was implemented. The work over the introduction of a new view for the data coming from the nagios component will be completed in the coming months. Other planned activities concern the creation of a status page for the Service Management Tools, the improvement of the notification mechanism for the services that change status, and a simplification of the way for deploying new probes on ARGO.

5.2.1 Implemented in 2021

Summary	New metrics / probes
Details and impact	Probes for fetching OIDC token, monitoring lifetime of refresh token, internal ARGO components and HTCondor-CE host certificate implemented. Provided support in developing probes for monitoring IGTF bundle on ARC-CE services and Check-in service.
Summary	Service Profile - Insights on Service Trends
Details and impact	Added detection of service flapping - situation when a service or host changes state too frequently. ARGO analyzes the status timelines in different levels, detects the flapping patterns and creates a report.
Summary	New View for raw data from nagios
Details and impact	Initial work on internal components finalized and view deployed on the development instance.
Summary	Harmonise metric names
Details and impact	This task will harmonise the metrics names so that they follow the same pattern and it is easier to understand what they do and what is their origin. Initial metric names schema is finalized.
Summary	Documentation
Details and impact	Documentation for ARGO Monitoring Service is continuously updated. UI and backend components were extensively documented in 2021. Documentation is provided here: https://argoeu.github.io/argo-monitoring .

5.2.2 Planned for 2022-2023

Summary	New metrics / probes	Expected date of completion	ongoing
Motivation and details	Providing support in developing new metrics and probes is continuous activity.		
Summary	Service Profile - Insights on Service Trends	Expected date of completion	March 2022
Motivation and details	In order to analyze trends and provide ranks we need to group some services based on the tags functionality of POEM ⁶ . These tags should be passed on to the analytics engine. The tags will be stored to the WEB API to store the tags in the metric profile.		
Summary	Status and incident communication tool / Set the problem	Expected date of completion	February /2022
Motivation and details	Status page summarising the failures and downtimes of a specific set of services. Downtime communication is only part of the story when it comes to status pages, the user may be informed about the main problem and let the service team focus on fixing the problem at hand. Incidents will be prominently displayed at the top of the status page and the status of all central services will also be displayed.		
Summary	Alerts to endpoints owners	Expected date of completion	July 2022
Motivation and details	Until today we are sending alerts to site owners. This means an admin is alerted when the group of endpoints changes status. We need to add an alert mechanism on each endpoint so as to inform the admin owner when the endpoint changes status.		
Summary	New view for raw data from nagios	Expected date of completion	February 2022
Motivation and details	Providing a view with all the data from Nagios is already in progress. We are planning to finalize it and move from development to production.		
Summary	Breaking the chains of RPM	Expected date of completion	July 2022
Motivation and details	Define a more flexible mechanism for probe deployment.		
Summary	Harmonise metric names	Expected date of completion	March 2022

⁶ https://poem.egi.eu/ui/public_home

Motivation and details	Initial metric names schema is finalized. We are planning to migrate existing metrics to a new schema and advertise it to all interested parties.		
Summary	Documentation	Expected date of completion	March 2022
Motivation and details	Finalize documentation of all ARGO Monitoring Service components.		

6 GGUS Helpdesk Service

6.1 Functional overview

The EGI Helpdesk is a distributed tool with central coordination, which provides the information and support needed to troubleshoot product and service problems. Users can report incidents, bugs or request changes using the GGUS system⁷ via its graphical interface as shown in Fig. 5.



Figure 5. The GGUS Helpdesk

⁷ <https://ggus.eu/>

The ticket will be routed to the related Support Unit (SU) assigned to handle the service support request. More information about the EGI Helpdesk are available in the EGI documentation⁸.

6.2 Roadmap

With the regular releases of GGUS over the last year several improvements to the users experience were introduced concerning the creation and the modification of a ticket, and the browsing of the tickets history. The most relevant feature request at the moment for the coming months is the improvement of the process to open a ticket to multiple sites at once.

6.2.1 Implemented in 2021

Summary	<i>Modified GUI layout</i>
Details and impact	Re-arranged some fields for better overview; added show-collapse feature to ticket history and some other areas. This improved the user's interaction with the tickets.
Summary	<i>Replaced plain text boxes by CKEditor</i>
Details and impact	CKEditor is a widely used editor that offers a lot of formatting features, e.g. source code, markdown and others. This improved the user's interaction with the tickets.
Summary	<i>Removed 4000 characters limit from diary fields, description and solution field</i>
Details and impact	In CKEditor the number of characters for payload may be heavily decreased as all formatting tags are considered by the character count.

