

## 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 (2002). All Rights Reserved.

<b>1. ABSTRACT.....</b>	<b>2</b>
<b>2. INTRODUCTION.....</b>	<b>2</b>
<b>3. BASE PROPERTIES.....</b>	<b>2</b>
3.1 CHARGE .....	3
3.2 STATUS .....	3
3.3 JOBNAME.....	3
3.4 NETWORK .....	4
3.5 DISK.....	4
3.6 MEMORY .....	4
3.7 WALLDURATION.....	5
3.8 CPUDURATION.....	5
3.9 NODECOUNT .....	5
3.10 PROCESSORS .....	6
3.11 END TIME.....	6
3.12 START TIME .....	7
3.13 MACHINENAME .....	7
3.14 HOST .....	7
3.15 SUBMIT HOST.....	7
3.16 QUEUE .....	8
3.17 LOCALJOBID .....	8
3.18 PROCESSID.....	8
3.19 GLOBALJOBID.....	8
3.20 PROJECT NAME.....	9
3.21 EXTENSION .....	9
3.22 LOCALUSERID .....	9
3.23 GLOBALUSERNAME .....	9
3.24 SWAP .....	10
3.25 SERVICE LEVEL.....	10
3.26 RECORD IDENTITY .....	10
<b>4. META PROPERTIES .....</b>	<b>11</b>
4.1 DESCRIPTION.....	11
4.2 UNITS DEFINITION .....	11
4.2.1 Volume Units .....	11
4.2.2 Phase Units.....	12
4.3 METRIC .....	12
<b>5. JOB LEVEL ACCOUNTING.....</b>	<b>12</b>

<b>6.</b>	<b>AGGREGATE ACCOUNTING .....</b>	<b>13</b>
<b>7.</b>	<b>SUPPORTED DATA TYPES .....</b>	<b>13</b>
7.1	STRING .....	13
7.2	INTEGER .....	13
7.3	POSITIVE INTEGER.....	13
7.4	BOOLEAN .....	13
7.5	FLOAT .....	13
7.6	TIMESTAMP .....	13
7.7	DOMAINNAME .....	13
<b>8.</b>	<b>SECURITY REQUIREMENTS .....</b>	<b>13</b>
<b>9.</b>	<b>CONTRIBUTORS .....</b>	<b>14</b>
<b>10.</b>	<b>INTELLECTUAL PROPERTY STATEMENT .....</b>	<b>14</b>
<b>11.</b>	<b>FULL COPYRIGHT NOTICE.....</b>	<b>15</b>
<b>12.</b>	<b>REFERENCES .....</b>	<b>15</b>

## 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. Introduction

The natural language document will focus on the representation of resource consumption data. The purpose of this document is to outline the basic building blocks of the accounting record, and how to properly represent them. All other tangential concerns such as the use, transport mechanism, and security are out of scope for this representation layer.

## 3. 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 7. 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”, “description”, and “units”. Common meta-properties are addressed more fully in section 4.

### 3.1 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.
- Meta-properties
  - Units MAY be specified.
  - Description MAY be specified.
  - Formula MAY be specified that describes how the charge was arrived at. There is no required format for the formula.

### 3.2 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.
- This property MUST support the following values:
  1. aborted – A policy or human intervention caused the job to cease execution.
  2. completed – The execution completed.
  3. failed – Execution halted without external intervention.
  4. held – Execution is held at the time this usage record was generated.
  5. queued – Execution was queued at the time this usage record was generated.
  6. started – Execution started at the time this usage record was generated.
  7. suspended – Execution was suspended at the time this usage record was generated.

### 3.3 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.
- Meta-properties

- Description MAY be specified

### 3.4 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. The metrics that MUST be supported are:
    1. average – the average flow rate over the entire usage window.
    2. total – rate of transfer in the specified unit. This is the default.
    3. min – minimum flow rate in the specified.
    4. max – flow rate in the specified unit.

### 3.5 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 MAY be specified
  - Description MAY be specified
  - Type MAY be specified. The types that MUST be supported are:
    1. scratch
    2. temp
  - Metric MAY be specified. The metrics that MUST be supported are:
    1. average
    2. total
    3. min
    4. max

### 3.6 Memory

The amount of 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.
- Description **MAY** be specified.
- Metric **MAY** be specified. The metrics that **MUST** be supported are:
  1. average
  2. total
  3. max
  4. min
- Type **MAY** be specified. The types that **MUST** be supported are:
  1. shared
  2. physical
  3. dedicated

### 3.7 WallDuration

Wall clock time that elapsed while the job was running.

- This property **SHOULD** be referred to as wall.
- This property **MUST** contain data of type timestamp.
- This property is optional.
- Meta-properties
  - Description **MAY** be specified.

### 3.8 CpuDuration

CPU time used, summed over all processes in the job.

This property **SHOULD** be referred to as cpuduration.

- This property **MUST** contain data of type timestamp.
- This property is optional.
- Meta-properties.
  - Duration **MAY** be specified.
  - Type **MAY** be specified. The types that **MUST** be supported are:
    1. user
    2. system

### 3.9 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.
- Meta-properties
  - Description MAY be specified.
  - Metric MAY be specified. The metrics that MUST be supported are:
    1. average
    2. total
    3. min
    4. max

### 3.10 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
  - Description MAY be specified.
  - Metric MAY be specified. The metrics that MUST be supported are:
    1. average
    2. total
    3. min
    4. max
  - 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.

### 3.11 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
- Meta-properties
  - Description MAY be specified

### 3.12 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
- Meta-properties
  - Description MAY be specified

### 3.13 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.
- Meta-properties
  - Description MAY be specified.

### 3.14 Host

The system hostname on which the job ran.

- This property SHOULD be referred to as host.
- This property MUST contain data of type domain name.
- This property is optional.
- Meta-properties
  - Description MAY be specified.
  - 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.

### 3.15 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 domain name.
- This property is optional.
- Meta-properties
  - Description MAY be specified.

### 3.16 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.
- Meta-properties
  - Description MAY be specified.

### 3.17 LocalJobId

The local job identifier as assigned by the batch queue

- This property SHOULD be referred to as localjobid.
- This property MUST have data of type string.
- This property is optional.
- Meta-properties
  - Description MAY be specified.

### 3.18 ProcessId

The process id of the jobs (PID).

- This property SHOULD be referred to as processid.
- This property MUST have data of type integer.
- This property is optional.
- Meta-properties
  - Description MAY be specified.

### 3.19 GlobalJobId

The global job identifier as assigned by a metascheduler or federation scheduler.

- This property SHOULD be referred to as globaljobid.
- This property MUST have data of type string.
- This property is optional.
- Meta-properties
  - Description MAY be specified.



### 3.20 Project name

The project associated with the resource usage reported with this record. Some accounting systems define this is the ACID. 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.
- Meta-properties
  - Description MAY be specified.

### 3.21 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.

### 3.22 LocalUserId

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 localuserid.
- This property MUST have data of type string.
- This property is optional.

### 3.23 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

### 3.24 Swap

This property specifies the swap usage

- This property SHOULD be referred to as swap.
- This property MUST have a data of type integer.
- This property is optional.
- Meta-properties
  - Description MAY be specified.

### 3.25 Service Level

This property identifies the quality of service associated with the resource consumption. For example, service level may represent a priority associated with the usage.

- This property SHOULD be referred to as record identity.
- This property MUST have a data of type string.
- This property is optional.

### 3.26 Record Identity

A record identity uniquely defines a record in the usage record.

- This property SHOULD be referred to as recordidentity.
- This property MUST have a data of type string.
- This property MUST exist.
- This property MUST be unique.
- Meta-properties
  - Create time of the record MUST be specified.

## 4. 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.**

### 4.1 Description

The description provides a mechanism for additional, optional information to be attached to a Usage Record base property. The value of this meta-property MAY provide clues to the semantic context to use while interpreting or examining the value of the owning base property.

### 4.2 Units Definition

The following section details the supported units for volume and phase units that apply to the properties that have those units. Some properties MAY have both volume and phase units. For example, megabytes per second. Phase and unit definitions MUST supply both the volume unit, and the phase unit separately.

#### 4.2.1 Volume Units

The units for storage and memory use volume units. The legal values for volume units are below. Each value combines a prefix that identifies a multiplier and a suffix that identifies the base quantity. For example, the suffix -B represents volume in bytes and the prefix Ki- implies a kibi multiplier. Therefore, KiB as the designated unit of measure value MUST be interpreted as kibibytes.

**Table 1 Volume Units**

Abbreviation	Definition	Quantity
b	bit	1 bit
B	byte	1 byte
KB	kilobyte	1000 bytes
KiB	kibibyte	1024 bytes
MB	megabyte	1000 KB
MiB	mibibyte	1024 KiB
GB	gigabyte	1000 MB
GiB	gibibyte	1024 MiB
TB	terabyte	1000 GB
TiB	tibibyte	1024 GiB
PB	petabyte	1000 TB
PiB	pibibyte	1024 TiB
EB	exabyte	1000 PB
EiB	exibyte	1024 PiB
Kb	kilobit	1000 bits

Kib	kibibit	1024 bits
Mb	megabit	1000 Kb
Mib	mibibit	1024 Kib
Gb	gigabit	1000 Mb
Gib	gibibit	1024 Mib
Tb	terabit	1000 Gb
Tib	tibibit	1024 Gib
Pb	petabit	1000 Tb
Pib	pibibit	1024 Tib
Eb	exabit	1000 Pb
Eib	exibit	1024 Pib

#### 4.2.2 Phase Units

Phase units represents the duration of time that is relevant to the usage reported. The lexical representation for **duration** is the [\[ISO 8601\]](#) extended format  $PnYnMnDTnHnMnS$ , where  $nY$  represents the number of years,  $nM$  the number of months,  $nD$  the number of days, 'T' is the date/time separator,  $nH$  the number of hours,  $nM$  the number of minutes and  $nS$  the number of seconds. The number of seconds can include decimal digits to arbitrary precision.

#### 4.3 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**

### 5. 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

## 6. 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.

## 7. Supported Data Types

### 7.1 String

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

### 7.2 Integer

### 7.3 Positive Integer

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

### 7.4 Boolean

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

### 7.5 Float

### 7.6 Timestamp

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

### 7.7 DomainName

Data of this type must comply with RFC 1034 format for fully qualified domain names. The constraints are a maximum 255 characters long, containing only alphabetic and numeric characters, the “-“, and the “.” characters.

## 8. Security Requirements

There may be security concerns that may be of concern with the usage data. Possible security concerns are:

1. Non-repudiation
2. Confidentiality of certain elements
3. Integrity
4. Secure Transport

There is not a required solution for these security concerns as it is out of scope for this layer. Another layer can provide these security requirements if appropriate.

## 9. 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

Scott M. Jackson  
Pacific Northwest National Laboratory  
[Scott.Jackson@pnl.gov](mailto:Scott.Jackson@pnl.gov)  
509-376-2205

### 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

## 10. 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.

## 11.Full Copyright Notice

Copyright (C) Global Grid Forum (1/18/2003). 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."

## 12.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

Mi Young Koo  
<http://www.psc.edu/~lfm/Grid/UR-WG/>