Enabling Grids for E-science in Europe

Accounting Update

Dave Kant, John Gordon
RAL
Javier Lopez, Pablo Rey Mayo
CESGA

Overview

Enabling Grids for E-science in Europe

- Status of Storage Accounting
- Status of APEL
- Status of UserLevel Accounting
- Status of DGAS2APEL
- Status of new SAM test
- Time scales

Storage Accounting Display

- Visualisation of Storage Used per VO for Disk and Tape
- http://goc02.grid-support.ac.uk/accountingDisplay/view.php
  - Select Resources via a Tree
  - Select time interval (last year, last month, last week, last day)

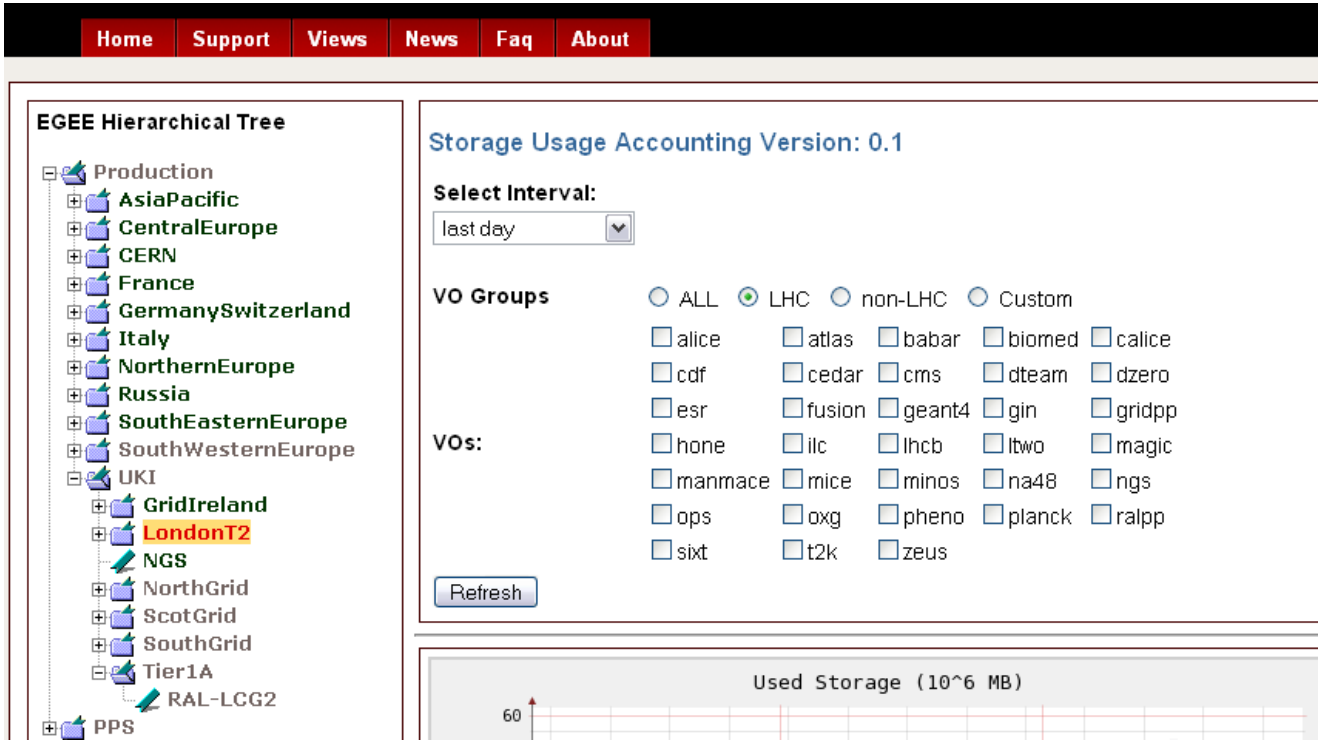

Storage Accounting Display

- Looking at data for RAL-LCG2
- Storage units are 1TB = 10^6 MB
- Tape Used + Disk Used = Total

