

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.

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

6.	<u>AGGREGATE ACCOUNTING</u>	12
7.	<u>SUPPORTED DATA TYPES</u>	13
7.1	<u>STRING</u>	13
7.2	<u>INTEGER</u>	13
7.3	<u>POSITIVE INTEGER</u>	13
7.4	<u>BOOLEAN</u>	13
7.5	<u>FLOAT</u>	13
7.6	<u>TIMESTAMP</u>	13
7.7	<u>DOMAINNAME</u>	13
8.	<u>SECURITY REQUIREMENTS</u>	13
9.	<u>APPENDIX</u>	13
9.1	<u>SURVEY BACKGROUND</u>	13
9.2	<u>SURVEY RESULTS</u>	14
9.3	<u>COMMON USAGE RECORD FIELDS</u>	19
10.	<u>CONTRIBUTORS</u>	20
11.	<u>INTELLECTUAL PROPERTY STATEMENT</u>	21
12.	<u>FULL COPYRIGHT NOTICE</u>	21
13.	<u>REFERENCES</u>	21

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 2 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. Appendix

9.1 Survey Background

In order to allow users to execute application on remote computer systems using some mechanism established between sites, sites should be able to provide the resource usage information in commonly defined record format. Currently, each site has it's own accounting record and resource attributes to describe the resources that are being tracked at individual sites. Thus, the common terms to describe a resource attributes must be identified for all resources that are being tracked at sites to exchange resource usage information between sites. To accomplish this goal, survey was conducted to collect all

the resources that are being tracked and terms that are used to describe resources at different sites. Based on responses received from sites (NASA, NPACI, PNNL, NCSA, and ANL), resources are grouped together to show the terms used to describe the same resources in following tabularized "Survey Results" section.

Based on survey results, a proposed common set of Usage Record fields are identified with the common terms to describe the resources. This proposed set doesn't limit the extension of the set as the new and/or additional resource attributes are encountered. An individual sites is not required to report all the resource attributes listed in the proposed set but should report resources that are currently being tracked at the individual site. These survey results are intended for use by other groups in the Grid community to reference resources that are currently being tracked and indicate common terms used to describe such resources.

9.2 Survey Results

Terms used at different sites	Globus Resource Specification Language(RSL)	Datatype	Reference/Description
LOGIN_ NAME, user, userName, username, UserName, USERNAME		Text/ char	User's login name corresponding to user Id in /etc/ passwd
uid, userId		Int	User identification number from the /etc/ passwd file
Type		String	This field indicates the type of transaction being recorded doWithdrawal, doDeposit, doTransfer, modifyAllocation etc.
AuthName		String	Authorized userid performing the transaction
ACCOUNT, project, AccountName, projectName, GROUPNAME	project	Text/ char	Users account name where usage will be charged
projid, projectName		int	The account ID
JOB_ ID, jid, jid_ num, jobId, jobid		Number/ int	Job id where job was submitted to the batch queue.
Session_ id		Number	session id from the originating system
Id		char	Identifier indicating the job_ id, session_ id, reservation allocation_ id, etc according to context

Terms used at different sites	Globus Resource Specification Language(RSL)	Datatype	Reference/Description
-------------------------------	---	----------	-----------------------

	RSL)		
Pid		Number/ int	The process identifier assigned by the operating system to the process.
client, hostname		Text/ char	Name of the system job was executed.
Machine		String	Machine name (This could be a list of machines (system spans clusters and each machine could be a composite of the host, partition, cluster, site, and/ or enterprise)
QUEUE_ ID		Text	Queueing system identification code. NQS id, LoadLeveler LSF id
QUEUE, qname, queue, Queue, queueName	queue	Text/ char	Queue name where job was executed. (LSF - job was submitted to)
QWAIT		Number	Queue wait time for batch jobs
QUEUE_ DATE, submitTime, QueueTime		Number/ long	The date the job was queued to the batch system. Number of the Epoch in GMT.
JOB_ QUEUE_ DATE		Date	This is the date the job was submitted in the date format.
START_ DATE, start_ time, beginTime, StartTime		Number/ long	The date the job was started by the system. Number of the Epoch in GMT
JOB_ START_ DATE		Date	This is the date the job started running in date format.
END_ DATE, end_ time, Event Time, EndTime, finishTime		Number/ long	The date the job was completed by the system. Number of the Epoch in GMT
JOB_ END_ DATE		Date	This is the date the job ended in date format.
REQUESTED_ PROCS, ncpus, Processors, limitNpe, numProcessors	count	Number/ int	Number of processors requested at job submission time.

