

## Accounting Interchange Natural Language Description (Requirements)

### Status of This Memo

This memo provides information to the Grid community regarding the Usage Record format requirements. Distribution is unlimited.

### Copyright Notice

Copyright © Global Grid Forum (2003). All Rights Reserved.

<b>1. ABSTRACT.....</b>	<b>2</b>
<b>2. BASE PROPERTIES.....</b>	<b>2</b>
2.1 CHARGE .....	2
2.2 STATUS .....	3
2.3 JOBNAME .....	3
2.4 NETWORK .....	3
2.5 DISK.....	3
2.6 MEMORY .....	3
2.7 WALLTIME .....	4
2.8 CPUTIME .....	4
2.9 NODECOUNT .....	4
2.10 PROCESSORS .....	4
2.11 NONCOMPUTE.....	5
2.12 END TIME.....	5
2.13 START TIME .....	5
2.14 MACHINENAME .....	6
2.15 HOST .....	6
2.16 SUBMIT HOST.....	6
2.17 QUEUE .....	6
2.18 LOCALJOBID .....	6
2.19 GLOBALJOBID.....	7
2.20 PROJECT NAME.....	7
2.21 EXTENSION .....	7
2.22 LOCALUSERNAME.....	7
2.23 GLOBALUSERNAME .....	8
<b>3. META PROPERTIES .....</b>	<b>8</b>
3.1 UNITS.....	8
3.2 METRIC .....	9
<b>4. JOB LEVEL ACCOUNTING.....</b>	<b>9</b>
<b>5. AGGREGATE ACCOUNTING.....</b>	<b>10</b>
<b>6. SUPPORTED DATA TYPES .....</b>	<b>10</b>
6.1 STRING.....	10
6.2 INTEGER.....	10
6.3 POSITIVE INTEGER.....	10
6.4 BOOLEAN .....	10
6.5 FLOAT .....	10

6.6	TIMESTAMP .....	10
7.	SECURITY REQUIREMENTS .....	10
8.	INTELLECTUAL PROPERTY STATEMENT .....	10
9.	FULL COPYRIGHT NOTICE.....	11
10.	REFERENCES .....	12

## 1. Abstract

For resources to be shared, sites must be able to exchange basic accounting and usage data in a common format. This format must encompass both job level accounting, and aggregate accounting. This document describes the requirements in natural language form for a Usage Record standard.

## 2. Base Properties

The following is a list of base properties that define common usage record requirements for both job level and aggregate properties. The definitions for the type of data (integer, string) are the base data types, as defined in section 6. Any additional restrictions on the legal values for a specific property are noted within the base property definition. The naming convention and naming style of these properties are implementation specific, and should use the referring name as a guideline for any naming recommendations.

Properties may carry certain meta-properties that can assist in establishing semantic meaning from the actual property value, such as “metric” and “units”. Common meta-properties are addressed more fully in section 3.

### 2.1 Charge

This property represents the total charge of the job in the system’s allocation unit. For example 100, 200, or 3000. The meaning of this charge will be site dependent. The value for this property MAY include premiums or discounts assessed on the actual usage represented within this record. Therefore, the reported charge may not be directly reconstructed from the specific usage reported.

- This property SHOULD be referred to as charge
- It MUST contain data of type integer
- This property is optional

## 2.2 Status

This property will represent the completion status of the job. For example, this may represent the exit status of an interactive running process or the exit status from the batch queuing system's accounting record. The semantic meaning of status is site dependent

- This property SHOULD be referred to as status
- This property MUST contain data of type integer
- This property is optional

## 2.3 Jobname

The job or application name. For example, this could be the name of the executable that ran, or the name of the batch queuing system's name for the job

- This property SHOULD be referred to as jobname
- This property MUST contain data of type string
- This property is optional

## 2.4 Network

The network used by the job.

- This property SHOULD be referred to as network
- This property MUST contain data of type integer
- This property is optional
- Meta-properties
  - Units SHOULD be specified
  - Metric MAY be specified

## 2.5 Disk

Disk storage used

- This property SHOULD be referred to as disk
- This property MUST contain data of type integer
- This property is optional
- Meta-properties
  - Units MUST be specified
  - Metric MAY be specified

## 2.6 Memory

The maximum amount of virtual memory used by all concurrent processes in the job

- This property SHOULD be referred to as memory
- This property MUST contain data of type integer
- This property is optional

- Meta-properties
  - Units MUST be specified
  - Metric MAY be specified

## 2.7 Walltime

Wall clock time that elapsed while the job was running

- This property SHOULD be referred to as walltime
- This property MUST contain data of type positive integer
- This property is optional
- Meta-properties
  - Units MUST be specified

## 2.8 Cputime

CPU time used, summed over all processes in the job

This property SHOULD be referred to as cputime

- This property MUST contain data of type positive integer
- This property is optional
- Meta-properties
  - Units MUST be specified

## 2.9 NodeCount

Number of nodes used. A node definition may be dependent on the architecture, but typically a node is a physical machine. For example a cluster of 16 physical machines with each machine having one processor each is a 16 “node” machine, each with one “processor”. A 16 processor SMP machine however, is 1 physical node (machine) with 16 processors.

- This property SHOULD be referred to as nodecount
- This property MUST contain data of type positive integer
- This property is optional

## 2.10 Processors

The number of processors used or requested. A processor definition may be dependent on the machine architecture. Typically processor is equivalent to the number of physical CPU’s used. For example, if a job uses two cluster “nodes”, each node having 16 cpu’s each, the number of total processors would be 32.

- This property SHOULD be referred to as processors
- This property MUST contain data of type positive integer
- This property is optional

- Meta-properties
  - ConsumptionRate
    - This meta property defines the consumption rate for this usage.
    - This meta-property MUST contain data of type float.
    - This meta-property MAY be specified.

## 2.11 NonCompute

A noncomputing resource used by a job. For example, consulting hours are considered a non-compute resource

- This property SHOULD be referred to as noncompute
- This property MUST contain data of type integer
- This property is optional
- Meta-properties
  - Units MAY be specified
  - Metric MAY be specified
  - Name
    - This meta-property should be referred to as name
    - This meta-property MUST contain data of type string
    - This meta property is optional
    - The value of this meta-property should contain an explanation of the non-compute resource used.

## 2.12 End Time

The time at which the job completed. The value of this property may depend on the selected queue system. For example some systems include time to stage files, others do not.

- This property SHOULD be referred to as endtime
- This property MUST contain data of type timestamp
- This property is optional

## 2.13 Start Time

The time at which the job started. The value of this property may depend on the selected queue system. For example some systems include time to stage files, others do not.

- This property SHOULD be referred to as starttime.
- This property MUST contain data of type timestamp
- This property is optional

### 2.14 MachineName

A descriptive name of the machine on which the job ran. This may be a system hostname, or may be a sites name for a cluster of machines.

- This property SHOULD be referred to as machinename
- This property MUST contain data of type string
- This property is optional

### 2.15 Host

The system hostname on which the job ran.

- This property SHOULD be referred to as host
- This property MUST contain data of type string
- This property is optional
- Meta-properties
  - Primary
    - This meta-property indicates whether the indicated host acted as the primary host for the execution of the job
    - The meta-property MUST contain data of type Boolean.
    - The default value for this meta-property is false.

### 2.16 Submit Host

The system hostname from which the job was submitted.

- This property SHOULD be referred to as submithost
- This property MUST contain data of type string
- This property is optional

### 2.17 Queue

The name of the queue from which the job was executed or submitted.

- This property SHOULD be referred to as queue
- This property MUST have data of type string
- This property is optional

### 2.18 LocalJobId

The local job identifier as assigned by the batch queue, or PID of the job for interactive jobs.

- This property SHOULD be referred to as localjobid
- This property MUST have data of type integer
- This property is optional

### 2.19 GlobalJobId

The global job identifier as assigned by the executing site.

- This property SHOULD be referred to as globaljobid
- This property MUST have data of type string
- This property is optional

### 2.20 Project name

The project associated with the resource usage reported with this record. Some accounting systems define this is the ACID, which is often the same as the users Unix group from /etc/passwd. The project is also referred to as the effective gid under which the job consumed resources on some systems.

- This property SHOULD be referred to as projectname
- This property MUST have data of type string
- This property is optional
- There SHOULD be no restriction on the length or available characters

### 2.21 Extension

For sites that may want to exchange data that is not one of the base properties (for example, perhaps grid telescope power ) the Extension property can be used to encode any type of usage information. The sites can agree on the meta properties supported for each extension.

- This property SHOULD be referred to as extension
- This property MUST have data of type string
- This property is optional
- Meta-properties
  - Units may be supported
  - Metric may be supported
  - Name may be supported
    - This meta-property must have data of type string.

### 2.22 LocalUsername

The local identity of the user associated with the resource consumption reported in this Usage Record. This user is often referred to as the requesting user. For example, the value may be the user's login name corresponding to the users uid in the /etc/passwd file on Unix systems.

- This property SHOULD be referred to as localusername
- This property MUST have data of type string
- This property is optional

### 2.23 GlobalUsername

The global identity of the user associated with the resource consumption reported in this Usage Record. For example, the value may be the distinguished name from the users certificate.

- This property SHOULD be referred to as globalusername
- This property MAY have data of type string.
- This property is optional

## 3. Meta Properties

Meta properties are associated with individual base properties to provide additional information and semantic meaning of the value for a base property. The meta properties outlined below are commonly encountered and should be supported for the indicated base properties.

### 3.1 Units

This meta-property represents the unit of measure that should be applied to the value of its owning base property. As an example, memory can be expressed in bytes, kilobytes, megabytes and gigabytes. The representation of these units of measure are: B, KB, MB and GB respectively. If a base property supports this meta-property, but it is omitted, the listed default value may be assumed.

- This meta-property MUST be called units
- This meta-property MUST have data of type string
- The supported values for each base property that MAY contain this meta-property are summarized in Table 1.

**Table 1 Supported and Default Values for Units Meta-property**

Property Name	Supported Units	Default Value
Network	B, KB, MB, GB	MB
Disk	B, KB, MB, GB	MB
Memory	B, KB, MB, GB	MB
Walltime	Seconds, minutes, hours	Seconds
Cputime	Seconds, minutes, hours	Seconds
EndTime	Seconds, minutes, hours	Seconds
StartTime	Seconds, minutes, hours	Seconds



### 3.2 Metric

This meta-property identifies the type of measurement used for quantifying the associated resource consumption if there are multiple methods to measure resource usage. As an example, disk usage may be measured as total, average, minimum or maximum usage. However, even if pertinent to the assessed charge, this meta-property does not attempt to differentiate between requested and utilized quantities of resource usage. If a base property supports this meta-property, but it is omitted, the listed default value may be assumed.

- This meta-property **MUST** be called metric
- This meta-property **MUST** have data of type string
- The supported values for each base property that **MAY** contain this meta-property are summarized in Table 1.

**Table 2 Supported and Default Values for Metric Meta-property**

Property Name	Supported Metric	Default Metric
Network	Average, Total, Max, Min, Requested	Total
Disk	Average, Total, Max, Min, Requested	Total
Memory	Average, Total, Max, Min, Requested	Total

## 4. Job Level Accounting

Job level accounting reports accounting data at the job level. PBS and LoadLeveller for example, use job level accounting where each job is assigned a unique id and the accounting for that job is reported for each id. The properties listed in this document are not required to appear within the usage record structure in the order in which they are defined.

- This type of usage record **MAY** contain any of the base properties.
- This type of usage record **MUST** contain at least one of the following properties:
  - LocalJobId
  - GlobalJobId
- The Host property **MAY** repeat within the record. Each occurrence of the Host property **SHOULD** contain a unique value.
- Job Level properties **SHOULD** be contained within a “job” property for each job

## 5. Aggregate Accounting

Aggregate accounting reports the accounting data in aggregate (summarized form). The properties listed in this document are not required to appear within the usage record structure in the order in which they are define.

- Aggregate accounting MAY contain any of base properties listed in this document, which appear exactly once per job.
- Aggregate properties SHOULD be contained within an “aggregate” property.

## 6. Supported Data Types

### 6.1 String

Data of this type has no required restrictions on the length or available characters.

### 6.2 Integer

### 6.3 Positive Integer

Data of this type must have a value of zero or greater.

### 6.4 Boolean

Data of this type may have a value of either true or false.

### 6.5 Float

### 6.6 Timestamp

Data of this type must comply with the UTC time zone format specified in ISO 8601

## 7. Security Requirements

There are no explicit security considerations for these requirements.

## 8. Contributors

### Authors

Rodney Mach  
University of Michigan  
[rmach@umich.edu](mailto:rmach@umich.edu)  
734-936-1985

Rebekah Lepro  
NASA Ames Research Center  
[rlepro@mail.arc.nasa.gov](mailto:rlepro@mail.arc.nasa.gov)  
650-604-4359

GWD-R

Scott M. Jackson  
Pacific Northwest National Laboratory  
Scott.Jackson@pnl.gov  
509-376-2205

### **Editor**

Rodney Mach  
University of Michigan  
[rmach@umich.edu](mailto:rmach@umich.edu)  
734-936-1985

### **Contributors**

Laura McGinnis  
Pittsburgh Supercomputing Center  
[lfm@psc.edu](mailto:lfm@psc.edu)  
412-268-5642

Jian Zhang  
Xcerla Corporation  
[jian@xcerla.com](mailto:jian@xcerla.com)  
408-807-9069

## **9. Intellectual Property Statement**

The GGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the GGF Secretariat.

The GGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the GGF Executive Director.

## **10.Full Copyright Notice**

Copyright (C) Global Grid Forum (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation

may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the GGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the GGF Document process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the GGF or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE GLOBAL GRID FORUM DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

## 11. References

### KEYWORDS

[RFC 2119](#). *Key words for use in RFCs to Indicate Requirement Levels*. S. Bradner. March 1997.  
<http://www.ietf.org/rfc/rfc2119.txt>

**Resource Survey** [Editor note: survey draft currently in GGF document review pipeline].