A Grid storage accounting system based on DGAS and HLRmon

2012 J. Phys.: Conf. Ser. 396 032029

(http://iopscience.iop.org/1742-6596/396/3/032029)

Journal of Physics: Conference Series **396** (2012) 032029 doi:10.1088/1742-6596/396/3/032029

A Grid storage accounting system based on DGAS and HLRmon

A. Cristofori[1], E. Fattibene[1], L. Gaido[2], A. Guarise[2], P. Veronesi[1]

[1] INFN-CNAF - Ist. Naz. di Fisica Nucleare Viale Berti Pichat, 6/2 40126 Bologna (Italy)

[2] INFN-TORINO - Ist. Naz. di Fisica Nucleare Viale Pietro Giuria, 1 10125 Torino (Italy)

E-mail: dgas-support@to.infn.it

Abstract. Accounting in a production-level Grid infrastructure is of paramount importance in order to measure the utilization of the available resources. While several CPU accounting systems are deployed within the European Grid Infrastructure (EGI), storage accounting systems, stable enough to be adopted in a production environment are not yet available. As a consequence, there is a growing interest in storage accounting and work on this is being carried out in the Open Grid Forum (OGF) where a Usage Record (UR) definition suitable for storage resources has been proposed for standardization.
In this paper we present a storage accounting system which is composed of three parts: a sensor layer, a data repository with a transport layer (Distributed Grid Accounting System - DGAS) and a web portal providing graphical and tabular reports (HLRmon). The sensor layer is responsible for the creation of URs according to the schema (described in this paper) that is currently being discussed within OGF. DGAS is one of the CPU accounting systems used within EGI, in particular by the Italian Grid Infrastructure (IGI) and some other National Grid Initiatives (NGIs) and projects. DGAS architecture is evolving in order to collect Usage Records for different types of resources. This improvement allows DGAS to be used as a 'general' data repository and transport layer. HLRmon is the web portal acting as an interface to DGAS. It has been improved to retrieve storage accounting data from the DGAS repository and create reports in an easy way. This is very useful not only for the Grid users and administrators but also for the stakeholders.

1. Introduction

A production Grid infrastructure consists of heterogeneous and geographically distributed resources that function together to satisfy the computational requirements of modern scientific experiments. Many research communities, working on different scientific disciplines - such as high energy physics, biology, astronomy, earth science, life sciences, etc. - benefit from the usage of computational Grids. In order to understand how the computing and storage resources have been used by those different communities, the Grid middleware needs to be integrated with scalable accounting services and intuitive interfaces offering reports about resource consumption.

EGI [1] is a federation of many National Grid Initiatives (NGIs) and two European Intergovernmental Research Organisations (EIROs). Each NGI provides its own computing and storage resources and Grid services to the leading edge computing infrastructure needed by European researchers. Accounting is provided by services integrated in the Grid middleware. Computing accounting systems deployed within EGI are considered sufficiently mature and stable.. Unfortunately this is not true for the accounting of storage resources although there is a growing interest in expanding accounting to include this type of resource.

IGI [2] has developed a storage accounting solution that is suitable not only for a Grid infrastructure but also for other different environments. Whatever the environment, the only requirement is that storage metrics are available and can be expressed with the same fields as in the system that will described in this paper.

This paper presents the storage accounting solution developed by IGI. Section 2 gives an overview of the accounting system and Section 3 describes in detail the storage UR. Section 4 is dedicated to the architecture and the implementation and Section 5 provides some general comments to this activity

2. The IGI accounting solution

The service responsible for the accounting task within IGI is DGAS [3][4][5]. It interoperates with the middleware stack provided by the European Middleware Initiative (EMI) [6]. It has been used for the accounting of computing resources in the IGI infrastructure since 2005. Figure 1 shows the accounting layout within IGI.

