EGI-Engage

First Data Accounting Prototype

D3.8

Date	08 March 2017
Activity	WP3
Lead Partner	STFC
Document Status	FINAL
Document Link	https://documents.egi.eu/document/2968

Abstract

This report documents the release of the first prototype for dataset accounting during EGI-Engage, focused on dataset usage, which will be run as a test bed by the EGI Accounting Repository team for further improvements that will form the second prototype. A dataset is defined as a logical set of files which may exist in several places at once and to which it is possible to assign some form of persistent unique identifier. This report looks at the metrics and architecture of this prototype, which uses modified software from the APEL project that can pull usage metrics from a REST endpoint, and the testing of it against the EGI DataHub. Lastly dissemination, exploitation and future plans are presented.

The EGI-Engage project is co-funded by the European Union (EU) Horizon 2020 program under Grant number 654142 http://go.egi.eu/eng

COPYRIGHT NOTICE

The EGI-Engage project is co-funded by the European Union Horizon 2020 programme under grant number 654142.

DELIVERY SLIP

	Name	*Partner/Activity*	*Date*
From:	A. Coveney, G. Corbett	STFC / JRA1	2017-01-24
Moderated by:	Małgorzata Krakowian	EGI Foundation /WP1	
Reviewed by	Kostas Koumantaros Matthew Viljoen Peter Solagna	GRNET / PMB EGI Foundation /WP4 EGI Foundation /WP5	2017-02-02 2017-02-04 2017-02-05
Approved by:	PMB and AMB		2017-03-08

DOCUMENT LOG

Issue	*Date*	*Comment*	*Author/Partner*
v0.1	2016-10-15	Document creation	A. Coveney / STFC
v0.2	2016-10-26	Revised following internal review	A. Coveney / STFC
v0.3	2017-01-24	Revised following additional testing	A. Coveney / STFC
v0.4	2017-01-30	Internal review	D. Scardaci / EGI F. - INFN
v1.0	2017-02-01	Revised and submitted for external review	A. Coveney / STFC
v1.1	2017-02-19	Revised following external review	A. Coveney / STFC
V1.2	2017-02-27	Further reviewers’ comments addressed	A. Coveney / STFC D. Scardaci / EGI F. - INFN

TERMINOLOGY

A complete project glossary and acronyms are provided at the following pages:

- https://wiki.egi.eu/wiki/Glossary
- https://wiki.egi.eu/wiki/Acronyms

Contents

Executive Summary

This report documents the release of the first prototype for dataset accounting during EGI-Engage, focused on dataset usage, which will be run as a test bed by the EGI Accounting Repository team. A dataset is defined as a logical set of files that may exist in several places at once and to which it is possible to assign some form of persistent unique identifier and to perform dataset accounting it is assumed that this unique identifier is available. This differs from storage accounting which accounts for disk allocation and usage without concern over what data is stored or who accesses it and how often. This prototype uses software from the APEL[1] project to account for the usage of datasets. The EGI Accounting Repository runs using software from the APEL project, which collects accounting data from sites participating in the EGI infrastructure as well as from sites belonging to other Grid organisations that are collaborating with EGI.

The EGI Open Data Platform (which will provide capabilities to publish, use and reuse openly accessible data identified by PID) seemed to be the best candidate in this set and was investigated as a source of dataset accounting data. The planned design of this platform included the integration of current EGI storage services into the platform backend, which should help to hide the complexity caused by the wide range of storage systems adopted.

This prototype has the following high-level architecture: the server-side messaging component can pull dataset accounting records directly from storage systems that expose these records via a REST interface. Once the records are on the APEL server, they can be loaded and summarised. The software has the following dependency: storage providers expose a REST interface that can be queried by HTTP.

This release includes a number of changes such as an updated SSM to support fetching dataset usage records from a RESTful HTTP interface, a database schema to store dataset usage records, and support for loading dataset usage records. This satisfies a number of requirements from the EGI-Engage Accounting Repository roadmap.

The software was tested by running it and pointing a modified instance of SSM, designed to pull data from a REST endpoint, exposed by the EGI DataHub[2], to extract usage data, which demonstrated that the prototype is capable of extracting space metrics from the test space, parsing them into an OGF based message format, and then loading the data into a database.

