


Storage Accounting

John Gordon, STFC

GDB June 2012



Outline

- EMI StAR
 timeline
- EGI Profile
- Visualisation
- Other Data Accounting

EMI StAR

- EMI recognised the requirement for storage accounting
- EMI created a proposal for a storage accounting record (StAR)
- ‘Finalised’ in June last year
- Published version:
<http://cdsweb.cern.ch/record/1352472?ln=en>
- Proposed to OGF
- Being implemented by EMI storage providers

What StAR is not

- To be able to arrive to a useful record in time we had to limit the scope
- Many interesting storage-related features to record
 - I/O (like in Amazon S3)
 - Metadata access (stat,ls)
 - Available space
 - File name
 - ...
- Scope of StAR is limited to consumption of storage space
- I'll come back to the above later

What StAR is

- Definition of usage record for storage
- Inspired by OGF UR1.0
 - Current implementations are for cpu
- Summarises used space
- XML-based schema
- A single value measured over TimeDuration means the product is an integral
 - allows averaging over time within a storage system, and averaging records.

StAR structure

- Resource
 - Fields describing the system the resource was consumed on
 - `StorageSystem`, `StorageShare`
- Resource consumption
 - How much of the described resource has been used
 - `ResourceCapacityUsed`, `LogicalCapacityUsed`
- 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`, ...)
- Record identity and duration

EMI StAR timeline (1)

- **May 2011: StAR definition agreement**
- Oct. 2011: Each SE to provide an implementation plan
- May 2012: Each SE to implement accounting sensors
- June 2012: Accounting record publishers (APEL and/or DGAS) to publish storage records
- December 2012: Each SE to test and deploy their accounting sensors (for EMI 3 release)
- January-April 2013: Bug-fixing in EMI 3 RC's
- April 2013: Accounting sensors released in EMI 3

EMI StAR timeline (2)

- May 2011: StAR definition agreement
- May 2012: Each SE provided an implementation plan
- June 2012: Each SE to implement accounting sensors using SSM to publish
- June 2012: Accounting repository (APEL) ready to receive storage records
- September 2012: Each SE to test and deploy their accounting sensors (for EMI 3 release)
- January-April 2013: Bug-fixing in EMI 3 RC's
- April 2013: Accounting sensors released in EMI 3

Storage Solutions

- dCache (EMI)
- StoRM (EMI)
- DPM (EMI)
- Castor
- EOS
- ~~BestMan~~
- Gratia (already collecting)
- xrootd
- hadoop
- Gstat
- Cloud Storage
- (Unicore)
- (ARC)
- iRods
- ??????
- The middleware in scope for EGI and WLCG
- Anything that collects storage information could also publish usage records
- eg Julia Andreeva talk at CHEP

EGI Profile for StAR

- My Proposal
- WLCG Use Case
 - https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf
- Required fields/values
- **StorageSystem** = Storage Element FQDN
- **StorageShare** = StorageArea
 - use the corresponding GLUE2 StorageShare.ID attribute.
- The **SubjectIdentity** property **MUST** be present in the record
 - It **MUST** contain
 - **Group** = VO name
 - It **MAY** contain:
 - **UserIdentity** = X.509 UserDN of the owner of the data.
 - **GroupAttribute** = VOMS FQAN

StAR Issues

- **Allocated Space** – equivalent to wallclock time for cpu. Shows resource blocked by VO.

```
<ur:StorageResourceCapacityUsed>14728</ur:StorageResourceCapacityUsed>
```

```
<ur:StorageResourceCapacityAllocated>50000</ur:StorageResourceCapacityAllocated>
```

Discussed and agreed in OGF since EMI document published

- **SITE** – EGI and NGI deal with SITEs. Defined in GOCDDB and used in many other tools inc BDII for resource discovery.

I thought this had been included after previous discussions. Surprised to see not.

Very Important to EGI

What will be Published?

- First Pass – Total allocated and used per site, per VO, per month
- Sum over all storage for a VO.
- Not obvious how storage tokens can be accounted

How would you like to see it?

- CPU Portal shows a variety of variables
 - Njobs, raw cpu, wallclock, etc
 - Aggregated at any point in topology tree
 - T1, T2, NGI, Country, Grid
- What would you like to view for storage?
 - Min/max/average over time period?
 - Tables and csv downloads
 - Custom selection of VOs?
 - Bar charts of VOs?

What now?

- The EMI timeline of release in May 2013 is far off
- Does WLCG wish for an interim solution?
 - Just for RRB?
 - Based on gstat?

gstat

- gstat records all the required data on use of storage by VOs.
- For single VO storage areas the results look good.
 - Multi VO storage is a problem
- The information could be gathered into a report for the RRB.
- Once exposed to scrutiny the quality of the data would improve.

GStat 2.0 - Site View - FZK-LCG2 - Mozilla Firefox

File Edit View History Bookmarks Tools Help

GDB (11 January 2012) GStat 2.0 - Site View - FZK-LCG2

gstat-prod.cern.ch/gstat/site/FZK-LCG2/

WLCG SSC BBC EGI EGI Accounting Portal GDB Doodle BA Vidyo

08:00	09:00	10:00
PhysicalCPUs	NOW:2.5k MAX:2.5k AVG:2.5k	
LogicalCPUs	NOW:14.2k MAX:14.2k AVG:14.2k	

VO-based Resource Allocation and Usage

aerogrid	alice	astrogrid	atlas	auger	babar	bauvogrid	belle	bioif	bisgrid	biz2grid	bwgrid	c3grid	cdf
cms	compass	dech	dgcms	dgops	dgtest	dteam	dzero	education	enmr.eu	fingrid	gdigrid	geclipse	hepcg
ingrid	interloggrid	kerndgrid	lhcb	lifescience	mediagrid	medigrid	mosgrid	ops	partgrid	progrid	textgrid		
undefinedvo	wisnet	wisnet											

Grid Jobs			Online Storage Usage (GB)		Nearline Storage Usage (GB)	
Total	Running	Waiting	TotalSize	UsedSize	TotalSize	UsedSize
10,141	7,027	3,114	3,827,629	2,501,102	9,294,100	3,459,180

[Click Me](#)

Job Number (VO: atlas, Site: FZK-LCG2)

TotalJobs	NOW:10.6k	MAX:12.5k	AVG:11.8k
RunningJobs	NOW:7.4k	MAX:9.6k	AVG:8.8k
WaitingJobs	NOW:3.2k	MAX:3.2k	AVG:3.0k

[Click Me](#)

Online Storage Space (VO: atlas, Site: FZK-LCG2)

UsedOnlineSize	NOW:2.5P	MAX:2.5P	AVG:2.5P
TotalOnlineSize	NOW:3.8P	MAX:3.8P	AVG:3.8P

[Click Me](#)

Nearline Storage Space (VO: atlas, Site: FZK-LCG2)

UsedNearlineSize	NOW:3.5P	MAX:3.5P	AVG:3.5P
TotalNearlineSize	NOW:9.3P	MAX:9.3P	AVG:9.3P

Data Centre Accounting

- StAR is about storage allocated and used.
- What about who used the data? And how often?
- What would the big sites like to see about the use made of their data centre?
 - Bandwidth?
 - Source/destination?
 - Identity of users?

Summary

- Storage Accounting is coming
 - But slowly
- Need to make sure all storage solutions deliver.

Questions

- Are you happy with EGI Profile for StAR?
- What should be visualised?
- Data Centre Accounting?