

Table of Contents

Accounting Frequently Asked Questions.....	1
Naming conventions.....	1
Is there a mapping between naming schemes?.....	1
APEL New portal.....	1
Where is the WLCG view?.....	1
How can I compare my batch system accounting with the APEL portal?.....	1
SSB comparison.....	1
What are the SSB accounting comparison links?.....	1
How often is the SSB comparison dashboard updated?.....	2
How can I get the numbers?.....	2
Benchmarking publishing.....	2
Are there guidelines on how to publish the benchmarking publishing?.....	2
ALICE specific.....	2
ATLAS specific.....	2
Where are the problems with ATLAS sites followed?.....	2
How do the ATLAS wall clock time and the APEL wall clock time differ?.....	2
How are the ATLAS numbers in SSB calculated?.....	3
I don't know where to look is there a list of possible problems already?.....	3
Is there a way to see if the changes have had any effect?.....	4
What is the ATLAS benchmarking status?.....	4
CMS specific.....	5
LHCb specific.....	5

Accounting Frequently Asked Questions

Naming conventions

Is there a mapping between naming schemes?

New APEL Names	Old APEL Names	Rebus	ATLAS	Units
processors	processors	logical CPUs	cores	processors
Wall clock time	Sum elapsed * number of processors	N/A	Wallclock consumption all jobs	seconds (or hours)
Wall clock work	Normalised Sum elapsed * number of processors	N/A	N/A	HS06 seconds (or hours)
CPU time	Sum CPU	N/A	CPU consumption All Jobs	seconds (or hours)
CPU work	Normalised Sum CPU	N/A	N/A	HS06 seconds (or hours)
Power	SpecInt (HS06*250)	N/A	Coefficient	HS06
Delivered power	N/A	N/A	WallClock HEPSP06	Total Used HS06
Pledged power	N/A	Federation Pledges (HS06)	Federation Pledges (HS06)	Total Pledged HS06
Installed power	N/A	Installed capacity (HS06)	HEPSP06	Total Installed HS06

APEL New portal

Where is the WLCG view?

- The new portal can be found here <https://accounting-next.egi.eu/> the WLCG views are under "Research Infrastructure" tab at the top. However to makes things easier WLCG has now its own view that can be accessed here <https://accounting-next.egi.eu/wlcg>

How can I compare my batch system accounting with the APEL portal?

- There aren't scripts centrally produced, some sites have written their own scripts. Liverpool has kindly made theirs available online for other sites.
 - ◆ Here are their scripts for CREAM/Torque, ARC/HTCondor, VAC.

SSB comparison

What are the SSB accounting comparison links?

- This is the snapshot of the latest results for all the experiments . If you click on each column title you'll get to the history page for that column. It will point to the last 24h which are empty. It is recommended you check the last 3 months. The most important histories pages for each experiment

are below.

- ◆ **ATLAS** wall clock work APEL/ATLAS ratio [↗](#)
- ◆ **CMS** wall clock work APEL/CMS ratio [↗](#)
- ◆ **LHCb** wall clock work APEL/LHCb ratio [↗](#)

How often is the SSB comparison dashboard updated?

- Every month few days after the end of the previous month to allow for all the APEL numbers to appear in the portal

How can I get the numbers?

- The SSB binning is 1 month. If you click on each month you get a json for that month. You can ignore most of its fields but the "Status". The Status is composed by 4 numbers in the following order:

Experiment work, APEL work, work ratio, Experiment wall time, APEL wall time, wall time ra

. The ratio is what the determines the colour on the dashboard and is APEL/Experiment. Work and wall time are displayed on different pages and the links reported above point to the work ratio pages. The walltime ratio can be a secondary check when the work ratio is off, but only at sites that don't scale internally in the batch system. ATLAS sites can get the numbers for ATLAS directly using the script linked in this FAQ.

Benchmarking publishing

Are there guidelines on how to publish the benchmarking publishing?

- HEPiX starting page for how to run an HS06 benchmark [↗](#)
- The GridPP wiki has guidelines on how to run the benchmark and how to publish the results in the BDII [↗](#) for standard grid sites
- Resources that do not get included in the BDII publishing are not accounted for in REBUS and are not seen by ATLAS. ATLAS exposes the values in the dashboard for sites selection for example [↗](#).

ALICE specific

ATLAS specific

Where are the problems with ATLAS sites followed?

- The known ATLAS sites problems are followed in this ticket <https://its.cern.ch/jira/browse/ADCINFR-28> [↗](#). If other tickets in JIRA or in GGUS get opened, they'll be linked to this ticket.

How do the ATLAS wall clock time and the APEL wall clock time differ?

- **ATLAS wall clock time** is recorded in panda as the length of the payloads (end time-start time). It is raw wall clock time.
- **APEL wall clock time** is the pilot walltime recorded by the batch system in the log files which may or may not be scaled
 - ◆ If sites don't scale the only difference is the discrepancy between the payloads lifetimes and the pilot one. For single payload pilots the discrepancy due to this is overall negligible with APEL slightly bigger numbers.

What are the SSB accounting comparison links?

- ◆ If sites scale the discrepancy between ATLAS raw wall clock time and APEL wall clock time can be significant. Usually when the work is calculated things go back in line, and if they don't it is likely that there is a mistake along the way. See list of possible problems if a site forgets to set the scaling factor on a large portion of the nodes.

How are the ATLAS numbers in SSB calculated?

