Accounting Status and Plans

GDB

John Gordon, STFC-RAL

GDB meeting @CERN October 8h, 2008

Science & Technology Facilities Council

Outline

- **CPU**
  - **Issues**
  - **UserDN**
- **Storage**
  - **Issues**
  - **New GIPs**
- **Megatable**

CPU Accounting - now

- The APEL repository gathers information from :-
  - Sites running the APEL client
  - Sites running their own accounting who then publish to the APEL repository (CC-IN2P3, NIKHEF)
  - Grids who collect their own accounting and bulk publish to the APEL repository (INFN, OSG, NGDF T1)
- Reports
  - T1 report extracted monthly and sent to T1s for approval
  - T2 report extracted monthly and sent to CB for approval
    - Based on Tier2s who have signed MoU. List manually maintained
  - Portal allows tailored reports and .csv output to manipulate yourself.
- Njobs, Cpu time, wallclock, normalised cpu and wallclock
- Cpu efficiency (cpu/wallclock)


EGEE Accounting Portal - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://www3.egee.cesga.es/gridsite/accounting/CESGA/country_view.html
Go Links
Google Go Bookmarks 127 blocked Check AutoLink AutoFill Send to
EGEE ACCOUNTING PORTAL
GLOBAL View
VO MANAGER View
VO MEMBER View
SITE ADMIN View
USER View
REPORTS
Hierarchical Tree
Tier1
Tier2
Countries
Armenia
Australia
Austria
GUP-JKU
HEPHY-UIBK
Hephy-Vienna
Belarus
Belgium
Brazil
Bulgaria
Canada
China
Croatia
Cyprus
Czech Republic
Denmark
Estonia
EGEE View -> Countries
Print Pa
Data to graph: Norm. Sum CPU Normalised CPU time to a reference value of 1000 SpecInt2000
Period: Start year: 2007 Start month: 11 End year: 2008 End month: 10
Groupings: Show data for: COUNTRY as a function of: VO
VO Groups: LHC Official EGEE TOP 10 ALL Custom
Group the rest of VOs in a new category
Chart: Type: ACCUM BAR Scale: LINEAR
dteam VO: Exclude dteam and ops VOs jobs information
Refresh
COUNTRY Normalised CPU time by COUNTRY and VO.
EGEE VOs. November 2007 - October 2008.
Enabling Grids

WLCG Accounting Summary

August 2008

Centre Summary of CERN + Tier-1s

Please report accounting data in the shaded cells and return the report to lcg.office@cern.ch

MoU pledges	2007 jan-mar	2008 apr-aug	2008 sep-dec
CPU (KSI2K-years)	22,470	52,262	53,410
Disk (Tbytes)	8,520	25,143	25,780
Tape (Tbytes)	8,910	32,169	33,350

Standard efficiency factors	
Scheduled cpu used (Tier-0, -1)	85%
Chaotic cpu used (Tier-2)	60%
Disk utilisation	70%
Mass store utilisation	100%

CPU used - *KSI2K-days*

			Jan-08	Feb-08	Mar-08	Apr-08	May-08	Jun-08	Jul-08	Aug-08	Sep-08	Oct-08	Nov-08	Dec-08	aggregate 2008 to date*** MoU pledge	total usage	%MoU
CPU Grid	ALICE	cpu	41,747	15,280	43,319	58,104	54,974	94,428	123,897	129,111						560,860	
		wall	65,605	36,043	76,366	106,242	87,923	165,699	182,076	177,374						897,326	
	ATLAS	cpu	114,452	86,392	180,819	184,554	119,214	175,867	146,937	180,464						1,188,699	
		wall	140,282	111,694	211,698	203,755	146,822	195,474	180,282	204,044						1,394,051	
	CMS	cpu	47,415	99,562	51,612	118,093	218,234	135,745	86,230	61,733						818,624	
		wall	119,977	199,524	161,274	211,439	380,744	206,142	216,681	107,277						1,603,056	
	LHCb	cpu	5,332	12,513	15,518	18,172	24,270	11,808	12,723	4,530						104,866	
		wall	13,415	23,868	23,757	26,467	32,284	20,376	25,582	11,086						176,835	
	TOTAL	cpu	208,946	213,747	291,268	378,923	416,692	417,848	369,787	375,838	0	0	0	0	8,515,630	2,673,049	31%
		wall	339,279	371,129	473,095	547,903	647,773	587,691	604,621	499,781	0	0	0	0		4,071,272	n/a
CPU Non-Grid	ALICE	cpu	624	948	476	154	734	378	2,199	1,911						7,424	
		wall	1,690	3,403	2,058	1,700	2,130	1,128	13,394	11,667						37,170	
	ATLAS	cpu	31,469	34,735	39,921	41,289	60,063	49,507	57,459	55,672						370,115	
		wall	42,313	46,864	54,362	51,816	74,112	65,244	93,547	98,091						526,349	
	CMS	cpu	29,776	18,137	36,363	60,532	62,183	37,082	26,715	32,056						302,844	
		wall	69,655	32,126	68,687	79,609	94,391	64,728	58,417	71,528						539,141	
	LHCb	cpu	663	1,127	954	1,148	1,283	1,028	1,780	1,383						9,566	
		wall	1,318	1,375	1,918	1,782	2,670	1,942	6,842	5,113						22,960	
	TOTAL	cpu	62,752	54,947	77,714	103,123	124,263	87,995	88,153	91,022	0	0	0	0	8,515,630	689,969	8%
		wall	114,976	83,768	127,025	134,907	173,303	133,042	172,200	186,399	0	0	0	0		1,125,620	n/a
CPU Total	ALICE	cpu	42,371	16,228	43,795	58,258	55,708	94,806	126,096	131,022						568,284	
		wall	67,295	39,446	78,424	107,942	90,053	166,827	195,470	189,041						934,496	
	ATLAS	cpu	145,921	121,127	220,740	225,843	179,277	225,374	204,396	236,136						1,558,814	
		wall	182,595	158,558	266,060	255,571	220,934	260,718	273,829	302,135						1,920,400	
	CMS	cpu	77,191	117,699	87,975	178,625	280,417	172,827	112,945	93,789						1,121,468	
		wall	189,632	231,650	229,961	291,048	475,135	270,870	275,098	178,805						2,142,199	
	LHCb	cpu	6,215	13,640	16,472	19,320	25,553	12,836	14,503	5,913						114,452	
		wall	14,733	25,243	25,675	28,249	34,954	22,318	32,424	16,199						199,795	
	TOTAL	cpu	271,698	268,694	368,982	482,046	540,955	505,843	457,940	466,860	0	0	0	0	8,515,630	3,363,016	39%
		wall	454,255	454,897	600,120	682,810	821,076	720,733	776,821	686,180	0	0	0	0		5,196,892	n/a
	*installed capacity***		22,725	23,054	24,584	45,854	44,967	50,010	50,010	50,010							
	*cpu usage as % installed****		45%	49%	57%	41%	47%	40%	36%	37%	n/a	n/a	n/a	n/a			
	*MoU pledge**		592,085	534,786	592,085	1,332,681	1,377,104	1,332,681	1,377,104	1,377,104	1,361,955	1,407,354	1,361,955	1,407,354			

Notes - cpu tables

* **MoU pledge** - pledge in MoU for KSI2K-days of cpu time for current month (includes efficiency factor)
note - 2007 pledge to end March, 2008 pledge from April

** **installed capacity** - capacity installed (KSI2K)

*** **aggregate 2008 to date** - **total** - KSI2K-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

Adobe Reader - [Master%20accounting%20summaries_August2008.pdf]

File Edit View Document Tools Window Help

Save a Copy Search Select Object Data Tool 100% Help Search Web

Pages Attachments Comments

Summary of CERN + Tier-1s

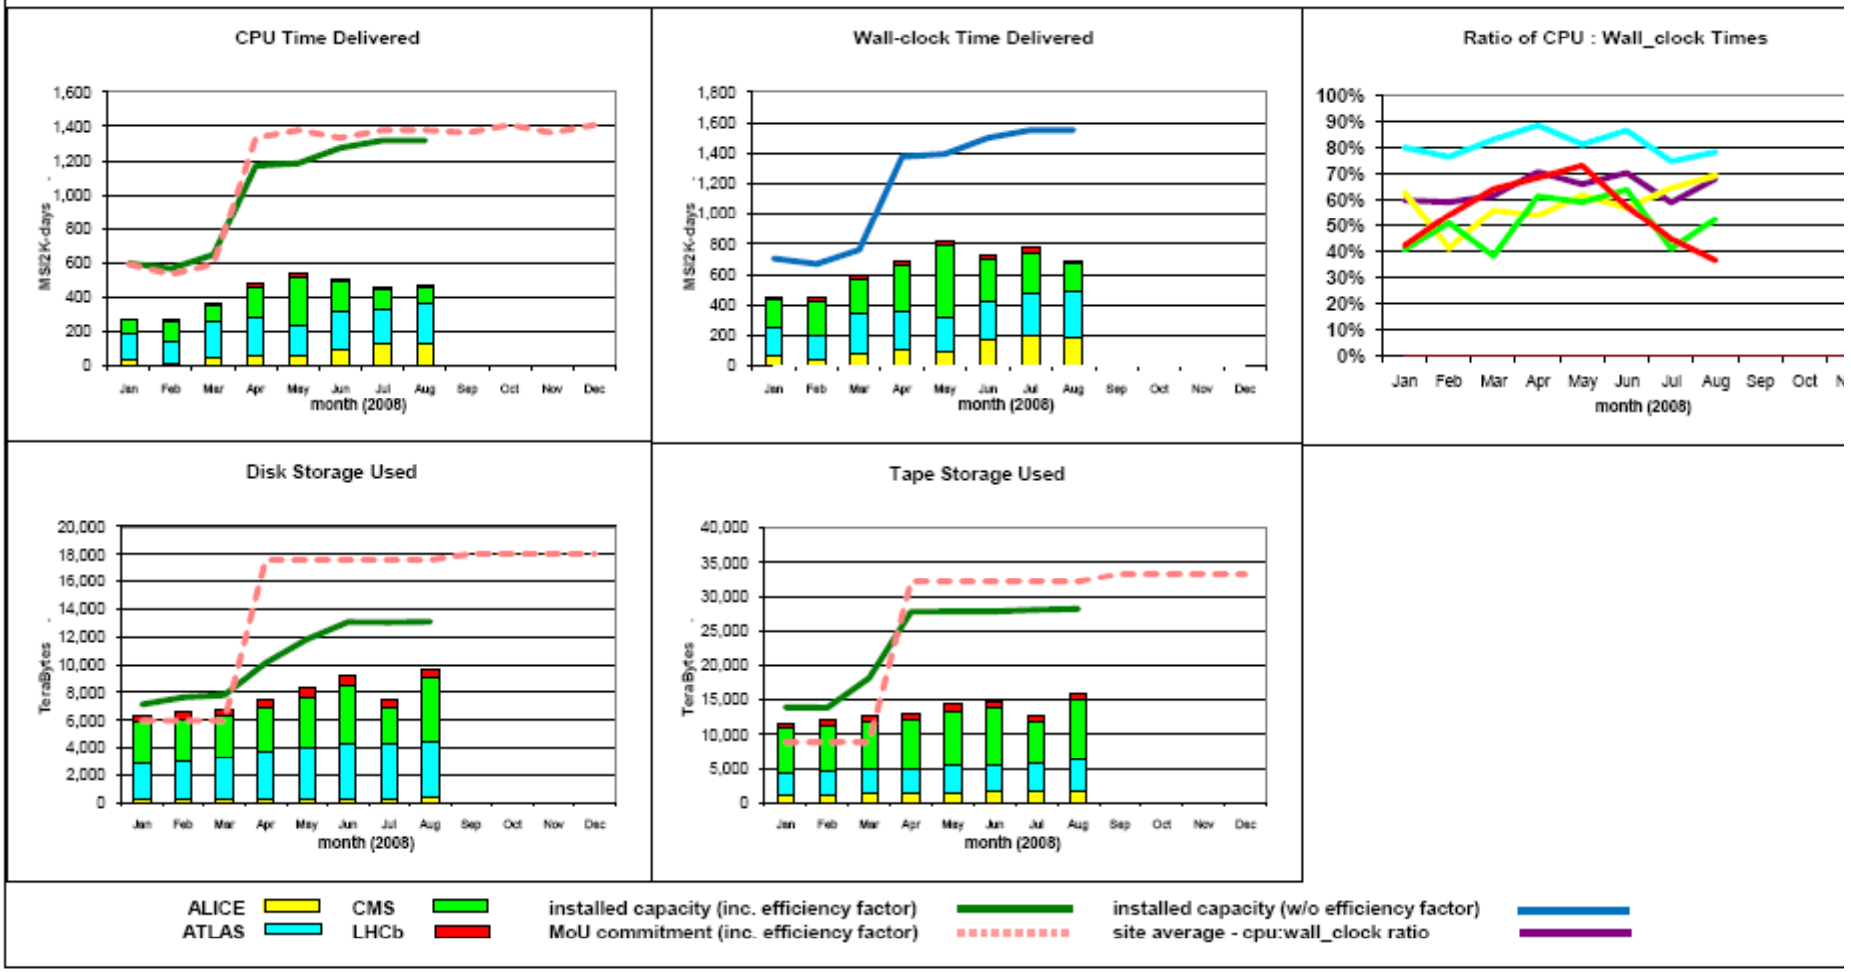

WLCG - Tier-2 Accounting Summary — July 2008

Efficiency factor for Tier-2 sites - utilisation 60% of pledge as specified in TDR

CPU usage in month (KSI2K-Hrs)

Federation - Accounting Name	2008 CPU Pledge (KSI2K)	pledge inc. efficiency (KSI2K-Hrs)	Site(s)	ALICE	ATLAS	CMS	LHCb	Total	used as % of pledge
Australia, University of Melbourne			Australia-ATLAS						
			Australia-UNIMELB-LCG2						
AU-ATLAS	150	66,960							
Austria, Austrian Tier-2 Federation			HEPHY-UIBK		17,473			17,473	
AT-HEPHY-VIENNA-UIBK	540	241,056			17,473			17,473	7%
Belgium, Belgian Tier-2 Federation			BEgrid-ULB-VUB						
			BelGrid-UCL						
BE-TIER2	1,050	468,720							
Canada-East Federation			TORONTO-LCG2						
CA-EAST-T2	200	89,280							
Canada-West Federation			ALBERTA-LCG2		45,794			45,794	
			SFU-LCG2		41,864			41,864	
			VICTORIA-LCG2		4,465			4,465	
CA-WEST-T2	300	133,920			92,123			92,123	69%
China, IHEP, Beijing			BEIJING-LCG2		10,036	15,384		25,420	
CN-IHEP	400	178,560			10,036	15,384		25,420	14%
Czech Rep., FZU AS, Prague			praguelcg2	37,283	11,470			48,753	
			prague_cesnet_lcg2		8,570			8,570	
CZ-Prague-T2	164	73,210		37,283	20,040			57,323	78%
Estonia, NICPB, Tallinn			T2_Estonia			25,519		25,519	
EE-NICPB	150	66,960				25,519		25,519	38%
Finland, NDGF/HIP Tier2			CSC						
FI-HIP-T2	564	251,770							
France, CC-IN2P3 AF			IN2P3-CC-T2	256,340	243,202	157,300	778	657,620	
FR-IN2P3-CC-T2	1,500	669,600		256,340	243,202	157,300	778	657,620	98%
France, GRIF, Paris			GRIF	341,664	171,920	113,501	122,574	749,659	
FR-GRIF	1,642	732,989		341,664	171,920	113,501	122,574	749,659	102%
France, LAPP, Annecy			IN2P3-LAPP		61,941		33,669	95,610	
FR-IN2P3-LAPP	600	267,840			61,941		33,669	95,610	36%
France, LPC, Clermont-Ferrand			IN2P3-LPC						
FR-IN2P3-LPC	800	357,120							
France, SUBATECH, Nantes			IN2P3-SUBATECH	209,209				209,209	

