



EOSC Hub Technical Specification

Federation Services

Accounting

Version:	1
Status:	EOSC-hub Proposal
Dissemination Level:	Public
Document Link:	https://confluence.egi.eu/display/EOSCDOC/Accounting

Abstract

The EOSC Accounting service collects, stores, aggregates, and displays usage information of HTC compute, storage space, cloud VM and data set resources. This usage data is collected from the Resource Centres of the EOSC infrastructure. Accounting information is gathered from distributed sensors into a central Accounting Repository where it is processed to generate summaries that are made available through an Accounting Portal. Depending on the use case the data may go via intermediate repositories that collate accounting data for particular regions, infrastructures or communities. EOSC resource centres can either directly publish accounting information to the EOSC central Accounting Repository or via an intermediate repository that can be related to an infrastructure (European, regional or thematic etc.). It is up to the infrastructure decide to have its own accounting infrastructure connected to the EOSC one or directly leverage the EOSC accounting infrastructure.



COPYRIGHT NOTICE



This work by Parties of the EOSC-hub Consortium is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>). The EOSC-hub project is co-funded by the European Union Horizon 2020 programme under grant number 777536.

DELIVERY SLIP

<i>Date</i>	<i>Name</i>	<i>Partner/Activity</i>
From:	Adrian Coveney Ivan Diaz Alvarez Pavel Weber Diego Scardaci	STFC CESGA KIT EGI Foundation
Reviewed by:	Heinrich Widmann William V. Karageorgos	DKRZ IASA

DOCUMENT LOG

<i>Issue</i>	<i>Date</i>	<i>Comment</i>	<i>Author</i>
v.1	27/01/2020	First release ready for public consultation	Adrian Coveney (STFC), Ivan Diaz Alvarez (CESGA), Pavel Weber (KIT), Diego Scardaci (EGI F.)

TERMINOLOGY

<https://wiki.eosc-hub.eu/display/EOSC/EOSC-hub+Glossary>

<i>Terminology/Acronym</i>	<i>Definition</i>

Contents

1	Introduction	4
2	High-level Service Architecture	4
3	Adopted standards.....	6
4	Interoperability guidelines	7
5	Examples of solutions implementing this specification	8
5.1	Procedure to integrate a service with the EOSC Hub Accounting	9

1 Introduction

The EOSC Accounting service collects, stores, aggregates, and displays usage information of HTC compute, storage space, cloud VM and data set resources. This usage data is collected from the Resource Centres of the EOSC infrastructure.

Accounting information is gathered from distributed sensors into a central Accounting Repository where it is processed to generate summaries that are made available through an 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 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. For example, by grouping resource centres in a country on specific time intervals a customized view can be generated and displayed. The databases are organized into resources record database (e.g. CPU, storage, dataset, etc), a User record database, and a topology database.

The main features of the EOSC Accounting can be grouped by two target groups.

Main features offered to the user are:

- Aggregated views of their usage wherever that usage occurred.
- Views that allow usage to be checked against allocation.

Features for resource providers:

- Provider-centric views of resource usage by users.
- Views that allow comparisons to be made between resource providers within and between regions and communities.

2 High-level Service Architecture

Describe the architecture (commented diagram) of the building block highlighting the interfaces towards the other services.

The architecture should be generic. Please, do not refer to specific service.

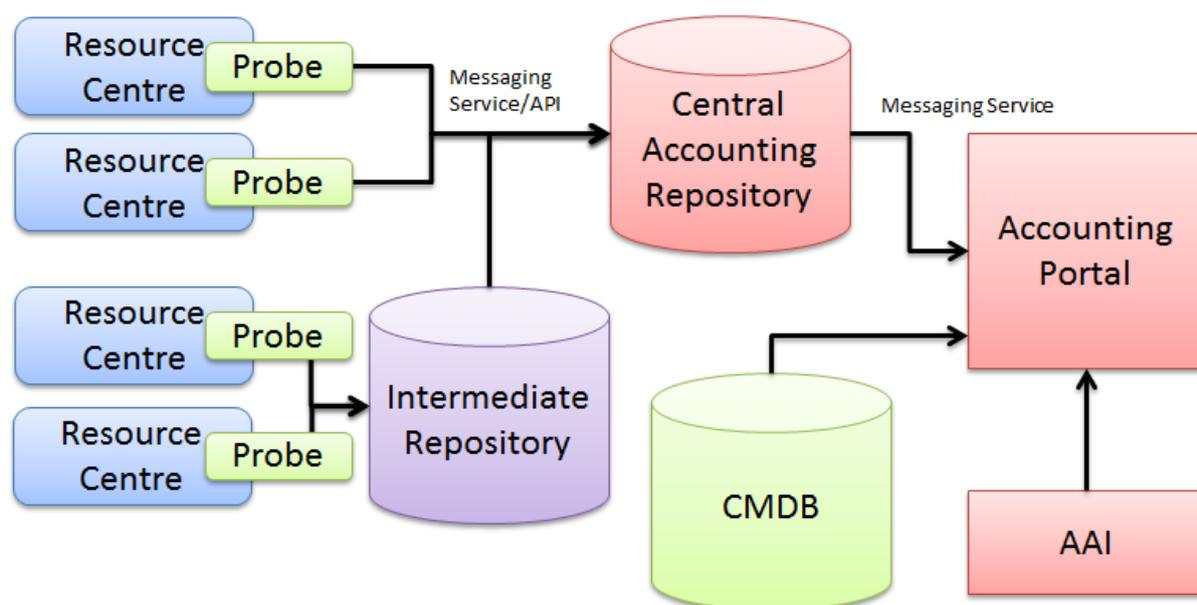


Figure 1. Components of the EOSC Accounting and their interactions

Resource centres that are providing compute or storage to the EOSC infrastructure have to implement a collector (a stand-alone script or program, or a built in function of their resource system) that gathers accounting metrics formatted into a standardised record format (see next section for details). These metrics are then transferred either via a messaging service or by being retrieved from an API to the Accounting Repository, which stores and processes the data to produce aggregations that are then sent to the Accounting Portal for display.

The Accounting Portal retrieves topology information on how resource centres relate to national infrastructures and regions from a configuration management database (CMDB) and community affiliation from the AAI service to properly organise the accounting data. Information related to groups or VOs should contain also information about scientific disciplines to allow the portal to properly classify the resource usage.

EOSC resource centres can either directly publish accounting information to the EOSC central Accounting Repository or via an intermediate repository that can be related to an infrastructure (European, regional or thematic etc.). It is up to each infrastructure to decide whether to have its own accounting infrastructure connected to the EOSC one or directly leverage the EOSC accounting infrastructure.

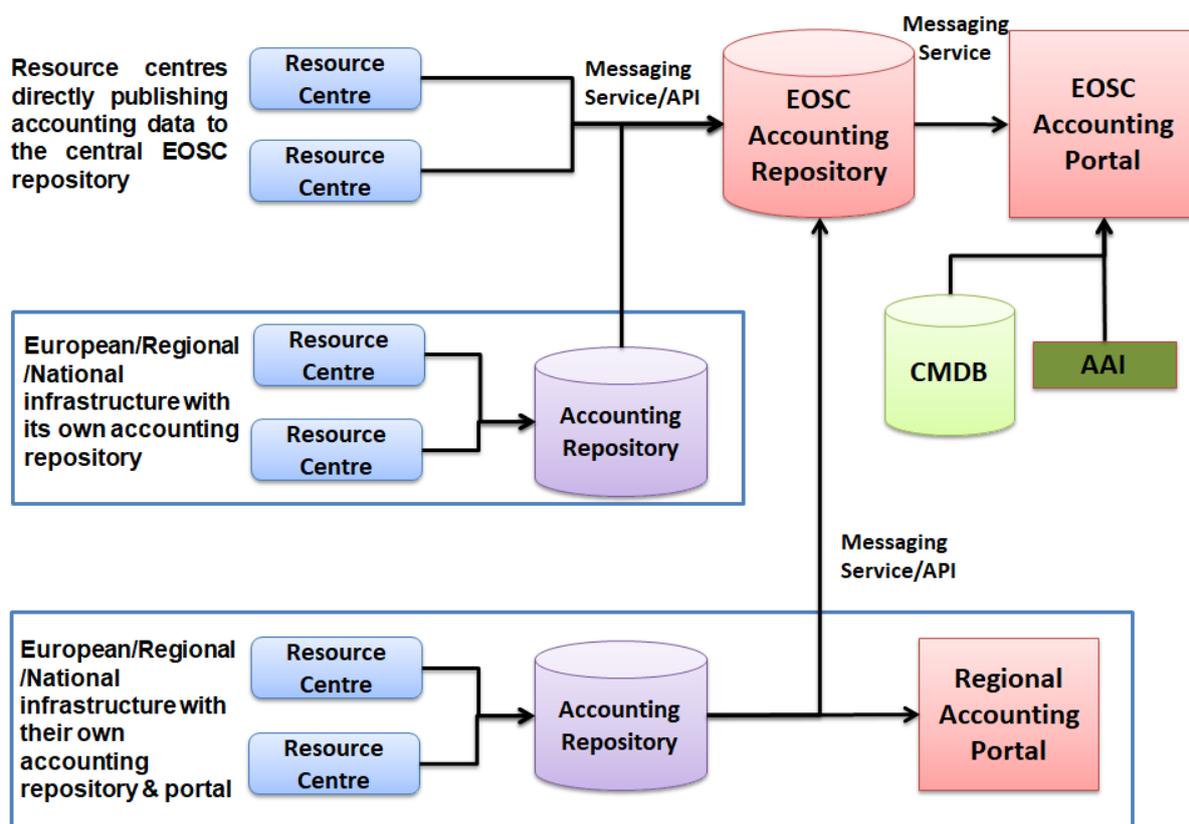


Figure 2. High-level architecture of the EOSC Accounting

3 Adopted standards

Standard	Short description	References
APEL Grid Job Usage Record	Standard used within WLCG and EGI for exchanging grid accounting metrics for individual grid jobs.	https://wiki.egi.eu/wiki/APEL/MessagingFormat#Job_Records
APEL Grid Summary Job Record	Standard used within WLCG and EGI for exchanging grid accounting metrics for aggregations of grid jobs.	https://wiki.egi.eu/wiki/APEL/MessagingFormat#Summary_Job_Records
Cloud VM Usage Record	Standard adopted by the EGI Federated Cloud for exchanging cloud accounting metrics.	https://wiki.egi.eu/wiki/FederatedCloudArchitecture#Cloud_Usage_Record
OGF StAR	Open Grid Forum standard for Storage Accounting Records, used to exchange storage space usage data.	http://www.ogf.org/documents/GFD.201.pdf

GOCDB Grid Topology	The GOCDB domain model closely resembles a subset of the GLUE 2 Grid model with additional entities.	https://wiki.egi.eu/w/images/d/d3/GOCDB5_Grid_Topology_Information_System.pdf
ARGO Messaging Service (AMS)	A Publish/Subscribe Service, which implements the Google PubSub protocol. It provides an HTTP API that enables Users/Systems to implement message oriented service using the Publish/Subscribe Model over plain HTTP.	http://argoeu.github.io/messaging/v1/

4 Interoperability guidelines

The following interoperability guidelines should be followed to connect an accounting infrastructure to the EOSC accounting infrastructure:

- Standard usage records; to be able to merge accounting data we need to have similar accounting information from the system. The table above lists the standards used.
- Either push to the ARGO Messaging Service (AMS) (<https://argoeu.github.io/guides/messaging/>) or provide an agreed HTTP API through which accounting data can be gathered.
- Topology information should follow GOCDB guidelines (https://wiki.egi.eu/w/images/d/d3/GOCDB5_Grid_Topology_Information_System.pdf), which allows other infrastructures (e.g. [OSG](#)) to coexist with separate topologies. The simplest way would be for infrastructures to register in [GOCDB](#) or even [REBUS](#), the WLCG topology platform. If this is not possible, a GOCDB compatible topology should be provided, with resource centres defined by a definite region, possible subregions and a numeric path identifier that should be consecutive, non-assigned integers, separated by dots (e.g. 1.2.3). An interface to extract the topology information should be provided.
- Metrics and units need to have a compliant format, not only in the datatype, but also the semantics must be commensurable, and the units clear. Before integrating an accounting infrastructure to the EOSC one, the provider of this infrastructure should send to the EOSC Accounting team a list of metrics and related descriptions to be published in the central EOSC Accounting Repository.
- If some of the fields contain URL pointers to metadata, these URLs must be of public access to unprivileged users, at least in a minimal form that can optionally obscure privileged information. In this way meaningful linking from the Portal is allowed.
- AAI should express group membership in a standard way following the EOSC Hub AAI interoperability guidelines (derived from the [AARC](#) guidelines).

The EOSC Accounting Repository can accept records produced by any service so long as they are in the correct format and are sent via AMS. Resource providers need to be registered in a configuration management database (e.g. GOCDB) or be individually authorised to publish via AMS.

5 Examples of solutions implementing this specification

APEL

APEL is an accounting tool that collects accounting data from sites participating in the EOSC infrastructure as well as from sites belonging to other organisations that are collaborating with EOSC. The accounting information is gathered from different collectors into a central accounting repository where it is processed to generate statistical summaries that are available through the EOSC Accounting Portal.

