



EGI-Engage

Report on Data Accounting

D3.14

Date	06 May 2017
Activity	WP3
Lead Partner	STFC
Document Status	FINAL
Document Link	https://documents.egi.eu/document/3025

Abstract

This report looks at the current state of dataset accounting within EGI-Engage, a new feature of the EGI accounting system that will enable storing information on dataset usage such as who has accessed datasets, how often a dataset is accessed, the data transfers, etc. This feature should enable Resource Centre and research community administrators to make decisions about the location and storage of data sets to make more efficient use of the infrastructure, to report on data usage to the data owners, data providers and funding agencies, and to assist scientists in assessing the impact of their work. The design of this new feature has been led by the users' requirements collected in the first part of the project, shortly summarised in this document, along with the accounting metrics that were derived from those requirements. In the first experiments, which culminated with the release of the first data accounting prototype, the EGI Accounting Repository has been integrated with the data provider Onedata, the underlying technology powering the EGI Open Data platform and EGI DataHub, as exemplar of a generic data provider.



This material by Parties of the EGI-Engage Consortium is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

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



This work by Parties of the EGI-Engage Consortium is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>). The EGI-Engage project is co-funded by the European Union Horizon 2020 programme under grant number 654142.

DELIVERY SLIP

	<i>Name</i>	<i>Partner/Activity</i>	<i>Date</i>
From:	A. Coveney	STFC / JRA1	2017-01-24
Moderated by:	M. Krakowian	EGI Foundation / NA1	
Reviewed by	J. Marco I. Haggstrom M. Viljoen	CSIC / SA2 EISCAT / SA2 EGI Foundation/ JRA2	2017-03-17
Approved by:	AMB and PMB		2017-05-03

DOCUMENT LOG

<i>Issue</i>	<i>Date</i>	<i>Comment</i>	<i>Author/Partner</i>
v0.1	2017-01-17	Document creation	A. Coveney / STFC
v0.2	2017-02-23	Revision	A. Coveney / STFC
v1.0	2017-03-10	Final revision	A. Coveney / STFC D. Scardaci / EGI F. -INFN
v1.1	2017-03-27	External review revision	A. Coveney / STFC
v1.3	2017-04-13	Further external review revisions	A. Coveney / STFC

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

1	Introduction.....	6
2	User Requirements.....	7
2.1	WLCG requirements.....	7
2.2	Example use case.....	8
2.3	Summary.....	8
3	Dataset Usage Metrics.....	9
4	Supported Storage Solutions.....	11
4.1	Onedata.....	11
4.1.1	Integration with the Accounting Repository.....	11
4.1.2	Metric collection example.....	12
4.2	Other data management systems.....	13
4.2.1	EUDAT.....	13
4.2.2	ICAT.....	14
4.2.3	iRODS.....	14
5	Plan for Exploitation and Dissemination.....	15
6	Future Work.....	16

Executive summary

This report looks at the current state of dataset accounting within EGI-Engage, a new feature of the EGI accounting system that will enable storing information on data set usage such as who has accessed a dataset, how often it is accessed, the data transfers, etc. 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 uses it and how often. Dataset accounting should enable site and experiment administrators to make decisions about the location and storage of datasets allowing more efficient use of the infrastructure, to report on data usage to data owners, data providers and funding agencies, to assist scientists in assessing the impact of their work and indirectly, to promote a culture of research data sharing and the development of altmetrics.

The design of this new feature has been led by user requirements collected in the first part of the project. From an initial survey¹, it was shown that some data centres already use persistent identifiers (PID). It identified the most important attributes needed for meaningful dataset accounting are requested to report on how often a dataset is accessed, who accessed it, and what data transfers occurred. Following the survey, a proposal (detailed in this report) was created for the metrics that could be extracted to perform dataset usage accounting.