What goes wrong?

- Sites don't publish
  - Old software. MON box not updated
  - EGEE now has a SAM test (name?) which becomes critical if a site doesn't publish for a month. COD in EGEE will raise tickets against a site.
  - APEL client configuration may not be straightforward due to local batch configurations. See the APEL wiki, Raise a GGUS ticket.
- Tier2 sites change (join, leave, change name) a federation
  - List is maintained by lcg.office@cern.ch
  - Tier2 sites are a subset of EGEE/OSG/NorduGrid/etc. Cannot automatically track changes.
- CPU Usage is compared with WLCG Pledges.
  - Not with actual installed capacity

UserDN/FQAN

- APEL publishes UserDN and FQAN (VOMS proxy)
  - if configured
  - Not default
  - UserDN encrypted
- In September 08,
  - 45% of sites (51% of jobs) publish UserDN
  - 80% of sites (74 of jobs) publish FQAN
- VOs would like to know this information
  - Only approved people can view it
  - Please configure it at your site.
  - The site needs to have the Savannah patch #898 installed in the CE so that it creates accounting log files.
  - Then they need to set the APEL parser configuration file to use the BlahdLogProcessor instead of the GkLogProcessor (BlahdLogProcessor is the default mode now).

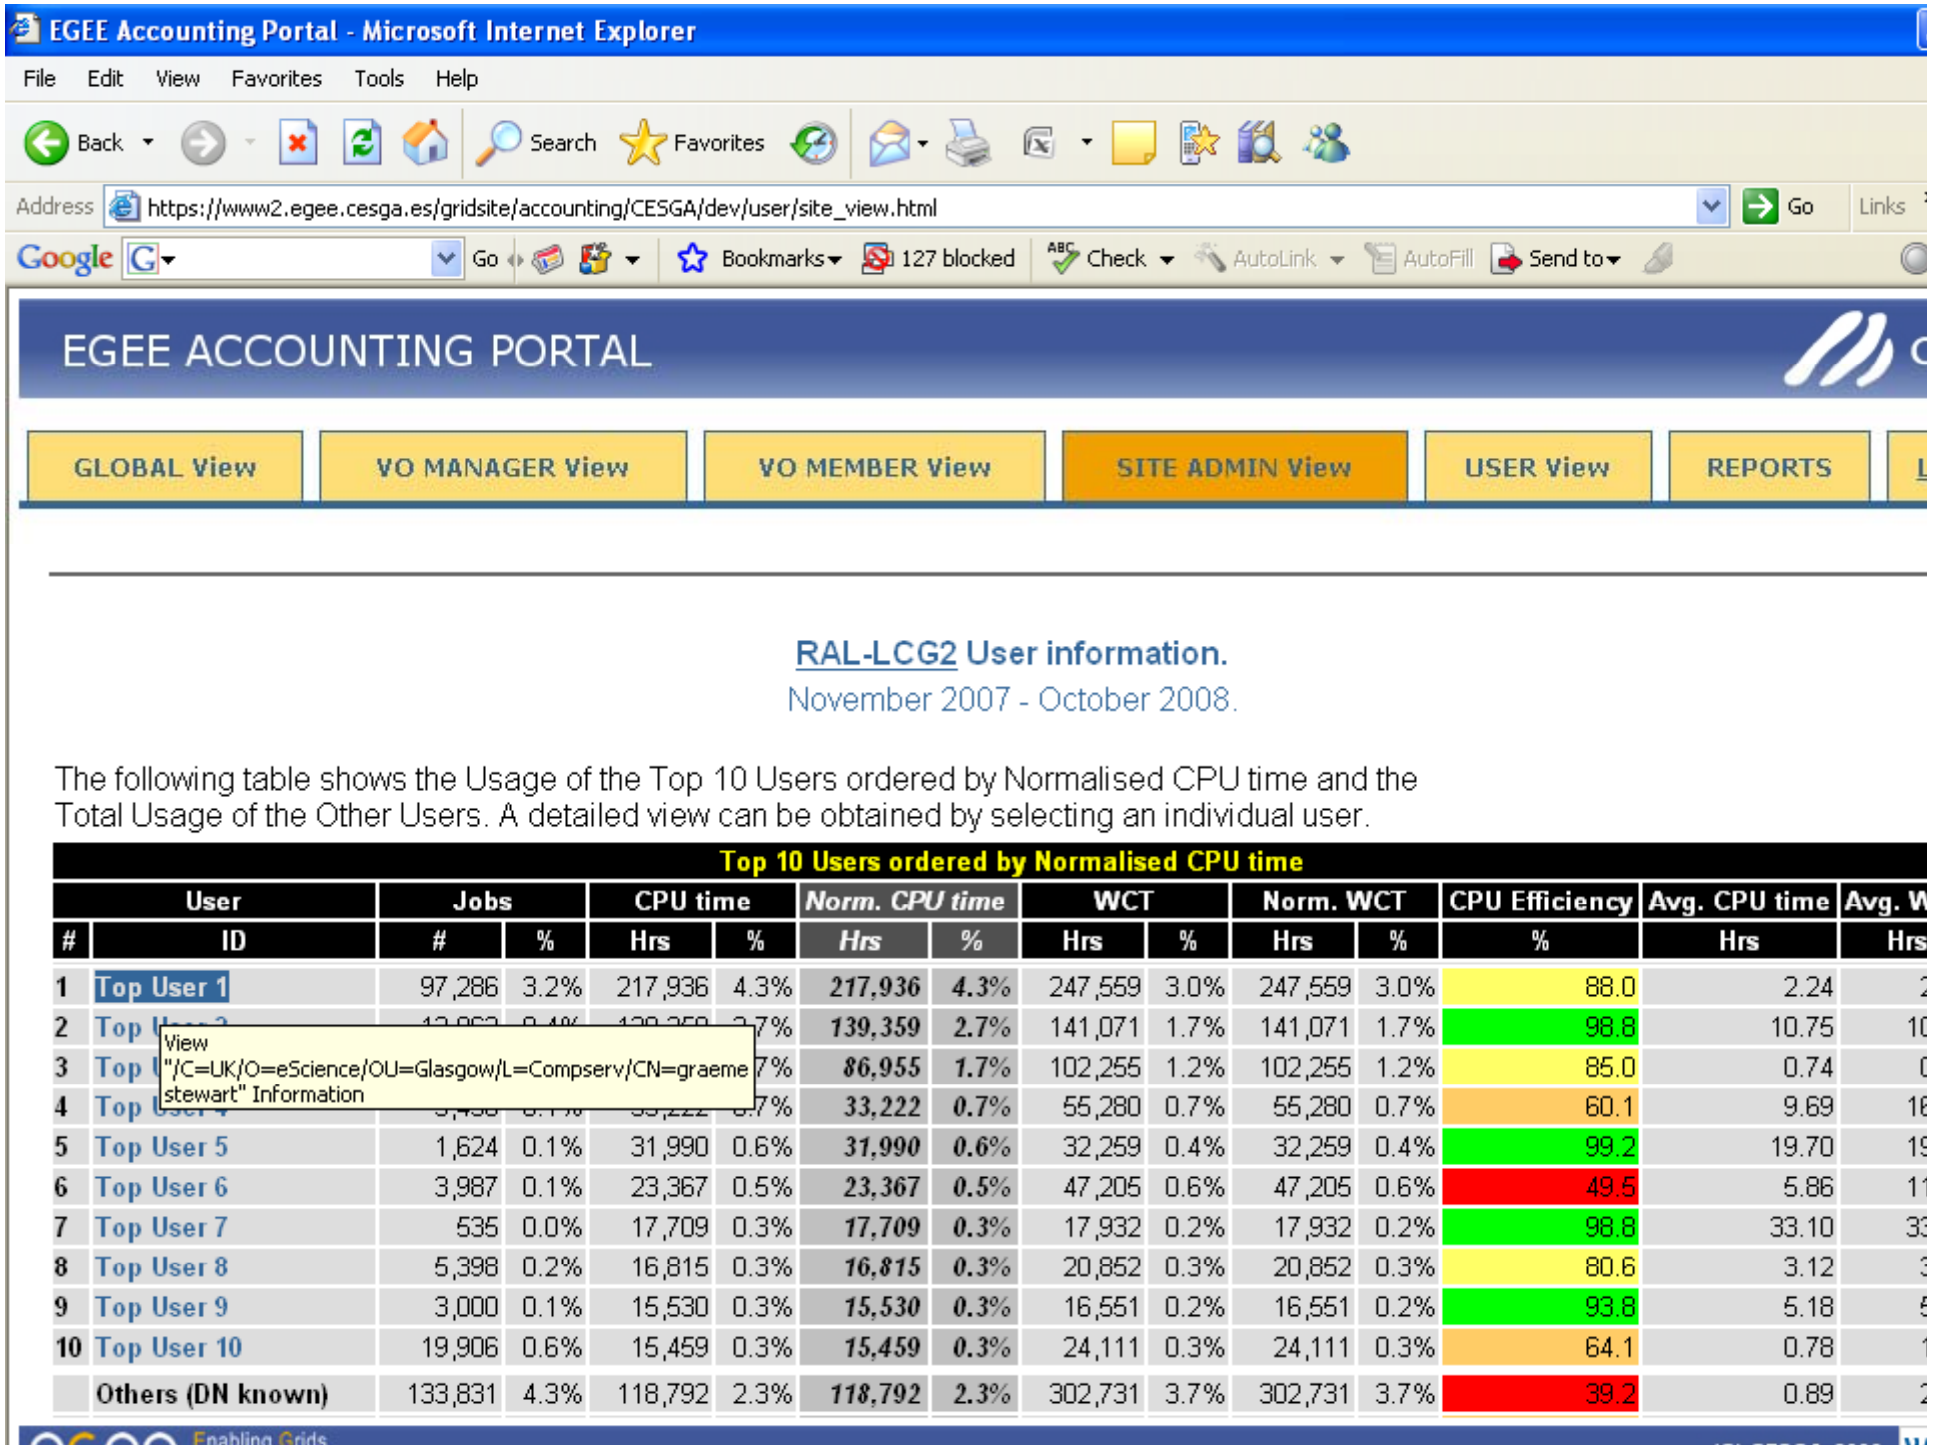
EGEE Accounting Portal - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address https://www2.egee.cesga.es/gridsite/accounting/CESGA/dev/user/site_view.html
Go Links
Google Go Bookmarks 127 blocked Check AutoLink AutoFill Send to
EGEE ACCOUNTING PORTAL
GLOBAL View
VO MANAGER View
VO MEMBER View
SITE ADMIN View
USER View
REPORTS
RAL-LCG2 User information.
November 2007 - October 2008.
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 Top User 1 97,286 3.2% 217,936 4.3% 217,936 4.3% 247,559 3.0% 247,559 3.0% 88.0 2.24
