GFD-I.201
Usage Record-WG
ur-wg@ogf.org

H. T. Jensen, Nordic Data Grid Facility
J. K. Nilsen, Univ. of Oslo
P. Millar, DESY
R. Müller-Pfefferkorn, Technische Universität Dresden
Zs. Molnar, CERN
R. Zappi, INFN
January 2013

EMI StAR – Definition of a Storage Accounting Record

Status of This Document

This document provides information to the Grid community. It does not define any standards or technical recommendations. Distribution is unlimited.

Abstract

In this document the EU-project European Middleware Initiative (EMI) describes a storage accounting record (StAR), defined to reflect practical, financial and legal requirements of storage location, usage and space and data flow. The defined record might be the base for a standardized schema or an extension of an existing record like the OGF UR and this document is intended as information to be taken as input for incorporating storage resources into the OGF UR.

The described definition is agreed upon and will be implemented by the storage elements within EMI.

Contents

H. T. Jensen, Nordic Data Grid Facility
J. K. Nilsen, Univ. of Oslo
P. Millar, DESY
R. Müller-Pfefferkorn, Technische Universität Dresden
Zs. Molnar, CERN
R. Zappi, INFN

1 Introduction

Measuring consumption of storage resources is distinctively different from the measurement of the consumption of computing resources in batch jobs. On a computing resource it is quite easy to continuously collect CPU usage information with a high time resolution - the operating system itself does it and an accounting system can access these data. Additionally, the resource usage data are usually strictly increasing, e.g. CPU time or wall clock time. A collection of storage resource consumption is much harder to achieve. The usage of each participant will vary over time and only few systems do permanently record the storage use of the single users or groups with a fine granularity in time. Usage snapshot are taken from time to time only resulting in rough consumption estimation especially in environments with a high fluctuation rate. Nevertheless, there is a need to do storage accounting both for the provider of storage resources and the users. It is the base for billing and for the development of a storage infrastructure.

1.1 Record Structure & Content

The structure of the format described in this document can be split into logical parts, each describing an aspect of the resource consumption. The parts are
Resource: Fields describing the system the resource was consumed on. They can specify a certain subsystem of the storage system.
Consumption Details: Fields describing what the data is consuming, e.g. storage classes, directory path, etc.
Identity: Describes the person or group accountable for the resource consumption.
Resource Consumption: Fields describing how much of the described resource has been used.

Please note, that these logical sections are not necessarily directly reflected in record format. They are merely a good mental model to have in mind.

Most of the information in the record is common to all files, e.g. resource or identity description. Certain fields are aggregates over the consumed resources. This includes the consumption itself, but can also be consumption details, such as number of files.

The record is not intended to be used for describing intricate information about the consumption. E.g., filenames, per-file data or application metadata should not be included in the record. Such details are out of scope for the record and are not important for accounting of resource consumption.
To see examples of records, see section 8.

1.2 Including Additional Information

It is allowed to add additional fields with information in the record, e.g. for a more accurate description of the data that has consumed the resources.

If any user or group information is added it must be added under the SubjectIdentity block. This makes it possible to automatically remove user and group information.

1.3 Purpose

Standardization is one of the main goals of EMI. The definition of a standard storage accounting record (StAR) is a first step towards enabling a shared storage infrastructure. StAR allows for accounting and reporting of the resources consumed by persons and groups in a common way.
The purpose of this document is to describe the format specification of StAR, which is planned to be implemented by the EMI storage providers and is proposed as input to future work on usage record standards in the OGF Usage Record WG. Thus this definition is intended to be aligned with the next version of the OGF usage record. Though this accounting record definition was developed in the context of distributed storage in a Grid environment it is not specific to such a storage infrastructure. Due to the generic approach it can be used in any storage environment.

1.4 About this Document

StAR tries to define a standardized way to exchange storage consumption data. This document does not address in detail how the records should be used, how the accounting data are aggregated, nor does it attempt to dictate the format in which the accounting records are stored at a local site. Furthermore, nothing is said regarding the communication mechanisms employed to exchange the records, i.e. transport layer, framing, authentication, integrity, etc.

2 Notational Conventions

The key words 'MUST," "MUST NOT," "REQUIRED," "SHALL," "SHALL NOT," "SHOULD," "SHOULD NOT," "RECOMMENDED," "MAY," and "OPTIONAL" are to be interpreted as described in RFC 2119 [BRADNER], except that the words do not appear in uppercase.

