



WLCG Accounting Task Force Update

Julia Andreeva CERN

GDB, 8th of June , 2016

Proposal to setup an Accounting Task Force

- Main area of activity is CPU/Wall Clock accounting. The highest priority task.
- Storage Space accounting should be also addressed in particular of what concerns understanding of the requirements of the experiments and sites.
- Stakeholders whose needs to be considered:
 - LHC experiments
 - WLCG sites
 - WLCG management
 - Funding agencies

Mandate

CPU/Wall clock accounting

- to agree with all stakeholders on the content of the accounting portal and accounting reports (metrics, units, aggregation in terms of time and set of attributes like site, VO, country, etc...);
- coordinate with APEL and accounting portal development team and follow up on the progress
- validate the new version of the portal, ideally a dedicated customized UI for WLCG, and make sure that the requirements of the stakeholders are addressed
- validate accounting data, ensure that it is correct in the APEL repository and manual fixes could be avoided
- follow up of on whether/ how we enable injection of the opportunistic resources to APEL and accounting portal
- coordinate with sites which provide their resources through cloud interfaces making sure that their resources are properly published in APEL. Debug corner cases when they are not properly published

Storage space accounting

- agree with the experiments and sites on the functionality which is required for the WLCG storage space accounting system considering the use case of the high level overview of the total and available space

Membership

Representatives of the

- LHC experiments
- WLCG sites
- WLCG operations
- APEL and EGI accounting portal development teams

Time line (1)

June

- **Fixing accounting reports**
 - *Exclude installed (available) capacities from the accounting reports. They should still be recorded in REBUS and visualized by the REBUS UI*
 - *Consolidate the CPU accounting data in the reports. Use normalized wall clock metric instead of normalized CPU for comparison with pledges. Wall clock should be expressed in HS06-days both for T1 and T2 reports.*
- **Check with the sites which provide their resources through private clouds that their usage is properly published in APEL**

June-July

- **Validate data in the EGI accounting portal by comparing it with the experiment-specific accounting systems. Debug sites which show big discrepancy**
 - *As a first step we need to understand whether we have consistent data for non-normalized wall clock multiplied by number of cores (June)*
 - *If at this level data is consistent, the problem might be caused by the normalization. Proceed with comparing data using normalized wall clock and CPU. Understand which normalization mechanisms are used by the experiments and sites, what can be improved, etc... (July)*
 - *If the problem is proven to be caused by normalization, follow up with the benchmarking working group (August)*

Time line (2)

Beginning-mid of June

- Confirm with the sites, experiments and WLCG management the set of metrics and views in the accounting reports

Mid of June – end of July

- Evaluate the new version of the EGI accounting portal
- Coordinate with the developers of the portal to make sure that all requirements are addressed and possible bugs are fixed

August

- Final validation and coordination with stakeholders to make sure that their concerns are addressed and the new version of the portal satisfies their needs

Time line (3)

June-July

- In parallel with the CPU/Wall clock accounting tasks get an agreement with the experiments and sites on the functionality which is required for the WLCG storage space accounting (the high level overview of the total and available space)

August

- Propose next steps for implementation of the WLCG storage space accounting system

Meetings and organization

- First meeting will take place on the 9th of June
<https://indico.cern.ch/event/539179/>
- Mailing list : wlcg-accounting@cern.ch
- Twiki page :
<https://twiki.cern.ch/twiki/bin/view/LCG/AccountingTaskForce>

If you would like to take part in the WLCG Accounting Task Force,
you are welcome!

Please contact Julia (Julia.Andreeva@cern.ch)

First results of data validation

- First metric to be compared is non-normalized wall clock multiplied by number of cores
- Monthly per-site metric is retrieved through the API from the new accounting portal and from ATLAS and CMS Dashboards.
Credits to Elena Tikhonenko from JINR who performed comparison.
- For ALICE and LHCb we currently do not have a convenient way to get the same data, but we made a request to the VO developers
- We plan to enable automatic checks with recording of the results of comparison in the WLCG instance of the Site Status Board

ATLAS results (1)

- Pretty good agreement, for many sites the difference is on the level of couple of percents (pilot overhead).
- However some sites show big discrepancy (up to 3 times), among them one T1 site.
- In case these sites are also used by CMS, in most cases we see that CMS results confirm the same problem. So the problem might be in data reporting to APEL.

ATLAS results (2)

January-April 2016

| Jan-16 | | | Feb-16 | | | Mar-16 | | | Apr-16 | | | Jan-Apr-2016 | | |
|----------|----------|----------------|----------|----------|----------------|----------|----------|----------------|----------|----------|----------------|--------------|----------|----------------|
| EGI | Dashb | EGI/Dashb in % | EGI | Dashb | EGI/Dashb in % |
| 3.49E+06 | 3.44E+06 | 101.45% | 3.22E+06 | 3.07E+06 | 104.89% | 3.23E+06 | 3.20E+06 | 100.94% | 3.41E+06 | 3.37E+06 | 101.19% | 1.33E+07 | 13087360 | 101.90% |
| 2.64E+06 | 5.18E+06 | 50.97% | 3.18E+06 | 5.17E+06 | 61.51% | 7.42E+06 | 5.45E+06 | 136.15% | 5.21E+06 | 5.39E+06 | 96.66% | 1.85E+07 | 21185200 | 87.12% |
| 3.70E+06 | 3.68E+06 | 100.54% | 3.90E+06 | 3.72E+06 | 104.84% | 2.80E+06 | 2.96E+06 | 94.59% | 3.99E+06 | 3.97E+06 | 100.50% | 1.44E+07 | 14336100 | 100.37% |
| 1.60E+06 | 1.56E+06 | 102.56% | 1.73E+06 | 1.58E+06 | 109.49% | 1.41E+06 | 1.27E+06 | 111.02% | 1.72E+06 | 1.54E+06 | 111.69% | 6.46E+06 | 5946550 | 108.67% |
| 4.73E+06 | 4.72E+06 | 100.21% | 4.91E+06 | 4.70E+06 | 104.47% | 4.02E+06 | 4.02E+06 | 100% | 4.74E+06 | 4.61E+06 | 102.82% | 1.84E+07 | 18054200 | 101.96% |
| 2.87E+06 | 3.03E+06 | 94.72% | 3.02E+06 | 2.78E+06 | 108.63% | 2.10E+06 | 1.94E+06 | 108.25% | 3.64E+06 | 3.42E+06 | 106.43% | 1.16E+07 | 11169070 | 104.12% |
| 2.24E+06 | 2.10E+06 | 106.67% | 1.68E+06 | 1.56E+06 | 107.69% | 1.58E+06 | 1.52E+06 | 103.95% | 1.63E+06 | 1.65E+06 | 98.79% | 7.13E+06 | 6838990 | 104.22% |
| 3474940 | 3.46E+06 | 100.43% | 2704690 | 2.49E+06 | 108.62% | 2222910 | 2.23E+06 | 99.68% | 3163050 | 3.16E+06 | 100.10% | 11565600 | 11342860 | 101.96% |
| 2.65E+06 | 2.76E+06 | 96.01% | 1.88E+06 | 1.93E+06 | 97.41% | 1.52E+06 | 1.60E+06 | 95% | 2.07E+06 | 2.14E+06 | 96.73% | 8.13E+06 | 8428150 | 96.41% |
| 1.68E+06 | 1.65E+06 | 101.82% | 1.32E+06 | 1.25E+06 | 105.60% | 955222 | 864014 | 110.56% | 1.75E+06 | 1.93E+06 | 90.67% | 5.71E+06 | 5691044 | 100.31% |
| 1.76E+07 | 6.50E+06 | 270.77% | 1.46E+07 | 5.39E+06 | 270.87% | 1.14E+07 | 4.44E+06 | 256.76% | 1.11E+07 | 4.35E+06 | 255.17% | 5.48E+07 | 20675690 | 264.94% |
| 1.07E+07 | 1.38E+07 | 77.54% | 9.76E+06 | 1.17E+07 | 83.42% | 1.05E+07 | 1.25E+07 | 80% | 1.07E+07 | 1.03E+07 | 103.88% | 4.12E+07 | 48304200 | 85.31% |
| 4.52E+06 | 4.48E+06 | 100.82% | 4.03E+06 | 3.85E+06 | 104.76% | 3.19E+06 | 3.10E+06 | 103.06% | 4.51E+06 | 4.49E+06 | 100.52% | 1.62E+07 | 15911810 | 102.12% |
| 613773 | 573116 | 107.09% | 626784 | 570024 | 109.96% | 596100 | 550496 | 108.28% | 809416 | 732032 | 110.57% | 2.65E+06 | 2425668 | 109.09% |
| 563066 | 574430 | 98.02% | 415652 | 517648 | 80.30% | 260733 | 505920 | 51.54% | 319279 | 863356 | 36.98% | 1.56E+06 | 2461354 | 63.33% |
| 2.03E+06 | 2.03E+06 | 100.41% | 2.70E+06 | 2.58E+06 | 104.95% | 2.38E+06 | 2.32E+06 | 102.32% | 2.88E+06 | 2.86E+06 | 100.66% | 9.99E+06 | 9785280 | 102.13% |
| 305983 | 382266 | 80.04% | 270157 | 383559 | 70.43% | 245607 | 331330 | 74.13% | 296978 | 497745 | 59.66% | 1.12E+06 | 1594900 | 70.14% |
| 707858 | 707094 | 100.11% | 700198 | 649043 | 107.88% | 502706 | 424438 | 118.44% | 732506 | 731568 | 100.13% | 2.64E+06 | 2512143 | 105.22% |
| 804308 | 786696 | 102.24% | 652141 | 596968 | 109.24% | 750538 | 718175 | 104.51% | 773634 | 758537 | 101.99% | 2.98E+06 | 2860376 | 104.20% |
| 939884 | 1.05E+06 | 89.63% | 653104 | 795391 | 82.11% | 700964 | 757279 | 92.56% | 294550 | 1.41E+06 | 20.87% | 2.59E+06 | 4012730 | 64.51% |
| 338703 | 895507 | 37.82% | 540391 | 1.19E+06 | 45.45% | 421960 | 1.17E+06 | 36.20% | 421030 | 1.13E+06 | 37.16% | 1.72E+06 | 4383207 | 39.29% |
| 6.36E+06 | 3.15E+06 | 202.19% | 6.29E+06 | 2.95E+06 | 213.47% | 6.62E+06 | 3.29E+06 | 201.46% | 7.87E+06 | 3.92E+06 | 200.59% | 2.71E+07 | 13300330 | 204.04% |
| 2.66E+06 | 881537 | 301.36% | 2.57E+06 | 825902 | 311.10% | 2.68E+06 | 898156 | 298.39% | 3.29E+06 | 1.09E+06 | 302.63% | 1.12E+07 | 3692655 | 303.19% |
| 247684 | 247687 | 100.00% | 199832 | 190066 | 105.14% | 221321 | 218114 | 101.47% | 243130 | 239911 | 101.34% | 911967 | 895778 | 101.81% |
| 305073 | 304629 | 100.15% | 268445 | 256801 | 104.53% | 243521 | 240119 | 101.42% | 270371 | 260999 | 103.59% | 1.09E+06 | 1062548 | 102.34% |
| 971085 | 985118 | 98.58% | 886102 | 854108 | 103.75% | 964232 | 685209 | 101.32% | 690762 | 678765 | 101.77% | 3.24E+06 | 3203200 | 101.22% |
| 3.51E+06 | 3293044 | 106.57% | 3.07E+06 | 2897267 | 106.09% | 2.65E+06 | 2611636 | 101.61% | 3.53E+06 | 3673468 | 96.14% | 1.28E+07 | 12475415 | 102.35% |
| 21222.5 | 20210.3 | 105.01% | 22869.7 | 22355.7 | 102.30% | 7874.93 | 7194.53 | 109.46% | 3927.03 | 3248.43 | 120.89% | 55894.1 | 53009 | 105.44% |
| 614397 | 565194 | 108.71% | 643547 | 568139 | 113.27% | 785099 | 735924 | 106.68% | 775569 | 663425 | 116.90% | 2.82E+06 | 2532682 | 111.29% |
| 220477 | 218401 | 100.95% | 237941 | 211899 | 112.29% | 214854 | 211696 | 101.49% | 294728 | 285487 | 103.24% | 968000 | 927483 | 104.37% |
| 458650 | 437178 | 104.91% | 563244 | 496997 | 113.33% | 646159 | 631559 | 102.31% | 542682 | 509379 | 106.54% | 2.21E+06 | 2075113 | 106.54% |
| 1.65E+06 | 962426 | 171.39% | 1.78E+06 | 959503 | 185.33% | 1.59E+06 | 864894 | 183.35% | 1.85E+06 | 1.02E+06 | 181.44% | 6.86E+06 | 3805863 | 180.31% |