2 Top User 2 139,359 2.7% 141,071 1.7% 141,071 1.7% 98.8 10.75
3 Top User 3 86,955 1.7% 102,255 1.2% 102,255 1.2% 85.0 0.74
4 Top User 4 33,222 0.7% 55,280 0.7% 55,280 0.7% 60.1 9.69
5 Top User 5 1,624 0.1% 31,990 0.6% 31,990 0.6% 32,259 0.4% 32,259 0.4% 99.2 19.70
6 Top User 6 3,987 0.1% 23,367 0.5% 23,367 0.5% 47,205 0.6% 47,205 0.6% 49.5 5.86
7 Top User 7 535 0.0% 17,709 0.3% 17,709 0.3% 17,932 0.2% 17,932 0.2% 98.8 33.10
8 Top User 8 5,398 0.2% 16,815 0.3% 16,815 0.3% 20,852 0.3% 20,852 0.3% 80.6 3.12
9 Top User 9 3,000 0.1% 15,530 0.3% 15,530 0.3% 16,551 0.2% 16,551 0.2% 93.8 5.18
10 Top User 10 19,906 0.6% 15,459 0.3% 15,459 0.3% 24,111 0.3% 24,111 0.3% 64.1 0.78
Others (DN known) 133,831 4.3% 118,792 2.3% 118,792 2.3% 302,731 3.7% 302,731 3.7% 39.2 0.89
View "/C=UK/O=eScience/OU=Glasgow/L=Compserv/CN=graeme stewart" Information
Enabling Grids

Storage

- APEL is a client that runs at a site and publishes cpu usage to a central location.
- Currently no equivalent client for storage systems
- The GOC Storage Portal attempts to harvest storage information from the GLUE Schema in the Information Service.

- BUT problems with the quality of the data
  - Some sites don't publish, some publish in the wrong units.
  - (up to factor of 10**9 out),
  - Some only publish

  - Information not yet useful.

Installed Capacity

- Accounting data needs putting into context of installed capacity
- Flavia Donno leading a group looking at this.
- Also looking at improving storage accounting

Installed Capacity WG

- A draft document with all details about the usage of the Glue Schema 1.3 to publish the computing and storage installed capacity is available:
  - https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.1.pdf
- Storage resource description has been agreed.
- For computing resources the situation is not so clear.
  - Please, check Steve Traylen's (Worker Nodes WG) proposal

Science & Technology Facilities Council

Storage Capacity Status

- The current CASTOR information provider (v1.1.0) is deployed at RAL.
  - Not fully compliant with the spec. However, only minor changes needed.
- dCache Information Providers will be available with dCache v.1.9.1 (20th of October 2008).
  - No installation yet available to validate the output of the information providers. Input waited from dCache developers for description of special pools (READ-only, WAN/LAN, etc.)
- DPM information providers installed at some sites in France and UK.
  - Not fully compliant with the spec. Minor changes needed.
  - The packaging of the DPM specific information providers and the yaim configuration generating component are available. In certification.
- StoRM information providers will probably be available around the middle of November with StoRM v. 1.4.0.
  - No installation yet available to validate the output of the information providers.

Science & Technology Facilities Council

Computing Capacity Status

- The needed information are retrieved by the Subluster/Host Glue Classes.
- WMS limits us to publishing one SubCluster per Cluster (basically the current WMS can't tell the LRMS which SubCluster it wants to use). The result is that sites have to publish some kind of average/minimum WN specification.
- The proposal is to split the system and have separate queues (CEs), e.g. for large memory nodes. Splitting the system to have separate homogeneous queues is the solution that is proposed to be adopted in WLCG.
  - See plan proposed by the Worker Nodes WG.
- The situation will change with the adoption of CREAM.

Science & Technology Facilities Council

Summary

- WLCG wants cpu and storage accounting for its VOs.
  - Please make sure yours is working.
  - Check your site in the portal
- New information providers coming for your SEs
  - Please install them once announced
- Megatable will estimate your installed capacity
  - We will be asking you to check this.
- UserDN, FAQN – please switch on.