In the first experiments, which culminated in the release of the first data accounting prototype, the EGI Accounting Repository has been integrated with the data provider Onedata, the underlying technology powering the EGI Open Data platform and DataHub², as an example of a generic data provider. Onedata is a global data management system, providing easy access to distributed storage resources, supporting a wide range of use cases. It provides a REST API which can be used to extract space and user metrics. Currently, it is not possible to extract metrics based on PIDs such as digital object identifiers (DOI), but this is a feature that is being added to support the Open Data Platform. To account for the usage of datasets, the prototype uses software from the APEL³ project, the underlying technology of the EGI Accounting Repository which collects accounting data from sites participating in EGI and other infrastructures. The software has been modified to support fetching dataset usage records from the REST API that Onedata provides.

Future work includes ironing out any problems with the integration and making use of PIDs once they are available in Onedata so that a second prototype could be produced, and improving the definition of the dataset usage metrics so that they could be standardised. The data accounting record is independent from the storage technology used to host the data. For this reason, the data

¹ <https://documents.egi.eu/document/2674>

² <https://datahub.egi.eu/>

³ <http://apel.github.io/>

accounting record can be used in a infrastructure environment with heterogeneous storage systems.

1 Introduction

This report looks at the current state of dataset usage accounting within EGI-Engage. 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 uses it and how often. Storage accounting is supported by a separate part of the Accounting Repository and is almost at the production level.

Accounting for dataset usage should enable site and experiment administrators to make 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.

After a first phase of the activity focused on requirements gathering (described in the next section) and on the definition of a data accounting usage record, a prototype dataset accounting system was created using software from the APEL project. APEL is an accounting tool that collects accounting data from sites participating in the EGI and WLCG infrastructures 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.

This report summarises the status of the activities, presenting: the user requirements that were captured from an initial questionnaire and interviews with relevant communities, the dataset metrics that were developed from those requirements, and the currently supported storage systems and how support is integrated into APEL and the EGI Federation. Finally exploitation, dissemination and future plans are shown.

2 User Requirements

The accounting team published a questionnaire to gather feedback from stakeholders on how best to implement a prototype system. In addition, communities that expressed the most interest in this activity were selected for interviews to clarify their needs. In general, the needs for this new feature for the EGI accounting system clearly emerged from the communities. There was a wide range of storage systems that were possible sources of information for dataset accounting, and with no consistent approach across the whole infrastructure to recording dataset usage. A de-facto data set storage technology or access method has not emerged; nor has a de-facto accounting record for datasets. Such an outcome was considered a good motivation to develop the data accounting prototype. More details on such analyses are available in M3.2⁴.

Considering the need, identified in the 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 DataCite⁵, ePIC, and Handle, as well as Uniform Resource Identifiers (URI) and persistent Uniform Resource Locators (URL). The survey identified the most important attributes needed for meaningful dataset accounting as

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

Other high priority data fields that should be included are: user identification (in whatever different forms are used in EGI, such as an x.509 certificate Distinguished Name (DN), an eduPersonPrincipleName (ePPN) attribute from a security realm, or the EGI unique identifier (UID) released by the CheckIn service); 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.

2.1 WLCG requirements

In the area of data accounting, the WLCG is mainly concerned with the optimisation of storage space and minimising the storage of data that is infrequently read, and so is looking at doing data

⁴ <https://documents.egi.eu/document/2674>

⁵ <https://www.datacite.org/>

accounting in a way that is much more customised for their use case. The WLCG Computing Resources Scrutiny Group asks all the experiments for data popularity information in a common form to monitor the efficiency of disk use. They split data into datasets according to creation data, to identify them as a group, and count the number of accesses over various periods of time (the previous 3, 6 and 12 months). The reporting aggregates data from a large number of sites, although the data available to sites is more fine-grained.

2.2 Example use case

The use cases for dataset accounting foreseen by one community, ELIXIR, cover both the optimisation and impact aspects of data accounting. Their use cases are: making replication decisions, i.e. moving data closer to where it is used and assessing dataset impact by measuring how many times datasets are reused and thus how popular they are. Therefore, the users of dataset accounting data would be site administrators, virtual organisation (VO) administrators, and dataset administrators.

In this case, Onedata⁶ would provide access to the storage resources where a user can upload and publish datasets. The dataset accounting system will collect usage information for those datasets in the background. Then the dataset usage can then be inspected by an administrator accessing a dataset accounting portal to see which datasets are being accessed most and if additional replicas are required to provide a sufficient level of access.

2.3 Summary

The space optimisation and impact assessment types of data accounting, contrast to a certain degree. Both need to store information on access events, but they differ in how they group and identify data. For example, the impact assessment types of data accounting, which looks at the reuse of datasets, is a more externally facing type of accounting and relies on some form of persistent unique identifier. Supporting both types of data accounting simultaneously does require more effort, but would be of benefit to many different communities.

⁶ <https://onedata.org/>

3 Dataset Usage Metrics

Following the questionnaire, a proposal was created for the metrics that could be extracted to perform dataset usage accounting. Table 1 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)⁷: the first two blocks are based directly on the OGF Usage Record, with the addition of an ORCID, and the dataset usage block is an extension. 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. Additionally there is currently no agreed way for the individual records to be aggregated into larger summaries and these metrics will need to be refined as dataset accounting is developed through the prototypes and integrations with storage systems.

Table 1 - Dataset accounting metrics

	Key	Type	Description
Record Identity Block	Resource provider	string	Resource provider at which the resource is located (e.g. GOCDDB site name)
Subject Identity Block	GlobalUserID	string	e.g. X.509 certificate DN / EGI unique ID (from Checkin service)
	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	End time of transfer

⁷ <https://www.ogf.org/documents/GFD.204.pdf>

		timestamp	
	TransferSize	integer	Bytes transferred
	HostType	string	Storage system Type
	FileCount	integer	Number of files accessed
	Status	string	Success / failure / partial transfer

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. These metrics should either be available from a REST endpoint that can be queried or it should be possible to send these metrics to the Accounting Repository via the EGI message brokers.

4 Supported Storage Solutions

4.1 Onedata

Onedata is a global data management system, providing easy access to distributed storage resources, supporting a wide range of use cases from personal data management to data-intensive scientific computations. With Onedata, users can access, store, process and publish data using global data storage backed by computing centres and storage providers worldwide. It is the underlying technology powering the EGI Open Data platform and DataHub⁹.

Onedata provides a REST API which can be used to extract space and user metrics. At the moment, it is not possible to extract metrics based on persistent identifiers (PIDs), such as DOIs, but this is a feature that is being added to support the EGI Open Data Platform and will be available shortly.

Additionally, the Onedata REST API provides metrics in a format that does not directly map onto the proposed metrics present in the previous section so some compromise will need to be found between the two. Also, since a single dataset can be divided between several storage providers, consideration should be made about how the metrics for a dataset can be collated from the data retrieved from disparate providers.

4.1.1 Integration with the Accounting Repository

The APEL software was modified to support the loading of dataset usage records into a specifically designed database schema, and the APEL Secure Stomp Messenger (SSM) component was modified to support fetching dataset usage records from a REST interface (as opposed to sending messages via the EGI Message Brokers), which is the method that Onedata provides access to monitoring data. Currently, the prototype uses a simple REST “puller”, similar to the receiver used to retrieve messages from the EGI Message Brokers. However, the updated SSM with REST support is being developed with the new the ARGO Messaging Service (AMS) in mind, as that service will also make use of HTTP interfaces. Although AMS can certainly be used to communicate with the Accounting Portal in the future, it is still not clear whether the messaging infrastructure can be used for communication with Onedata. If it cannot, then that would mean that effort would need to be spent supporting an additional interface to the Accounting Repository, although the added flexibility may be beneficial.

The prototype works by using the Python urllib2 library to query the REST interface. The returned data is then parsed into the new message format based on the OGF Usage Record. This message is then saved for future loading, as currently happens with all other messages received via the message broker.

⁹ <https://datahub.egi.eu/>

The APEL software has been modified to support loading of this new format into a database by starting a separate loader process with its own configuration file. This means that the prototype is capable of extracting the space metrics of the test space, parsing them into the OGF message format then loading the data into the database.

A lot of the metrics proposed are available internally to Onedata, but not all of them are exposed by the REST API and the ones that are use different keys, and some are not yet implemented (mainly PIDs, ORCIDIDs, and specific metrics about transfers). Additional modifications to the software are thus required to convert the data retrieved from the Onedata API into a format suitable for ingestion by the Accounting Repository and further collaboration will be required between the Onedata and APEL developers to ensure all the right metrics are exposed.

4.1.2 Metric collection example

The integration with Onedata has been tested by running the accounting software and pointing an instance of the SSM, modified to allow interaction with a REST endpoint, at an instance of Onedata to extract usage data from a test space containing an image. The known “spaceID” was used to query the Onedata REST API¹⁰ hosted at the EGI DataHub for the quota metric of the space (e.g. https://datahub.plgrid.pl/api/v3/oneprovider/metrics/space/1I8DOQUXXiezOAcTpAewz40HVNzy-Sr2mIBZZtEmpA?metric=storage_quota&step=1m) and a JSON response like the following was received:

```
{
  "rrd":{
    "meta":{
      "step":2678400,
      "start":1454371200,
      "legend":[
        "space ZgaqavtmWKV805-KXrpfxCXD9UvL_wbvEe0VgeZM1I; metric
storage_quota; oneprovider ID H0e-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":"H0e-D_aZvrpggVyhnTkmBc9czucj1nQ3z-NOJQKew"
}
```

