



# Accounting Update

John Gordon, CCLRC-RAL  
WLCG Workshop, CERN  
24/1/2007



# Outline


- This is not a technical explanation of WLCG accounting
- It is to bring the wider community up to date with the direction the management Board and Grid Deployment Board (GDB) are hoping accounting will progress this year and with some decisions the MB has made.
  - They will affect you
    - Automated reporting
    - Tier2 Accounting
    - User level Accounting
    - Storage Accounting
- Recently there have been presentations and discussions at the GDB and MB in December and January
  - See references at end




# Current Accounting

- Many sensors exist:- APEL, DGAS, SGAS, Gratia, local databases
- A central repository exists to hold information from multiple sources and produce integrated reports.
  - The APEL Repository,
  - Sometimes referred to as GOCDB (not strictly correct)

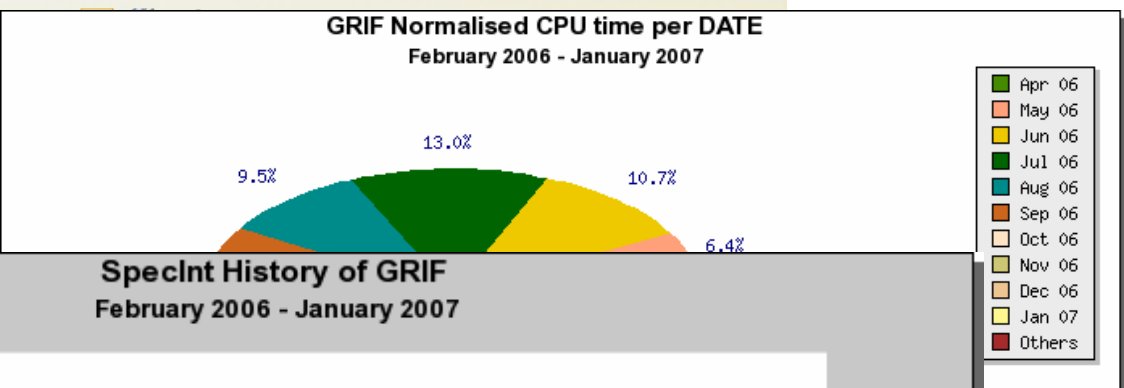

GRIF Normalised CPU time by February 2006 - January 2007



GRIF Cumulative Normalised CPU time (in hours) February 2006 - January 2007



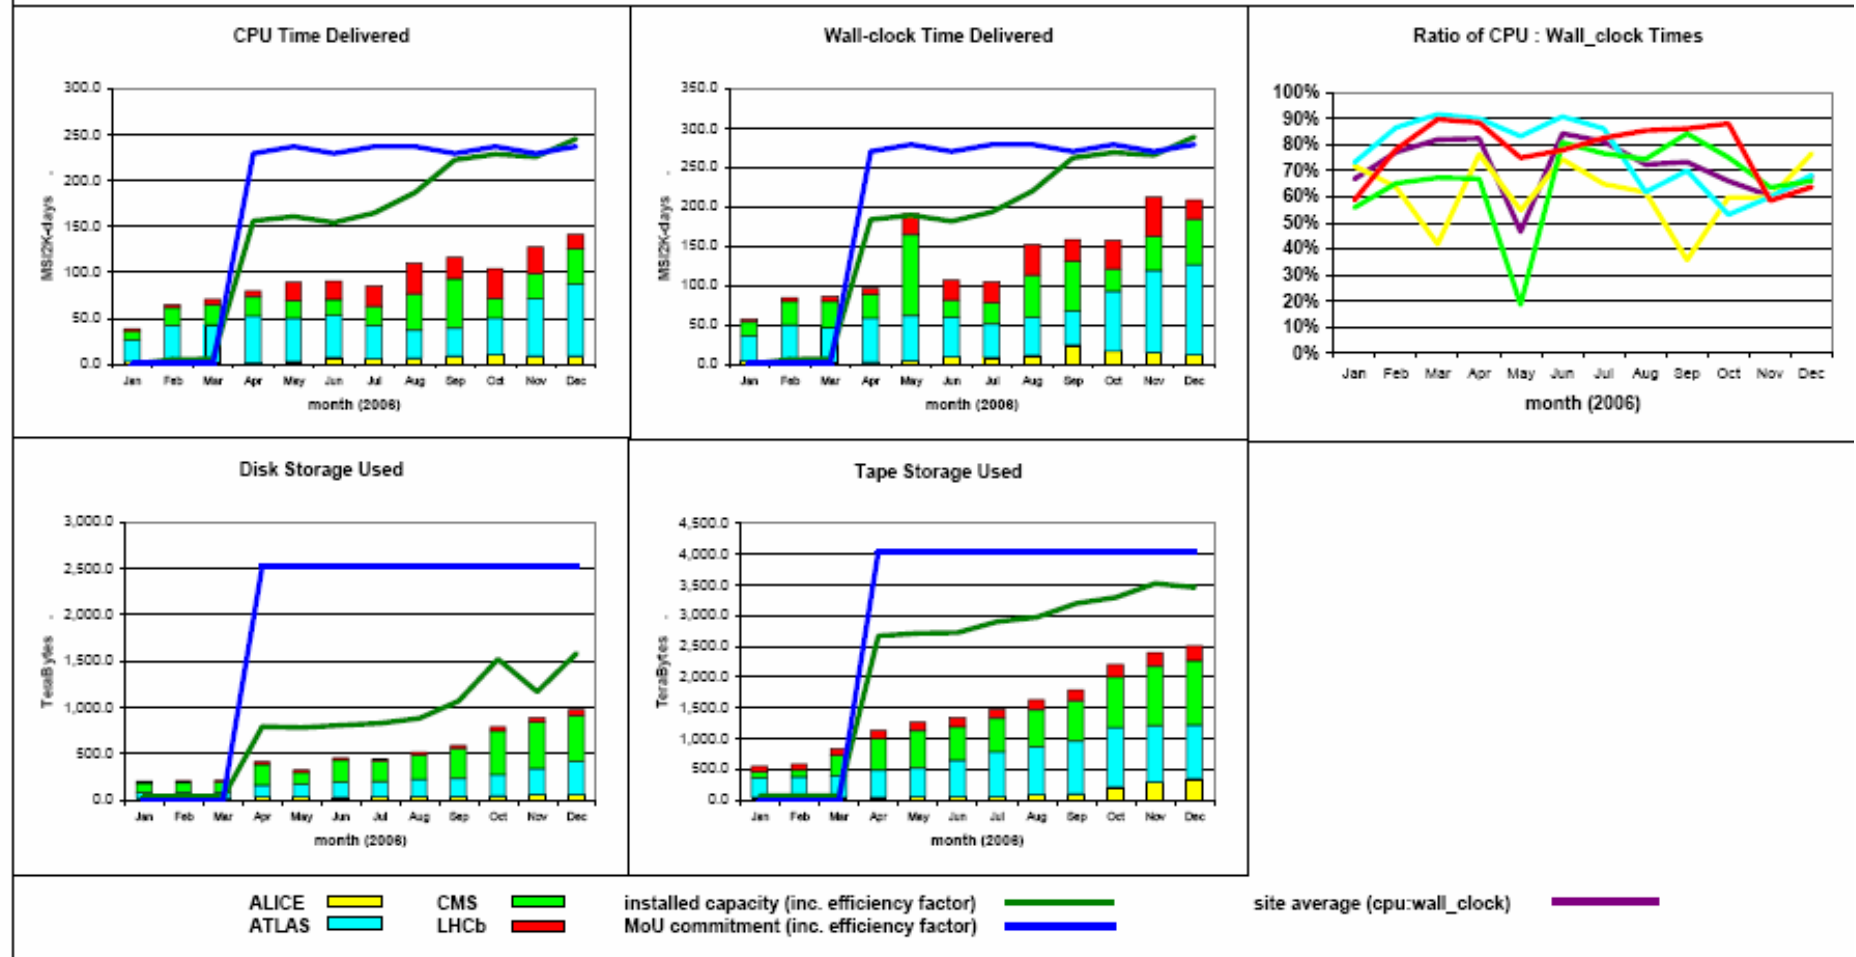
SpecInt History of GRIF February 2006 - January 2007



Accounting Data for GRIF February 2006 - January 2007

SiteName	VO	2006												
		Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	
GRIF	alice													
	atlas													
	biomed													
	calice													
	cms													
	cppm													
	daphnia													
	dteam													
	dzero													
	egeode													
	esr													
	fusion													
	geant4													
	grif													
	hone													
	ilc													
	lal													
	lhcb													
	ops													
	planck													
	proactive													

## Summary of Tier-1s



	total	112,000	124,064	147,140	164,400	204,114	300,000	404,440	420,100	440,000	460,000	471,140	494,000	total	4,001,011	total
installed capacity**		225	225	8,405	8,385	8,336	8,529	9,377	11,017	11,004	10,027	11,215				
cpu usage as % installed****	n/a	1819%	1792%	55%	66%	58%	61%	68%	60%	57%	60%	65%				

### Notes - cpu tables

- \* MoU monthly average - commitment in MoU for KB2K-days of cpu time per month (Includes efficiency factor)
- \*\* installed capacity - capacity installed (KB2K)
- \*\*\* aggregate 2008 to date - total - KB2K-days delivered in calendar year to current date  
- % MoU - percentage of MoU commitment in calendar year to current date
- \*\*\*\* cpu usage as % installed - includes efficiency factor



# Automate Reporting

Accounting should move to automatically-generated reports as soon as possible.

- a) The APEL repository will be the source of data
- b) Data will be extracted using a standard query early each month, on the previous month's data
- c) The target is 1<sup>st</sup> March 2007 but sites should publish data back to 1/1/2007 if possible.
- d) Sites can send an additional manual report until the end of May
- e) If sites wish their non-Grid usage to be reported to CRRB then they should use the APEL publishing tools to add this data to the repository.



# Tier2 Accounting

1. Reports to C-RRB should also include Tier2 data.
2. The source of the data will be the APEL repository.
3. Data will be reported from 1 April 07.
4. All T2s should be publishing by 1 September to enable report to C-RRB in April 2008 with 6 months data..



# User Level Accounting

1. APEL in gLite 3.0.1u10 has encrypted DN included in UR but not published by default.
  - FQAN also included
2. Sites should install this release and configure it to publish encrypted user data. Target for T1s 1 June but everyone welcome to publish
3. No DNs will be viewable until policy agreed.





# VO-Resource Manager

- 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	<a href="#">2daf892f6a1f68d0</a>	15,629	25.2%	120	24.3%	88	22.6%	1,314	16.5%	990	16.0%	9.1	0.01
2	<a href="#">671e5b277fa6d1ac</a>	20	0.0%	98	19.8%	87	22.3%	103	1.3%	91	1.5%	95.1	4.90
3	<a href="#">4cd78d6b029f7050</a>	7,773	12.5%	80	16.2%	55	14.1%	924	11.6%	765	12.4%	8.7	0.01
4	<a href="#">5d61bd2201bec5f9</a>	7,816	12.6%	52	10.5%	41	10.5%	331	4.2%	250	4.0%	15.7	0.01
5	<a href="#">2a153e141e98f06a</a>	1,950	3.1%	36	7.3%	36	9.2%	45	0.6%	46	0.7%	80.0	0.02
6	<a href="#">62d3866c3c8260d6</a>	39	0.1%	27	5.5%	16	4.1%	31	0.4%	18	0.3%	87.1	0.69
7	<a href="#">2756dfcb65975a47</a>	95	0.2%	19	3.8%	15	3.8%	22	0.3%	18	0.3%	86.4	0.20
8	<a href="#">748206ea352cce31</a>	467	0.8%	12	2.4%	14	3.6%	15	0.2%	16	0.3%	80.0	0.03
9	<a href="#">2d04c1be5d64c1b8</a>	3	0.0%	11	2.2%	13	3.3%	26	0.3%	21	0.3%	42.3	3.67
10	<a href="#">2d72edb26620a697</a>	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
<b>Total</b>		<b>61,969</b>		<b>494</b>		<b>390</b>		<b>7,974</b>		<b>6,177</b>		<b>6.2</b>	<b>0.01</b>