This initial prototype will be improved by using feedback following this release to ensure it will be meet user requirements. Once Onedata has made support for unique dataset identifiers available in the EGI DataHub, the prototype will be updated to query those metrics. These improvements will then be integrated into the release of the second data accounting prototype. The optimum balance between accounting granularity and data volume will be investigated. A method for

[1] http://apel.github.io/

[2] https://datahub.egi.eu/

summarising the data and sending it to the Accounting Portal will be developed for the next prototype alongside defining Portal views on that data.

1 Introduction

This report documents the release of the first prototype for dataset usage accounting. Here, a dataset is defined as a logical set of files which may exist in several places at once and to which it is possible to assign some form of persistent unique identifier and to perform dataset accounting it is assumed that this unique identifier is available. This differs from storage accounting which accounts for disk allocation and usage without concern over what data is stored or who accesses it and how often. This prototype uses software from the APEL project to account for the usage of datasets. Storage accounting is also supported by APEL and is almost at the production level.

APEL is an accounting tool that collects accounting data from sites participating in the EGI infrastructure as well as from sites belonging to other Grid organisations that are collaborating with EGI, including OSG and NorduGrid. The accounting information is gathered from different sensors into a central accounting repository where it is processed to generate statistical summaries that are available through the EGI Accounting Portal. Statistics are available for view in different detail by users, Virtual Organisation (VO) managers, site administrators and anonymous users according to well-defined access rights.

Table 1 provides a summary of the tool covered in this release.

Table 1 - APEL tool summary

Tool name	APEL – Dataset accounting feature
Tool URL	http://apel.github.io/
Tool wiki page	https://wiki.egi.eu/wiki/Accounting_Repository
Description	EGI Core Service – The Accounting Repository collects and stores user accounting records from various services offered by EGI.
Value proposition	Support for dataset usage accounting can aid site and experiment administrators in making decisions about the location and storage of datasets to make more efficient use of the infrastructure, and to assist scientists in assessing the impact of their work.
Customer of the tool	EGI
User of the service	EGI Accounting Repository
User Documentation	https://twiki.cern.ch/twiki/bin/view/EMI/EMI3APELClient
Technical Documentation	https://twiki.cern.ch/twiki/bin/view/EMI/EMI3APELClient
Product team	STFC
License	Apache License, Version 2.0
Source code	https://github.com/apel

	https://github.com/gregcorbett/ssm/tree/receive_via_rest

The accounting team published a questionnaire to gather feedback from stakeholders on how best to implement a prototype system. In addition, communities that expressed great interest in this activity were selected for interviews to clarify their needs[3]. Considering the need, identified in our preliminary analysis, for a persistent identifier (PID) management system to implement a data accounting feature, special attention was devoted on gathering information about current usage of digital object identifiers (DOI) from DataCite, ePIC, and Handle, as well as Uniform Resource Identifiers (URI) and persistent Uniform Resource Locators (URL).

The EGI Open Data Platform (which will provide capabilities to publish, use and reuse openly accessible data identified by PID) seemed to be the best candidate in this set and was investigated as a source of dataset accounting data. The planned design of this platform included the integration of current EGI storage services into the platform backend, which should help to hide the complexity caused by the wide range of storage systems adopted.

The survey identified the most important attributes needed for meaningful dataset accounting as

- how often a dataset is accessed,
- who accessed them, and
- what transfers of the dataset occurred.

Other high priority data fields that should be included are: the different forms of user identification (such as an x.509 certificate Distinguished Name (DN) or an eduPersonPrincipleName (ePPN) attribute from a security realm); user groupings such as VO, or home-site; number of store and retrieve operations; number of files transferred; success or failure of the transfer; and the dataset identifier.

Other, medium priority data fields which should probably be accounted for include: storage system implementation, i.e. the type of storage system this data was extracted from; transfer start time and end time or duration; the source and destination IP address; and the volume of data transferred.

The outline of this deliverable is as follows: first we provide a short introduction to the components provided by the APEL project as part of this prototype. Then the high-level architecture of the tool and its components are described, along with the integrations and dependencies it has. Release notes and the results of testing for this release are provided, followed by a dissemination and exploitation plan. Finally, a selection of future developments is shown.

[3] https://documents.egi.eu/document/2674

2 Service architecture

2.1 High-Level Service architecture

Currently, APEL is capable of accounting for Virtual Machine (VM) usage and non-GPGPU grid computing usage, and storage system usage is under development. Under the current model, a provider-side parser creates APEL accounting records in a message format suitable for transmission. These records are then sent by a provider-side instance of the APEL Secure Stomp Messenger[4] (SSM) to the EGI Message Broker Network. A server-side SSM then pulls down the data from the Message Broker Network, where it is loaded into the APEL server, summarised and finally these summaries are sent to the EGI Accounting Portal, again via SSM and the Message Broker. This is shown in Figure 1.

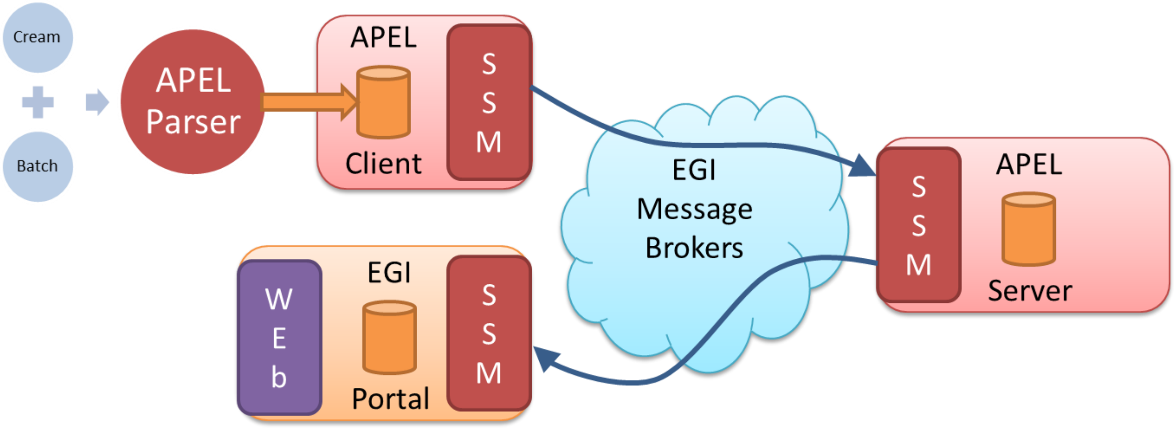

Figure 1 - APEL components and their interactions. Components in red are provided by the APEL project

For dataset usage accounting, the SSM component was updated to add support for retrieving records from a RESTful HTTP interface. This means that the server-side SSM instance can pull dataset accounting records directly from storage systems that expose these records via a REST interface. This new architecture is shown in Figure 2. This allowed the APEL software to integrate with EGI DataHub[5] which only exposes a REST interface for obtaining accounting metrics and so does not require the use of a message broker. Transfers are performed over HTTPS to ensure the integrity and security of the data.

[4] https://github.com/apel/ssm
[5] https://datahub.egi.eu/

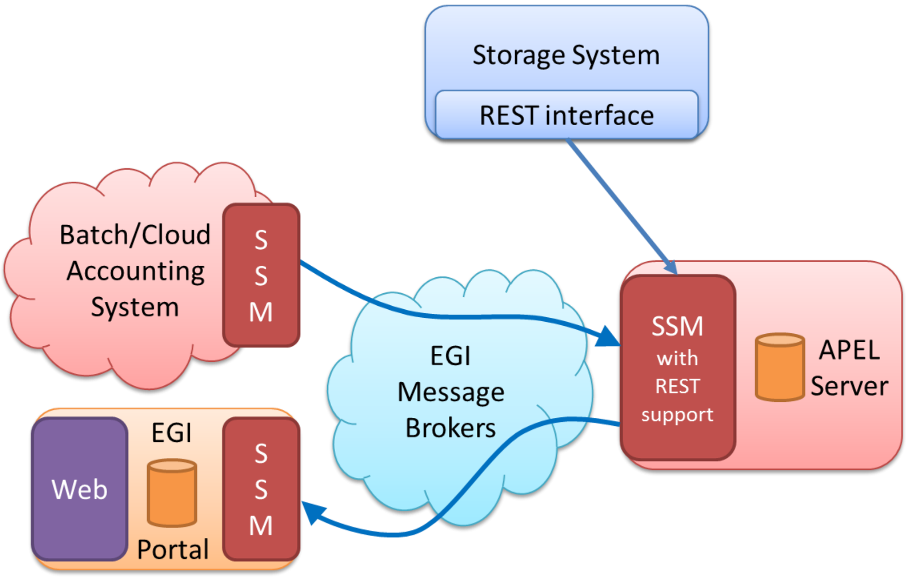

Figure 2 - New APEL architecture that supports pulling data from a REST interface.

However, the updated SSM with REST support is being developed with the new the ARGO Messaging Service[6] (AMS) in mind, as it will also make use of HTTP interfaces. The prototype SSM provides functions to pull data from (and push data to) any given HTTP endpoint.

Research into how APEL will interact with the AMS is still ongoing, with the possibility of using an off the shelf module for communication. If this happens, the existing methods in the dataset accounting SSM prototype could be transferred to the parser, making it more akin to the cloud accounting collectors: cASO[7] and OneacctExport[8].

Separate extensions to the SSM software will be needed for each use case (i.e. pulling data from Onedata[9] based systems like the EGI DataHub, or messages from AMS) as Onedata and other storage systems will possibly differ in how they present dataset accounting metrics. Pulling messages from the AMS will require parsing the received message to extract the accounting records.

Pulling the records directly will require the REST interface to be available at the same time as the SSM initiates the pull. This differs from sending records via the Message Broker, which acts as a buffer, storing records for multiple days, allowing asynchronous publishing and consuming.

Once the records are on the APEL server, they will be loaded and summarised. The portal is then updated via the Message Broker as happens now.

[6] http://argoeu-devel.github.io/messaging/v1/
[7] https://caso.readthedocs.io/
[8] https://github.com/EGI-FCTF/oneacct_export/
[9] https://onedata.org/

2.2 Integration and dependencies

For existing types of accounting, communication between clients and the central APEL server is via the EGI Message Broker network using the APEL SSM package. The SSM can be configured to send or receive messages. Where the messages are destined for is controlled by the queue, which is set in the SSM configuration.

For this dataset usage accounting prototype, the central APEL server uses an updated SSM with support for pulling data directly from RESTful HTTP interfaces so that it can interact with EGI DataHub. However, the next release of the prototype will provide again dataset accounting records via the messaging system, profiting on the new HTTP based interface of the messaging infrastructure.

The central APEL server can use the EGI service registry (GOCDB[10]) to get a list of endpoints so that only data from endpoints correctly defined in GOCDB is processed. This feature is not currently used by this prototype, but it is considered a mandatory requirement to integrate the new feature into production in the future.

SSM can be configured to get a list of message brokers from the EGI information system (querying a BDII) or it can be pointed directly at a message broker.

2.3 Record metrics

Table 2 shows an outline of the metrics that were proposed for performing dataset usage accounting. They are intended as an extension to the Open Grid Forum (OGF) Usage Record version 2 (UR-2.0)[11], but the final implementation for Dataset Usage Accounting may need to be quite different once more feedback is received and use cases developed. For example, the format of the Dataset field should probably not be mandated in the accounting record, and should just provide optional fields to specify which replica the record refers to. This prototype supports a subset of these metrics but can easily be extended.

Table 2 - Dataset accounting metrics

	Key	*Type*	*Description*
Record Identity Block	Resource provider	string	Resource provider at which the resource is located (e.g. GOCDB sitename)
Subject Identity	GlobalUserID	string	e.g. X.509 certificate DN / EGI unique ID (from Checkin service)

[10] http://goc.egi.eu/

[11] https://www.ogf.org/documents/GFD.204.pdf

Block	GlobalGroupId	string	e.g. VO
	GlobalGroupAttribute	string	e.g. VO Group and/or Role
	ORCID	string	ORCID iD of the user
Dataset Usage Block	Dataset	string	unique identifier such as a PID / DOI
	AccessEvents	integer	Number of read and write operations
	Source	IP address / other	Source of transfer at resource provider
	Destination	IP address / other	Destination of transfer
	StartTime	ISO 8601 timestamp	Start time of transfer
	Duration	ISO 8601 duration	Duration of transfer
	EndTime	ISO 8601 timestamp	End time of transfer
	TransferSize	integer	Bytes transferred
	HostType	string	Storage system Type
	FileCount	integer	Number of files accessed
	Status	string	Success / failure / partial transfer

The association of the ORCID of a user with a dataset usage record was not covered in the initial survey and was suggested at a later stage of the development process with the aim to help linking datasets to research publications. However, additional personal information should only be collected if there is a clear need and there is agreement between stakeholders. For next release, we will consider the option to remove the ORCID identifier from the basic dataset usage record and gather it from third party services (e.g. from the CheckIn service or directly from ORCID) only when needed.

At a minimum, the number of times that a dataset is accessed and who that dataset belongs to should be recorded by the storage system so that the accounting system can retrieve that information and fill in the Record and Subject Identity blocks, and the "Dataset" and "AccessEvents" fields. If possible, the origin of these access events (including who performed it) would also be recorded. This would then cover the attributes that were considered most important in the survey. It is yet to be decided if this would form the set of mandatory fields. Specific use cases would help to clarify this.

Depending on how detailed the accounting data is, a method for aggregating this information should be created so that the volume of accounting data does not become unmanageable. This applies especially to the rest of the fields that relate to transfer operations as getting information for these fields could require quite a fine-grained approach to the usage data.

3 Release notes

These are the changes included in this prototype release of the APEL software.

- Updated SSM to support fetching dataset usage records from a REST interface.
- Added a database schema to store dataset usage records.
- Added support for loading dataset usage records.

3.1 Requirements covered in the release

Requirements from the EGI Engage Accounting Repository roadmap[12] that are covered in this release include the following.

- 3.1.5 – Data Accounting: proof of concept – Feedback and requirements were captured using the survey reported in M3.2[13]. These were used to produce a draft set of metrics that could be extracted from a storage system to perform dataset accounting. This draft was then used to gather further feedback.
- 3.1.9 – Data Accounting: minimal requirements implemented – All the feedback and requirements collected in the earlier stages were used to produce a prototype version of a dataset accounting system, the release of which is presented in this report.

4 Result of testing

As mentioned previously, the EGI Open Data Platform seemed to be the best source of dataset accounting metrics as it is planned to be integrated with the current EGI storage services, which reduces the number of storage systems that the dataset accounting system needs to integrate with.

The software was tested by running it and pointing a modified instance of SSM at an instance of Onedata to extract usage data. A test space was created by an EGI DataHub developer containing a test image. During the test, this known “spaceID” was used to query the Onedata REST API[14] hosted at the EGI DataHub for the quota metric of the space (e.g. https://datahub.plgrid.pl/api/v3/oneprovider/metrics/space/1I8DOQUXXiezOAcTpAewz40HVNzy-Sr2mlBZZtEmpA?metric=storage_quota&step=1m) and a response like the following was received:

[12] https://wiki.egi.eu/wiki/TASK_JRA1.3_Accounting#Accounting_Repository

[13] https://documents.egi.eu/document/2674

[14] https://onedata.org/docs/doc/advanced/rest/index.html

```
{
"rrd":{
    "meta":{
        "step":2678400,
        "start":1454371200,
        "legend":[
            "space ZgaqavtmWKV8O5-KXrpfxCXD9UvL_wbvEe0VgeZM1I; metric
storage_quota; oneprovider ID HOe-D_aZvrpggVyhnTkmBc9czucj19nQ3z-
NOJQKew; storage_quota[bytes]"
        ],
        "end":1486512000
    },
    "data":[[null],[null],[null],[null],[null],[null],
            [null],[null],[null],[null],[null],[null]],
    "about":"RRDtool graph JSON output"
},
"providerId":"HOe-D_aZvrpggVyhSnTkmBc9czucj1nQ3z-NOJQKew"
}
```

In future, this would be extended to first determine an exhaustive list of “spaceID”s and then to query each one programmatically. Although this means that the central server for the Accounting Repository has to do more work, rather than passively accepting records, it allows control of when data is retrieved to lie with the central server. More testing would be required to determine if this would be a scalable solution for the potential size of the EGI DataHub.