3 Context

The specifications that are made in the following are based on a context that the reader needs to comprehend.

A Storage resource is a logical resource (either local or distributed) that allows an individual user or a group of users to store data. Such a system can contain single disks or can be created by pooling together physical storage media. This is transparent to the user and does not need to be considered when accounting for resource consumption.
Storage accounting is the recording and summarizing of the consumption of a storage resource by an individual user or a group of users in a specified time frame.

4 Releated Work

The record format described in this document is clearly related to the usage record (UR) format recommendation of the OGF 98 standard[1], as it tries to achieve a shared record format for accounting of consumed resources. Furthermore it shares several element names and semantics of the fields.
While efforts have been made to keep the StAR format close to the OGF UR format, the OGF UR does not allow for extensions and StAR and OGF UR have a limited number of properties in common. Hence, it has been decided to define a separate storage record format at this point in time.

[1] http://ogf.org/documents/GFD.98.pdf

5 Record Properties

This section describes the record properties and their fields and attributes. A summary of the fields is given in section 7, while examples of using the fields are given in section 8.

The format of the record is XML, using QNames. The currently defined name space is http://eu-emi.eu/namespaces/2011/02/storagerecord, denoted as "sr" in this document. All time and duration formats are ISO types[2]. These design choices are made in order to keep the format as close as possible to OGF usage record format.

Many of the properties presented in this section are optional, however a few are not. For the required properties, it is explicitly listed that they must be present in the record. None of the properties are allowed to be repeated.

A record should only represent a single identity. This identity can either be a person or a group of users. If a record contains both user and group information, the implementation should assume that the resources have been consumed by the user in the context of the group information.

5.1 StorageUsageRecord

This is the top container property of the record format when recording a single storage usage record.

- StorageUsageRecord MUST be present in the record.
- StorageUsageRecord MUST be either the top container property of the document or the direct child property of a StorageUsageRecords.
- StorageUsageRecord MUST NOT have a value.

Example

```
<sr:StorageUsageRecord>
<!—Record properties go in here -->
</sr:StorageUsageRecord>
```

5.2 StorageUsageRecords

This property can hold a number of StorageUsageRecord properties, i.e., act as a container for several storage usage records.

- StorageUsageRecords MUST only contain StorageUsageRecord elements.
- If present, StorageUsageRecords MUST be the root element of the record.
- StorageUsageRecords MUST NOT have a value.

Example

```
<sr:StorageUsageRecords>
<sr:StorageUsageRecord>
<!—Record properties go in here -->
</sr:StorageUsageRecord>
<sr:StorageUsageRecord>
<!—Record properties go in here -->
</sr:StorageUsageRecord>
</sr:StorageUsageRecords>
```

[2] http://www.iso.org/iso/catalogue_detail?csnumber=40874

5.3 RecordIdentity

This property describes the identity of the record. The field has two attributes: recordId and createTime. The recordId should be constructed in such a way that it is globally unique and records with the same value is not generated accidentally. Hence this field can be used to identify the record, and be used for duplicate detection. The createTime attribute describes when the record was created. It might differ from the time when the resource usage was measured (see section 4.16)

The field is similar to the field with the same name in the Usage Record standard.

- The RecordIdentity property MUST be present in the record.
- The RecordIdentity field MUST NOT have a value.
- The recordId attribute MUST be present in the record.
- The recordId attribute MUST have the type string.
- The createTime attribute MUST be present in the record.
- The createTime attribute MUST be an ISO timestamp.

Example

```
<sr:RecordIdentity
sr:createTime="2010-11-09T09:06:52Z"
sr:recordId="host.example.org/sr/87912469269276"/>
```

5.4 StorageSystem

This property describes the storage system on which the resources have been consumed. This value should be chosen in such a way that it globally identifies the storage system, on which resources are being consumed. E.g., the FQDN of the storage system could be used.

In Grid terms, this would be a storage element.

- The StorageSystem property MUST be present in the record.
- The StorageSystem field MUST have the type string.
- The StorageSystem value SHOULD be constructed in such a way, that it globally identifies the storage system.

Example

```
<sr:StorageSystem>host.example.org</sr:StorageSystem>
```

5.5 Site

This property describes the site at which the resource is located. This property should contain a descriptive name of the group of resources containing the storage system which is accounted for in the record. The Site value should be constructed in such a way that it is unique within the context where it is used.

- The Site field type MUST be a string.