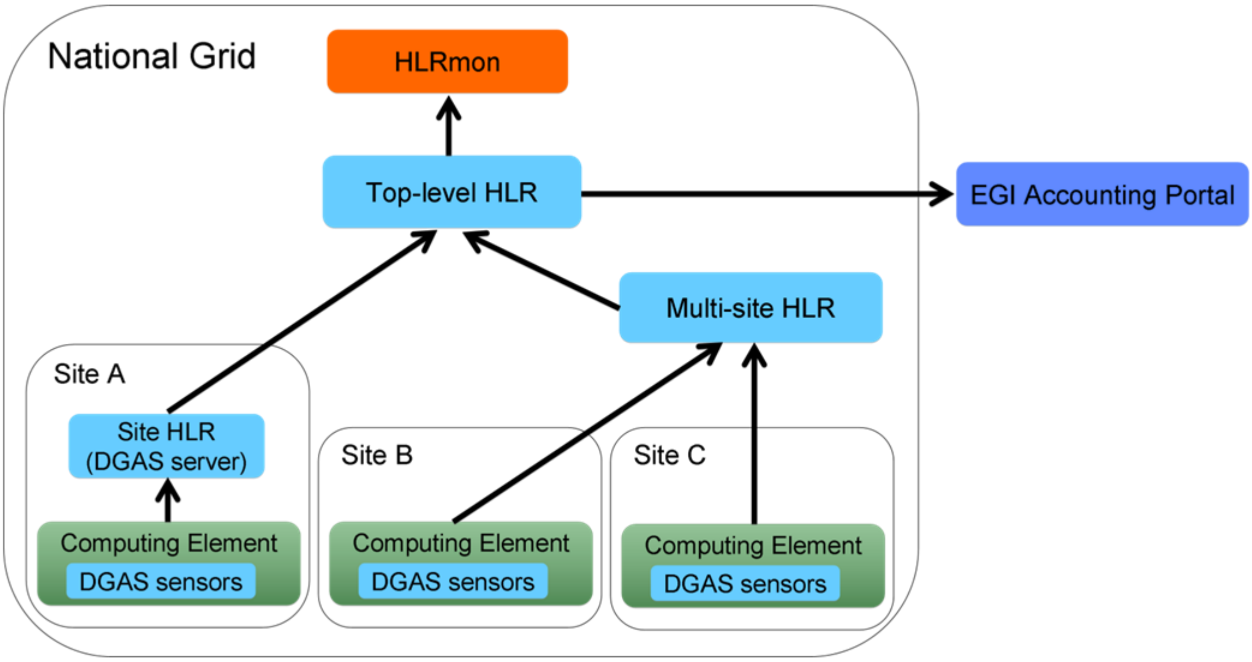

Figure 1. IGI accounting layout

At Grid site level, DGAS sensors retrieve accounting information and produce a UR for each Grid job; the UR contains the relevant information about the job resources consumption, such as the values of the CPU Time, Wall Clock Time and memory used during the job execution. The DGAS server is termed the Home Location Register (HLR). The URs are stored in the HLRs that may either be dedicated to a single site or multi-site. A top level HLR collects all the URs from the distributed HLRs and acts as the IGI accounting repository. On top of this service, the HLRmon web portal [7][8]

aggregates the large amount of computing accounting data collected by DGAS and provides multiple reports in both graphic and tabular formats.
Since IGI is part of EGI, the accounting data aggregated by Grid sites and by Virtual Organizations (VOs) - the communities of users - are sent from the top level HLR to the EGI accounting repository that collects accounting data from all the NGIs and provides reports through the EGI central accounting portal.

DGAS and HLRmon have been adopted by some national production Grids - HellasGrid in Greece and D-Grid in Germany - as well as by We-NMR and EUMed projects.

3. The storage Usage Record

A storage accounting system can be defined as a collection of applications that allows the retrieval and recording of the resources used and the interfaces necessary for billing and pricing, the processes in which the consumption of the resources are accounted and a price is calculated depending on specific rules. It might also allow the aggregation of URs using different criteria, e.g.: users, groups, VOs, Storage Areas (SAs) - the SA is a portion of storage space that can be defined on the Storage Element (SE), that provides uniform access to data storage resources -, etc.

