



EGI-Engage

User Requirements on Data Accounting

M3.2

Date	03 March 2016
Activity	WP3
Lead Partner	STFC
Document Status	FINAL
Document Link	https://documents.egi.eu/document/2674

Abstract

The APEL accounting team is proposing to extend accounting to account for usage of data sets. Here, we define a data set 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. Accounting for data set usage should enable site and experiment administrators to make decisions about the location and storage of data sets to make more efficient use of the infrastructure and assist scientists in assessing the impact of their work. Before embarking on software development, a questionnaire was used to solicit interest and requirements. It reveals a clear interest from the surveyed communities. The survey reveals a number of core requirements for data set accounting, including user-access, data set PID logging, and recording transfer operations. There is a wide range of data set storage systems in use with no consistent approach to recording data set usage; a de-facto accounting record for data sets has not yet emerged. In lieu of a de-facto accounting record for data sets, we propose to develop a data set accounting record based on the OGF UR2 record format, which third parties can implement, and to develop a central accounting repository which can receive and process this data.



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:	Stuart Pullinger	STFC/JRA1	03-03-2016
Moderated by:	Małgorzata Krakowian	EGI.eu/NA1	
Reviewed by	Tiziana Ferrari Alexandre Bonvin	EGI.eu/NA1 BCBR-UU	03-03-2016 21-02-2016
Approved by:	AMB and PMB		03-03-2016

DOCUMENT LOG

<i>Issue</i>	<i>Date</i>	<i>Comment</i>	<i>Author/Partner</i>
v.1	02/02/2016	Initial version	Stuart Pullinger/STFC
v.2	16/02/16	Internal review	Stuart Pullinger/STFC
V3	03/03/16	Input from external reviews	Stuart Pullinger/STFC
FINAL	03/03/16	Final version	Stuart Pullinger/STFC

TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/about/glossary/>

Contents

1	Introduction.....	5
2	Methodology	6
2.1	Questionnaire.....	6
2.2	Interviews.....	6
3	Questionnaire Responses	7
3.1	Respondents.....	7
3.2	Data Set Accounting	7
3.3	Storage Systems	8
3.4	Measures.....	9
4	Interviews	11
4.1.1	ELIXIR.....	11
4.1.2	ESA's Geohazards and Hydrology thematic exploitation platforms	11
4.1.3	CESNET Data Storage Department.....	12
4.1.4	EPOS	12
4.1.5	D4Science	12
5	Data Accounting Systems Already in Use	14
5.1	Data Accounting System at CESNET Data Storage Department	14
5.2	WLCG Computing Resources Scrutiny Group.....	14
6	Analysis and Outcomes.....	15
7	Conclusions and next steps.....	17
	Appendix I. The Data Accounting Questionnaire	18
7.1	You	18
7.2	Data Accounting	18
7.3	Storage Systems	18
7.4	Measures.....	18
7.5	Technical	19

Executive summary

EGI-Engage supports the development of the extensions to accounting needed to account for usage of data sets; this is a necessary extension needed to the current accounting infrastructure in order to support Data as a Service (DaaS), intended as the on demand provisioning of data to the user regardless of geographic or organizational separation of provider and consumer.

Here, we define a data set 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 differs from storage accounting which accounts for storage (e.g. disk or tape) allocation and usage without concern over what data is stored or who uses it and how often. Accounting for data set usage will include recording who accesses 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, report on data usage to the data owners, data providers and funding agencies, and assist scientists in assessing the impact of their work.

Before embarking on software development, a questionnaire was used to solicit interest and requirements; and to gather feedback from stakeholders on how best to implement a prototype system. It reveals a clear interest from the 14 surveyed communities. In addition, communities that expressed great interest in this activity were selected for interviews to clarify their needs.

The survey and the interviews reveal a number of core requirements for data set accounting, including user-access auditing (who and how often), data set source/sink/PID logging, and recording transfer/copy operations (store/retrieve, bytes transferred, success/fail).

There is a wide range of data set storage systems in use with no consistent approach to recording data set usage. A de-facto data set storage technology or access method has not emerged; nor has a de-facto accounting record for data sets. This makes the creation of a single data set accounting tool difficult. In lieu of this, we propose to develop a data set accounting record based on the OGF Usage Record version 2 XML format, which third parties can implement, and to develop a central accounting repository which can receive and process this data.

Furthermore, the design of this new accounting feature will consider the EGI Open Data Platform (WP4) as a possible source of data set accounting data. Indeed, this new EGI services, which provides capabilities to publish, use and reuse openly accessible data identified via a persistent unique identifier (PID), looks promising to become widely adopted, guarantees the use of a handle (the PID) to trace data set access through the whole infrastructure and hide the complexity due to the wide range of storage systems adopted.