Example

```
<sr:Site>ACME-University</sr:Site>
```

5.6 StorageShare

This property describes the part of of the storage system which is accounted for in the record. For a storage system, which is split into several logical parts, this can be used to account for consumption on each of these parts. The value should be able to identity the share of the storage system, given the storage system property.

- The StorageShare field type MUST be a string.

Example

```
<sr:StorageShare>pool-003</sr:StorageShare>
```

5.7 StorageMedia

This property describes the media type of storage that is accounted for in the record, e.g. “disk” or “tape”. This allows for accounting of different backend storage types.

- The StorageMedia field type MUST be a string.

Example

```
<sr:StorageMedia>disk</sr:StorageMedia>
```

5.8 StorageClass

This property describes the class of the stored data, e.g. ”pinned”, ”replicated”, ”precious”. This is a descriptive value, which allows the storage system to provide details about the stored data.

- The StorageClass field type MUST be a string.

Example

```
<sr:StorageClass>replicated</sr:StorageClass>
```

5.9 FileCount

This property describes the number of files which are accounted for in the record.

- The FileCount field type MUST be a positive non-zero integer.

Example

```
<sr:FileCount>42</sr:FileCount>
```

5.10 DirectoryPath

This property describes the directory path being accounted for. If the property is included in the record, the record should account for all usage in the directory and only that directory.

- The DirectoryPath field type MUST be a string.

Example

```
<sr:DirectoryPath>/projectA</sr:DirectoryPath>
```

5.11 SubjectIdentity

This property is a container for all user and group properties. Its purpose is to clearly mark one or more properties describing a user or group, i.e. the entity accountable for the resource consumption. The property is similar to the UserIdentity block in the OGF Usage Record format, but it can also be used for describing group affiliations.

- The SubjectIdentity property SHOULD be present in the record.
- The SubjectIdentity property SHOULD include at least one sub element.
- The SubjectIdentity field MUST NOT have a value.

Example

```
<sr:SubjectIdentity>
<!-- Various identity fields go in here -->
</sr:SubjectIdentity>
```

5.12 LocalUser

This property describes the local user name of the person accountable for the resource consumption on the storage system. It can be defined on the operating system level or as an internal user name in the storage system.

- If included, the LocalUser property MUST be under the SubjectIdentity.
- The LocalUser field type MUST be a string.

Example

```
<sr:LocalUser>johndoe</sr:LocalUser>
```

5.13 LocalGroup

This property describes the local user group accountable for the resource consumption on the storage system. It can be defined on the operating system level or as an internal group in the storage system.

- If included, the LocalGroup property MUST be under the SubjectIdentity.
- The LocalGroup field type MUST be a string.

Example

```
<sr:LocalGroup>binarydataproject</sr:LocalGroup>
```

5.14 UserIdentity

This property describes the global identity of the user accountable for the resource consumption. The property should identify the user globally, such that clashes do not happen accidentally, e.g. it could be an X509 identity.

- If included, the UserIdentity property MUST be under the SubjectIdentity.
- The UserIdentity field type MUST be a string.

Example

```
<sr:UserIdentity>/O=Grid/OU=example.org/CN=John Doe
```

```
</sr:UserIdentity>
```

5.15 Group

This property describes the global group accountable for the resource consumption. The property should identify the group globally, such that clashes do not happen accidentally, e.g. using a FQDN to construct it. In Grid terms, this would typically be the VO name.

- If included, the Group property MUST be under the SubjectIdentity.
- The Group field type MUST be a string.

Example

```
<sr:Group>binarydataproject.example.org</sr:Group>
```

5.16 GroupAttribute

This property describes supplemental traits of the group property, e.g., sub-groups, role or authority. This makes it possible to account for segments of a group, while still being able to account for the group as a whole. The property consists of a type which denotes the type of attribute and an actual value for the attribute.

- If included, the GroupAttribute property MUST be under the SubjectIdentity.
- The GroupAttribute property can be repeated.
- The Group property MUST exist in the record if GroupAttribute is specified.
- The GroupAttribute type and field values MUST exist.
- The GroupAttribute type MUST have the type string.
- The GroupAttribute field type MUST be a string.

Examples

```
<sr:GroupAttribute
sr:attributeType="role">production</sr:GroupAttribute>
<sr:GroupAttribute
sr:attributeType="subgroup">analysis</sr:GroupAttribute>
<sr:GroupAttribute sr:attributeType="authority">
/O=Grid/OU=example.org/CN=host/auth.example.org
</sr:GroupAttribute>
```

5.17 StartTime

This property describes a timestamp indicating the time at which the measured resource consumption started. Together with EndTime (see section 5.18) this defines a period over which the resource has been consumed.

- The StartTime field MUST be present in the record.
- The StartTime field type MUST be an ISO timestamp.
- The time zone may be specified as Z (UTC) or (+|-)hh:mm. Time zones that aren't specified are considered undetermined.

Example

```
<sr:StartTime>2010-10-11T09:31:40Z</sr:StartTime>
```

Implementation Note

The period defined by EndTime-StartTime could define the period over which the storage system measured a storage integral. In a simpler case, EndTime is the current time of mesasurment and StartTime the preceding measurement.

5.18 EndTime

This property describes a timestamp indicating the time at which the measured resource consumption ended. Together with StartTime (see section 5.17) this defines a period over which the resource has been consumed.

- The EndTime field MUST be present in the record.
- The EndTime field type MUST be an ISO timestamp.
- The time zone may be specified as Z (UTC) or (+|-)hh:mm. Time zones that aren't specified are considered undetermined.

Example

```
<sr:EndTime>2010-10-12T09:29:42Z</sr:EndTime>
```

Implementation Note

Same as for StartTime property (see section 5.17).

5.19 ResourceCapacityUsed

This property describes the number of bytes used on the storage system. This is the main metric for measuring resource consumption. It should include all resources for which the identity of the record is accountable for.

ResourceCapacityUsed can include reserved space, file metadata, space used for redundancy in RAID setups, tape holes, or similar. The decision about including such “additional” space is left to the resource owner but should be made known to the user e.g. via the usage policy. In contrary the LogicalCapacityUsed (see section 5.20) denotes the pure file size.

- The ResourceCapacityUsed property MUST be present in the record.
- The ResourceCapacityUsed field type MUST be a nonnegative integer.
- ResourceCapacityUsed SHOULD include all resources that are used to store the files. The value MAY also include resources that are no longer in use but are unavailable for reuse, as documented in the appropriate SLA or usage policy documents.

Example

```
<sr:ResourceCapacityUsed>14728</sr:ResourceCapacityUsed>
```

Implementation Note

Using bytes saves us from the argument of discussing if 1000 or 1024 should be used as a base. However, this also means that the number reported can be very large. Therefore any implementation should use at least a 64-bit integer to hold this variable (a signed 64-bit integer will overflow at 8 Exabytes).

5.20 LogicalCapacityUsed

This property describes an integer denoting the number of “logical” bytes used on the storage system by the identity of the record. By ‘logical” is meant the sum of bytes of the files stored, i.e.

excluding reservation and any underlying replicas of files (see ResourceCapacityUsed in section 5.19).

- The LogicalCapacityUsed field type MUST be a nonnegative integer.

Example

```
<sr:LogicalCapacityUsed>13617</sr:LogicalCapacityUsed>
```

Implementation Note

Same as for ResourceCapacityUsed property (see section 5.19).

5.21 ResourceCapacityAllocated

This property describes the number of bytes allocated on the storage system or storage share where appropriate. Depending on implementation this property may be equal to ResourceCapacityUsed (see section 2.4.19), however this property should only take into account space allocated to the entity described in the record, not resources used for redundancy in RAID setups, tape holes, or similar.

- The ResourceCapacityAllocated field type MUST be a nonnegative integer.

Example

```
<sr:ResourceCapacityAllocated>14624</sr:ResourceCapacityAllocated>
```

Implementation Note

Same as for ResourceCapacityUsed property (see section 5.19).

6 Intentionally Left Out Properties

In the preparation phase yielding the current draft of the record a number of properties were in discussion to be included. They are listed here for the sake of completeness.

FileNames Providing a list of file names in the record would allow per-file accounting, or allow certain files to be accounted for separately. However, the available properties like path, storage system, group, and user provide a sufficient definition of the resource consumption.

File Metadata The focus of the record is accounting. Any per-file data or metadata are out of scope of the record.

SpaceAvailable This property would describe how much space is available for the identity on the storage system. This property would however not report any form of consumption and is often difficult to determine.

Transfer Information The original suggestion included properties for describing how much data has been transferred. There are however a range of issues with this: A. network resources are not storage resources. B: Users transferring data are not necessarily the user owning the data. Thus, the accounting of the network usage should be defined elsewhere.

7 Field Summary

Field	Short Description	Field Type	Requirement
StorageUsageRecord	Top container		REQUIRED
StorageUsageRecords	Container for grouping multiple		OPTIONAL

	records		
RecordIdentity Attributes:	Identity of the record		
recordID	Global unique record ID	String	REQUIRED
createTime	Record creation time	ISO timestamp	REQUIRED
StorageSystem	Storage system description	String	REQUIRED
Site	Site where the storage system is located	String	OPTIONAL
StorageShare	Part of the storage system accounted for	String	OPTIONAL
StorageMedia	Media type	String	OPTIONAL
StorageClass	Class of stored data	String	OPTIONAL
FileCount	Number of files accounted for	String	OPTIONAL
DirectoryPath	Directory path accounted for	String	OPTIONAL
SubjectIdentity Elements:	Container for user and group properties		RECOMMENDED
LocalUser	User name on the storage system	String	OPTIONAL
LocalGroup	Group name on the storage system	String	OPTIONAL
UserIdentity	Global user ID	String	OPTIONAL
Group	Global group	String	OPTIONAL
GroupAttribute	Group attribute	String	OPTIONAL
attributeType	Type of attribute	String	REQUIRED if GroupAttribute is defined
StartTime	Start time of measurement	ISO timestamp	REQUIRED
EndTime	End time for measurement	ISO timestamp	REQUIRED
ResourceCapacityUsed	Number of bytes used on the storage system	Nonnegative Integer	REQUIRED
LogicalCapacityUsed	Number of “logical” bytes used on the storage system	Nonnegative Integer	OPTIONAL
ResourceCapacityAllocated	Number of bytes allocated on the storage system	Nonnegative Integer	OPTIONAL

Table 1: Summary of the fields of StAR

8 Record Examples

8.1 Minimal Example

Minimal record that is actually useful. There is no identity block, which should be interpreted as the record accounts for all usage on the storage system.

```
<sr:StorageUsageRecord
 xmlns:sr="http://eu-emi.eu/namespaces/2011/02/storagerecord">
  <sr:RecordIdentity sr:createTime="2010-11-09T09:06:52Z"
                     sr:recordId="host.example.org/sr/87912469269276"/>
  <sr:StorageSystem>host.example.org</sr:StorageSystem>
  <sr:StartTime>2010-10-11T09:31:40Z</sr:StartTime>
  <sr:EndTime>2010-10-12T09:29:42Z</sr:EndTime>
  <sr:ResourceCapacityUsed>13617</sr:ResourceCapacityUsed>
</sr:StorageUsageRecord>
```

8.2 Local Usage Example

Example of how a record accounting for a local user could look like.

```
<sr:StorageUsageRecord
 xmlns:sr="http://eu-emi.eu/namespaces/2011/02/storagerecord">
  <sr:RecordIdentity sr:createTime="2010-11-09T09:06:52Z"
                     sr:recordId="host.example.org/sr/87912469269276"/>
  <sr:StorageSystem>host.example.org</sr:StorageSystem>
  <sr:SubjectIdentity>
    <sr:LocalUser>johndoe</sr:LocalUser>
  </sr:SubjectIdentity>
  <sr:StorageMedia>tape</sr:StorageMedia>
  <sr:FileCount>42</sr:FileCount>
  <sr:StartTime>2010-10-11T09:31:40Z</sr:StartTime>
  <sr:EndTime>2010-10-12T09:29:42Z</sr:EndTime>
  <sr:ResourceCapacityUsed>913617</sr:ResourceCapacityUsed>
</sr:StorageUsageRecord>
```

8.3 Grid Usage Example

Example of how a record accounting for Grid usage could look like.

```
<sr:StorageUsageRecord
 xmlns:sr="http://eu-emi.eu/namespaces/2011/02/storagerecord">
  <sr:RecordIdentity sr:createTime="2010-11-09T09:06:52Z"
                     sr:recordId="host.example.org/sr/87912469269276"/>
```

```
   <sr:StorageSystem>host.example.org</sr:StorageSystem>
   <sr:Site>ACME-University</sr:Site>
   <sr:StorageShare>pool-003</sr:StorageShare>
   <sr:SubjectIdentity>
     <sr:Group>binarydataproject.example.org</sr:Group>
     <sr:GroupAttribute
sr:attributeType="subgroup">ukusers</sr:GroupAttribute>
   </sr:SubjectIdentity>
   <sr:StorageMedia>disk</sr:StorageMedia>
   <sr:FileCount>42</sr:FileCount>
   <sr:StartTime>2010-10-11T09:31:40Z</sr:StartTime>
   <sr:EndTime>2010-10-12T09:29:42Z</sr:EndTime>
   <sr:ResourceCapacityUsed>14728</sr:ResourceCapacityUsed>
   <sr:LogicalCapacityUsed>13617</sr:LogicalCapacityUsed>
</sr:StorageUsageRecord>
```

8.4 Full Example

Example using all fields.

```
<sr:StorageUsageRecords
      xmlns:sr="http://eu-emi.eu/namespaces/2011/02/storagerecord">
   <sr:StorageUsageRecord>
     <sr:RecordIdentity sr:createTime="2010-11-09T09:06:52Z"

     sr:recordId="host.example.org/sr/87912469269276"/>
     <sr:StorageSystem>host.example.org</sr:StorageSystem>
     <sr:Site>ACME-University</sr:Site>
     <sr:StorageShare>pool-003</sr:StorageShare>
     <sr:StorageMedia>disk</sr:StorageMedia>
     <sr:StorageClass>replicated</sr:StorageClass>
     <sr:FileCount>42</sr:FileCount>
     <sr:DirectoryPath>/home/projectA</sr:DirectoryPath>
     <sr:SubjectIdentity>
       <sr:LocalUser>johndoe</sr:LocalUser>
       <sr:LocalGroup>projectA</sr:LocalGroup>
       <sr:UserIdentity>/O=Grid/OU=example.org/CN=John
Doe</sr:UserIdentity>
       <sr:Group>binarydataproject.example.org</sr:Group>
       <sr:GroupAttribute
sr:attributeType="subgroup">ukusers</sr:GroupAttribute>
```

```
      </sr:SubjectIdentity>
      <sr:StartTime>2010-10-11T09:31:40Z</sr:StartTime>
      <sr:EndTime>2010-10-12T09:29:42Z</sr:EndTime>
      <sr:ResourceCapacityUsed>14728</sr:ResourceCapacityUsed>
      <sr:LogicalCapacityUsed>13617</sr:LogicalCapacityUsed>
      <sr:ResourceCapacityAllocated>14624</sr:ResourceCapacityAllocated>
   </sr:StorageUsageRecord>
</sr:StorageUsageRecords>
```

9 Security Considerations

There may be security concerns that should be addressed with respect to usage data. Possible security issues might include:

- Non-repudiation
- Confidentiality of certain elements
- Integrity
- Secure Transport

Recommendation of required solutions for these security concerns is out of scope for this layer. Another layer should address the necessary security requirements.

10 Glossary

StAR	**St**orage **A**ccounting **R**ecord
FQDN	Fully Qualified Domain Name
ISO	International Organization for Standardization
SLA	Service Level Agreement
NDGF	Nordic Data Grid Facility
EGI	European Grid Infrastructure
OSG	Open Science Grid
INFN	Istituto Nazionale di Fisica Nucleare
CMS	Compact Muon Solenoid
ATLAS	A Toroid LHC Aparatus

11 Contributors

Henrik Thostrup Jensen
Nordic Data Grid Facilty / NORDUnet
htj@nordu.net

Paul Millar
DESY
paul.millar@desy.de

Ralph Müller-Pfefferkorn
Technische Universität Dresden
ralph.mueller-pfefferkorn@tu-dresden.de

Zsolt Molnár
CERN
zsolt.molnar@cern.ch

Jon Kerr Nilsen (Corresponding Author)
University of Oslo
j.k.nilsen@fys.uio.no

Riccardo Zappi
INFN
riccardo.zappi@cnaf.infn.it

12 Acknowledgments

During the making of the document, contacts have been made with organizations and potential user groups, such as NDGF, OGF, EGI, OSG, INFN, CMS and ATLAS, and internally in EMI with JRA1.4 (security) and JRA1.6 (standardization) resulting in valuable feedback and active cooperation. Special gratitude goes to Brian Bockleman (OSG), Andrea Cristofori (INFN) and John Alan Kennedy (RZG/EUDAT) for valuable comments.

13 Intellectual Property Statement

14 Disclaimer

of the information herein will not infringe any rights or any implied warranties of merchantability or fitness for a particular purpose.

15 Full Copyright Notice

16 References

[BRADNER] Scott Bradner. Key Words for Use in RFCs to Indicate Requirement Levels, RFC 2119. The Internet Society. March 1997. http://tools.ietf.org/html/rfc2026