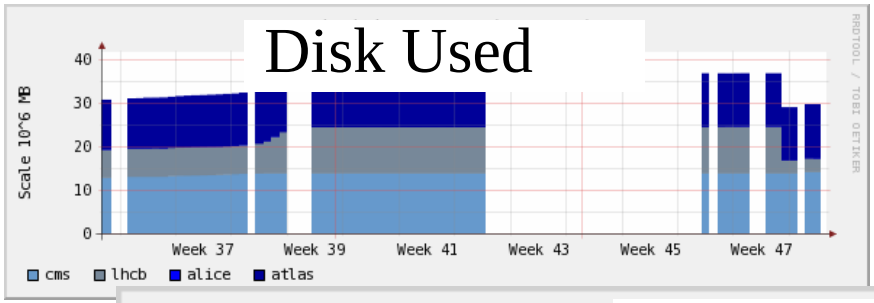

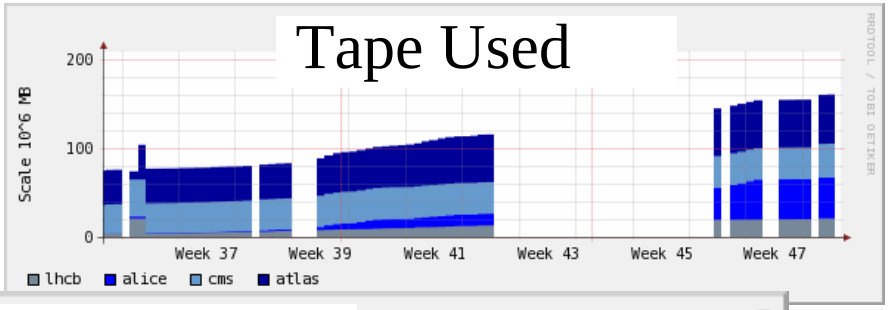

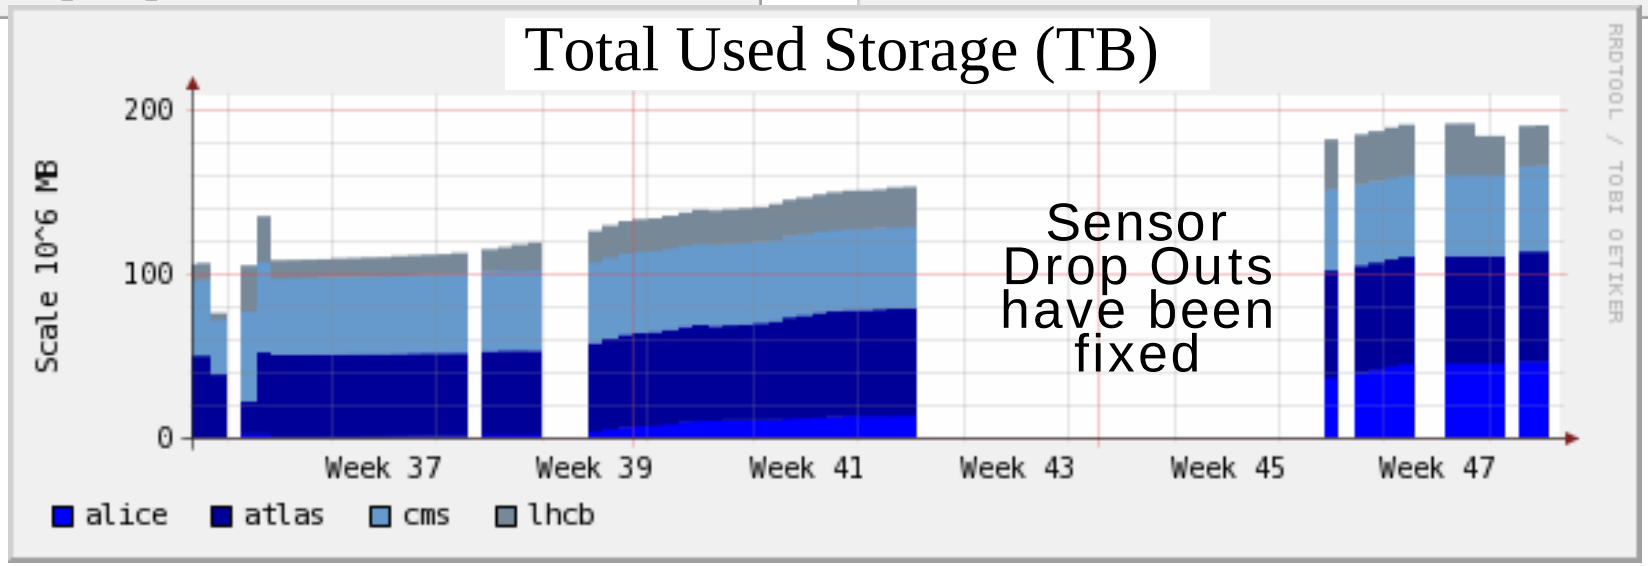

What's Next

Enabling Grids for E-science in Europe

- Currently, the Implementation is based on querying the GridPP sites in the information system. We will extend this to cover all sites EGEE/WLCG
- As with Job Accounting, we are planning to have a Tier-1 and Country View of Disk and Tape storage used per VO
- Bearing in mind that we are limited to showing information that is published via Glue into the information system, what other information should we display?
- Glue 1.3
  - Intra-VO reporting via Glue 1.3
  - More state information: UsedOnlineSpace, ReservedNearlineSpace

APEL Status

- APEL2 has been released in PPS in gLite 3.0.2 Update 10
- Main features are:
  - More reliable publisher which can handle tcp connection timeouts with the archiver.
  - Encryption of UserDN using a 1024-bit RSA key
  - Support for the Blah accounting file on the gLiteCE

DGAS2APEL Status

- Data from DGAS HLR has been inserted into the GOC accounting database using DGAS2APEL
- Tests performed by Rosario Piro (INFN)
  - ~ 2,000 Test records published for INFN-TORINO-DEV
  - To be deployed at Torino site first, then other Tier-2s
- Main issue is that Blah accounting continues to have problems (Savannah bug # 17591)
  - Blah prints out mangled information when job queues are full.

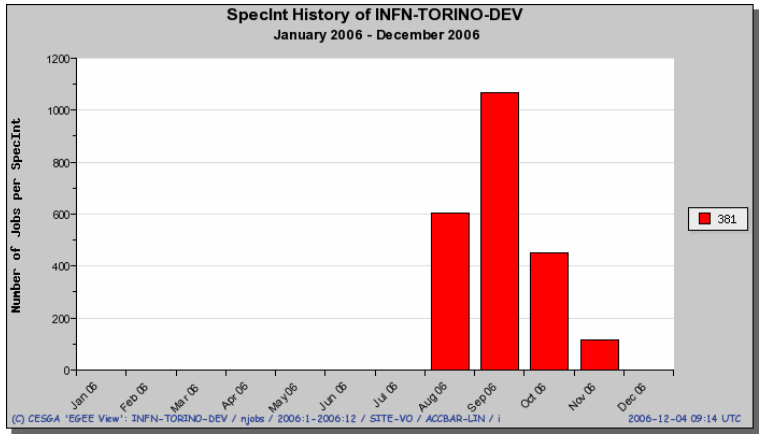

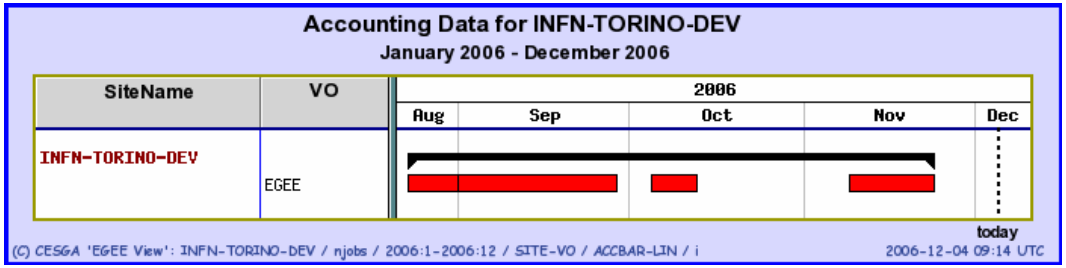

User-Level Accounting

Enabling Grids for E-science in Europe

- Development of a prototype User-level reporting display based on the “Five Actors” described:
  - VO Resource Manager
  - VO Member
  - User
  - Site Administrator
  - GOC Developer.
- Screen shots demonstrate this in action

VO-Resource Manager

Enabling Grids for E-science in Europe

- Table shows CPU, WCT and Job Eff. of the Top 10 Anonymised Users
- This example shows that the largest WCT User has a job efficiency of 10%…clearly the VO Manager may wish to contact this person?

EGEE View | VO MANAGER View | VO MEMBER View | SITE ADMIN View | USER View

January 2006 - December 2006.

The following table shows the Usage of the Top 10 Users ordered by Normalised CPU time and the Total Usage of the Other Users. A detailed view can be obtained by selecting an individual user.

Top 10 Users ordered by Normalised CPU time

User		Jobs		CPU time		*Norm. CPU time*		WCT		Norm. WCT		CPU Efficiency	Avg. CPU time
#	ID	#	%	Hrs	%	*Hrs*	*%*	Hrs	%	Hrs	%	%	Hrs
1	2daf892f6a1f68d0	15,629	25.2%	120	24.3%	***88***	***22.6%***	1,314	16.5%	990	16.0%	9.1	0.01
2	671e5b277fa6d1ac	20	0.0%	98	19.8%	***87***	***22.3%***	103	1.3%	91	1.5%	95.1	4.90
3	4cd78d6b029f7050	7,773	12.5%	80	16.2%	***55***	***14.1%***	924	11.6%	765	12.4%	8.7	0.01
4	5d61bd2201bec5f9	7,816	12.6%	52	10.5%	***41***	***10.5%***	331	4.2%	250	4.0%	15.7	0.01
5	2a153e141e98f06a	1,950	3.1%	36	7.3%	***36***	***9.2%***	45	0.6%	46	0.7%	80.0	0.02
6	62d3866c3c8260d6	39	0.1%	27	5.5%	***16***	***4.1%***	31	0.4%	18	0.3%	87.1	0.69
7	2756dfcb65975a47	95	0.2%	19	3.8%	***15***	***3.8%***	22	0.3%	18	0.3%	86.4	0.20
8	748206ea352cce31	467	0.8%	12	2.4%	***14***	***3.6%***	15	0.2%	16	0.3%	80.0	0.03
9	2d04c1be5d64c1b8	3	0.0%	11	2.2%	***13***	***3.3%***	26	0.3%	21	0.3%	42.3	3.67
10	2d72edb26620a697	83	0.1%	9	1.8%	***6***	***1.5%***	125	1.6%	79	1.3%	7.2	0.11
	Others (DN known)	25,073	40.5%	18	3.6%	***10***	***2.6%***	4,695	58.9%	3,527	57.1%	0.4	0.00
	Others (DN unknown)	3,021	4.9%	12	2.4%	***9***	***2.3%***	343	4.3%	356	5.8%	3.5	0.00
	Total	**61,969**		**494**		***390***		**7,974**		**6,177**		**6.2**	**0.01**

Click here for a csv dump of this table

Key: 0% <= eff < 75%; 75% <= eff < 90%; 90% <= eff < 100%; eff >= 100% (parallel jobs)