1 Introduction

The EGI has a production accounting infrastructure for grid and cloud systems that records CPU usage and a number of other usage metrics. A proposal to extend the accounting system to include the usage of data sets by different sites is a likely future requirement.

Data set accounting would include:

- Usage by a project across many sites
- Usage by a user within a project or individually across many sites
- Usage of a Resource Centre by many projects
- Management view of usage by all projects of all Resource Centres.

There are a variety of clients including those developed by APEL and site-specific clients which all use a common EGI Messaging infrastructure to send usage records (UR) to a central APEL repository. At the central repository, they are loaded into the relevant database and summarised daily on various fields. Summaries are subsequently sent to the EGI Accounting Portal where they can be visualized dynamically.

The APEL accounting team are proposing to extend accounting to account for usage of data sets. Here, we define a data set 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 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.

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. The results of that questionnaire, interviews and a brief analysis are given below.

2 Methodology

In order to establish the best way for the accounting team to proceed, a questionnaire was formulated and distributed for response. Interviews with selected communities were organised to better refine the collected requirements.

2.1 Questionnaire

The questionnaire aimed to address the following:

- Establish if data accounting was already being done and if there was interest in EGI doing it
- Establish which technologies were in use for providing data sets
- Establish which measures should be accounted for
- Create a group of volunteers who are willing to collaborate with software testing

The questionnaire was structured with a section to correspond to each of these areas. The full questionnaire can be found at the end of this document. The questionnaire was created on Survey Monkey and the link was distributed to the mailing lists representing interested communities: site administrators and resource providers; virtual organisation administrators and representatives of research communities. The responses are collated in the next section.

2.2 Interviews

A small number of communities that expressed specific interest in this activity were selected for interviews to better understand their needs. The results from these interviews are given below.

3 Questionnaire Responses

The questionnaire responses from the 14 respondents are split into the categories from the questionnaire.

3.1 Respondents

The respondents represented a number of Resource Centres (also known as “sites”), Virtual Organisations (VOs) and research communities. Some respondents represented several of these at the same time. The complete list is:

- Sites: CESNET (2 respondents), IFCA, CSIC-UC, DASSH: the UK Archive for Marine Species and Habitats Data, CNR
- VOs: enmr.eu, vo.nbis.se, excelerate, ESA's Geohazards and Hydrology thematic exploitation platforms, LifeWatch, D4Science
- Research Communities: WeNMR, Instruct (West-Life), ELIXIR, NBIS, Biodiversity & Ecosystems (LifeWatch), EISCAT: the European Incoherent Scatter Scientific Association, iMarine, The European Plate Observing System (EPOS), ESA's Geohazards and Hydrology thematic exploitation platforms

1. Do you represent...		
a Site providing resources?	a Virtual Organisation (VO)?	a research community?
6	7	8

3.2 Data Set Accounting

This section aims to establish the interest in data set accounting and to assess whether solutions already exist. The respondents expressed strong interest in data set accounting. 2 respondents stated that they had an existing method of data set accounting and 4 respondents described partial solutions in place. None of the systems described in the responses were applicable across the infrastructure as proposed for data set accounting here. For example, one solution used Google analytics to record accesses to web pages; another used a custom metadata catalogue. Use of PIDs was low or incomplete and, where used, used a variety of different schemes.

1. Are you interested in accounting for usage of data sets?		
Yes	No	
12	1	
2. Do you have an existing method for accounting for usage of data sets? If so, please describe it.		
Yes	Partial	No
2	4	4
3. Do you use Global Identifiers/URIs/PIDs to identify your data sets? If so, which scheme do you use?		
Yes	Partial	No
4	3	4
URIs, Persistent URLs, DOIs from DataCite	DOI, EPIC, Handle	
4. Which of the following features are you interested for? Multiple answers are possible.		
Record how often a data set is accessed	10	
Record who accesses data sets	9	
Record the data transfers	9	

3.3 Storage Systems

A wide variety of storage systems are in use for providing data sets and these are accessed by a diverse range of methods. No single combination of storage system and access method is established as the most popular.

1. Which storage system(s) do you use?	
Other (MongoDB + Apache Jackrabbit, S3, HDFS, gCube-based solutions, STORM, OneData)	5
GPFS	4
Local storage/POSIX	3
Ceph	2
NFS	1
2. How is data on your storage system accessed? Is there an interface to users? Is there a programmatic interface?	
gCube-based, custom tool, custom api, ownCloud, Globus Online,	5

Aspera Web	
Web GUI	4
local filesystem, NFS	3
REST/SOAP Web Service	2
WebDAV, scp/sftp	2
dCache	1
xrootd	1

3.4 Measures

The survey established the metrics which are important in data set accounting: who accessed a data set and how often; and logging of data set transfers (copies) between a source URI and a sink URI. This includes a request to include meta-data to describe ‘why’ the data was copied. Other high priority data fields that should be definitely included are: the different forms of user ID (such as an x509 certificate DN and/or an eduPersonPrincipleName (eppn) attribute from a FEM security realm); user groupings such as VO, home-site; store and retrieve operations; files transferred; success or failure; and the data set Identifier

Other, medium priority data fields which should probably be accounted for include: storage system implementation ie. 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. Finally, the lower priority field which may be excluded is the sub-unit of the storage system.

1. Which of the following data fields would you be interested in accounting for? Please rate them:			
Data Field	high priority	interested	not interested
User Group/Virtual Organisation	5	6	1
Store Operations	5	6	1
Retrieve Operations	5	6	1
User	5	5	1
Data Set (some kind of unique identifier such as a PID)	4	7	1
Destination of transfer (IP address or something else?)	3	7	0
Amount of data transferred (which units?)	3	6	1
File(s) accessed	3	6	1
Success and Failure ie. do you want to account for failed	3	6	1

and/or partial transfers?			
Site	2	9	0
Sub-unit of Storage System	2	3	5
Source of transfer (IP address or something else?)	1	9	0
Storage System	1	7	3
Start time of transfer	1	7	2
Duration of transfer	1	7	2
VO Group and/or Role	1	6	4
End time of transfer	1	6	3

4 Interviews

Communities that expressed great interest in data accounting were selected for interviews to better clarify their requirements. This section summarises the outcome of these interviews.

4.1.1 ELIXIR

The ELIXIR research community representative felt that they had requirements which did not suit the format of the questionnaire so requested a meeting to further explain their requirements. A member of the accounting team attended the EGI-Engage ELIXIR Competence Centre meeting to acquire more detailed information. The results of this meeting are summarised below.

In the questionnaire, the ELIXIR respondent expressed interest in accounting for store and retrieve operations at a site level as well as recording the bytes transferred. At the meeting, it emerged that the more useful metrics for this community were related to per-data set measurements rather than per-site or per-user totals of bytes transferred. They are already looking into this in the metrics work package of the ELIXIR EXCELERATE project.

The use cases for data set accounting foreseen by this community are for replication decisions ie. moving data closer to where it is used and assessing data set impact by measuring which data sets are most popular. Therefore, the users of data set accounting data would be site administrators, VO administrators and data set administrators.

The technologies in use by this community for providing data are diverse. Data is provided through locally attached filesystems as well as storage systems. Technologies such as dCache, xrootd and iRods, popular in the grid computing world, are not commonly used in the life sciences. There is a central installation of Aspera – a commercial product – which is used for data transfers to/from the EBI repository. Elsewhere in the community, open-source solutions are in use such as GridFTP. Data sets are generally released on a quarterly/monthly timescale but there is no single PID system for identifying data sets in use.

4.1.2 ESA's Geohazards and Hydrology thematic exploitation platforms

The ESA's Geohazards and Hydrology thematic exploitation platforms (TEPs) community is interested to integrate the whole EGI accounting system into the TEPs. They want to use the APEL-based accounting system to trace the consumption of resources they provide which do not belong to the EGI federation since the TEPs will exploit EGI and non-EGI sites for their computational needs.

It has been agreed to rely on the EGI accounting system for accounting data related to the federation and integrate a new APEL server in the TEPs architecture to retrieve accounting information from non EGI sites, which will setup proper APEL probes. The APEL server dedicated to

the TEPs will fetch accounting data from the EGI accounting repository and combine them with the other collected data to have an integrated view from all the TEPs resource providers.

This community needs the accounting information to measure the cost of the resources consumed by each user; access to datasets is one of the variables to take into account to compute this cost. They expressed interest in data accounting and are willing to provide requirements for the usage record definition in the coming months.

TEPs will process satellite data images. Each of these images has to be considered as a separated product with a PID assigned. A PID management system for the satellite datasets is not available yet.

4.1.3 CESNET Data Storage Department

CESNET is interested in the amount of data stored by users and VOs. The current CESNET accounting system provides a storage accounting feature. It traces the amount of data transferred but no information about the specific file accessed is stored.