The returned data was then parsed into a message format based on the OGF Usage Record to give the following:

```
<?xml version="1.0" encoding="UTF-8"?><ur:UsageRecords
xmlns:ur="http://eu-emi.eu/namespaces/2017/01/datasetrecord">
<ur:UsageRecord>
      <ur:RecordIdentityBlock>
            <ur:RecordId>"host.example.org/ur/1485278427</ur:RecordId>
            <ur:CreateTime>2017-01-24T17:20:27.942629</ur:CreateTime>
            <ur:ResourceProvider>HOe-D_aZvrpggVyhSnTkmBc9czucj1nQ3z-
NOJQKew</ur:ResourceProvider>
      </ur:RecordIdentityBlock>
      <ur:SubjectIdentityBlock>
      </ur:SubjectIdentityBlock>
      <ur:DataSetUsageBlock>
            <ur:StartTime>2016-02-02 00:00:00</ur:StartTime>
            <ur:EndTime>2017-02-08 00:00:00</ur:EndTime>
      </ur:DataSetUsageBlock>
</ur:UsageRecord>
</ur:UsageRecords>
```

The message was then saved for future loading, as currently also happens with messages received via the message broker network. The message was then loaded into a database by starting a separate loader process with its own configuration file, modified to support the loading of this new format.

Onedata has only just developed support for unique dataset identifiers, so that feature was not ready for testing. However, it was still possible to test the new method of obtaining accounting information via a REST interface and the method to transform the received data into a format suitable for loading into a database. This test demonstrated that the prototype is capable of extracting space metrics from the test space, parsing them into an OGF based message format, and then loading the data into a database.

5 Dissemination and exploitation plan

Name of the result	Prototype dataset usage accounting system
DEFINITION	
Category of result	Software & service innovation
Description of the result	This prototype system extends the types of usage accounting that the EGI Accounting Repository can perform by adding features to support dataset usage accounting.
EXPLOITATION	
Target group(s)	RIs, international research collaborations, storage providers
Needs	Provide sufficient information about the location and storage of datasets to make more efficient use of computing infrastructures. Enable scientists to assess the impact of their work.
How the target groups will use the result?	With the right information on dataset usage, a dataset provider (i.e. an infrastructure provider like EGI) could create multiple replicas of a dataset if it is requested many times, and a scientist can know how many people have accessed the dataset created with their research.
Benefits	Demonstrate the potential that dataset usage accounting has to aid in fulfilling the needs above. Allow the Accounting Repository team to gather more specific feedback on dataset accounting and to identify any potential issues that will need to be overcome in future.
How will you protect the	Open source license (Apache License, Version 2.0)

results?	
Actions for exploitation	The prototype will be run by the Accounting Repository team as a test bed for future developments of dataset usage accounting. Selected resource providers will be asked to make REST endpoints available so that the Accounting Repository can extract dataset accounting records to further test the prototype. Feedback will be solicited from potential users of dataset accounting on how useful the current prototype is and what new features they would like to be included in the future. The software will be made available in a public repository.
URL to project result	https://github.com/gregcorbett/apel/tree/dataset_accounting https://github.com/gregcorbett/ssm/tree/dataset_accounting
Success criteria	Feedback from the customers that will become inputs for a more developed second prototype that regularly collects dataset accounting
DISSEMINATION	
Key messages	Test system for dataset accounting can be made available so that feedback can be gathered.
Channels	Operations Management Board meetings, EGI Engagement channels, Competence Centres
Actions for dissemination	Present results at an OMB and solicit feedback on prototype
Cost	N/K
Evaluation	Quality of feedback

6 Future plans

This initial prototype will be improved by using feedback following this release to ensure it will meet user requirements. Once Onedata has made support for unique dataset identifiers available in the EGI DataHub, the prototype will be updated to query those metrics. These improvements will then be integrated into the release of the second data accounting prototype (D3.15) as shown in the EGI Engage Accounting Repository roadmap[15]. Additionally there will be a report on data accounting (D3.13) to record what has been learnt from the development of, and feedback received for, this data accounting prototype.

As discussed in the record metrics section, there are different levels of detail that could be specified for dataset accounting records, depending on what action trigger the update of the

[15] https://wiki.egi.eu/wiki/TASK_JRA1.3_Accounting#Accounting_Repository

dataset metrics. This should be investigated to determine the optimum balance between accounting granularity and data volume.

Although there is no dependency on the EGI Message Broker network for fetching data with the current prototype, changes to the SSM will be required to integrate the software with the new ARGO Messaging Service for sending data to the Accounting Portal. Summarising the data and sending it to the Accounting Portal is currently not supported, so this should be investigated for the next prototype alongside working with the Portal to define views on that data.