

Accounting Overview:

From main page:

Sum Normalized Elapsed time * Number Processor: why in Tier1/s but not in main?

Snapshot of main page

The screenshot shows the EGI Accounting Portal interface. The top navigation bar includes links for GLOBAL View, VO MANAGER View, VO MEMBER View, SITE ADMIN View, USER View, REPORTS, and METRICS PORTAL. The main content area is titled 'EGI View --> Production'. On the left is a 'Hierarchical Tree' with a tree view showing various categories like Tier1, Tier2, Countries, EMI3(WLCG), EGI, OSG, UNREGISTERED, VO_Discipline, VO_Metrics, CUSTOM_view, Cloud, Cloud Tier1, and Cloud Tier2. The central panel displays a configuration window for 'Data to graph' with a dropdown menu open, showing options like 'Norm. Sum CPU (kSI2K-hours)', 'Norm. Sum CPU (HEPSPEC06-hours)', 'Sum CPU', 'Norm. Sum Elapsed (kSI2K-hours)', and 'Norm. Sum Elapsed (HEPSPEC06-hours)'. The 'Period' is set to 5, 'End year' to 2016, and 'End month' to 4. The 'Groupings' are set to 'as a function of: VO'. The 'Chart' is set to 'Sum Elapsed CPU Efficiency'. The 'dteam VO' is set to 'Exclude dteam and ops VOs jobs information'. The 'Local Jobs' are set to 'Grid Jobs Only'. A 'Refresh' button is at the bottom. Below the configuration window, the text reads: 'Normalised CPU time (kSI2K) by REGION and VO. LHC VOs. May 2015 - April 2016.'

Vs snapshot of Tier1 page

The screenshot shows the EGI Accounting Portal interface for the 'Tier1' view. The top navigation bar is the same as the main page. The main content area is titled 'EGI View --> Tier1'. The 'Hierarchical Tree' on the left is the same. The central panel displays a configuration window for 'Data to graph' with a dropdown menu open, showing options like 'Sum Normalised Elapsed time * number Processors', 'Sum CPU time', 'Sum Elapsed time', 'Sum Normalised CPU time', 'Sum Normalised Elapsed time', 'Computation Monetary Cost', 'Estimated Monetary Cost', and 'CPU Efficiency'. The 'Period' is set to 5, 'End year' to 2016, and 'End month' to 4. The 'Groupings' are set to 'as a function of: VO'. The 'Chart' is set to 'Sum Normalised Elapsed time * number Processors'. The 'dteam VO' is set to 'Exclude dteam and ops VOs jobs information'. The 'Local Jobs' are set to 'Grid Jobs Only'. A 'Refresh' button is at the bottom. Below the configuration window, a table of VOs is displayed with columns for VO name and checkboxes for selection. The table includes VOs like 50051, 50052, 5300, 5800, alic, alicsgm, ams, apesci, asci, astro.vo.eu-egee.org, astron, atcan, atlas-prd, atlas-user, atlaspil, atlaspl, atpr, atpilot, auger013, auger015, auger024, babar, bbmri.nl, beapps, belle, bfactory, bio, biogrid, biomed, biowur, bmed, bwgrid, calice, cdf, cedar, chem.biggrid.nl, cmpd, cms, cmsgm, cmsprd, cmssgm, compass, compassit, d0, and d4sci.

why Sum Normalised Elapsed time * number Processors is available only in that view?

CPU Efficiency (CPU/WALL*processors) in Tier1/2 view

It seems OK in Tier1 view

<http://accounting.egi.eu/tier1.php?query=cpueff&startYear=2015&startMonth=5&endYear=2016&endMonth=4&yrange=TIER1&xrange=VO&groupVO=lhc&chart=GRBAR&scale=LIN&localJobs=onlygridjobs>

VIEW

VO MANAGER VIEW

VO MEMBER VIEW

SITE ADMIN VIEW

USER VIEW

REPORTS

METRICS PORTAL

LINK

Local Jobs:

☒ Grid Jobs Only
 ☐ Grid Jobs and Local Jobs
 ☐ Local Jobs Only

Refresh

TIER1 CPU Efficiency by TIER1 and VO.

LHC VOs. May 2015 - April 2016.

Warning:

For metrics including number of processors we count records with 0 processors as having 1 processor in order to improve reporting.

The following table shows the distribution of CPU Efficiency grouped by TIER1 and VO (only information about LHC VOs is returned).

