

Towards Usage Record Format 2.0


Usage Record (UR) Working Group
Working Group
Morris Riedel et al.

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




Outline



- Welcome & Administrative Aspects
- Goals and Visions
- Discussions from OGF31
- Survey of ongoing activities
- Storage Accounting Records
- Next Steps

Welcome & Administrative Aspects



New Chair of Working Group




- Jon Kerr Nilsen is new official group co-chair
 - University of Oslo
 - Member of the European Middleware Initiative (EMI)
 - Driver of storage accounting record activities within the community
 - Couldn't come to OGF this time
 - (Morris chairs session only on behalf of Jon and EMI)

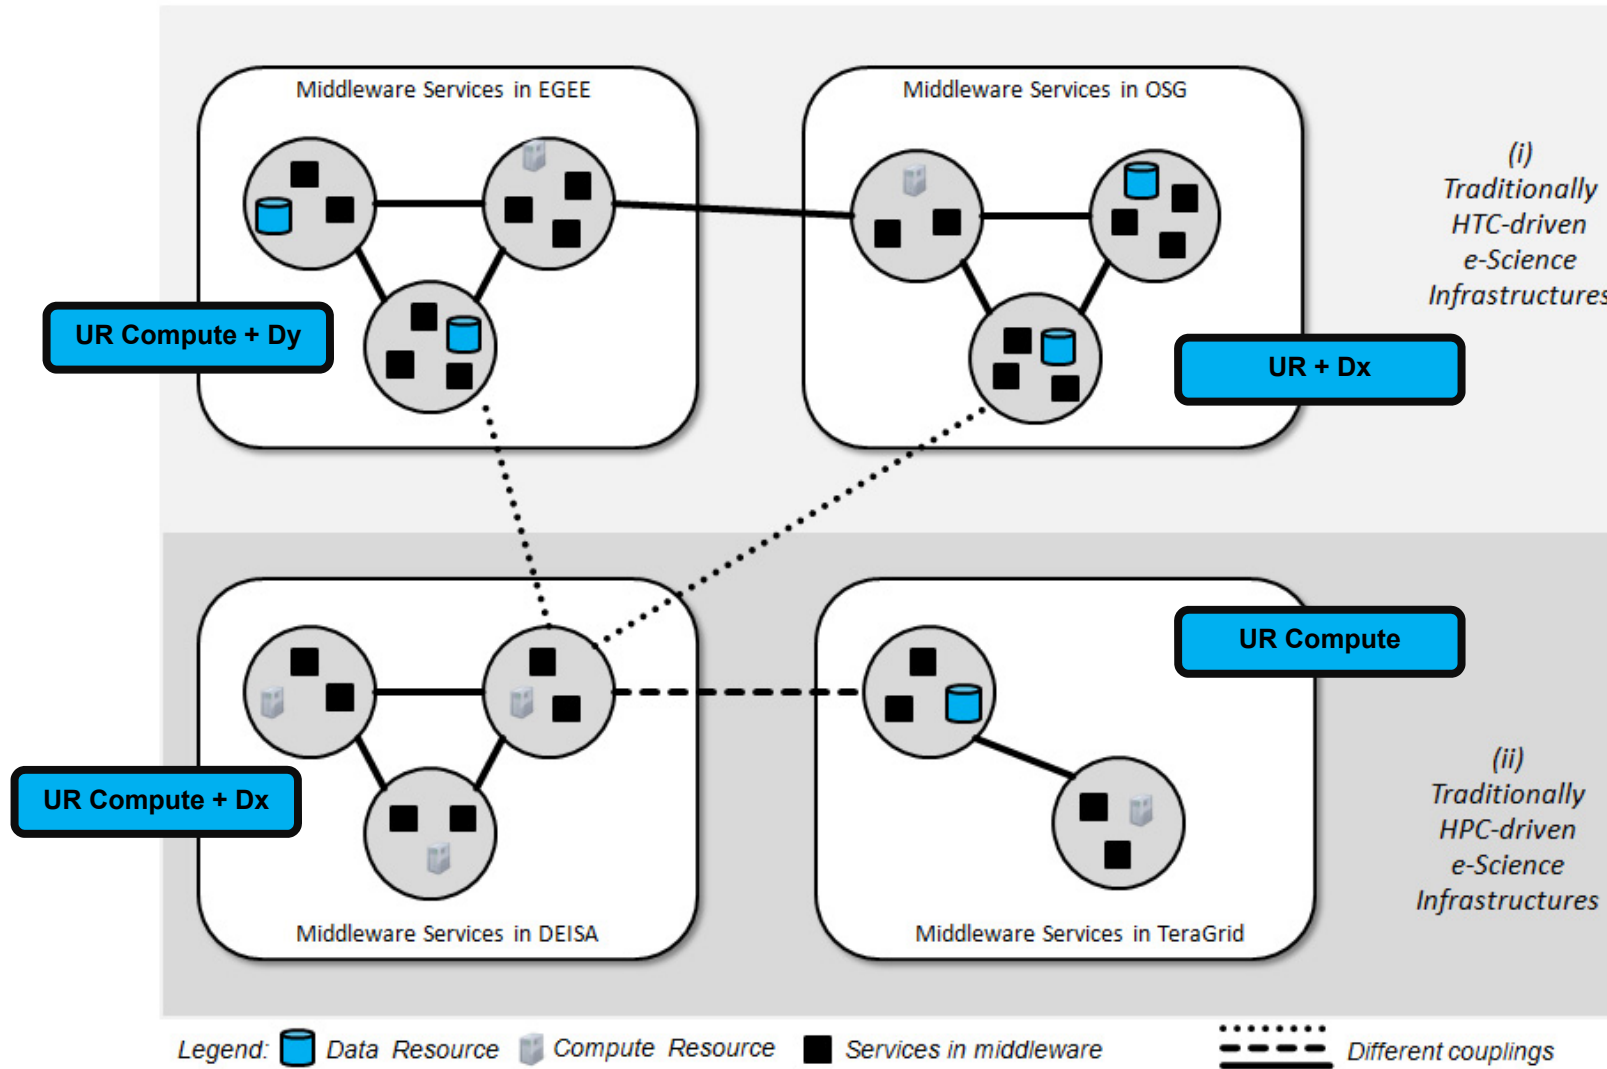


- New co-chair will lead to new activities around UR 2.0
 - Also in contact with the other chairs to understand their interest