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

¹⁰ <https://onedata.org/docs/doc/advanced/rest/index.html>

```
<?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.

Due to the limitations of the current Onedata implementation, it was not possible to extract most of the metrics that were rated as a priority in the survey and so the output was limited to resource provider metrics. Extracting more detailed metrics will be achieved in the coming months during the development of the second data accounting prototype.

4.2 Other data management systems

4.2.1 EUDAT

B2SHARE¹¹ is the service in the [EUDAT](https://www.eudat.eu/) Collaborative Data Infrastructure¹² that appears to be appropriate for dataset accounting. B2SHARE can receive requests for a digital object's PID and use another service called B2HANDLE¹³ to look up the replicas that exist in the data resources B2SAFE¹⁴ or iRODS (described later), and then deliver the data from one of those sources to the user. B2SHARE currently only logs storage accounting information but they see the need for dataset accounting and have plans to work on usage statistics.

¹¹ <https://www.eudat.eu/services/b2share>

¹² <https://www.eudat.eu/eudat-cdi>

¹³ <https://www.eudat.eu/services/b2handle>

¹⁴ <https://www.eudat.eu/services/b2safe>

There is potential for a future collaboration with APEL to ensure that what they develop meets the EGI-Engage requirements, and the technical capabilities of the APEL software.

4.2.2 ICAT

ICAT¹⁵ is a metadata catalogue, originally developed by the STFC, but now used at a number of neutron and photon sources in Europe and the US through the Photon and Neutron data infrastructure initiative¹⁶ (PANdata) FP7 project. It keeps information on an ‘investigation’ which is typically a series of measurements by a team on one instrument at a large facility. It tracks the investigation from project submission, approval, and data-taking, through to scientific publications. The relevant part for dataset accounting is its knowledge of the raw and derived data stored as part of the investigation. This is not only of use by the participating scientists but is also used by publications to point to the data for subsequent use.

In the ISIS¹⁷ neutron source, the investigation itself has a single DOI in DataCite. Other models are available; for example the Diamond Light Source¹⁸ has a DOI for each individual dataset. Both models could benefit from dataset accounting. ICAT and its associated web interface TopCat look worthy of further investigation.

4.2.3 iRODS

The integrated Rule-Oriented Data System¹⁹ (iRODS) is open source data management software, provided by the iRODS Consortium, used by many scientific sites and collaborations, and is the backend of the B2SAFE service of EUDAT. Although it is a stand-alone system with its own catalogues, metadata, collections, and data transfer mechanisms, it is also integrated into EUDAT services.

It is worth further investigation, both in terms of usage within the EGI communities and technically that it keeps sufficient logging to report what dataset accounting needs.

¹⁵ <https://icatproject.org/>

¹⁶ <http://pan-data.eu/>

¹⁷ <http://www.isis.stfc.ac.uk/about/aboutisis.html>

¹⁸ <http://www.diamond.ac.uk/Home/About.html>

¹⁹ <https://irods.org/>

5 Plan for Exploitation and Dissemination

Name of the result	Knowledge of current state of dataset accounting
DEFINITION	
Category of result	Technical input to standards and know-how
Description of the result	This report gives an overview of the current state of dataset accounting in EGI-Engage and looks at some of the areas that will need future development.
EXPLOITATION	
Target group(s)	Service providers, standardisation bodies
Needs	Knowledge of what the required metrics to perform dataset accounting are and how they might be integrated into the EGI accounting system.
How the target groups will use the result?	Get agreement on a standard for dataset usage accounting that is useful for users and that can be reused for other storage systems and infrastructures.
Benefits	Greater understanding between the service providers and accounting team. Starting point for future discussion on standardisation of dataset accounting.
How will you protect the results?	All software and standards released under an open source licence.
Actions for exploitation	Further collaboration between EGI accounting team and Onedata. Solicit feedback on prototype and proposed metrics.
URL to result	
Success criteria	A more developed second prototype that regularly collects dataset accounting and an agreed standard for collecting dataset accounting metrics.
DISSEMINATION	
Key messages	Replicate your datasets according to real user needs. Identify the datasets that must be on-line and those that can be archived or even deleted.
Channels	EGI conferences and meetings, direct communications with collaborators.
Actions for dissemination	Further collaboration with Onedata, discussion with developers of other storage systems that support unique IDs (e.g. EUDAT), solicit feedback on prototype and metrics, track developments in the WLCG Data Management group.
Cost	N/K
Evaluation	Quality of feedback and discussions.

6 Future Work

Having produced an initial prototype, work is needed for the second prototype to move it towards a production quality service. This includes fixing any problems with the integration with Onedata and making use of unique dataset identifiers once they are available. Work will also be needed to modify the software to use the Onedata API to retrieve the list of space IDs and metrics programmatically. 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 roadmap²⁰. Summarising the data 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 for this data.

The initial strategy was to use the EGI DataHub and to integrate the prototype with the Onedata API hosted there as this was the EGI provided solution. However, considering the heterogeneous nature of the storage infrastructure involved, and the need to develop a system that is as generic as possible, other data management systems will be analysed. A variety of other storage systems have already been reviewed – B2SHARE, ICAT, and iRODS – although even these do not yet provide all the features required for dataset accounting. This means that work is needed to collaborate with the developers of these systems to add the required features. With this in mind, a check list of features should be produced to list what a project or storage system needs to provide to enable dataset usage accounting.

Overall, there is now a good starting point for dataset accounting in the form of a prototype schema, which may need some tweaking, and a method for storing those records. However, the more challenging problem is the lack of data tools in the EGI communities that actually store the right information to do dataset accounting, and the fact that the current prototype does not really cover the requirements of WLCG. WLCG is only just starting to define its requirements in this area and it is currently focussed on optimising storage space usage – storage is file-based and the data is not currently stored with an associated PID, although EUDAT is working towards providing PID-aware services that could be used by WLCG. It would be of benefit if WLCG could collect this data through APEL. The Accounting Repository team will follow along with these developments, and increasing the amount of engagement with WLCG should enable the second prototype to be increasingly aligned with the WLCG and other community requirements. Having sorted out some of the basic requirements for dataset accounting, they are now in a stronger position to collaborate with WLCG, Onedata, EUDAT, and others.

²⁰ https://wiki.egi.eu/wiki/TASK_JRA1.3_Accounting#Accounting_Repository