[Click here for a csv dump of this table](#)

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



# Storage Accounting

- A Storage Accounting schema has been defined
- A single storage accounting sensor at Edinburgh queries the BDII once per day to gather information on current storage usage from all EGEE sites
- This sensor populates an R-GMA archiver with the gathered information
  - StorageRecords table
- The information is aggregated per VO/Site/Month as in APEL
  - Value stored is average over the number of measurements
  - $\text{Sum}(\text{measurements})/\text{N}(\text{measurements})$
- APEL-like visualisation interface has been developed
  - Grids, countries, regions
  - <http://goc02.grid-support.ac.uk/accountingDisplay/>
  - Only Storage Used is displayed to date
  - Storage Installed will be displayed soon (data is stored)

### EGEE Hierarchical Tree

- Production
  - AsiaPacific
  - CentralEurope
  - CERN
  - France
  - GermanySwitzerland
  - Italy
  - NorthernEurope
  - Russia
  - SouthEasternEurope
  - SouthWesternEurope
  - UKI
    - GridIreland
    - LondonT2
    - NGS
    - NorthGrid
      - UKI-NORTHGRID-LANCS-HEP
      - UKI-NORTHGRID-LIV-HEP
      - UKI-NORTHGRID-MAN-HEP
      - UKI-NORTHGRID-SHEF-HEP
    - ScotGrid
      - ScotGRID-Edinburgh
      - scotgrid-gla
      - UKI-SCOTGRID-DURHAM
      - UKI-ScotGrid-Glasgow
    - SouthGrid
    - Tier1A
    - PPS

### Storage Accounting Display (Version 0.2)

#### Select Interval:

last month

#### VO Groups

ALL  LHC  non-LHC  Custom

alice  atlas  babar  biomed  cdf

cedar  cms  dteam  dzero  esr

#### VOs:


fusion  geant4  hone  ilc  lhcb

magic  mariachi  minos  ops  pheno

t2k  zeus

Refresh

#### Displaying the Total






#### Displaying Disk



# Storage Accounting Display

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





# Storage Issues

- Data not complete. Still only a prototype so sometimes doesn't run.
  - Hence the gaps in the data
- Need to check correctness of data
  - Can sites look at their published data and see if it makes sense?
  - Is it ready for automatic reporting?
- UK using new GIPs for DPM and dCache which report more accurate VO usage.
  - These improve quality of reporting but will need rewriting for SRM2.2 so not worth general release.
- More work to be done on display
  - New release including installed storage in a couple of weeks
  - Comments welcome



# Actions

- T1s compare recent APEL cpu data for their site with their recent manual reports
- T1s consider whether they should report their non-Grid use.
- T2s check they have been publishing into APEL reliably
  - T1s help their T2s with this.
- T1s start publishing user data
  - T2s follow suit if no problems.
- OSG to pass T2 cpu data to APEL
- OSG to publish user DN information
- All sites to check their storage accounting results



# References

- GDB December  
<http://indico.cern.ch/getFile.py/access?contribId=13&resId=0&materialId=slides&confId=a057712>
- MB December 19  
<http://indico.cern.ch/conferenceDisplay.py?confId=a063279>
- GDB January  
<http://indico.cern.ch/getFile.py/access?contribId=0&resId=0&materialId=slides&confId=8468>
- MB January 9