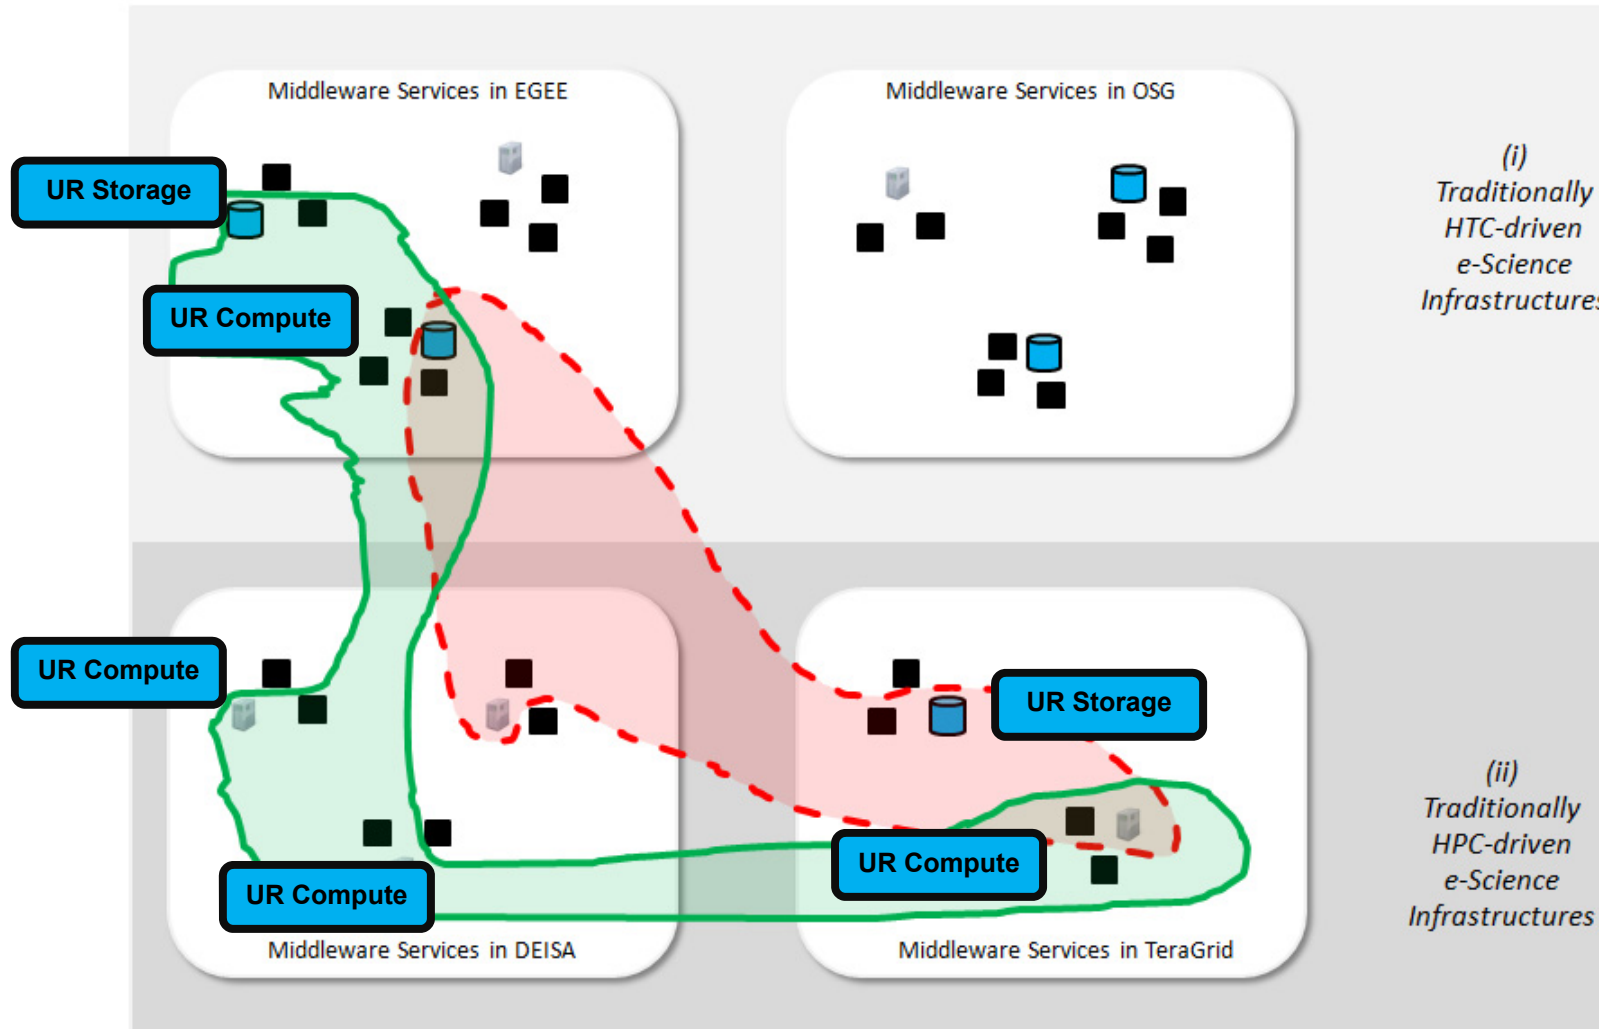
Goals



Goals: Situation Today







Goals: Resource Usage Tracking



(i)
Traditionally
HTC-driven
e-Science
Infrastructures

(ii)
Traditionally
HPC-driven
e-Science
Infrastructures

Legend:  Data Resource  Compute Resource  Services in middleware

 Individual Infrastructures

Discussions in context



- Can we deliver a higher vision that integrates all of these approaches, including...
- Service Usage Tracking
- License Usage Tracking
 - MatLAB license, etc.
 - FlexLM, etc.
- Network Usage Tracking (Network Measurement)

More active members



- Jon Kerr Nilsen
- (INFN)
- John Gordon + Dx → STFC
- XSEDE...Vanamala (UVA),
- JP??

Motivation: UR Storage

- Note: Storage resources are shared between VO's, so more than one owner sometimes
- Challenges of developing a storage infrastructure
 - We need to know how much storage space is used
 - Which group/user uses it
 - Which storage media is used
 - Motivation: To know where to put the money when increasing the storage space and to know who to ask for the money to increase the storage space
- Basis for billing (in the near future?)
 - Storage is expensive and 'data wave' foresees more data coming
 - Non-academic resource owners not like to give it away for free
 - Possibly even academic institutions in future need to 're-think'

Helping Directions



- How would you charge for it – Money!
- Starting from this...

- Use cases GLUE2 in terms of accounting (Jens)
 - If UR covers this functionality
 - Reasons and to follow-up

Vision




- We need to track the resource usage in storage and compute (and network → OGF network groups)
- One comprehensive Usage Record Format
 - Covers basic compute AND storage aspects
 - Not an academic exercise → Production needs matter!
 - Extensible format to allow for profiling of specific extensions
 - Complementary to OGF networking activities in context
- Backwards compatibility
 - Usage Record Format 1.0 is broadly used in different infrastructures, so backwards compatibility is desired
 - **Computing straightforward; Storage records not existed in the past**
- Common format across infrastructures
 - Not a profile for each infrastructure (e.g. HPC vs. HTC)
 - Works together with complementary networking activities

UR Network

UR Storage

UR Compute

Discussions from OGF31



Common Understanding?



- Accounting and/or resource usage tracking means...
- Accounting is the recording and summarizing of the consumption of a resource by an individual user or a group of users (aggregated resource records)
- Typically used to find out who uses how much of a set of resources
- Typically not used to find out how individual resource components are used
 - Perhaps this would make sense?
 - Example: Typical batch system submission or middleware used?
 - Hence: Add a little hint how it was used
- Timeframe is important in context – needs to be explored

What we have vs. what we want

UR Compute

- Job accounting
- Job accounting is the recording and summarizing of the consumption of computing resources
- Recording wall-time, cpu-time, etc. – strictly increasing values.
- When a job finishes it stays finished
- Typically measured per job – information readily available from local batch system
- Standardized through UR 1.0
- (Several +Dx exists)

UR Storage

- Storage accounting
- Storage accounting is the recording and summarizing of the consumption of a storage resource
- Storage usage varies in time – recorded usage is only relevant in a given time frame
- Data is rarely in finished state and can live for a long time
- Typically recorded through (in)frequent usage snapshots – high fluctuation rate
- No standardized record

Commonalities?




- Both job records and storage records have a subject ID
- URF 1.0 has no way to include VO info
 - This information is crucial for many production infrastructures and technology providers
 - URF 2.0 should support VOs
- None of them care about network
 - Complementary activities in the networking groups of OGF, but more cooperation needed

Type of specifications



- Joint specifications or separate ones?
- UR1 not very much extensible and is missing VO information, etc. (+Dx elements)
- No draft for service record in UR 2.0 yet
- Implementations need to be done in parallel to ensure a format that is useful (no academic exercise)
- One of the reasons why the StAR format is close to OGF UR 1.0, but still separate
- Merge with UR 2.0 when available should be discussed

Survey of ongoing activities



Initial Survey performed



- Surveys and discussions performed within EMI with other activities
- Available at...
 - <https://twiki.cern.ch/twiki/bin/view/EMI/StorageAccounting>

EMI Project



- Jon represents the EMI project
- EMI Description of Work
 - *“EMI will address known UR issues and extend accounting records to include VO-aware storage usage accounting”*
 - *“The refined and extended standard UR format, service interface and data transport protocol will be implemented in gLite, ARC, dCache and UNICORE”*
 - *EMI will thus contribute to the definition of UR 2.0 of OGF*
- Worked together so far with many infrastructures
- Three major involvements
 1. Chairing of OGF activities → starting telcons soon, etc.
 2. Contribute with EMI STAR proposal to the standards process
 3. Contribute with compute production experience to UR 2.0




Other projects



- TBD: Audience?

EMI Star Proposal



EMI Activities



- EMI is finalized a proposal for a storage accounting record (StAR)
- Contacted and received input from organizations and potential user groups
 - NDGF, OGF, EGI, OSG, INFN, CMS, ATLAS, ...
 - Not only an EMI in-house production
- Will be implemented by EMI storage providers
- Work already started outside EMI
 - Swedish Grid Accounting System (SGAS)
- Link to activities:
 - <https://twiki.cern.ch/twiki/bin/view/EMI/StorageAccounting>

STAR Structure



- Resource
 - Fields describing the system the resource was consumed on
 - **StorageSystem, StorageShare**
- Consumption details
 - Fields describing what is consumed
 - **StorageMedia, StorageClass, DirectoryPath, ...**
- Identity
 - Describes the person or group accountable for the consumption
 - **SubjectIdentity** block (**UserIdentity, Group, GroupAttribute, ...**)
- Resource consumption
 - How much of the described resource has been used
 - **ResourceCapacityUsed, LogicalCapacityUsed**
- Record identity and duration

