

Improving Published Capacity WLCG MB - September - 2009



Steve Traylen
CERN/IT

Tier 2 CPU Capacity

- A report is being produced every month:
 - <https://twiki.cern.ch/twiki/bin/view/LCG/SamMbReports>
 - e.g. Tier2_Reliab_200908.pdf
- All software for Installed Capacity is now:
 - Released
 - Info providers for lcg-CE and cream.
 - YAIM for DPM, lcg-CE and Cream.
 - Documented.
 - Detailed Docs - Installed Capacity Doc
 - YAIM documentation.
 - Instructions for running benchmarks
- All sites can now publish CPU installed capacity.



CPU data Major Omissions

- Major Omissions, i.e zero or null capacity:
- OSG sites all missing.
 - OSG have provided data via MyOSG:
 - <http://tinyurl.com/y9hfl5s>
 - Gridview team contacted to include in report.
 - Easy now report as switched to installed capacity.
- AU-ATLAS - Australia-UNIMELB-LCG2
 - Site closed for good, should drop out next month.
- IL-HEPTier-2 - IL-TAU-HEP
 - Completely unclear why this Null - gridview team contacted.
- UK-London - UKI-LT2-IC-LESC
 - Site is closed in GOCDB.



Random Tier 2 CPU Anomalies(1)

- HEPHY-Vienna 416*phy, 416*log
 - Probably wrong unless really 1 core processors.
 - gstat2 shows 10 errors, e.g:
 - <http://gstat-prod.cern.ch/gstat/site/Hephy-Vienna/ce/>
 - ERROR: hephygr.oeaw.ac.at:2119/jobmanager-lcgpbs-dteam, **GlueCEPolicyAssignedJobSlots has negative or null value**
 - ERROR: hephygr.oeaw.ac.at, **The Cores format is wrong, Cores not set**
 - ERROR: hephygr.oeaw.ac.at, **The Benchmark format is wrong, Benchmark not set**
 - i.e site is not publishing:
 - GlueHostProcessorOtherDescription:
Cores=<N>, Benchmark=<X>-HEP-SPEC06
 - If it were $\langle \text{Cores} \rangle * \langle \text{Physical} \rangle = \langle \text{Logical} \rangle$ could be checked.
 - i.e gstat2 will perform a cross check of these 3 values.



Random Tier 2 CPU Anomalies(2)

- UKI-NORTHGRID-LANCS-HEP
 - 1 Logical CPU spread over 656 physical CPUs.
 - Obviously wrong.
 - 5 errors reported by GStat2
- Finland - CSC
 - 4 Logical CPUs spread over 64 physical CPUs.
 - 6 errors reported by GStat2.
- There are clearly inconsistencies.
 - GStat2 is detecting them though.
 - Sites and respective Tier1s please review GStat2 results.
- GStat2 can never determine a correct capacity.
 - Only a sane one.
 - Read the report or gstat2 for instantaneous values.