6.2.2 Planned for 2022-2023

Summary	<i>Improvements on multiple sites tickets</i>	Expected date of completion	March 2022
Motivation and details	Sophisticated selection of sites grouped by several attributes (e.g. EGI site, OIM site, <voabc> site, etc.) will be introduced when opening a ticket to multiple sites.		

⁸ <https://docs.egi.eu/internal/helpdesk/>

7 GOCDB Configuration database

7.1 Functional overview

The EGI Configuration Database (GOCDB)⁹ is a central registry that records topology information about all sites and services participating in the EGI infrastructure. The Configuration Database also provides different rules and grouping mechanisms for filtering and managing the information associated with the resources. This can include entities such as operations and resource centres, service endpoints and their downtimes, contact information and roles of staff responsible for operations at different levels.

The Configuration Database is used by all the actors (end-users, site managers, NGI managers, support teams, and VO managers), by other tools, and by third parties middleware to discover information about the infrastructure topology.

The tool provides a web portal for inserting/editing information and a REST style programmatic interface (API) for querying data in XML (Fig. 6). Relationships between different objects are defined using a well constrained relational schema that closely resembles a subset of the GLUE2 information model¹⁰. A comprehensive role-based permissions model controls user permissions.

⁹ <https://goc.egi.eu/>

¹⁰ <http://www.ogf.org/documents/GFD.147.pdf>

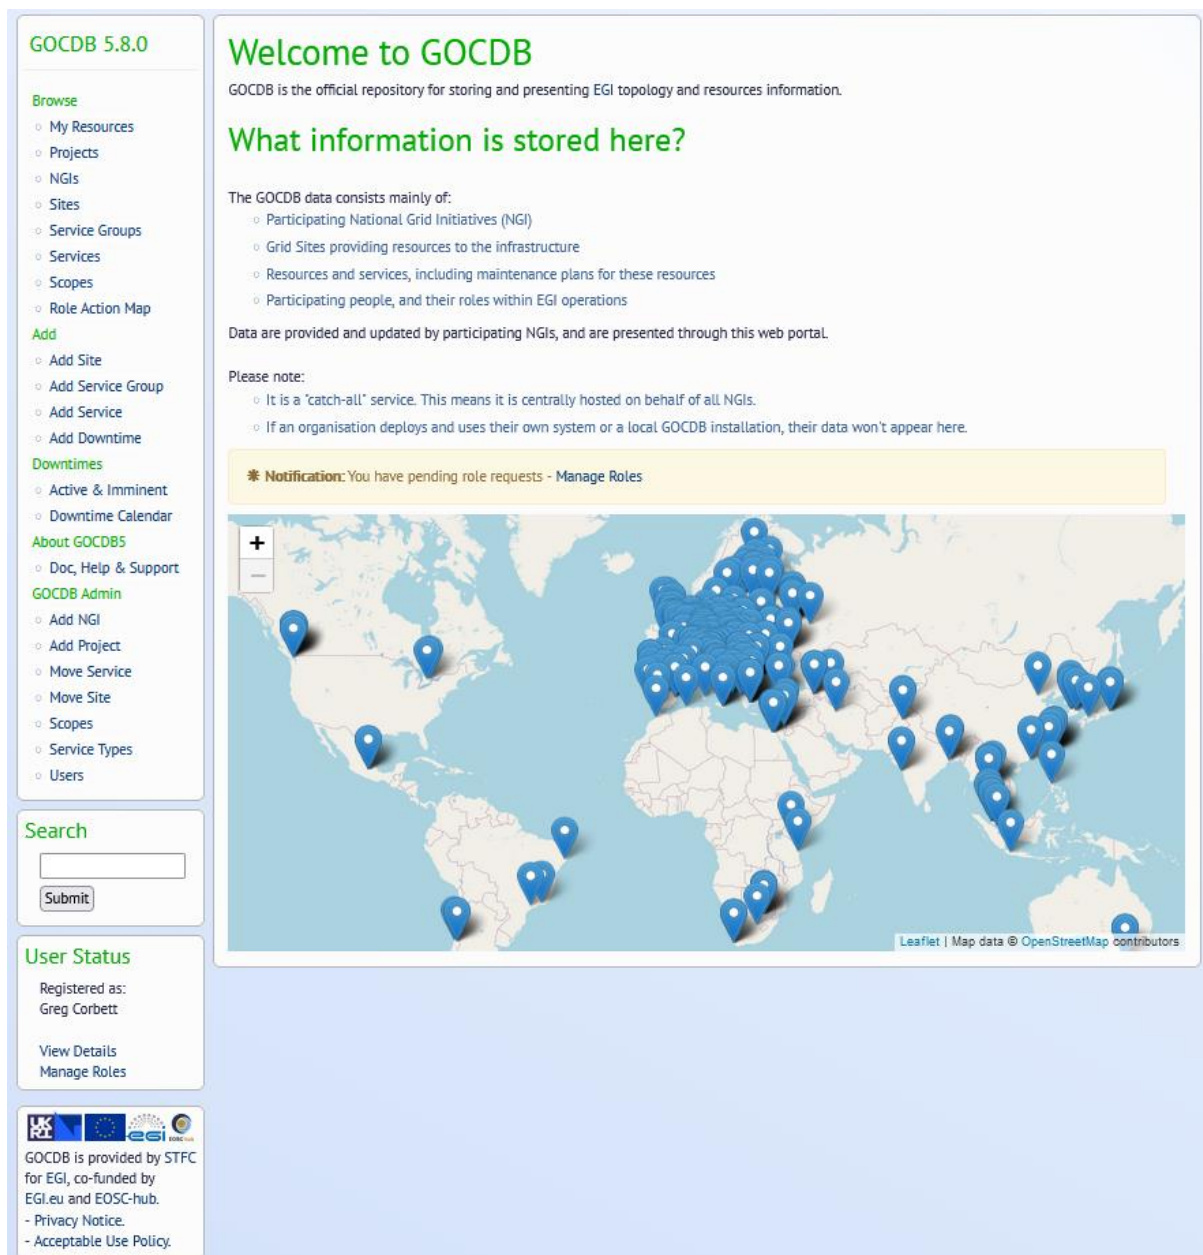


Figure 6. GOCDB web GUI

A flexible tag-cloud mechanism allows objects to be tagged with one or more 'scope-tags'. This allows resources to be tagged and grouped into multiple categories without duplication of information – this is essential to maintain the integrity of topology information across different infrastructures and projects. Different scope tags can be defined when necessary. For example, tags can be used to reflect different projects, infrastructure groupings and sub-projects. Resources can be flexibly 'filtered-by-tag' when querying for data via the Application Programmatic Interface (API). Some tags may be 'reserved' which means they are protected as they are used to restrict tag usage and prevent non authorised sites/services from using tags not intended for them.

Core objects can also be extended using a powerful extensibility mechanism that allows custom key-value pairs to be added to objects. These objects can then be flexibly ‘filtered-by-custom-property’ when selecting / querying data.

7.2 Roadmap

Regular releases over the past year introduced new features that improved the users experience, preparing at the same time the road for a migration to a new database framework planned for the coming months. Furthermore, the support of technologies based on OIDC authentication mechanisms will be enabled, in addition to the Identity Linking feature for the users.

7.2.1 Implemented in 2021

Summary	<i>GOCDDB 5.7.6</i>	
Details and impact	<p>Bug fixes providing full compatibility between GOCDDB and MariaDB like back ends.</p> <p>This will allow the service to migrate to an open source database alternative and avoid having to pay Oracle license fees. It will also simplify internal development, release and deployment procedures, by allowing us to use a common database technology across these stages of software development.</p>	
Summary	<i>GOCDDB 5.8.0</i>	
Details and impact	<p>Large feature release, including but not limited to:</p> <ul style="list-style-type: none"> • Extensions to the WriteAPI at the request of WLCG to allow Service and Service Endpoints to be updated programmatically. • Allowing multiple DNs to be added as a single extension property separated by “<” and “>”, at the request of the CERN Argus team. • Improvements to the emails sent as part of a role request so they are more personalised to the relevant users and less generic. • Maintainability improvements to the code base, making it easier to extend in future. 	

7.2.2 Planned for 2022-2023

Summary	<i>Support for integration with Indigo IAM technologies via OIDC</i>	Expected date of completion	April 2022
Motivation and details	<p>Motivated by both EGI’s and WLCG’s aim of using OIDC-based end user authentication mechanisms across the core services. This work will involve reworking the existing portal integration with the EGI Check-In service¹¹. It will also allow access to the API via non-X.509 credentials for the first time.</p>		

¹¹ <https://www.egi.eu/internal-services/check-in/>

Summary	<i>Identity linking within GOCDB across multiple IdPs.</i>	Expected date of completion	April 2022
Motivation and details	<p>Currently, each IdP a user uses to log into GOCDB produces a new account, i.e. any GOCDB roles a user has with a X.509 based account are completely separate from any GOCDB roles they have with a EGI Check-In based account. As GOCDB integrates with more IdP proxies, this problem becomes more unwieldy and unmanageable for GOCDB's users and service staff.</p> <p>As such, we are adding a self service process to associate multiple identifiers with the same GOCDB account.</p>		
Summary	<i>Database Migration to GaleraDB</i>	Expected date of completion	April 2022
Motivation and details	<p>GOCDB currently has an ORACLE database as it's backend. This is a risk to the long-term sustainability of the service due to the license fees required. The GOCDB team has avoided paying this in the past due to other larger Oracle use-cases in our organisation bearing the cost. This will not continue however as these use cases are reaching end of life.</p>		
Summary	<i>OIDC Based integration with multiple IdPs</i>	Expected date of completion	January 2023
Motivation and details	<p>Motivated by both EGI's and WLCG's aim of using OIDC-based end user authentication mechanisms across the core services. This work will allow users to log in with WLCG IAMs and other IdPs as deemed appropriate, such as the EOSC AAI.</p>		
Summary	<i>Improvements to User privacy</i>	Expected date of completion	January 2023
Motivation and details	<p>We will further restrict access to personal information within GOCDB.</p>		

8 Operations Portal

8.1 Functional overview

The Operations Portal provides VO management functions and other capabilities which support the daily operations of EGI (Fig. 7). It is a central portal for the operations community that offers a bundle of different capabilities, such as the broadcast tool, VO management facilities, a security dashboard and an operations dashboard that is used to display information about failing monitoring probes and to open tickets to the Resource Centres affected. The dashboard also supports the infrastructure oversight activities. It is fully interfaced with the EGI Helpdesk and the monitoring system through messaging service. It is a critical component as it is used by all EGI Operations Centres to provide support to the respective Resource Centres. The Operations Portal has been splitted into 2 instances in February 2021 : one instance is dedicated to EGI ACE¹² the other one is dedicated to EOSC Future¹³

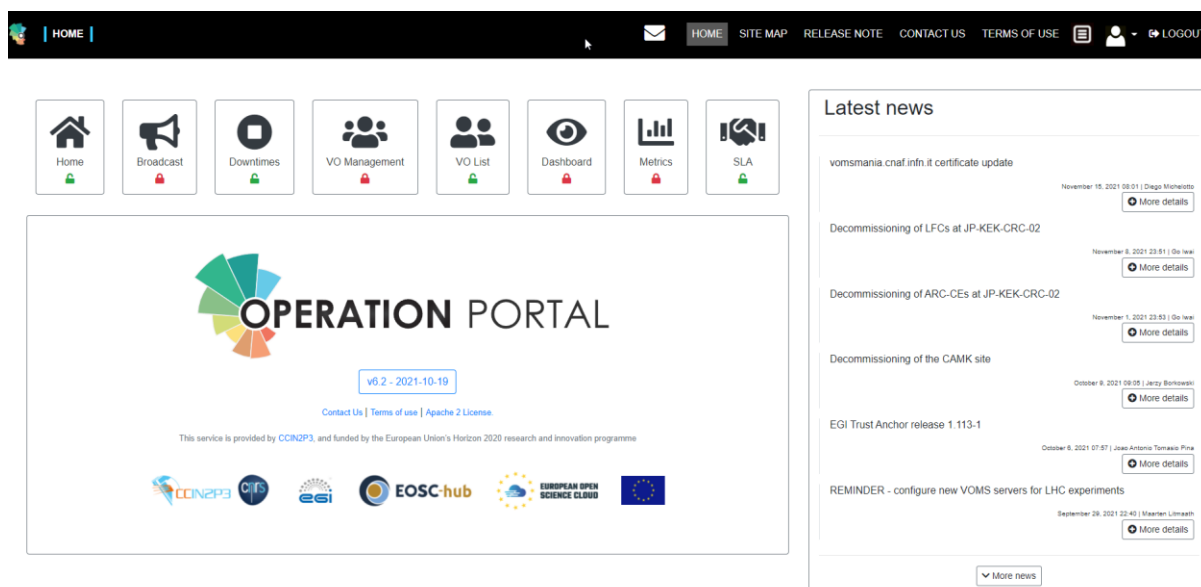


Figure 7. The Operations Portal landing page.