| 499445 | 492011 | 101.51% | 470668 | 434567 | 108.31% | 431874 | 425329 | 101.54% | 639399 | 631915 | 101.18% | 2.04E+06 | 1983822 | 102.90% |
| 925808 | 904463 | 102.36% | 877783 | 847712 | 103.55% | 688780 | 676519 | 101.81% | 541637 | 540783 | 100.16% | 3.03E+06 | 2969477 | 102.17% |
| 521982 | 796474 | 65.54% | 701804 | 671743 | 104.48% | 760592 | 751848 | 101.16% | 781362 | 700790 | 111.50% | 2.77E+06 | 2920855 | 94.69% |
| 702302 | 615594 | 114.09% | 406697 | 392551 | 103.60% | 480561 | 460473 | 104.36% | 678298 | 669562 | 101.31% | 2.27E+06 | 2131880 | 106.07% |
| 605690 | 586354 | 103.30% | 561454 | 535939 | 104.76% | 566152 | 555793 | 101.86% | 625397 | 622271 | 100.50% | 2.36E+06 | 2300357 | 102.54% |
| 1.10E+06 | 1.07E+06 | 102.37% | 950397 | 907300 | 104.75% | 1.06E+06 | 1.05E+06 | 101.49% | 1.10E+06 | 1.07E+06 | 102.83% | 4.21E+06 | 4099040 | 102.79% |
| 1.33E+06 | 1.31E+06 | 102.13% | 1.17E+06 | 1.12E+06 | 104.30% | 1.07E+06 | 1.06E+06 | 100.79% | 1.27E+06 | 1.23E+06 | 102.57% | 4.84E+06 | 4726700 | 102.46% |
| 2.07E+06 | 2.11E+06 | 98.36% | 2.10E+06 | 2.09E+06 | 100.33% | 2.16E+06 | 2.17E+06 | 99.15% | 2.31E+06 | 2.33E+06 | 99.02% | 8.63E+06 | 8701560 | 99.21% |
| 715402 | 708873 | 100.92% | 951330 | 908909 | 104.67% | 947539 | 941051 | 100.69% | 1.02E+06 | 954555 | 107.34% | 3.64E+06 | 3513388 | 103.57% |
| 6670.82 | 3419.58 | 195.08% | 6803.18 | 4342.09 | 156.68% | 7163.94 | 2909.4 | 246.23% | 10787.1 | 9310.91 | 115.85% | 31425 | 19982 | 157.27% |
| 690985 | 683466 | 101.10% | 437796 | 419243 | 104.43% | 369552 | 363909 | 101.55% | 407432 | 404526 | 100.72% | 1.91E+06 | 1871144 | 101.85% |
| 2.61E+06 | 2.78E+06 | 93.88% | 3.24E+06 | 3.17E+06 | 102.03% | 2.14E+06 | 2.14E+06 | 100.30% | 4.64E+06 | 4.50E+06 | 103.04% | 1.26E+07 | 12595970 | 100.30% |
| 195394 | 1.29E+06 | 15.19% | 139647 | 1.45E+06 | 9.60% | 96540.1 | 1.16E+06 | 8.36% | 121104 | 1.43E+06 | 8.44% | 552685 | 5330330 | 10.37% |
| 8.01E+06 | 7.62E+06 | 105.05% | 8.35E+06 | 7.63E+06 | 109.47% | 7.69E+06 | 7.19E+06 | 106.97% | 8.63E+06 | 8.15E+06 | 105.92% | 3.27E+07 | 30587310 | 106.84% |
| 351554 | 351564 | 100.00% | 198446 | 194471 | 102.04% | 312439 | 311738 | 100.23% | 352288 | 337898 | 104.29% | 1.21E+06 | 1195671 | 101.60% |
| 585575 | 580890 | 100.81% | 403007 | 378131 | 106.58% | 429455 | 380951 | 112.73% | 341835 | 342110 | 99.92% | 1.76E+06 | 1682082 | 104.63% |
| 2.91E+06 | 1.24E+06 | 235.53% | 2.86E+06 | 1.18E+06 | 242.56% | 3.00E+06 | 1.27E+06 | 236.80% | 3.31E+06 | 1.40E+06 | 236.14% | 1.21E+07 | 5087330 | 237.65% |

Up to 2-3 times difference

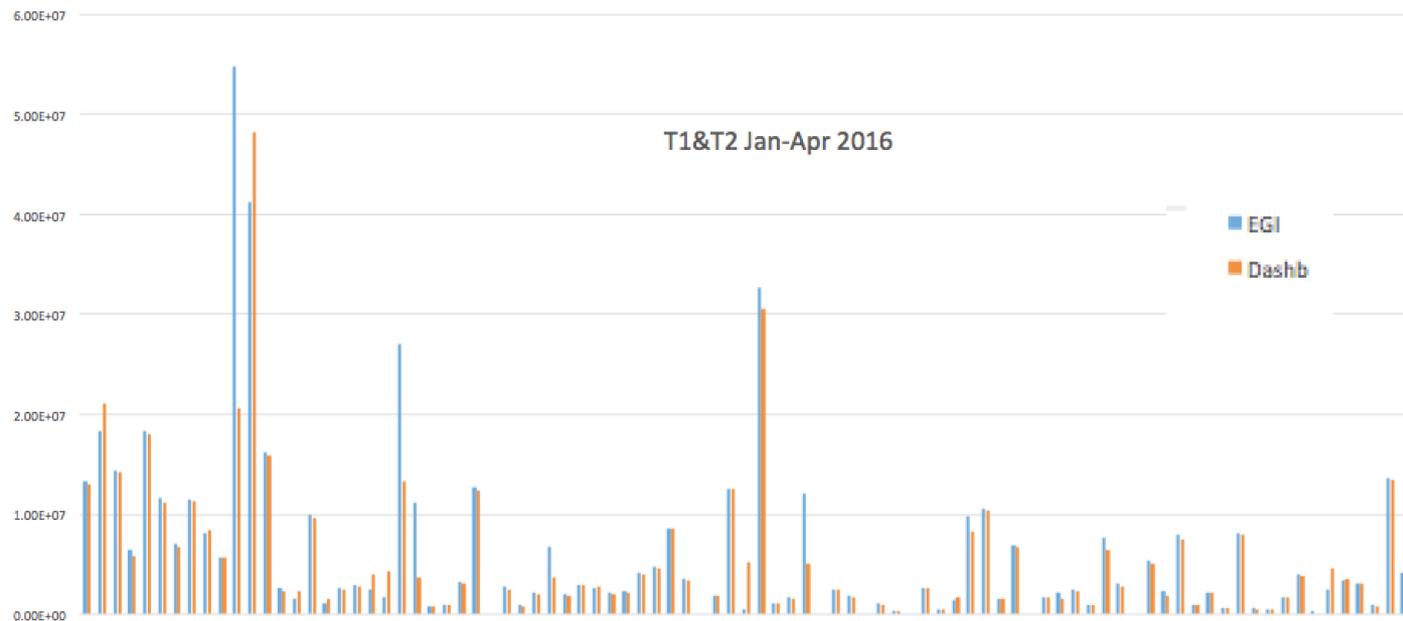
Dashboard shows higher statistics than EGI

The problem persists through all months



ATLAS results (3)