Although a standard storage UR has not yet been defined, the need for a storage accounting system has led communities and projects such as OGF [9] and EMI to start the work on its definition. As a result of this work, EMI has defined the Storage Accounting Record (StAR) [10] and OGF has drafted a document on the OGF UR 2.0 proposing a standardisation of the UR definition that also includes storage. While EMI's StAR definition is focused on the storage, OGF UR has a broader scope and also encompasses computing, memory, etc.

The following schema of the UR has been defined during the development of the IGI storage accounting system and is part of our contribution to the definition of the OGF standard UR. It has many similarities with the StAR UR definition and many of its features are now in OGF UR 2.0. Table 1 contains the list of attributes of this UR definition. The UR can refer to a single file or to an aggregate, i.e. the files stored in a SA; in case of aggregate the fields that refer to a file should have NULL value.

Attribute name	**Note**
RecordIdentity	a unique identifier for the UR
GlobalFileId	a string that globally identifies the file
LocalFileId	full path and file name on the storage
GlobalGroup	the group that is globally unique and whom the file belongs to
GlobalUsername	the user that is globally unique and whom the file belongs to
LocalUserId	a local user identifier
Charge	the charge associated with the storage utilization
Status	status of the request to the storage system
Host	the hostname of the machine that hosts the file
SubmitHost	the hostname of the machine that submits the file to the storage system

ProjectName	the name of the project, such as an experiment or a group of users, whom the file is associated to
ProjectPartition	a logic partition of the project
StorageType	the implementation of the storage system used
ProtocolType	the protocol used in the event
OperationType	the operation performed on the file
Network	the amount of bytes transferred through the network for the operation
Disk	depending on the OperationType performed could have different meanings, i.e.: CREATE - total file size; DELETE - size of the deleted file; READ - the amount of bytes read from the file
TimeDuration	the time used for the operation (if specified) or the amount of time that the file(s) space has been occupied on the storage
TimeInstant	the timestamp of the beginning of the operation
ServiceLevel	service level associated with the file; it depends on the storage implementation

Table 1. Storage UR attributes

4. The storage accounting system

The architecture of the DGAS storage accounting system is shown in Figure 2.