VO-Resource Manager

- Accumulative CPU of the Top 10

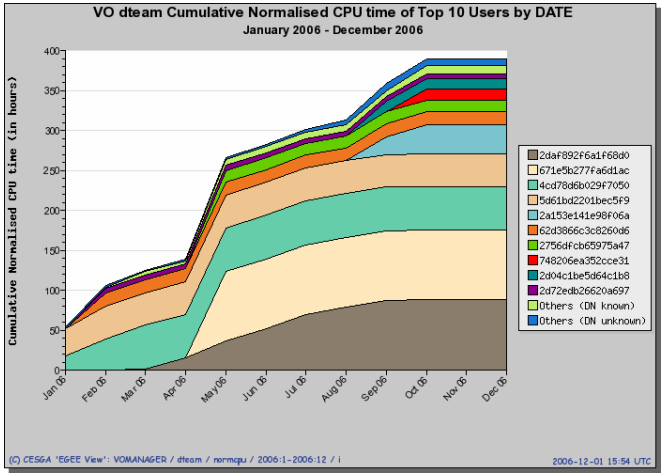

- Relative Share of Top 10 compared to the VO Total

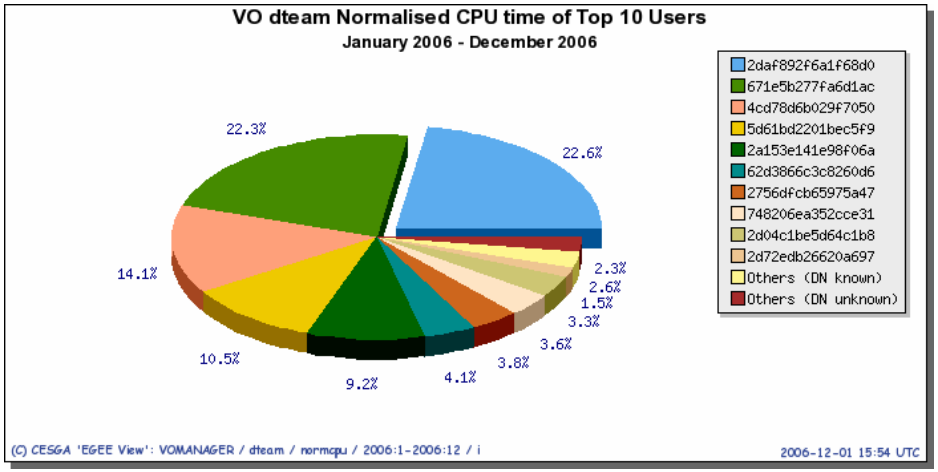

Site Admin View

- The Site Administrator can view usage of anonymous grid users who executed jobs at the site.

EGEE View | VO MANAGER View | VO MEMBER View | SITE ADMIN View | USER View

CESGA-EGEE User information.

January 2006 - December 2006.

The following table shows the Usage of the Top 10 Users ordered by Normalised CPU time and the Total Usage of the Other Users. A detailed view can be obtained by selecting an individual user.

Top 10 Users ordered by Normalised CPU time

	User	Jobs		CPU time		*Norm. CPU time*		WCT		Norm. WCT		CPU Efficiency	Avg. CPU time	Avg. WCT
#	ID	#	%	Hrs	%	*Hrs*	*%*	Hrs	%	Hrs	%	%	Hrs	Hrs
1	007c482b7a509753	335	0.5%	6,069	30.9%	***2,313***	***30.9%***	6,204	19.6%	2,364	19.6%	97.8	18.12	18.52
2	006f8bf719df068f	49,214	68.6%	2,769	14.1%	***1,054***	***14.1%***	6,955	21.9%	2,648	21.9%	39.8	0.06	0.14
3	57684c0c3d621a53	1,598	2.2%	1,717	8.7%	***653***	***8.7%***	2,052	6.5%	781	6.5%	83.7	1.07	1.28
4	43289fd45f650e5e	101	0.1%	1,616	8.2%	***616***	***8.2%***	1,981	6.2%	754	6.2%	81.6	16.00	19.61
5	11e9316e4987c00c	541	0.8%	1,295	6.6%	***493***	***6.6%***	1,337	4.2%	510	4.2%	96.9	2.39	2.47
6	6b641d611ad9af1e	101	0.1%	1,148	5.8%	***437***	***5.8%***	2,824	8.9%	1,076	8.9%	40.7	11.37	27.96
7	2b87e56d01c14620	186	0.3%	787	4.0%	***300***	***4.0%***	886	2.8%	338	2.8%	88.8	4.23	4.76
8	398ccd8b5d6c6e30	51	0.1%	678	3.5%	***258***	***3.4%***	788	2.5%	300	2.5%	86.0	13.29	15.45
9	0e2564274a92c273	1,098	1.5%	608	3.1%	***232***	***3.1%***	637	2.0%	243	2.0%	95.4	0.55	0.58
10	77f727593fa532af	625	0.9%	401	2.0%	***154***	***2.1%***	413	1.3%	157	1.3%	97.1	0.64	0.66
	Others (DN known)	17,930	25.0%	2,553	13.0%	***969***	***13.0%***	7,643	24.1%	2,903	24.0%	33.4	0.14	0.43
	Others (DN unknown)	0	0.0%	0	0.0%	***0***	***0.0%***	0	0.0%	0	0.0%		0.00	0.00
	Total	**71,780**		**19,641**		*7,479*		**31,720**		**12,074**		**61.9**	**0.27**	**0.44**

Click here for a csv dump of this table

Key: 0% <= eff < 75%; 75% <= eff < 90%; 90% <= eff < 100%; eff >= 100% (parallel jobs)

User

- Each Grid User can interrogate their own accounting data
  - Tables showing what they did and when
  - Number of Jobs, CPU and WCT per Month (per VO)
  - Average Job Efficiency per VO
  - Accumulative Njobs, CPU and WCT per VO
  - The sites which executed the jobs, and when they were done

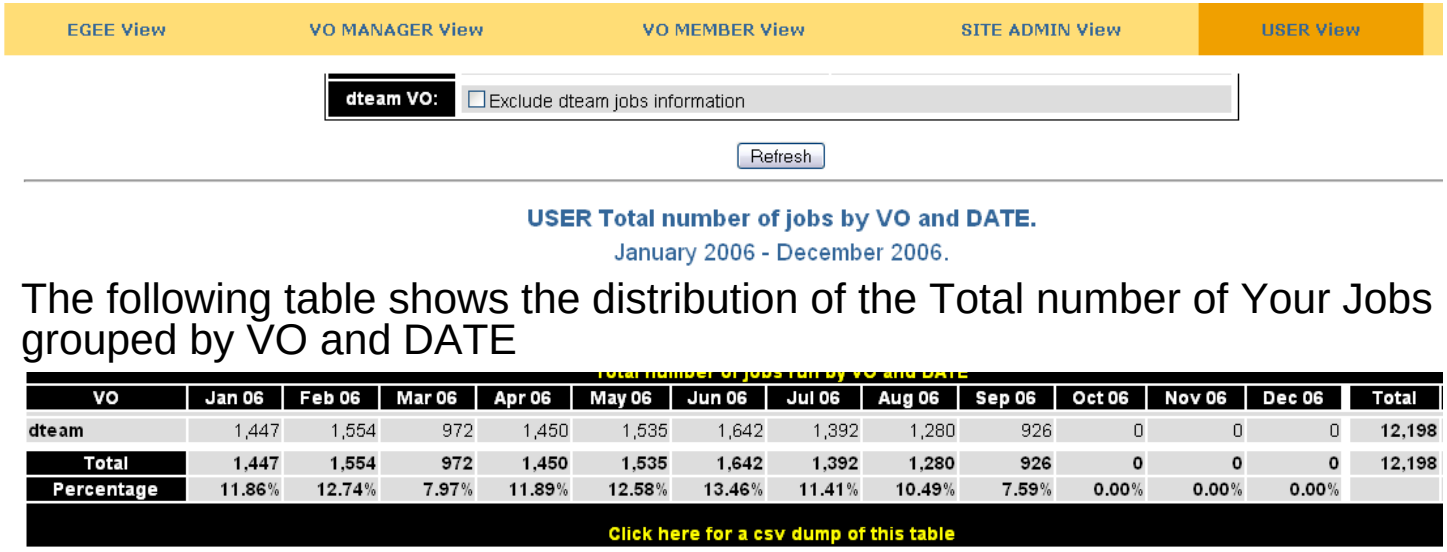

VO	Jan 06	Feb 06	Mar 06	Apr 06	May 06	Jun 06	Jul 06	Aug 06	Sep 06	Oct 06	Nov 06	Dec 06	Total
dteam	1,447	1,554	972	1,450	1,535	1,642	1,392	1,280	926	0	0	0	12,198
Total	**1,447**	**1,554**	**972**	**1,450**	**1,535**	**1,642**	**1,392**	**1,280**	**926**	**0**	**0**	**0**	**12,198**
Percentage	**11.86%**	**12.74%**	**7.97%**	**11.89%**	**12.58%**	**13.46%**	**11.41%**	**10.49%**	**7.59%**	**0.00%**	**0.00%**	**0.00%**	