CPU Efficiency (%) by TIER1 and VO

TIER1	alice	atlas	cms	lhcb	Total
CA-TRIUMF		86.1			86.1
CH-CERN	68.2	83.6	71.9	91.6	78.8
DE-KIT	89.7	92.7	79.3	99.1	90.2
ES-PIC		83.6	62.2	98.4	81.4
FR-CCIN2P3	80.7	87.1	62.8	96.1	81.7
IT-INFN-CNAF	81.4	83.6	70.1	94.1	82.3
KR-KISTI-GSDC	77.2				77.2
NDGF	75.5	75.9			75.7
NL-T1	87.6	91.6		97.7	92.3
NRC-KI-T1	83.5	86.7		97.2	89.1
RU-JINR-T1			60.7		60.7
TW-ASGC		88.1			88.1
UK-T1-RAL	81.7	85.2	62.8	94.5	81.1
US-FNAL-CMS		81.5	67.8		74.7
US-T1-BNL		76.7	0.5		38.6
Total	80.6	84.7	58.3	96.3	80.0

Click here for a CSV dump of this table

Click here for an Extended CSV dump of this table

Click here for XML encoded data

Key:

0% <= eff < 50%; 50% <= eff < 60%; 60% <= eff < 75%; 75% <= eff < 90%; 90% <= eff < 100%; eff >= 100% (parallel jobs)

The information in the previous table is also shown in the following graph.

The information in the previous table is also shown in the following graph.

CPU/WALL >>100% in country view

But it is NOT ok in Country view.

<http://accounting.egi.eu/country.php?query=cpueff&startYear=2015&startMonth=5&endYear=2016&endMonth=4&yRange=COUNTRY&xRange=VO&voGroup=lhc&chart=GRBAR&scale=LIN&localJobs=onlygridjobs>

accounting.egi.eu/country.php?query=cpueff&startYear=2015&startMonth=5&endYear=2016&endMonth=4

EGI ACCOUNTING PORTAL

GLOBAL View VO MANAGER View VO MEMBER View SITE ADMIN View USER View REPORTS METRICS PORTAL LINK

EGI View --> Country

Tier1
Tier2
Countries
EMI3(WLCG)
EGI
OSG
UNREGISTERED
VO_Discipline
VO_Metrics
CUSTOM_view
Cloud
Cloud Tier1
Cloud Tier2

Data to graph: CPU Efficiency overall efficiency = SUM (CPU times) / SUM (wall times)

Period: Start year: 2015 Start month: 5 End year: 2016 End month: 4

Groupings: Show data for: COUNTRY as a function of: VO

VO Groups: LHC TOP 10 ALL Custom

Chart: Type: GROUP BAR Scale: LINEAR

dteam VO: Exclude dteam and ops VOs jobs information

Local Jobs: Grid Jobs Only Grid Jobs and Local Jobs Local Jobs Only

Refresh

COUNTRY CPU Efficiency by COUNTRY and VO.
LHC VOs. May 2015 - April 2016.



Developed by CES-A

The following table shows the distribution of CPU Efficiency grouped by COUNTRY and VO (only information about LHC VOs is returned).

COUNTRY	CPU Efficiency (%) by COUNTRY and VO					Total
	alice	atlas	cms	lhcb		
Armenia	81.6	63.2				69.3
Australia		123.4				123.4
Austria		82.6	77.8			78.1
Belgium			76.7			76.7
Brazil	84.0			97.3		85.8
Bulgaria			27.8	95.8		86.1
Canada		161.5				161.5
Chile		138.0				138.0
China		107.0	79.5			94.7
Croatia			16.4			16.4
Czech Republic	78.9	136.7				102.3
Estonia			59.5			59.5
Finland	40.0		69.2			69.2
France	83.7	175.0	107.0	96.4		115.6

And not even in the devel portal.

<http://accounting-devel.egi.eu/country.php?query=cpueff&startYear=2015&startMonth=5&endYear=2016&endMonth=4&yrange=COUNTRY&xrange=VO&groupVO=lhc&chart=GRBAR&scale=LIN&localJobs=onlygridjobs>

Hierarchical Tree EGI View --> Country

- Tier1
- Tier2
- Countries
- EMI3
- EGI
- OSG
- UNREGISTERED
- VO_Discipline
- VO_Metrics
- CUSTOM_view
- Cloud
- Cloud Tier1
- Cloud Tier2



Developed by CES-A

Data to graph: CPU Efficiency overall efficiency = SUM (CPU times) / SUM (wall times)