Local Example



```
<sr:StorageUsageRecord
  xmlns:sr="http://eu-emi.eu/namespaces/2011/02/storagerecord">
  <sr:RecordIdentity sr:createTime="2010-11-09T09:06:52Z"
    sr:recordId="host.example.org/sr/87912469269276"/>
  <sr:StorageSystem>host.example.org</sr:StorageSystem>
  <sr:SubjectIdentity>
    <sr:LocalUser>johndoe</sr:LocalUser>
  </sr:SubjectIdentity>
  <sr:StorageMedia>tape</sr:StorageMedia>
  <sr:FileCount>42</sr:FileCount>
  <sr:MeasureTime>2010-10-11T09:31:40Z</sr:MeasureTime>
  <sr:ValidDuration>PT3600S</sr:ValidDuration>
  <sr:ResourceCapacityUsed>913617</sr:ResourceCapacityUsed>
</sr:StorageUsageRecord>
```

Grid Example



```
<sr:StorageUsageRecord
  xmlns:sr="http://eu-emi.eu/namespaces/2011/02/storagerecord">
  <sr:RecordIdentity sr:createTime="2010-11-09T09:06:52Z"
    sr:recordId="host.example.org/sr/87912469269276"/>
  <sr:StorageSystem>host.example.org</sr:StorageSystem>
  <sr:SubjectIdentity>
    <sr:Group>binarydataproject.example.org</sr:Group>
    <sr:GroupAttribute sr:attributeType="subgroup">ukusers
    </sr:GroupAttribute>
  </sr:SubjectIdentity>
  <sr:StorageMedia>tape</sr:StorageMedia>
  <sr:FileCount>42</sr:FileCount>
  <sr:MeasureTime>2010-10-11T09:31:40Z</sr:MeasureTime>
  <sr:ValidDuration>PT3600S</sr:ValidDuration>
  <sr:ResourceCapacityUsed>913617</sr:ResourceCapacityUsed>
</sr:StorageUsageRecord>
```

Full Example

Why Files?


Starttime and Endtime, rather measurement now

Not correct



```
<sr:StorageUsageRecords xmlns:sr="http://eu-emi.eu/namespaces/2011/02/storagerecord">
  <sr:StorageUsageRecord>
    <sr:RecordIdentity sr:createTime="2010-11-09T09:06:52Z"
      sr:recordId="host.example.org/sr/87912469269276"/>
    <sr:StorageSystem>host.example.org</sr:StorageSystem>
    <sr:StorageShare>pool-003</sr:StorageShare>
    <sr:StorageMedia>disk</sr:StorageMedia>
    <sr:StorageClass>replicated</sr:StorageClass>
    <sr:FileCount>42</sr:FileCount>
    <sr:DirectoryPath>/home/projectA</sr:DirectoryPath>
    <sr:SubjectIdentity>
      <sr:LocalUser>johndoe</sr:LocalUser>
      <sr:LocalGroup>projectA</sr:LocalGroup>
      <sr:UserIdentity>/O=Grid/OU=example.org/CN=John Doe</sr:UserIdentity>
      <sr:Group>binarydatapoint.example.org</sr:Group>
      <sr:GroupAttribute sr:attributeType="subgroup">ukusers</sr:GroupAttribute>
    </sr:SubjectIdentity>
    <sr:MeasureTime>2010-10-11T09:31:40Z</sr:MeasureTime>
    <sr:ValidDuration>PT3600S</sr:ValidDuration>
    <sr:ResourceCapacityUsed>14728</sr:ResourceCapacityUsed>
    <sr:LogicalCapacityUsed>13617</sr:LogicalCapacityUsed>
  </sr:StorageUsageRecord>
</sr:StorageUsageRecords>
```

Next Steps



Next Steps

- **Use of normal mailing list (no sub-lists)**
 - Subject indicating storage, etc.
- Starting telcons (e.g. bi-weekly)
 - Checking people with real interest and commitment
 - Discussions of how STAR and URF 1.0 +Dx can be joined or specified as separate complementary specs
- Setting Milestones
 - E.g. First draft expected at next OGF for discussions
 - E.g. OGF Draft finalized towards X-MAS
- Interop demonstrations of early adopters
 - In collaboration with GIN at Supercomputing (maybe)
- **TBD: WG-ALL mailing list post with UR...**
- **TBD: Audience suggestions / commitments?**

Next Steps



- **Cloud Accounting?**
 - Matteo Turilli
 - Knowledge from this aspects also relevant

Put it in GridForge

Full Copyright Notice



Copyright (C) Open Grid Forum (2011). 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.

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