What's Next

- Implement a proper security model
  - Ongoing in 2007
  - Currently, the prototype uses HTTPS (Apache/mod_SSL / Certificates) for Authentication, and a database table to describe the UserRoles for Authorisation.
    - The Site Administrator roles should be taken from GOCDB.
    - The VO Resource Members and VO Members should be taken from VOMS.
  - Can a short-lived VOMS certificate be place in the browser?
- Would Like Tier-1s to consider using this portal
  - Sometime in 2007
  - When APEL2 has been released into production
  - When the accounting policy document is ready to be signed
- Implications of glexec?
  - Investigate possibility of extracting user information from LCMAPS logs and joining with blah log

SAM Tests

- Synchronisation Test
  - Checks that the site accounting database and GOC are in agreement
  - Two RGMA tables:
    - one is built from site data (LcgRecordsSync_v1) every-time the site publishes
    - the other is built from job records at the GOC (LcgRecordsSyncGoc_v1)
  - Compare job numbers: LcgRecordsSync_v1 with LcgRecordsSyncGoc_v1
    - They must be equal for each Monthly aggregate
- Site APEL DB Up-to-date Test
  - Checks that the site accounting database contains recent data.
  - In LcgRecordsSync_v1: Compare RecordEnd with MeasurementDate
  - If site is publishing and is up-to-date, these will agree
  - If site is publishing but hasn't run any jobs for sometime, RecordEnd<MeasurementDate
- Is Site Publishing Test
  - Very similar to the existing SFT test
  - In LcgRecordSync_v1: MeasurementDate should be not older than 3 days.

Future Work / Wish List

- Can we correlate with the Resource Broker?
  - Given GlobalJobId, can we get L&B data?
  - Capture events from RGMA?
    - Difficult because we need to capture data from all RBs!
    - Clearly useful if you can divide resource shares according to RB:
      - Top 10 RBs, Efficiency of RB
- Is it sufficient to Aggregate at the per RB level?
  - Global WCT for all jobs through each RB
  - Compare with RBs estimate via L&B
  - Can the RB people provide a table that we could query?
  - Average values per RB per CE per Week
    - Waiting time, Wall Clock Time estimated,Number submitted / Number Done (= Efficiency of RB)
- Ask the RealTime monitor people for this!
- What happens if APEL doesn't contain RB Hostname?
  - Aggregate the RB data across all RBs and look at what happened at the site.

What Now?

- Storage
  - new GIPs for certification etc
  - GOC will starting collecting from all EGEE sites.
  - T1s should check their results
- APEL
  - As we requested in Rome(?), sites should check their APEL accounts with their local records.
  - Sites should check their published SI2K values
  - Adopt new SAM test
- DGAS
  - Extend testing to more INFN sites, VOs

What Now 2

- User Level Accounting
  - Once APEL2 is released for production, Tier1s should install (can do so now) and accumulate records.
  - Sites can publish encrypted DNs, we will not implement user identification until policies in place. Anonymous information will be displayed.
  - Feedback on portal views please.
- Does GLEXEC change anything?
  - Currently jobs will be accounted under DN of pilot job.
  - May be possible to join job to actual user from LCMAPS

Other Issues

- Agreement between GOC and local accounting (Holger)
  - Partly processed job logs ignored unless config is correct. Jobs appear as ‘local’.
  - See wiki. http://goc.grid.sinica.edu.tw/gocwiki/apel_bug_isActiveLog
  - Fixed in 3.0.2u10 currently in PPS
- ATLAS Agreement
  - ATLAS job database (?) do not agree with GOC results.
  - Some sites have a good match, some are out by up to 50%
  - Some probably due to issue above
  - Published kSI2K also looks a candidate.

- http://www3.egee.cesga.es/gridsite/accounting/CESGA/tree_egee.php?ExecutingSite=FZK-LCG2
- http://www3.egee.cesga.es/gridsite/accounting/CESGA/tree_egee.php?ExecutingSite=CERN-PROD

Accounting

- Is CERN Publishing wrong specInt value?