Period: Start year: 2015 Start month: 5 End year: 2016 End month: 4

Groupings: Show data for: COUNTRY as a function of: VO

VO Groups: LHC TOP 10 ALL Custom

Chart: Type: GROUP BAR Scale: LINEAR

dteam VO: Exclude dteam and ops VOs jobs information

Local Jobs: Grid Jobs Only Grid Jobs and Local Jobs Local Jobs Only

Refresh

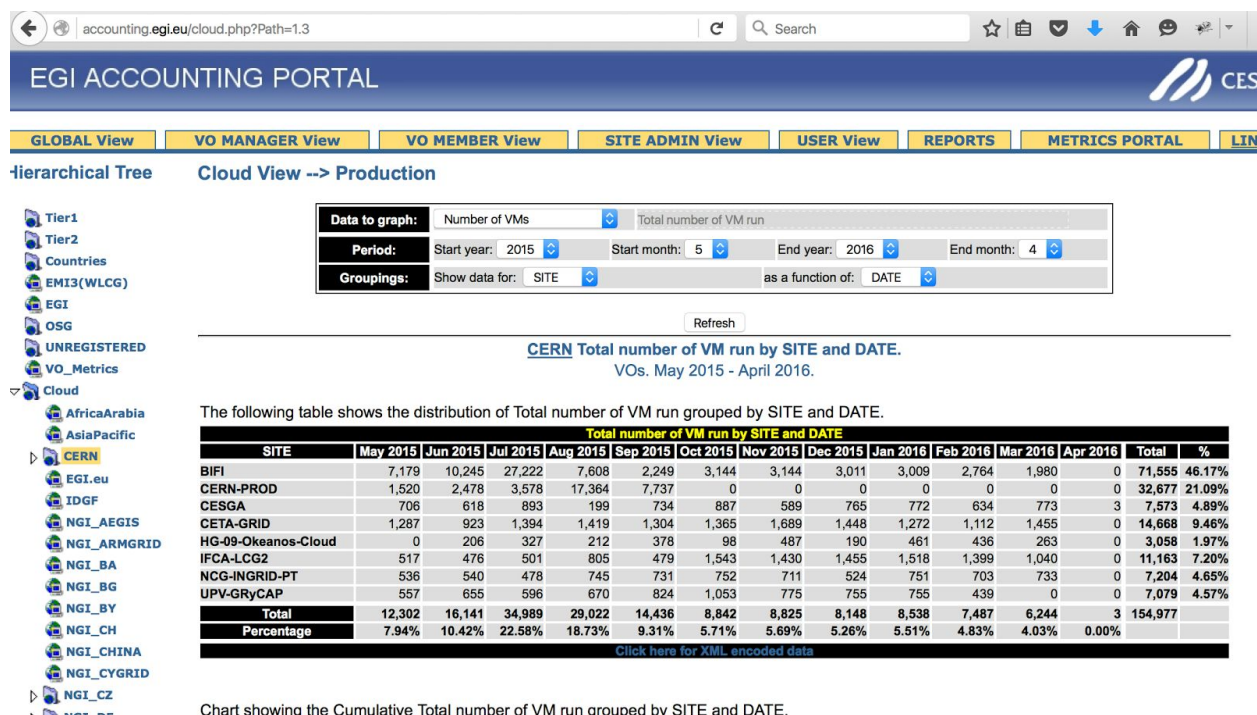
COUNTRY CPU Efficiency (%) by COUNTRY and VO.
LHC VOs. May 2015 - April 2016.

The following table shows the distribution of CPU Efficiency (%) grouped by COUNTRY and VO (only information about LHC VOs is returned).

CPU Efficiency (%) by COUNTRY and VO					
COUNTRY	alice	atlas	cms	lhcb	Total
Armenia	60.0	100.0			80.0
Australia					0.0
Austria		11.1			11.1
Belgium					0.0
Brazil				11.1	11.1
Bulgaria			266.7	14.3	90.0
Canada		166.7			166.7
Chile		20.0			20.0
China		50.0	22.2		35.3
Croatia			120.0		120.0
Czech Republic	100.0	33.3			50.0

Cloud view: what is this?

<http://accounting.egi.eu/cloud.php?Path=1.3>



Computation Monetary cost: what is this?

http://accounting.egi.eu/tier2.php?query=cost&startYear=2015&startMonth=5&endYear=2016&endMonth=4&yrange=COUNTRY_T2&xrange=VO&groupVO=lhc&chart=GRBAR&scale=LIN&localJobs=onlygridjobs

accounting.egi.eu/tier2.php?query=cost&startYear=2015&startMonth=5&endYear=2016&endMonth=
Search

EGI ACCOUNTING PORTAL
CESC

GLOBAL View
VO MANAGER View
VO MEMBER View
SITE ADMIN View
USER View
REPORTS
METRICS PORTAL
LINK

☐ vo.photonics.es-ngi.eu
☐ vo.sixt.cern.ch
☐ vo.ucad.sn
☐ vo_pheno
☐ xfelusr
☐ zh

☐ vo.plgrid.pl
☐ vo.slgrid.es-ngi.eu
☐ voce
☐ vo_scotgrid
☐ xray.vo.eu-egee.org

Chart:
Type:
GROUP BAR

dteam VO:
☐ Exclude dteam and ops VOs jobs information

Local Jobs:
☒ Grid Jobs Only

Refresh

TIER2 by COUNTRY_T2 and VO.

LHC VOs. May 2015 - April 2016.

Warning: For metrics including number of processors we count records with 0 processors as having 1 processor in order to improve reporting.

The following table shows the distribution of grouped by COUNTRY_T2 and VO (only information about LHC VOs is returned).

COUNTRY_T2	by COUNTRY_T2 and VO					Total	%
	alice	atlas	cms	lhcb			
Italy	35,548,308€	0	20,414,754€	4,346,865€		60,309,927€	76.48%
Poland	2,468,476€	5,207,500€	0	1,242,220€		8,918,196€	11.31%
Spain	0	0	2,117,257€	0		2,117,257€	2.68%
Switzerland	0	7,516,542€	0	0		7,516,542€	9.53%
Turkey	0	0	0	0		0€	0.00%
Total	38,016,784€	12,724,043€	22,532,011€	5,589,084€		78,861,923€	
Percentage	48.21%	16.13%	28.57%	7.09%			

[Click here for a CSV dump of this table](#)
[Click here for a Extended CSV dump of this table](#)
[Click here for XML encoded data](#)

Chart showing the Cumulative grouped by COUNTRY_T2 and VO (only information about LHC VOs is returned).

One other small point:

Why not having Elapsed time * Num Processor?

why not having also elapsed time * number of processors?

—> this would be actually the easiest more basic info that also experiment can crosscheck. The experiments don't know the normalisation applied by the site but they know how much time *processors they got from the site.

In the link here one example, TRIUMF.

http://accounting.egi.eu/tier1.php?Path=1.2&query=sum_normelap_nprocessors&startYear=2015&startMonth=5&endYear=2016&endMonth=4&yrange=SITE&xrange=DATE&groupVO=lhc&chart=GRBAR&scale=LIN&localJobs=onlygridjobs

...

...

In general on the portal:

Would be nice if from an accounting portal we have all and only the useful and validated information.

All the rest should be hidden/removed.

We also need to agree on which are the info that we really need.

General view on accounting for memory/cores

As agreed at the WLCG management board few months ago,

ATLAS is requesting the sites to provide 2GB physical memory per core.

ATLAS is also kindly requesting to some sites (the ones that can) to provide the possibility to run a fraction of the ATLAS jobs on high memory (≥ 4 GB memory per core) resources.

ATLAS is going to provide the memory requirements for all the jobs to allow the site to maximize the resources usage.

ATLAS suggest that we keep accounting as simple as possible.

Minor note

If the goodwilling sites providing HighMemory slots jobs need to improve their accounting for these slots and would need to advertise the cores “not fully used” we would need to know how they do this to make sure that our workload management system broker jobs accordingly.

Opportunistic resources:

•Please, provide feedback on how you would like opportunistic resources accounting to be performed and whether you consider this could be done through central accounting portal

It depends on the opportunistic resources.

we believe we cannot ask places like HPC/CloudProviders to publish into the accounting portal.

What we do have is our JobDashboard, and there for sure we want to have the overview.

If we want to use the accounting portal as the “WLCG overview”, then maybe we should ask at least the possibility to inject data a posteriori, in this case we could (experiments or opportunistic resource provider) inject “if needed”.

Space accounting

•Please, provide some details on how you currently perform space accounting

•Whether you have any needs from WLCG for this

•How you envisage to do this in a SRM-less world

ATLAS do direct checks each approx20mins, using gfal clients, on the space tokens the sites allocate to ATLAS.

ATLAS do have a plan for the SRMless world: DDM is proposing a structure (extension of what we had few years ago for gridftp only sites), and this could be then fetched and collected as today we collect the info from each storage. This is also discussed with CMS

REBUS

•Please, explain what REBUS functionality is used and whether any of the published data there is useful, otherwise, please describe which information or functionality is missing

•Please, state whether information about installed/available capacity is needed

•Available capacity (HW which is actually usable) as opposed to installed capacity (HW which is in place)

•Note that capacity definitions are being discussed!

ATLAS use REBUS as the place where the pledge are collected from the sites. we need it (or similar).

ATLAS use installed/available capacity to have a guess of the core power.

Would be useful to know how much a site can provide to each VO, e.g. a working installed capacity by site. This would most probably need to be extended because as of today this installed capacity per federation is used to match with pledges, while ATLAS would be interested to know how much it is really available, including over-the pledge and opportunistic resources.

Additional note: for accounting numbers for computing a +/- 10% can be considered ok.