Next-generation Usage Record

12 March 2013, OGF37, Charlottesville

Andrea Cristofori, Jon Kerr Nilsen

OGF IPR Policies Apply

- "I acknowledge that participation in this meeting is subject to the OGF Intellectual Property Policy."
- Intellectual Property Notices Note Well: All statements related to the activities of the OGF and addressed to the OGF are subject to all provisions of Appendix B of GFD-C.1, which grants to the OGF and its participants certain licenses and rights in such statements. Such statements include verbal statements in OGF meetings, as well as written and electronic communications made at any time or place, which are addressed to:
  - the OGF plenary session,
  - any OGF working group or portion thereof,
  - the OGF Board of Directors, the GFSG, or any member thereof on behalf of the OGF,
  - the ADCOM, or any member thereof on behalf of the ADCOM,
  - any OGF mailing list, including any group list, or any other list functioning under OGF auspices,
  - the OGF Editor or the document authoring and review process
- Statements made outside of a OGF meeting, mailing list or other function, that are clearly not intended to be input to an OGF activity, group or function, are not subject to these provisions.
- Excerpt from Appendix B of GFD-C.1: "Where the OGF knows of rights, or claimed rights, the OGF secretariat shall attempt to obtain from the claimant of such rights, a written assurance that upon approval by the GFSG of the relevant OGF document(s), any party will be able to obtain the right to implement, use and distribute the technology or works when implementing, using or distributing technology based upon the specific specification(s) under openly specified, reasonable, non-discriminatory terms. The working group or research group proposing the use of the technology with respect to which the proprietary rights are claimed may assist the OGF secretariat in this effort. The results of this procedure shall not affect advancement of document, except that the GFSG may defer approval where a delay may facilitate the obtaining of such assurances. The results will, however, be recorded by the OGF Secretariat, and made available. The GFSG may also direct that a summary of the results be included in any GFD published containing the specification."

- OGF Intellectual Property Policies are adapted from the IETF Intellectual Property Policies that support the Internet Standards Process.

Outline

- Charter.
- From UR1 to UR2.
- Background for UR2 format.
- Requirements and use cases.
- Format of UR2.
- Conclusion.

Proposed charter from OGF34 (14 March 2012)

Type: Recommendation Document

Milestone	Date (YYYY-MM)
First Draft	2012-09
Public Comment	2013-02 (started 2013-03-05)
Published	2013-04

Slightly optimistic?

Should we also give ourselves some internal deadlines? (Agreeing on use-cases/requirements etc.).

From UR1 to UR2

- UR1 focused on accounting for Grid jobs.
- Many non-compatible implementations of UR1 - many fields open for interpretation.
- During the years the need for accounting for other kinds of resources became apparent.
- Several new record formats started to emerge for use cases not covered by UR1:
  - EMI-StAR and SAI for storage accounting
  - EMI-CAR and others for (more) proper handling of VOs
  - EGI Cloud record for clouds
- Some GLUE fields could be useful for usage records but not used in UR1 (e.g., HostType).

Background for UR2 format

- The initial work on UR2 was due to the need for storage accounting.
- From the preliminary discussions it became clear that accounting of other resources would need to be added at some point in time.
- To achieve this, it was decided to organize the resources in blocks:
  - Inititially: RecordIdentityBlock, SubjectIdentityBlock, ComputeUsageBlock, JobUsageBlock and StorageUsageBlock
  - MemoryUsageBlock was initially part of ComputeUsageBlock, but as it is a separate resource it got its own block
  - Then came the EGI Cloud Task Force
  - To accommodate cloud resources we added CloudUsageBlock and NetworkUsageBlock (since network was included in the EGI cloud records)
- This way new blocks can be added easily to account for future resources.
- With the block separation changes in one resource type will (hopefully) not affect the other resource definitions.

Requirements and use cases

- Requirements and use cases were collected:

  https://redmine.ogf.org/projects/ur-wg/wiki/UseCases

- Half a year discussion of use cases on mailing list and bi-weekly meetings before even starting the document.

