Storage Accounting

John Gordon
STFC
GDB, Lyon
6th April2011

GDB January 2012

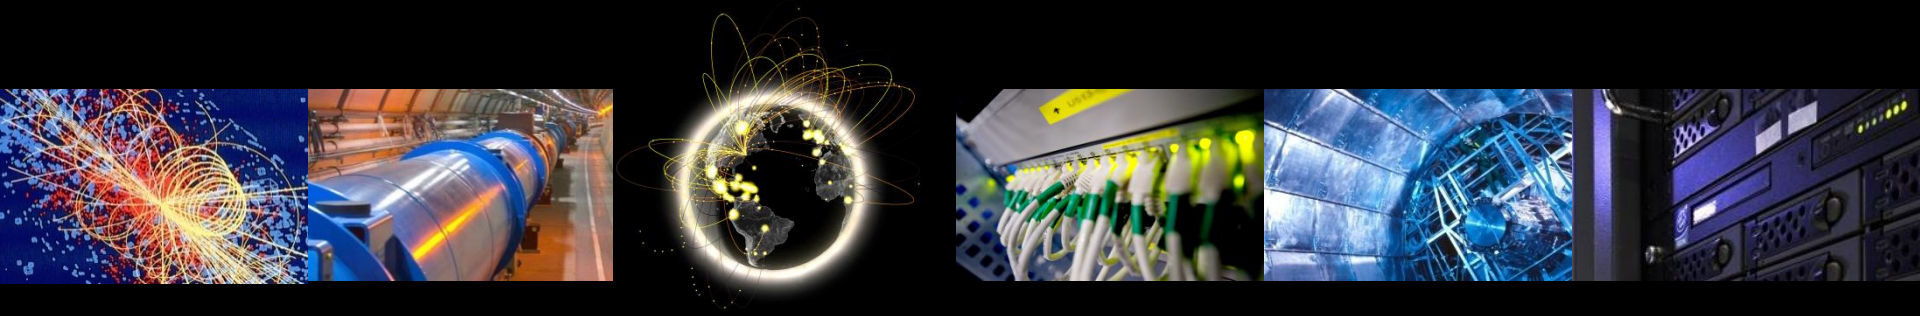

WLCG
Worldwide LHC Computing Grid

Storage accounting in EMI and OGF

Jon Kerr Nilsen
EMI JRA1 deputy
Dept. of Physics, Univ. of Oslo

EMI StAR

- EMI has created a proposal for a storage accounting record (StAR)
- Finalised in June last year
- Final version: http://cdsweb.cern.ch/record/1352472?ln=en
- Will be 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

What StAR is

- Definition of usage record for storage
- Inspired by OGF UR1.0
- Summarises used space
- XML-based schema
- Non-overlapping/non-contiguous records means no space used (ValidDuration configured by sys-admin)

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

- May 2011: StAR definition agreement
- Oct. 2011: Each SE to provide an implementation plan
  - Done, not documented
- 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

WLCG view of StAR

- As a standard document StAR allows some flexibility with optional fields and possibility of different semantics.
- For WLCG and EGI I intend to define a profile (to be agreed) that forces certain fields to be mandatory (eg project) and to have a certain semantic (eg VO name)
- Still a couple of outstanding issues around installed capacity and timeduration
- At some stage there will be a revision of the general UR and StAR will be reconciled with this.

EGI

- As I reported to the November GDB the EGI infrastructure is changing.
- When the server end is complete in a few months it will be ready to receive other type of accounting records including storage.
- The SSM transport level is available to EMI PTs. It follows EMI guidelines although they could write their own compatible one if they wish.

  Like Unicore

gstat

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

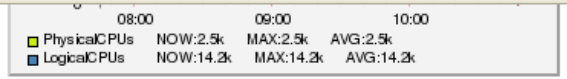

VO-based Resource Allocation and Usage

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 | wisent | 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

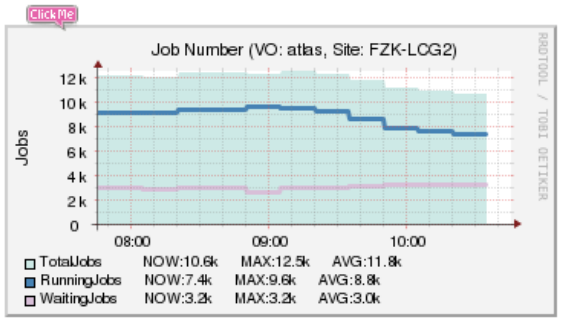

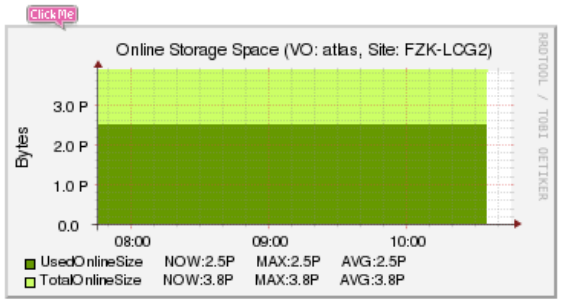

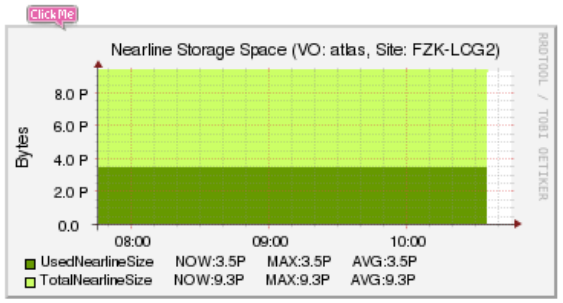

gstat

- 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
- Could the information be gathered into a report for the RRB?

Storage accounting approach

- Storage accounting system integrated in DGAS under development
- Non standard schema has been used to generate storage URs
  - no standard schema exists
  - we joined the Open Grid Forum (OGF) working group on the storage UR definition
- The information on the used space is collected from the BDII once a day
  - specific sensors can be developed to retrieve more accurate data making the system independent from the BDII
- Storage URs are stored on the HLR server just like the computing URs
- HLRmon is able to retrieve storage information directly from the desired HLR
- Current implementation of the sensor does not generate an XML (this will be changed in future implementations)

Storage accounting approach

- Some differences with the StAR (Storage Accounting Record) defined in EMI:
  - There is not a validity field but the TimeDuration that records the amount of time the storage has been used
  - The UR provides a field for storing the specific file to account, if a per file account is chosen

Storage report

- HLRmon shows the daily disk consumption in terms of used space
- Data are aggregated per site, SE and Storage Area (SA)
- Used space will be compared to the total available space (at SA, SE and site level)
- The plotting of these data allows to detect wrong values published in the BDII
  - Sites are notified in order to correct the values

Storage charts

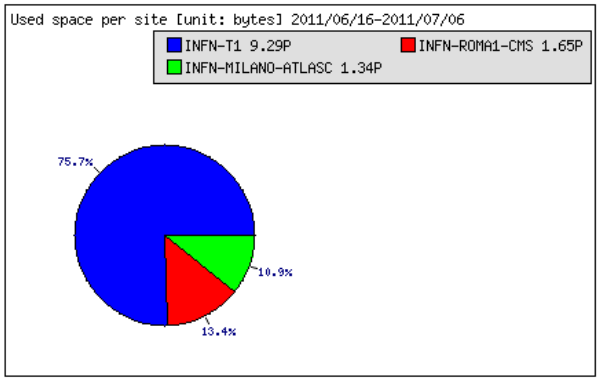

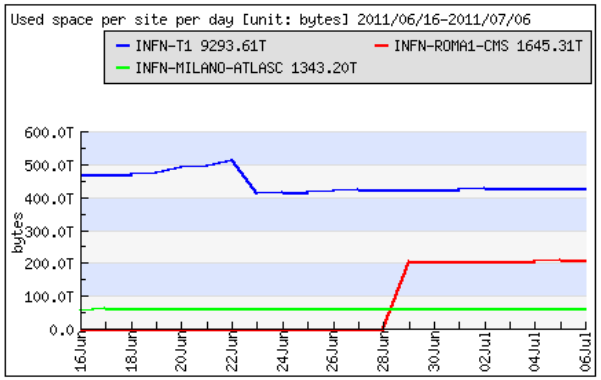

Used space per site
Used space per site per day

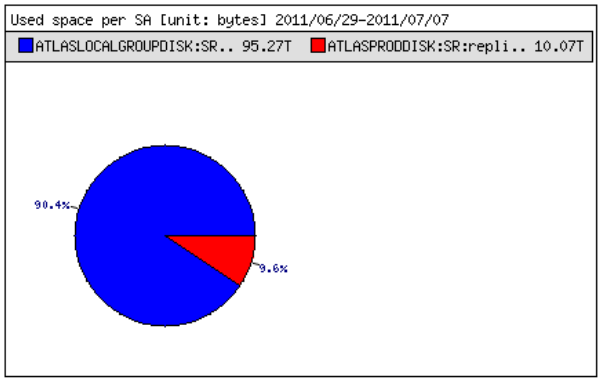

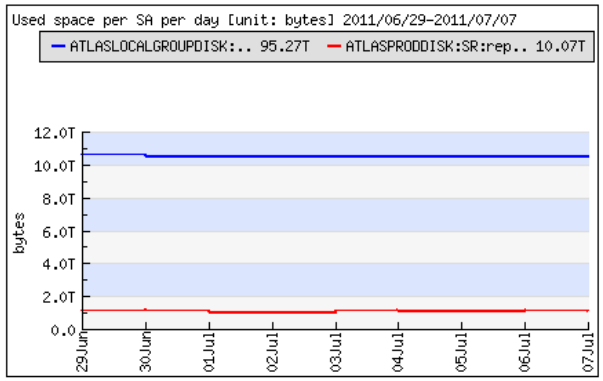

Used space per SA
Used per SA per day

OSG

- OSG already collects storage accounting through Gratia.
- Like CPU accounting they can massage their data into a common format and publish to RAL.

  Similar for any other infrastructures/sites

What now?

- The EMI timeline of release in May 2013 is far off
  - But perhaps it could be accelerated.
- Does WLCG wish for an interim solution?
  - Based on gstat or the IGI solution.