Terms used at different sites	Globus Resource Specification Language(RSL)	Datatype	Reference/ Description
nprocs, peakNpe, maxNumProcessors, MAXPROCS		Number/ int	The number of CPU's used.
MINPROCS		Number/ int	
NODES, Nodes		Number	Cumulative sum of all nodes allocated to the job. (number_of_nodes times cpu_per_node)
num_nodes		int	Number of nodes used. max((cputime + process_per_node * mb_per_node -1)/mb_per_node)
Nodemask		char	A hexadecimal string representing the bit mask specifying of processors) to be associated with this job.
NodeType		String	Type of node might factor into performance and charge
MAXPAR		Number	Maximum node partition. Largest number of processors parallel applications within the job. On all systems except systems this number will be the same as NODES. On C multiple parallel applications per job can be run, therefore describe the largest number of NODES allocated for the

Cpupercent		percent	The maximum percentage of a cpu which the job used. means 1 cpu. This cannot be set, it is only reported.
ProconsumptionRate		Number	Percentage of Total CPU used for prorating charge – a value between 0 and 1
CPU_ TIME, cput, connect_ time, cputime, CPUTime, cpuTime		Number/long	CPU time used by all processes of job

Terms used at different sites	Globus Resource Specification Language(RSL)	Datatype	Reference/Description
CONNECT_ TIME		Number	Connect time for interactive session
Pcput	max_cpu_time	long	Maximum amount of CPU time used by any single process
user_cpu, ru_utime		long/ double	The user CPU time in seconds
sys_cpu, ru_stime		long/ double	The system CPU time in seconds
interactive_cpu		double	Interactive cpus used (user_cpu + sys_cpu)
Batch_cpu		double	Batch cpus used (user_cpu + sys_cpu)
mt_user_cpu		double	The total user CPU time in seconds in multitasking (MT) mode
mt_sys_cpu		double	total system CPU time in seconds in MT (sys_cpu)
mt_connect		double	total connect time in MT queue
mt_nconnect		double	The sum of (connect_time * nprocs) in each of the CPU queues in MT queue.
mt_non- mt		double	number of seconds which are not multitasking in MT queue (user_cpu - mt_nconnect)
WALLCLOCK, walltime, Wallclock, runTime	max_wall_time	Number/long	Wall clock time which elapsed while the job was in running state. In clusters where a node is exclusively allocated the wallclock time is the number of processors yielding wallclock processor hours. In an IBM SP system this is actually the wallclock node hours divided by the number of cpus”

Terms used at different sites	Globus Resource Specification Language(RSL)	Datatype	Reference/Description
-------------------------------	---	----------	-----------------------

	RSL)		
REQUESTED_ TIME, limitRuntime		Number/ long	Amount of time requested at queue submission time for either wallclock time for parallel jobs or cpu time for ve
MAXMEMORY, high_ mem		Number /long	Memory high water mark for entire job
MEMORY		Number	Memory usage in Kcore- hours
REQUESTED_ MEM, limitMem		Number	Amount of memory requested at job submission time
Pmem		size	Maximum amount of virtual memory (workingset) used process of the job.
vmem, memory, maxRSwap, mem, Memory, peakMem	max_ memory	size	Maximum amount of virtual memory used by all concurr the job.
workingset		size	Maximum amount of physical memory used by any sing job.
maxRMem		size	Maximum amount of resident memory used by all proce
NUMMPPJOBS		Number	Number of parallel applications run in this job. On all sy T3E systems this number will be one. On Cray T3E sys parallel applications per job can be run, therefore, NUM describe the number of parallel applications run in this j
kword_ sec, kword- minutes		double	memory integral in seconds and in minutes(kword_ sec
I_ O, Mbytes I/ O, io_ kbytes, IO		Number/ double	IO usage in megabytes or kbytes transferred
IOread		number/ double	total number of bytes read by the job
IOwrite		Number/ double	total number of bytes written by the job

Terms used at different sites	Globus Resource Specification Language(RSL)	Datatype	Reference/ Description
Iobread		Number/ double	total number of bytes read by the job to block devices
Iobwrite		Number/ double	total number of bytes written by the job to block devices
io_ physreq(Physical I/ O)		double	The number of physical I/ O requests
DISK, Disk		Number	Disk storage used or Disk Charge in units defined by C other.
Network		int	Network used (withdrawals) or requested (reservations) AVG, TOT, or MAX]
EXPF		Number	Expansion factor. (QWAIT+ WALLCLOCK)/WALLC whether queue times are proportional to job size
File		Size	The largest size of any single file that may be created by
Fsblkused		long	The number of file system blocks consumed during the
Nice		int	The nice value under which the job is to be run.

PRIORITY		Number	Priority weight value.
JOB_COMP_STATUS, Status, jStatus		Number	Number representing completion status of the job.
ExitStatus		int	UNIX exit status of the job.
KillReason		Text/ char	if killed, reason the job was killed(npe, mem, cputime,
command, Executable		char	The name of the executable or system command
APP_NAME, jobName, JobName	executable	Text/ char	Job or Application name.

Terms used at different sites	Globus Resource Specification Language(RSL)	Datatype	Reference/ Description
Class		String	Class of job (batch, interactive, etc.)
JobType		String	Here you could distinguish between RMS job types, NQ etc.
QOS		String	Quality of Service
Total_charge, Charge		double, float	The total charge of the job in system's billing unit.. Am credited to account or allocation/ reservation/ quotation
SU		Number	Total charge for this job in System Billing Units. (second

Datatypes

The resource datatype corresponds to the following units.

- Date: Date in human readable format
- Number: specifies the maximum amount in terms or time period, integer, long integer, or double.
- Text: Specifies the character representation of string
- time: specifies a maximum time period the resource can be used. Time is expressed in seconds as an integer, or in the form
[[hours:] minutes:] seconds[. milliseconds]
- size: specifies the maximum amount in terms of bytes or words. It is expressed in the form integer[suffix]. The suffix is a multiplier defined in the following table, The size of a word is the word size on the execution host.
b or w bytes or words.
kb or kw Kilo (1024) bytes or words.
mb or mw Mega (1,048,576) bytes or words.
gb or gw Giga (1,073,741,824) bytes or words.

- unitary: The maximum amount of a resource which is expressed as a simple integer.
- int: specifies the numeric representation in integer
- long: specifies the numeric representation in long integer.
- double: specifies the numeric representation in float.
- char: specifies the character representation.
- String: specifies the character representation.
- percent: specifies the numeric representation in percentage (i. e 0- 100)

9.3 Common Usage Record Fields

Table 1: Usage Record Fields

Resource Name	Descriptions
Username	User's login name corresponding to user Id in /etc/ passwd file
ProjectName	Users account name where usage will be charged
JobId	Job id where job was submitted to the batch queue for batch jobs. Process Id for the interactive jobs
Queue	Queue name where job was executed or submitted depending on the batch system
GridId	User's global unique Id that identifies the user. Distinguish Name in the user's X509 certificate
FromHost	System name where job was submitted from
execHost	System name where job ran on.
StartTime	The date job started running in date time format (UTC timezone)
EndTime	The date job completed in date time format (UTC timezone)
Processors	Number of processors either used or requested that each center uses for billing purpose..
NumNodes	Number of nodes used.
Cputime	CPU time used, summed over all processes in the job.
Walltime	Wall clock time which elapsed while the job was in the running state.
Memory	Maximum amount of virtual memory used by all concurrent processes in the job.
Disk	Disk storage used or Disk Charge in units defined by CPU: disk blocks or other.
Network	Network used (withdrawals) or requested (reservations) by job [could be

	AVG, TOT, or MAX]
jobName	Job or Application name
Status	Number representing completion status of the job.
Charge	The total charge of the job in system's allocation unit

10. Contributors

Authors

Rodney Mach
University of Michigan
rmach@umich.edu
734-936-1985

Mi Young Koo
NASA Ames Research Center
mkoo@nas.nasa.gov
650-604-4528

Rebekah Lepro
NASA Ames Research Center
rlepro@mail.arc.nasa.gov
650-604-4359

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

Contributors

Laura McGinnis
Pittsburgh Supercomputing Center
lfm@psc.edu
412-268-5642

Jian Zhang
Xcerla Corporation
jian@xcerla.com
408-807-9069

11. 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.

12. 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."

13. 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>