Usage Record definitions

- Usage Record 2.0 format recommendation is written to be as compatible as possible with existing UR definitions (StAR, CAR, EGI Cloud TF). See "Attribute Matrix".
- The work done during the definition of UR 2.0 influenced also the definition of the previously mentioned URs.
  - E.g., StAR: MesureTime and ValidDuration -> StartTime and EndTime
- "Attribute Matrix" simplifies the transition between URs.

Block Summary

- RecordIdentityBlock
- SubjectIdentityBlock
- ComputeUsageBlock
- JobUsageBlock
- MemoryUsageBlock
- StorageUsageBlock
- CloudUsageBlock
- NetworkUsageBlock

General record info

- All elements are optional unless otherwise stated.
- The only mandatory block is the RecordIdentityBlock, all other blocks are optional.
  - Every implementation should follow a pre-defined profile suitable to their needs.
- These block must be used only once in each record:
  - RecordIdentity
  - SubjectIdentity
  - JobUsageBlock
  - CloudUsageBlock

Block combinations

Different combinations of blocks can be used for accounting of different resources. Following are some typical examples.

- Accounting of a grid job:
  - RecordIdentityBlock
  - SubjectIdentityBlock
  - ComputeUsageBlock
  - JobUsageBlock
  - (MemoryUsageBlock)
- Accounting of storage resources:
  - RecordIdentityBlock
  - SubjectIdentityBlock
  - StorageUsageBlock

Block combinations

- Accounting of a cloud virtual machine:
  - RecordIdentityBlock
  - SubjectIdentityBlock
  - ComputeUsageBlock
  - (MemoryUsageBlock)
  - (StorageUsageBlock)
  - CloudUsageBlock
  - (NetworkUsageBlock)
- Accounting of network resources:
  - RecordIdentityBlock
  - SubjectIdentityBlock
  - NetworkUsageBlock

RecordIdentityBlock

- Distinguish the record unequivocally.
- Gives information on:
  - the record ID
  - the creation time
  - the site that produced the record
  - the infrastructure
- This block must be present and can only be present once.

SubjectIdentityBlock

- Defines the user/group identity consuming the resources.
- Gives information on:
  - local/global user
  - local/global group
  - additional properties useful to identify the consumer
- If used this block can only be present once.

ComputeUsageBlock

- Contains fields for the compute usage.
- Gives information on:
  - CPU/Wall Time consumed
  - details on the resources used
  - beginning and end of the utilization
  - exit status of the application
  - charge
- If used this block can be present multiple times.

JobUsageBlock

- Identifies the details related to grid jobs.
- Gives information on:
  - identity of the job
  - the submission method
  - CE details
  - status of the job
- If used this block can only be present once.

MemoryUsageBlock

- Identifies the details related to memory consumption.
- Gives information on:
  - memory type
  - used/allocated/requested memory
  - beginning and end of the utilization
  - details on the resources used
  - charge
- If used this block can be present multiple times.

StorageUsageBlock

- Identifies the details related to storage consumption.
- Gives information on:
  - storage type
  - physical/logical used storage
  - resource capacity allocated
  - beginning and end of the utilization
  - details on the resources used
  - charge
- If used this block can be present multiple times.

CloudUsageBlock

- Identifies the details related to cloud usage.
- Gives information on:
  - identity of the VM
  - details on the VM (e.g., amount of time the VM has been suspended)
  - status of the VM
- If used this block can only be present once.

NetworkUsageBlock

- Identifies the details related to network usage.
- Gives information on:
  - type of network
  - inbound and outbound bytes
  - charge
- If used this block can be present multiple times.

The rest

- Attribute Matrix.
- Field summaries.
- Working examples.
- Schema.
- Security considerations.

Conclusions/outlook

- UR-2.0 entered the public comments period on the 5th of March 2013.
- Covers: compute, storage, memory, network, grid job and cloud usage.
- Next steps:
  - Ask for comments from EMI, EGI, OSG, Prace, ... (comments period ends on the 5th of June).
  - Reply to comments.
  - Encourage people to start using UR-2.0.