CESNET believes that the data accounting could be useful for a part of its community and it would be exploited to improve its policy for dataset replication management.

4.1.4 EPOS

The EPOS research infrastructure¹ is very interested in this topic and data accounting is a feature that has been envisaged on the design of its infrastructure. The EPOS implementation phase project started in the last months of 2015 and the EPOS collaboration is now working on collecting needs from the solid earth science communities. It is currently too early for EPOS to provide EGI with specific requirements on data accounting. However, they expressed their willingness to collaborate with EGI on the definition of the data accounting usage record and to evaluate APEL as a system for data accounting for their infrastructure.

EGI will contact again EPOS in the next months, when the activity to define the data accounting usage record will start.

4.1.5 D4Science

D4Science already has a data accounting system for its internal distributed storage based on MongoDB. Each dataset is seen as a file with associated a unique identifier that identifies the dataset within the whole D4Science infrastructure. Each operation to the datasets is recorded: read, update, delete, etc.

This data accounting system has been introduced in the D4Science infrastructure to manage quotas (e.g. a dataset can be accessed not more than one time each day). Now, D4Science is investigating on how to exploit the data collected (intra and inter sites) to proper manage replicas.

¹ <http://www.epos-eu.org/>

In particular, proper analytics will be designed, which will take as input data about the time needed to transfer data and the volumes transferred, to understand how to improve the datasets distribution in the infrastructure. Until now, the accounting data are only collected but not used for any specific aims. All the data transfers are done via HTTP or hidden behind an HTTP communication.

D4Science needs a programmatic interface towards the accounting repository to get the same accounting data related to their usage of the EGI Federated Cloud.

The D4Science infrastructure includes other data sources like relational, geographical and statistics databases. The data account feature is not available yet for these systems since data stored there has not associated to a unique identifier.

D4Science is willing to provide further requirements during the design of the EGI Data Accounting system and share its experience with the EGI accounting team.

5 Data Accounting Systems Already in Use

One of the aims of the survey has been the identification of data accounting solutions already in place that could be adopted into the EGI infrastructure or used as reference for the specific EGI developments.

Interviewed communities suggested analysing two systems that could fit with their requirements:

- Data Accounting system of the CESNET data storage department
- WLCG Computing Resources Scrutiny Group

An assessment of these solutions against the collected requirements is described in the follow.

5.1 Data Accounting System at CESNET Data Storage Department

The CESNET data storage department designed and developed an accounting system for its internal hierarchical storage to trace the amount of data stored by users and VOs. Internal redundancy is guaranteed through the features provided by hierarchical storage.

During the interview, it has been clarified that this accounting system logs the amount of transferred data but not information about the specific file/dataset accessed is saved. It means that this system is providing a storage accounting feature and not a data accounting one.

Currently, CESNET has not adopted any persistent unique identifier management system to identify its datasets, which are just files in the file system. The development of an open access data interface and the related adoption of a PID system will start in the next months in the context of a new project.

5.2 WLCG Computing Resources Scrutiny Group²

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 focus on data which remains on disk and is not accessed. They report on data with zero or low numbers of accesses for 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. The methods employed to achieve this deserve more study to see if they are applicable to data set accounting from research communities outside of the WLCG. It will be further examined during the design of the EGI Data Accounting prototype.

² <http://wlcg.web.cern.ch/collaboration/management/computing-resources-scrutiny-group>

6 Analysis and Outcomes

The survey was successful in establishing that there is clear interest in data set accounting; 13 out of the 14 respondents expressed an interest. There is a wide range of storage systems currently in use, with no consistent approach to recording data set usage or metrics across those systems.

Existing methods for data accounting appear to be varied and only partially supported, ranging from publishing DOIs on the Web to site-local log files. There appears to be no consistent approach to record these data across the sites surveyed. Global data set identifiers are used incompletely (four sites are using such IDs and three are using partial IDs). We can conclude that global identifiers should be included in any future data accounting record, especially given the increasing popularity of persistent identifier services (PID) such as SeeGrid³. The format of such an ID field should probably not be mandated in the data set accounting record, which should provide only (one-to-many) optional fields to specify replicas.

Data sets are accessed using a variety of methods including WEB UIs, REST/WS services, local-FS only, custom APIs, and Globus Online. Therefore, any future data set accounting record should be extensible and suitable for Web/Catalogue publishing (e.g. XSD Schema definition with XML/JSON renderings, or RDF). We recommend the format of the data set accounting is based on the OGF Usage Record version 2 (UR-2.0)⁴ - an XML format which could be extended as required.

The range of protocols in use for the location of data sets is wide and includes WebDav, SFTP/SCP, FTP, HTTP(S), and local file. Therefore, the fields for the data set source/sink for any future data set record should define simple URIs and not mandate particular protocols. Some further research would be needed to review relevant works such as data staging profiles for existing middleware (e.g. OGSA Data Movement Interface, JSDL data staging, or data copy documents).

The survey established the metrics which are important in data set accounting, with the most important being: who accessed a data set and how often; and logging of data set transfers (copies) between a source URI and a sink URI.

Through the interviews, requirements from the most interested communities have been further discussed and clarified and useful information has been collected to refine the survey outcome. Interviews have also been good a method to establish better connections with the communities and to involve them in the design phase.

Finally, the need for a PID management system to implement a data accounting feature clearly emerged in this analysis. EGI systems which expect to offer this feature should be properly taken into account during the design phase. The EGI Open Data Platform (which will provide capabilities to publish, use and reuse openly accessible data identified by PIDs) is the best candidate in this set

³ <https://www.seegrid.csiro.au/wiki/Siss/PIDService>

⁴ <https://www.ogf.org/documents/GFD.204.pdf>

and will be investigated as a possible source of data set accounting data. Furthermore, the design of this platform includes the integration of current EGI storage services into the platform backend. Therefore, it could also help to hide the complexity caused by the wide range of storage systems adopted.

7 Conclusions and next steps

The questionnaire and interviews reveal a clear interest in data set accounting from the surveyed communities. The questionnaire reveals a number of core requirements for data set accounting, including user-access, data set PID logging, and recording transfer operations. Interviews allowed the requirements to be refined and created collaboration links with the most interested communities.

There is a wide range of data set storage systems in use with no consistent approach to recording data set usage. A de-facto data set storage technology or access method has not emerged; nor has a de-facto accounting record for data sets. This makes the creation of a single data set accounting tool difficult. In lieu of this, we propose to develop a data set accounting record based on the OGF Usage Record version 2 XML format, which third parties can implement, and to develop a central accounting repository which can receive and process this data.

In the coming months, the APEL accounting team will further investigate those projects, indicated in the questionnaire and interview responses, which are developing or have developed data set accounting to see if there are lessons to be learned. This will include: the ELIXIR EXCELERATE project which is looking at data set accounting; the WLCG scrutiny group who have some data set accounting in place already; and the EGI Open Data Platform which will provide an access method for Open Data from the EGI infrastructure. Furthermore, a working group will be established from the technical contacts in the questionnaire responses. This working group will aim to seek agreement on a message format and field definitions to extend the OGF Usage Record version 2 to data set accounting.

Appendix I. The Data Accounting Questionnaire

7.1 You

1. Do you represent...
 - a. a Site providing resources? Which Site.....
 - b. A VO? Which VO.....
 - c. A research community? Please describe your community.....
 - d. Other. Please describe.....

7.2 Data Accounting

2. Are you interested in accounting for usage of data sets?
 - a. Yes/No
3. Do you have an existing method for accounting for usage of data sets? If so, please describe it.
 - a. Yes:...../No
4. Do you use Global Identifiers/URIs/PIDs to identify your data sets? If so, which scheme do you use?
 - a. Yes:/No

7.3 Storage Systems

5. Which storage system(s) do you use?
 - a.
 - b.
 - c.
6. How is data on your storage system accessed? Is there an interface to users? Is there a programmatic interface?

7.4 Measures

7. Which of the following data fields would you be interested in accounting for? Please rate them: 0=not interested; 1=interested; 2=high priority.
 - a. User (identified by X509 DN)

- b. VO
- c. VO Group and/or Role
- d. Site
- e. Storage System
- f. Sub-unit of a Storage System
- g. Store operations
- h. Retrieve operations
- i. Start time of transfer
- j. End time of transfer
- k. Duration of transfer
- l. Source of transfer (IP address or something else?)
- m. Destination of transfer (IP address or something else?)
- n. Amount of data transferred (which units?)
- o. File(s) accessed
- p. Data Set (some kind of unique identifier such as a PID)
- q. Success and Failure ie. Do you want to account for failed and/or partial transfers?

7.5 Technical

- 8. Are you willing to provide the APEL accounting team with a log file to enable them to test their prototype?
- 9. Are you willing to assist the APEL accounting team by testing the prototype accounting software on your storage system?
- 10. Are you able to provide a technical contact at your site to communicate with the APEL accounting team?

Thank you very much for taking the time to fill in this questionnaire.

Stuart Pullinger

APEL Accounting Team Leader