



Storage Accounting Records in EMI

EMI is partially funded by the European Commission under Grant Agreement RI-261611

Jon Kerr Nilsen j.k.nilsen@fys.uio.no

EMI Data Group
OGF30, Brussels

EMI Storage Usage Records



- From DoW:
 - 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 v2.0 of OGF
- SAR not just an academic exercise for EMI

Plan for storage accounting record in EMI

- Project Month 8 (**January 2011**):
Definition of SAR
- End of EMI: SAR implemented in DPM,
StoRM, dCache, gLite FTS, UNICORE FTS

What have been done



- Wiki page
 - <https://twiki.cern.ch/twiki/bin/view/EMI/StorageAccounting>
- Mailing list
 - <http://mail.eu-emi.eu/mailman/listinfo/emi-jra1-data-sar>

What have been done (contd.)

- Requesting requirements
 - Contacted
 - OGF, EGI, OSG, INFN, CMS, ATLAS, NDGF
 - EGI has contacted NGI's
 - A few people seem to get lots of forwarded request...
 - Asking for input/requirements for storage accounting
 - Only general requests to get in contact
 - Need something concrete to get more input (maybe from OGF30?)

What have been done (contd.)

- Received some feedback
 - INFN need some way to assess if new storage is needed (e.g., are there never-accessed files)
 - OSG extends OGF UR with a StorageElement and a StorageElementRecord – would like something similar (but standardized)
- Also looked at Amazon S3 and iRODS

What have been done (contd.)

- Drafting a record format
 - Henrik will tell more...

Amazon S3

- Amazon S3 accounting record, is written hourly
 - Service (= S3)
 - Operation
 - the operation that was done like listing, reading, writing objects (called buckets on S3)
 - UsageType (see <http://aws-portal.amazon.com/gp/aws/developer/common/amz-storage-usage-type-help.html>)
 - category of usage data
 - some examples:
 - TimedStorage-ByteHrs: The amount of storage you've used over time. This is the raw data used to calculate your storage charges, and is useful if you are interested in auditing your bill. If you just want to track how much storage you are using on a daily basis, you should look at the AverageStorage-Bytes usage type instead. At least twice a day, we check to see how much storage is used by all your Amazon S3 buckets. The result is multiplied by the amount of time passed since the last checkpoint. For example, if you used 1024 bytes for 12 hours, that is 12,288 byte-hours.
 - DataTransfer-In-Bytes: This contains records of network data transfer into your Amazon S3 account from the outside Internet. This is the same data used to calculate your charges for network use at the end of the billing period. Every time a request is received to put an object or list a bucket, the amount of network traffic involved in transmitting the object data, metadata, or keys is recorded here. You can view the network usage associated with one of these operations by specifying the operation of choice before generating the usage report.
 - Requests-Tier1: This contains records of the number of PUT and LIST requests made to your Amazon S3 account for the given month.
 - StartTime and EndTime
 - UsageValue in the according unit (bytes or counts)

iRODS

- iRODS

- Accounting is a way to document the properties of a collection
 - to track how much is stored
 - assert current status of a collection
 - track changes to the collection
 - validate whether the collection meets assessment criteria
- Files systems do not provide the context needed to validate assesment criteria - need to think in terms of collection
 - Accounting system for file system is useless to the data grid, digital library and preservation communities
- Within a collection iRODS need to track
 - desired properties
 - policies that enforce each property
 - procedures that implement the policy
 - system state information that is generated by applying the policies
 - assessment criteria (policies that are run periodically)
 - audit trails on every operation
- Accounting then corresponds to
 - queries on current state information for current status
 - parsing of audit trails to track compliance with policies over time
 - validation of assessment criteria
- To iRODS data grid, accounting is the set of procedures composed from micro-services that implement these three functions for each policy