- **Raw wall clock time:** the raw walltime in seconds in panda for each job. This corresponds to the lifetime of the payload (end_time-start_time). raw wallclock time is reported in the dashboard for example this is the json with the wallclock for all the jobs that run at UKI-NORTHGRID-MAN-HEP [↗](#) site for each month of Q3. The quantity to look at is SUM of course.
- **Wall clock work:** the wallclock time is then multiplied for the average power (HS06) of the site. The average power is calculated by dividing the total capacity by the number of logical CPUs using the numbers in REBUS [↗](#). REBUS get these numbers from the BDII for each site.

`wallclock work = raw wall clock time*average power = raw wall clock time*(total capacity/1`

- **Delivered power:** delivered power is the wall clock work divided by the number of seconds in the selected time interval (for example for august $31*24*60*60$). Derived power is what is reported as Wallclock HEPSPEC in the ATLAS dashboard. An example of json always using UKI-NORTHGRID-MAN-HEP is here [↗](#). The quantity to look at is SUM_HEPSPECWC.

`derived power = wall clock work/time interval seconds`

- **ATLAS wall clock work in the SSB:** ATLAS wall clock work in the SSB is obtained multiplying the derived power for the number of hours in the period selected. In that case the month of August $31*24$

`SSB ATLAS wall clock work = derived power*time interval hours`

- **ATLAS wall clock time in the SSB:** ATLAS wall clock time in the SSB is obtain converting the raw wall clock time in the dashboard in hours, i.e. dividing by 3600.

`SSB ATLAS wall clock time = raw wall clock time/3600`

I don't know where to look is there a list of possible problems already?

The numbers in the SSB are extracted from the ATLAS dashboard and the APEL accounting portal. These two rely on a host of other services to put the numbers together

- **ATLAS:** panda, dashboard summaries, REBUS, BDII
- **APEL:** batch systems, CEs, BDII, GOCDB, parsers or SSMsend (else?), APEL, summaries for portal

Problems encountered so far varied here is a list and how they affected the site. The effect of some of these problems are evident and lasting other times they cause a variation of colors on the SSB month after month because they affect only part of the resources.

- ARC/HTcondor parser: when flocking is active the current version doesn't see all the jobs anymore under reporting in APEL
- Wrong internal scaling on a big portion of resources over reporting in APEL
- Wrong HS06 in the BDII wrong reporting in APEL and ATLAS
- Resources not reported in the BDII work done disappears from ATLAS
 - ◆ For example sites with cloud resources or non traditional sites using VAC but also normal batch system resources not added for one reason or another.
 - ◇ VAC and traditional batch system are affected if they don't adjust the BDII to include them

◇ VAC only sites have no BDII work is not reported in ATLAS similar to the next problem

- BDII seemingly correctly publishing but REBUS doesn't see the site ATLAS doesn't record any work though the wall clock time is there
- Wrong DN/missing service in GOCDB missing resources in APEL
- Site capacity misreported in REBUS ATLAS numbers are affected
 - ◆ Typically sites on university clusters with highly variable usage
- Wrong benchmark: official benchmark is still HS06 **32bit not 64bit** overreporting both in ATLAS and APEL
 - ◆ The comparison ATLAS - APEL may be green but the accounting is still wrong
- Event Service in ATLAS (still under investigation) ATLAS under reporting
- APEL clients stops publishing under reporting in APEL
 - ◆ Most obvious of all problems but if you have multiple sources and one stops you may not notice.
- Resources declared separately in AGIS. If you setup a different site the resources will not be accounted for this will result in ATLAS under reporting.

Is there a way to see if the changes have had any effect?

- There is a python script to extract the same information from the ATLAS dashboard and the EGI portal as the SSB here. You will have to remove the .txt extension the twiki decided to add. It assumes the name of the site is the same in both portals which is true for most, but not all sites.

```
python get-atlas-accounting-data.py -s <SITE> -m 6
```

will return the last 6 months and the current one for site .

```
python get-atlas-accounting-data.py -s <SITE> -m 6 -b10
```

will return only the entries that have either the work or the wall clock time discrepancies bigger than 10%. If you apply this only to the last complete month you can create alerts in your system.

Script takes only few seconds to run here is an example of the output. Current month (in this case 2017-01) is always a bit off because of incomplete records and can be removed with the option -n

```
python get-atlas-accounting-data.py -s UKI-NORTHGRID-MAN-HEP -m3
Date,ATLAS work,EGI work, (wE-wA)*100/wA,ATLAS wc,EGI wc, (wcE-wcA)*100/wcA
2016-10,18245809,19316185,5.9%,1851162,1947290,5.2%
2016-11,22827795,22321641,-2.2%,2316036,2261186,-2.4%
2016-12,24884445,23447118,-5.8%,2524697,2369523,-6.1%
2017-01,14263609,11960116,-16.1%,1447140,1225311,-15.3%
```

The order of the columns is the same as the one in the SSB Status field explained here though instead of reporting the ratio the script reports the percentage of the discrepancy. For shortness wE=workEGI wA=workATLAS wcE=wallclockEGI wcA=wallclockATLAS. Negative numbers indicate ATLAS has recorded higher for that month.

What is the ATLAS benchmarking status?

- Benchmarking status is summarised in this twiki
<https://twiki.cern.ch/twiki/bin/view/AtlasComputing/ADCResourcesBenchmarking>

CMS specific

LHCb specific

This topic: LCG > AccountingFAQ

Topic revision: r19 - 2018-03-12 - AlessandraForti



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback