

Draft 1 of Cloud Accounting Usage Record.

The following are a set of assumptions I have made related to this draft Usage Record:

1. For each instantiated VM there is one record, this has implications for charging as there is one 'Charge' attribute per record it does not fit with a charging model where each of the elements of CPU, Memory, I/O etc. are charged for separately.
2. The fields for GlobalUserName, VirtualOrganization, Projectname and Group relate to an existing method of authorizing grid users, if this authorization mechanism were to be used for cloud usage as well then these would provide a method of associating VM usage with a user and then querying usage in a portal could be based on user, VO etc.
3. A VM might be in one of three states: started, stopped and suspended. I think we would expect to gather many VM records whilst a VM is running, so I propose using the existing UR "Status" property for this - it may be possible to extend the implementation of this property to include some VM-specific statuses if these existing ones are not deemed appropriate.
4. We need a method of identifying where a specific running VM is - there are many ways to do this, I suggest that the RecordIdentity might contain a concatenation of some of the other fields in the record - e.g. TimeStamp SiteName MachineName ImageID - other suggestions welcome.
5. We have job accounting and storage accounting systems and their related usage records already - we therefore account for jobs and storage using these existing records/methods.

The following table shows what might be used and some additional elements required for cloud accounting:

Current UR Property	Current Definition	Extended for Cloud Accounting
RecordIdentity	Unique identifier of a record (String)	(keep the same)
GlobalUserName	Global identity of user (certificate DN) (String)	(keep the same)
VirtualOrganization	All three may be retrieved from FQAN - use if VOs will be used as part of authorization mechanism	(keep the same)
ProjectName		
Group		
Charge	Total charge of job (Float)	
Status	Completion status (String), supported values are: aborted, completed, failed, held, queued, started, suspended.	For VM instance "completed", "started" and "suspended" could be used.
WallDuration	WallClock time	Is this the EndTime in the case of VM?
CPUDuration	CPU time consumed (Duration)	(keep the same)
EndTime	Completion time (Timestamp)	Set to NULL until Status = "completed"
StartTime	Start time of the job (Timestamp)	Must be set if Status = started
MachineName	Hostname of the LRMS (String)	VM Hostname
Network	Amount of network resource used by the job (+ integer)	Network resource used by the VM

Memory	Amount of physical memory used by the job (+ integer)	Memory associated with the VM
TimeDuration	Additional measure of time duration (Duration)	This could be used for suspended VM to record suspended-endtime
SiteName	GOCDB SiteName	Extend service types in GOCDB to include cloud resource providers
<b>Proposed New Property</b>	<b>Definition</b>	
ImageID	Every image has a unique ID associated with it	
CloudType	Cloud type	
StorageRecordIdentity	Link to associated storage record	