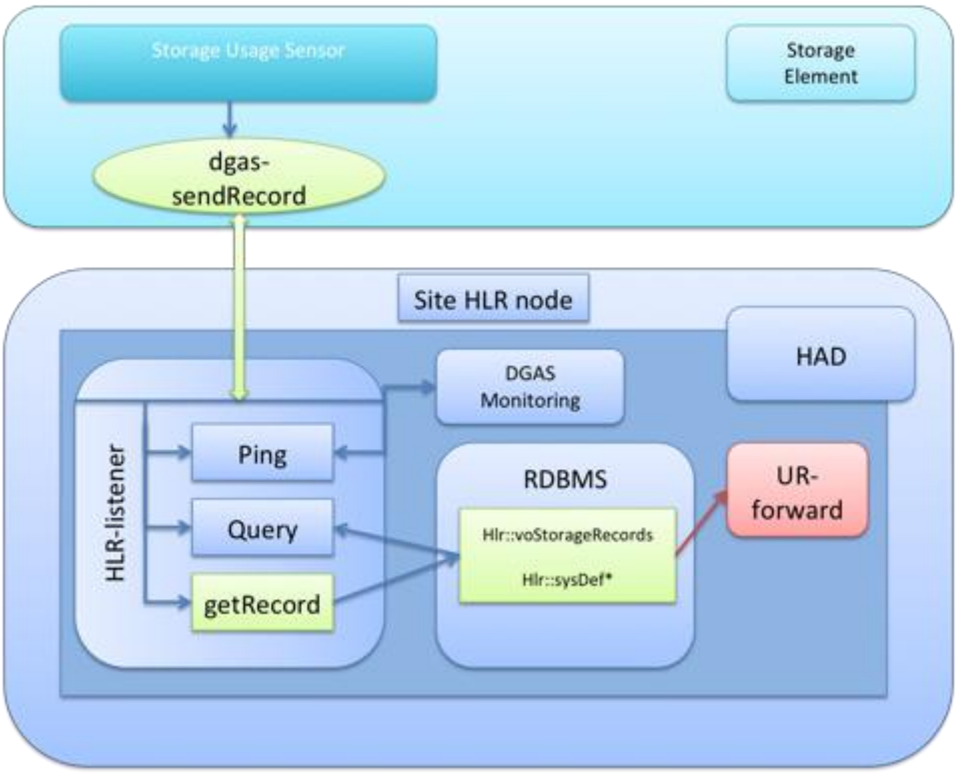

Figure 2. Storage accounting system architecture

In this architecture the sensor collects the storage usage data directly from the storage system and then a UR is produced according to the format described in the previous section. In the current implementation of the prototype, the sensor is a Python script that queries the Grid Information System (IS) to collect the available accounting information for the storage systems of the Italian Grid sites, as published according to the GLUE schema [11]. In particular it retrieves the usage metrics of each SE and of each SA if present.

As already described, the current implementation extracts storage accounting information from the IS. This is far from being optimal since it is affected by the inaccuracy of the current implementations of the storage information providers used in the EMI framework or by problems related to publishing to the IS. As soon as more accurate data is available, due either to an improvement of the storage information providers or to the development of appropriate sensors for storage resources, the DGAS storage accounting framework will easily be able to handle them.

Before being delivered to the HLR repository by the *dgas-sendRecord* tool URs are converted to a suitable format. The records collected by the HLR are stored in a MySQL database. A software component has been developed that allows HLRmon to retrieve summary aggregations of those records. An example of the graphical reports provided by HLRmon is shown in Figure 3. Through the web interface it is possible to get historical graphs aggregated for site, VO and SA.