| | | | | | | | | | | | | | | |
|----------|----------|----------|----------|----------|---------|----------|----------|---------|----------|----------|---------|----------|----------|---------|
| 1.96E+06 | 1.95E+06 | 100.40% | 2.20E+06 | 2.03E+06 | 108.40% | 1.70E+06 | 1.67E+06 | 101.27% | 2.23E+06 | 1.95E+06 | 114.42% | 8.08E+06 | 7599020 | 106.32% |
| 273839 | 266727 | 102.67% | 243997 | 231979 | 105.18% | 245525 | 237423 | 103.41% | 313090 | 309639 | 101.12% | 1.08E+06 | 1045768 | 102.93% |
| 316185 | 300150 | 105.34% | 375446 | 345641 | 108.62% | 789445 | 755370 | 104.51% | 784908 | 782937 | 100.25% | 2.27E+06 | 2184098 | 103.75% |
| 325332 | 317239 | 102.55% | 59515.9 | 53605.5 | 111.03% | 86079.4 | 81272 | 105.92% | 314038 | 297221 | 105.66% | 784965 | 749338 | 104.75% |
| 1.66E+06 | 1.63E+06 | 101.88% | 2.10E+06 | 2.03E+06 | 103.19% | 1.99E+06 | 1.97E+06 | 101.26% | 2.45E+06 | 2.34E+06 | 104.67% | 8.20E+06 | 7968980 | 102.88% |
| 174632 | 156169 | 111.82% | 213300 | 187235 | 113.92% | 120219 | 101827 | 118.06% | 139374 | 123725 | 112.65% | 647525 | 568956 | 113.81% |
| 180658 | 175934 | 102.69% | 154926 | 145712 | 106.32% | 126499 | 124376 | 101.71% | 136620 | 133707 | 102.18% | 598703 | 579729 | 103.27% |
| 565075 | 563996 | 100.19% | 520959 | 513059 | 101.54% | 360730 | 352835 | 102.24% | 302093 | 302382 | 99.90% | 1.75E+06 | 1732272 | 100.96% |
| 1.47E+06 | 1.45E+06 | 101.48% | 1.32E+06 | 1.27E+06 | 103.28% | 770098 | 762784 | 100.96% | 484127 | 484736 | 99.87% | 4.04E+06 | 3973940 | 101.76% |
| 12567.8 | 746.929 | 1682.60% | 108692 | 90197.6 | 120.50% | 148077 | 146348 | 101.18% | 99303.7 | 97625.6 | 101.72% | 368640 | 334918 | 110.07% |
| 806380 | 997801 | 80.82% | 1.19E+06 | 1.27E+06 | 93.46% | 584770 | 1.24E+06 | 47.04% | 0 | 1.11E+06 | 0% | 2.58E+06 | 4619801 | 55.82% |
| 1.07E+06 | 1.10E+06 | 97.84% | 738011 | 723815 | 101.96% | 794974 | 833549 | 95.37% | 928161 | 925756 | 100.26% | 3.54E+06 | 3581050 | 98.72% |
| 720407 | 713566 | 100.96% | 657758 | 629087 | 104.56% | 829612 | 824192 | 100.66% | 978276 | 977617 | 100.07% | 3.19E+06 | 3144462 | 101.32% |
| 560685 | 549055 | 102.12% | 296057 | 281233 | 105.27% | 118495 | 114293 | 103.68% | 6208.76 | 2215.15 | 280.29% | 981446 | 946796 | 103.66% |
| 3.38E+06 | 3.41E+06 | 99.05% | 3.86E+06 | 3.57E+06 | 107.90% | 2.96E+06 | 2.96E+06 | 99.97% | 3.52E+06 | 3.54E+06 | 99.59% | 1.37E+07 | 13478010 | 101.74% |
| 1.16E+06 | 1.17E+06 | 99.39% | 992020 | 961043 | 103.22% | 975434 | 990936 | 98.44% | 1.10E+06 | 1.09E+06 | 101.44% | 4.23E+06 | 4210159 | 100.57% |
| 1.42E+08 | 1.29E+08 | 110.09% | 1.35E+08 | 1.19E+08 | 113.45% | 1.23E+08 | 1.10E+08 | 111.82% | 1.42E+08 | 1.28E+08 | 110.95% | 5.42E+08 | 4.86E+08 | 111.53% |



Overall, the result is not bad. Difference between two information sources January- April 2016 is on the level of 11.5% All problematic cases to be investigated

CMS results

- For CMS discrepancy is higher than for ATLAS, overall on the level of 30%. This can be partially explained by the fact that CMS Dashboard is not based on the information of the central server, but based on the individual job reports and therefore less accurate. If job died not being able to report, or CPU was not reported correctly, CMS Dashboard does not consider such jobs.
- In most cases (not always) if the sites are used both by ATLAS and CMS, if the site looks suspicious for ATLAS, we see a similar problem for CMS

Portal validation

- Just started. The new version of the portal is now available for validation of the pilot users.
- Looks promising
- All information required for data validation can be retrieved using an API
- Very good collaboration with Ivan Diaz Alvarez (EGI accounting portal developer)

Summary

- WLCG accounting task force has been setup
- Data validation and new portal validation work has already started
- We will provide regular updates at the GDB and WLCG operations coordination meetings