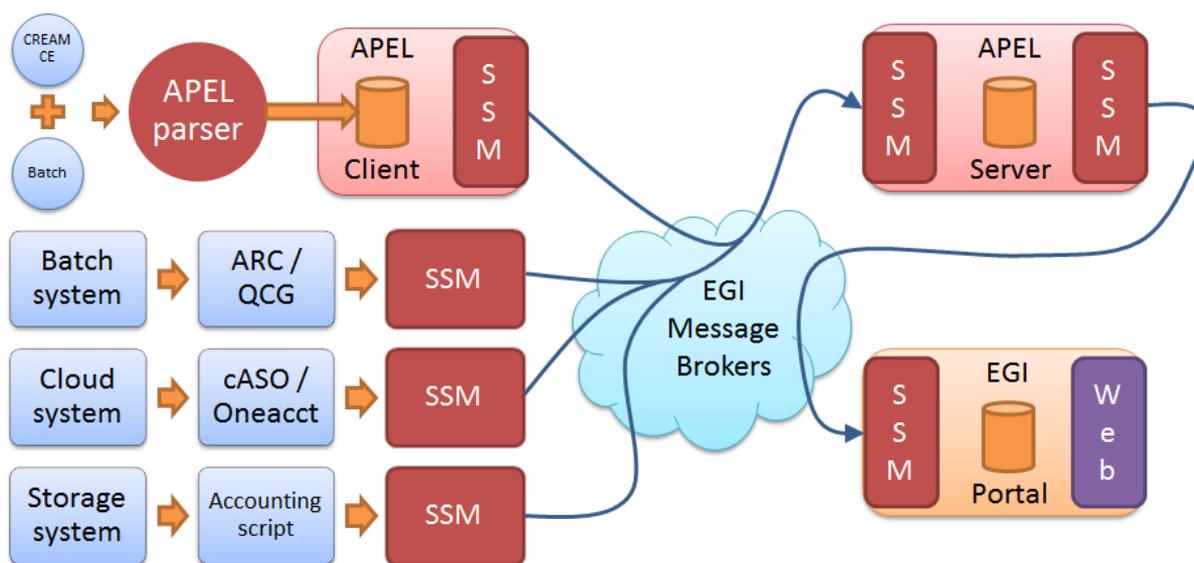


Figure 3. APEL architecture

APEL collects accounting information for compute, cloud and storage resources. Typically a site will deploy some form of accounting collector which will interact with the underlying resource provider and produce an accounting record in a supported format which is then sent to the APEL central repository via the Argo Messaging Service (AMS) and Secure STOMP Messenger (SSM, see <https://github.com/apel/ssm>). However, APEL is agnostic to the exact source of accounting data, so it is possible to set up regional APEL servers which receive the accounting data from national sites before sending a copy of the information on to the central server.

1. APEL clients (<https://github.com/apel/apel>) can run an APEL parser to extract data from a batch system and place it in their client database, or they can use third-party tools to extract

batch or cloud data. This data is then unloaded into a message format suitable for transmission.

2. APEL clients run a sending SSM to send these messages containing records via the EGI Message Brokers to the central APEL server. The messages can contain either Job Records or Summary records. This is configurable in the APEL client.
3. The central APEL server runs an instance of the SSM, which receives these messages and a “loader” processes the records in the messages and loads them into a MariaDB database.
4. A “summariser” process runs to create summaries of any Job Records received and load them in a “SuperSummaries” table along with any Summary records. This summariser runs as a cron job approximately once a day.
5. A database “unloader” process unloads the summary records into the message format to be sent on by the sending SSM via the EGI Message Brokers to the EGI Accounting Portal.

Accounting Portal

The Accounting Portal receives data from APEL and ultimately from resource centres participating in the EOSC infrastructure as well as from sites belonging to other organisations that are collaborating with EOSC. This is crossed with metadata from other sources to offer an integrated view of accounting data on the Infrastructure.

It is capable of:

- Accounting of CPU time (Normalized or not), Wall Time (Normalized or not), Number of jobs, and efficiency. Grid, Cloud and Storage support
- Grouping of resource usage by Country/Region, date, user group, usage by country.
- Grouping by infrastructure (e.g. WLCG, OSG, etc.)
- Discipline Views
- Views generated by User group Manager, User group Member, Site Admin or User views

DPMT

The Data Project Management Tool (DPMT) used within EUDAT provides an HTTP API that can be used to perform queries to retrieve usage metrics. This is done by the Accounting Repository which then stores the data alongside metrics from other systems. The Accounting Portal then transforms some of the data to enable it to be displayed in aggregated views.

5.1 Procedure to integrate a service with the EOSC Hub Accounting

The integration of a resource centre with the EOSC Accounting infrastructure requires two steps:

1. Probes to produce data in the correct format should be installed in the resource centre. The EOSC Accounting Repository will accept records formatting to the standards above. Ready-to-use probes for a large set of resource types are already available (<https://github.com/apel/apel/blob/dev/README.md#apel-parsers>).
2. Accounting records should be sent to the Accounting Repository. For sending accounting records it is recommended to use SSM to handle the interfacing with AMS, but if it is desired

to create a 3rd-party service, the ARGO AMS Library (<https://github.com/ARGOeu/argo-ams-library>) can be used to simplify the integration.

The EOSC Hub accounting team provides technical support to infrastructures and resources centre willing to be integrated in the EOSC Accounting system.