8.2 Roadmap

Over the last year several new features have been implemented to improve the users experience with the service and to support the infrastructure oversight activities. For the coming months there is a plan to upgrade the framework layer of the service and to improve some aspects of the VO registration process. Furthermore, the authenticated access to the service through EGI Check-in will be simplified by merging in only one profile the multiple accounts that users may own.

¹² <https://operations-portal.egi.eu/>

¹³ <https://opsportal.eosc-portal.eu/>

8.2.1 Implemented in 2021

Summary	<i>Operations portal v6.0, v6.1, v6.2</i>
Details and impact	<p>The service was migrated to the new web platform and it was improved the proxy to split accesses between the programmatic part and the user's interfaces; the access to the page listing the VOs was made public, making it easier for the users to browse this information. Some features of the VO ID Card were reviewed and the SAML protocol was replaced by Oauth for the authentication with EGI Check-in. Moreover, two new features were integrated in the Portal:</p> <ol style="list-style-type: none"> 1. Cloud badges: a page displaying reports of the availability and reliability of the EGI FedCloud sites over the last 3 months 2. Dashboard for SLA Violations: a page reporting the performance of the sites involved in the several SLAs and the related violations to the service level targets

8.2.2 Planned for 2022-2023

Summary	<i>Migration to symfony 5</i>	Expected date of completion	February 2002
Motivation and details	The framework used by the portal needs to be upgraded to a recent version. This part is also completed by a migration to the last version of Bootstrap (css framework).		
Summary	<i>Improve the registration mechanism for VO registries</i>	Expected date of completion	April 2022
Motivation and details	Since 5 years the VOs have multiplied the number of different registries: Voms, IAM, Perun, Co-Manage. We have tried to use the existing interface but this is not optimal and we need to review the process of registration into the VO ID Cards.		
Summary	<i>Improve the capture of users for VO registries</i>	Expected date of completion	June 2022
Motivation and details	Following the previous task we need also to complete the current capture based on VOMS with the new type of registries .		
Summary	<i>Manage multiple profiles</i>	Expected date of completion	September 2022
Motivation and details	With the use of the EGI Check-in system we have ensured the continuity of use of the authentication. Users in Check-in might have multiple profiles, so we have to design a new model to merge all the profiles a user owns into a single one.		

9 Software Provisioning Infrastructure

9.1 Functional overview

The EGI Software Provisioning Infrastructure is composed of services that support the management and provisioning of the software artefacts managed by the Software Provisioning team. Currently the supported repositories are the UMD (Unified Middleware Distribution), the CMD (Cloud Middleware Distribution), and the Community Repositories. The components of the infrastructure are the following:

- Repository back-end and Composer (admin-repo.egi.eu)
- Repository front-end (repository.egi.eu)
- UMD, CMD & Community repositories (repository.egi.eu)

The Repository back-end and Composer is the sub-component within the EGI Software Repository ecosystem responsible for the construction of UMD and CMD releases and their related repositories. The Repository front-end makes the produced repositories and all the required information to be available to the public (Fig. 8).

The EGI Software Provisioning Infrastructure is strongly integrated with the Application Database¹⁴ (AppDB). In this case, the AppDB acts as the backend “engine” for creating and managing the Community repositories populated through the EGI Software Provisioning Infrastructure: in this way it is possible to release with some quality control and validation tests software that is not distributed through the usual UMD channel and that is managed by third party providers.

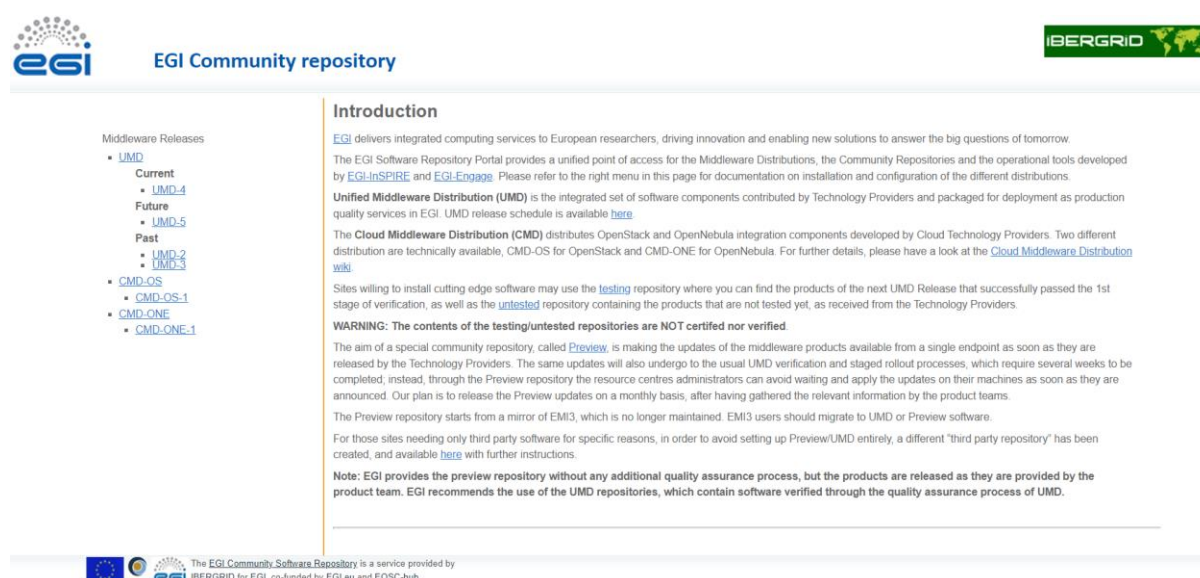


Figure 8. The Repository front-end landing page.

¹⁴ <https://appdb.egi.eu/>

9.2 Roadmap

Over the past year the migration of the service to the new provider was completed, introducing a number of changes that improved the provisioning of the service as a whole and its security. Besides, some preparatory work was performed for deploying the service on a new infrastructure planned for the coming months.

9.2.1 Implemented in 2021

Summary	<i>Resource Provider Migration</i>
Details and impact	<p>Migration of the service to the new service provider with containerization of the services.</p> <p>This migration allowed the automated deployment of the backend and frontend with the use of ansible-roles allowing operations teams to treat service provisioning and deployment in a similar fashion as developers manage the software code.</p>
Summary	<i>Front-end migration</i>
Details and impact	<p>The front-end services were based on word-press which made the component vulnerable to attacks and due to the criticality of the service a new front-end based on static pages was implemented.</p>
Summary	<i>New architecture of the service</i>
Details and impact	<p>The current production service was designed many years ago and does not support new types of artifacts like containers and has issues with integration with tools like jenkins that support continuous integration using a DevOps approach.</p> <p>So a new infrastructure based on Nexus Repository OSS it's currently being deployed together with a new front-end.</p> <p>This change will have a high impact on the development of a fully automated approach for Quality assurance allowing a faster release cycle and higher QoS.</p>
Summary	<i>Repository mirror service</i>
Details and impact	<p>Creation of a mirror for the current yum / deb repositories (mirror-repository.egi.eu)</p>

9.2.2 Planned for 2022-2023

Summary	<i>New production Infrastructure ready for testing</i>	Expected date of completion	February 2022
---------	--	-----------------------------	---------------

Motivation and details	<p>Replacement of the old infrastructure based on pearl scripts. The new infrastructure based on Nexus Repository OSS together with a new front-end will allow a fully automated approach for Quality assurance allowing a faster release cycle and higher QoS.</p> <p>Nexus Repository OSS besides the yum / deb support will also provide support to other artifacts like containers, pip, etc.</p> <p>Requirements gathering finalized for the Unified view of the front-end.</p>		
Summary	<i>First prototype of the unified view</i>	Expected date of completion	July 2022
Motivation and details	<p>The new front-end web pages will provide an unified view of all software available in the EGI Repository Infrastructure. This centralization of the information will provide easy access to users about distributed software and configuration recipes.</p>		
Summary	<i>First fully release</i>	Expected date of completion	September 2022
Motivation and details	<p>Final infrastructure tested and deployed into production.</p>		