

# Improving Published Capacity WLCG GDB - October - 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, cream, DPM, Castor, storm, dCache.
    - YAIM for DPM, lcg-CE and Cream.
  - Documented.
    - Detailed Docs - Installed Capacity Doc
    - YAIM documentation.
    - Instructions for running benchmarks
- All sites can and should now publish CPU installed capacity.

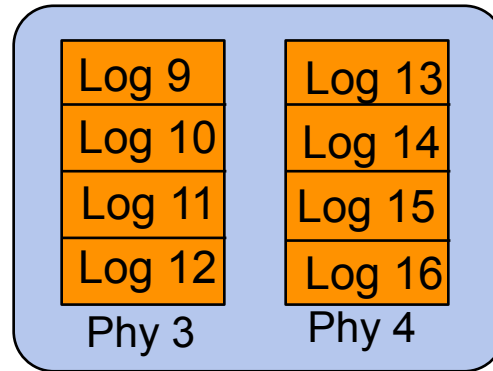


# Log/Phys/SMP/Cores Reminder

WN 1



WN 2



- 2 Node farm
- Dual 4 Core Nodes.

- **GlueSubClusterLogicalCPUs**
  - Total#Logicals = Total#Cores = 16
- **GlueSubClusterPhysicalCPUs**
  - Total#Physicals = #Chips or Sockets = 4
- **GlueHostProcessorOtherDescription: Cores=N**
  - Cores = #Cores/Physical = 4
- **GlueHostArchitectureSMPSize**
  - SMPSize = #Logicals/WN = 8



# CPU data Omissions from T2s

- Major Omissions, i.e zero or null capacity:
- OSG sites all missing.
  - OSG have provided data via MyOSG:
  - <http://tinyurl.com/y9hfl5s>
  - Gridview team working into report now.
    - Easy now report has switched to installed capacity.
    - Expected as draft with next months report.
- AU-ATLAS - Australia-UNIMELB-LCG2
  - Site closed for good, should drop out next month.
- IL-HEPTier-2 - IL-TAU-HEP
  - Site was down for all last month.
- UK-London - UKI-LT2-IC-LESC
  - Site is closed in GOCDB and will drop out.



# 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, 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
  - GStat does a sanity check across 3 numbers: Example:
    - GlueSubClusterUniqueID=svr026.gla.scotgrid.ac.uk: A value is incorrect, Cores \* PhysicalCPUs != LogicalCPUs



# 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 for summaries 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 - see next slide
- Gstat2 - Less Pretty Nagios Results
  - <https://gstat-prod.cern.ch/nagios/>
  - Search by site-bdii hostname.
  - Will be linked from pretty page at some point.
  - nagios results open to [dteam](#), [ops](#) and [lhc vos](#) members.
- Gstat2 view requests [project-grid-info-support@cern.ch](mailto:project-grid-info-support@cern.ch)
- New Gstat2 RC hoped end of month
  - easier to browse installed capacities.
- OAT Nagios, ROC or Site instance.
  - Both of these have same probes running.



# GStat2 and Site Objects

- Inconstancies in GStat2 summaries:

Country Name ^	Sites ^	Logical CPUs^	Physical CPUs^	Storage Space ^	Waiting Jobs
<a href="#">UK</a>	<a href="#">16</a>	<a href="#">13558</a>	<a href="#">8060</a>	70%	72%
<a href="#">Ukraine</a>	<a href="#">1</a>	<a href="#">54</a>	<a href="#">54</a>	0%	81%
<a href="#">United Kingdom</a>	<a href="#">2</a>	<a href="#">1965</a>	<a href="#">1965</a>	0%	39%
<a href="#">Scotland</a>	<a href="#">1</a>	<a href="#">1436</a>	<a href="#">460</a>	0%	0%

- GStat2 constructs summaries from GlueSite object.
- Population of GlueSite is well defined:
  - [http://goc.grid.sinica.edu.tw/gocwiki/How\\_to\\_publish\\_my\\_site\\_information](http://goc.grid.sinica.edu.tw/gocwiki/How_to_publish_my_site_information)
- Please sites review your site objects for consistency:
  - Country - Where is IN (India? No its Indiana)?
  - WLCG Tiers - There are currently only six tier1s?
  - ROCs - RDIG, RU-RDIG and Russia are all unique ROCs?
- We can/will add some sanity checks for this.
- Tickets will be sent to the really obvious ones now.






# Debugging Errors from GStat

- Error messages being added to GOC WIKI@Taiwan
  - <http://goc.grid.sinica.edu.tw/gocwiki/SiteProblemsFollowUpFaq>
  - [http://goc.grid.sinica.edu.tw/gocwiki/Cores\\_format\\_is\\_wrong](http://goc.grid.sinica.edu.tw/gocwiki/Cores_format_is_wrong)
  - [http://goc.grid.sinica.edu.tw/gocwiki/Benchmark\\_format\\_is\\_wrong](http://goc.grid.sinica.edu.tw/gocwiki/Benchmark_format_is_wrong)
  - [http://goc.grid.sinica.edu.tw/gocwiki/GlueCEPolicyAssignedJobSlots\\_has\\_negative\\_or\\_null\\_value](http://goc.grid.sinica.edu.tw/gocwiki/GlueCEPolicyAssignedJobSlots_has_negative_or_null_value)
  - [http://goc.grid.sinica.edu.tw/gocwiki/GlueHostProcessorOtherDescription\\_does\\_not\\_exist](http://goc.grid.sinica.edu.tw/gocwiki/GlueHostProcessorOtherDescription_does_not_exist)
- There are/may be bugs though:
  - e.g torque not publishing AssignedJobSlots
    - not relevant to installed capacity but adds noise.
    - <https://savannah.cern.ch/patch/?3320>
      - Was certified yesterday.
  - New ones to EGEE-OAT savannah.
    - <https://savannah.cern.ch/projects/sa1tools/>




# CE GStat 2 Results - 2 weeks old



- 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 - 2 weeks old



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



# HEPSpec Publication Status

- Number of sites per country that are publishing:
  - 13 - Italy
  - 13 - UK
  - 3 - Spain, Portugal
  - 2 - Country(?), Austria, Canada, CZ and Hungary
  - 1 - Australia, Brazil, Cyprus, France, Germany, Greece, Japan, Korea, Malaysia, Pakistan, Romania, Scotland, Taiwan, Thailand, The Netherlands.
- Some adoption issues from non-WLCG EGEE sites.
  - Some want to avoid to spending money.
  - Some reluctance on acceptance due HEP prefix.
  - Decisions from SA1 coordination meeting:
    - Any serious grid site should be able to spend this kind of money.
    - A rename (or dual name) is possible.



# Installed Capacity by OS

- Capacity is certainly moving to SL5.
  - From today and one month ago, value in ().

OS	#SubClusters	#Logical CPUs	#SI2000
AIX 5	1	95	105735
Debian 4	1(2)	8(2908)	14400(6684400)
Debian 5	1(0)	60(0)	120000(0)
RHEL3 Compat	5(10)	0	0
RHEL4 Compat	389(410)	97385(103513)	169328506 (181804949)
RHEL5 Compat	121(69)	67778(52625)	111400684(80644686)
Suse 10	3	1464	2051600
Ubuntu 8	2	944	1898080



# 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.
  - GRNET are now providing a review.
- OSG GIP plugin seems better or works
  - OSG publishes multiple non-overlapping SubClusters today.



# Summary & References

- Add OSG results to report.
- Sites can now view sanity errors via gstat2.
  - Review summary tables in gstat.
    - Correct site meta data.
  - More sites to measure and publish HEPSPREC.
  - 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](https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf)