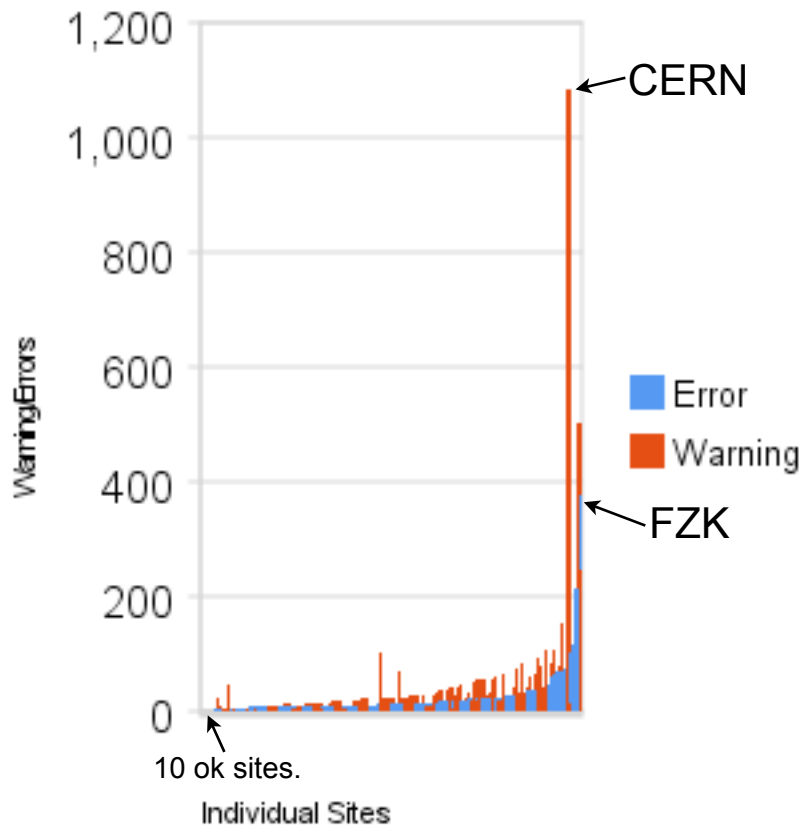


Reading/Browsing Gstat2 results

- GStat2 - Pretty Page
 - <http://gstat-prod.cern.ch/gstat/summary>
 - Search by site, region or tier.
- Gstat2 - Less Pretty Nagios Results
 - <https://gstat-dev.cern.ch/nagios/>
 - Search by site-bdii hostname.
 - Will be linked from pretty page at some point.
- Gstat2 view requests project-grid-info-support@cern.ch
- There are/may be bugs though:
 - <https://savannah.cern.ch/bugs/?55235>
 - New ones to EGEE-OAT savannah.
 - <https://savannah.cern.ch/projects/sa1tools/>
- OAT Nagios, ROC or Site instance.
 - Both of these have same probes running.



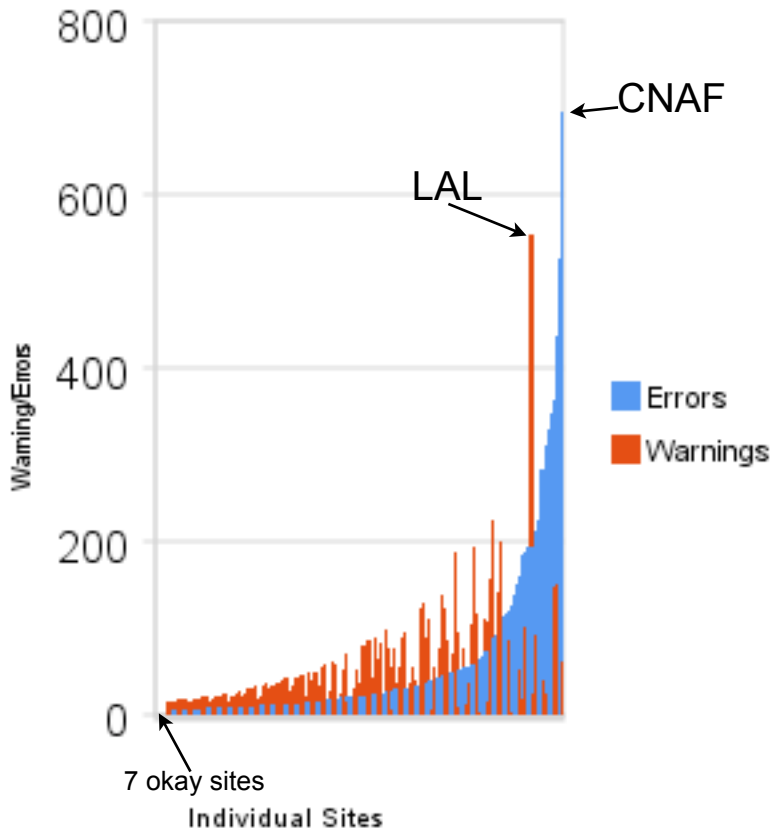
CE GStat 2 Results



- Shows #errors and #warnings from gstat2.
- Ordered by #errors
- Totals:
 - 268 Site BDIIs with CEs.
 - 6202 errors.
 - 9497 warnings.
- Many duplicates.
 - e.g Every CESEbind gives one error.



SE GStat2 Results.



- Shows #errors and #warnings from gstat2.
- Ordered by #errors
- Totals:
 - Site BDIIs with SEs: 268
 - Errors: 15522
 - Warnings: 12995



Related Work - SubClusters

- We must improve SubCluster publishing in EGEE
 - Requested by a number of regions at EGEE conference.
 - This does **NOT** stop correct publication of installed capacity.
 - It will make occupancy figures easier/possible.
- New code for YAIM exists but needs another testing round.
 - YAIM code has no semantics.
 - A bash representation of glue basically.
- OSG GIP plugin seems better and works
 - OSG publishes multiple non-overlapping SubClusters today.
 - Review OSG plugin for EGEE inclusion.
 - Can it be separated from everything else.
 - Is it worth YAIM wrapping the already simple .ini file.



Summary & References

- Add OSG results to report.
- Sites can now view sanity errors via gstat2.
 - T1s please follow up with sites.
 - GGUS or submit bugs where you disagree.
- Installed Capacity Documents
 - https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf