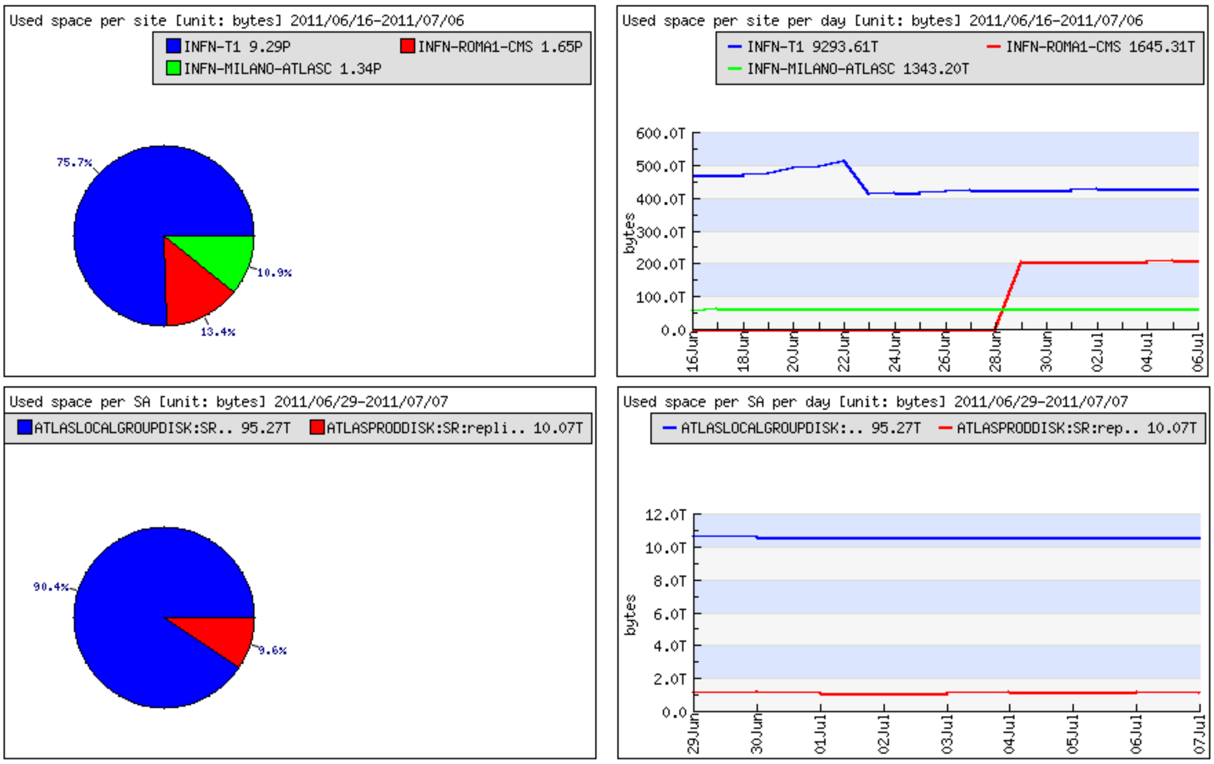

Figure 3. Sample screenshot from HLRmon storage accounting view

In the future the HLRmon plots will compare the used space of each storage system (SE or SA) with the overall amount of available space on those systems. This feature will provide an understanding of how much the storage resources are being used against their total usable space. The UR stored in the HLR database does not provide any information on the total available storage space, since this is simply a static administrative value. HLRmon will query the IS for the specific attribute defined within the GLUE schema.

Journal of Physics: Conference Series **396** (2012) 032029 doi:10.1088/1742-6596/396/3/032029

5. Conclusions

The system described in this paper already provides accounting for the computing resources. It has been improved to retrieve and report also storage accounting information. The storage accounting system relies on several components: the sensor that collects the usage information and produces the UR according to a schema that has been developed by IGI and has contributed to the OGF UR definition; DGAS that provides the accounting framework to collect and store the URs in a repository; HLRmon that has been improved to allow the visualization of the storage accounting data. The current implementation uses a sensor that retrieves the data from the Grid IS. A great improvement is foreseen as soon as specific sensors will be provided for the different storage systems, in order to have more accurate measurements of the accounting metrics. When the proposed OGF UR 2.0 standard has been approved our prototype will be easily adapted to adopt the OGF standard UR format for storage accounting.

References

[1] https://www.egi.eu/ European Grid Infrastructure
[2] http://www.italiangrid.it/ Italian Grid Infrastructure
[3] Guarise A; Piro R, M, (2009) Distributed Grid Accounting System, Market Oriented Grid and Utility Computing., R. Buyya and K. Bubendorfer (eds), ISBN: 978-0470287682, Wiley Press, Hoboken, New Jersey, USA
[4] Piro R M; , Guarise A; Patania G; Werbrouck, A (2008). Using historical accounting information to predict resource usage of grid jobs. Future Generation Computing System , Volume 25, Issue 5, May 2009, pp. 499-510, doi:10.1016/j.future.2008.11.003
[5] http://www.to.infn.it/dgas/ Distributed Grid Accounting System: project home page
[6] http://www.eu-emi.eu/ European Middleware Infrastructure
[7] Cristofori, A.; Fattibene, E.; Gaido, L.; Recent improvements in HLRmon, an accounting portal suitable for national Grids; in the Proc. of the 2011 International Conference on Grid Computing and Applications, Jul 2011, Las Vegas, USA (pp. 10-13 ISBN: 1-60132-181-3)
[8] Fattibene, E.; Ferrari, T.; Gaido, L.; Misurelli, G.; Solagna, P.; HLRmon, a Tool for Advanced Visualization of Resource Utilization. Use Cases and Successful Applications of Distributed Computing Infrastructures (ISGC 2010)
[9] http://www.ogf.org/ Open Grid Forum
[10] Definition of a Storage Accounting Record - Jensen, H. T. et al - EMI-JRA1-TD002
[11] Andreozzi, S.; Burke, S.; Donno, F. Et alt.; GLUE Schema Specification, version 1.3, Jan 2007.