

Installed Capacity Update GDB - July - 2009



Steve Traylen/Flavia Donno
CERN/IT

Computing Resources - Step 1

- Publish non-zero values for Installed Capacity.
 - #Logical CPUS * KS2K value summed over GlueSubClust'.
 - gstat1 is checking that sites do not publish zero.

Sites publishing Zero	CERN-ROC	North	AP	Italy	SE	SW	Russia
June GDB	1	6	2	6	3	2	2
July GDB	1	2	2	2	3	1	1

- gstat2 (nagios) is also checking these values.
 - It is a lot more thorough.

e.g Validating SMPSize, Logical and Physical Numbers



Computing Resources - Step 2

- HEPSpec 2006 Values.
 - Measuring - HEPSpec 2006
 - <https://twiki.cern.ch/twiki/bin/view/FIOgroup/TsiBenchHEPSPECWlcg>
 - Some sites are clearly doing this.
 - Support tickets, mails, blog entries exist on the matter.
 - **ALL sites should be doing this now.**
 - Publication
 - gLite 3.1 update 49 was released this Monday.
 - Allows the following to be set:
 - YAIM variable - **CE_OTHERDESCR**
 - Benchmark=<value>-HEP-SPEC06
 - 1 site publishing - **congratulations FZK**
 - YAIM variable - **CE_CAPABILITY**
 - CPUScalingReferenceSI00=<referenceCPU SI00>
 - 8 USA sites, 1 * Greece, Slovenia, Spain, UK, Latvia, Bulgaria, Russia are publishing.
- Checking for compliance (info now) is added to gstat2 probes.



Double use of SI00 and APEL.

- Previously GlueHostBenchmarkSI00 was used twice:
 - Represent the Speed of hosts in the SubCluster.
 - Represent the normalized CPU speeds.
 - The guaranteed speed that max queue times respect.
- YAIM now supports two variables.
 - GlueHostBenchmarkSI00 - Speed of SubCluster
 - CPUScalingFactorSI00 - Guaranteed CPU speed.
- APEL will respect these changes.
 - <https://savannah.cern.ch/bugs/?51176>
 - Ready to be submitted to certification following current update.
- Sites need this to fix their installed capacity.



OSG Resources.

- OSG has an XML feed of installed capacity per site.
 - CPU and Storage capacity.
 - Many sites are publishing 0 size.
 - OSG are awaiting a final format report.
- For OSG OIM numbers must be authoritative in reports.
 - A risk that OSG BDII numbers creep in.



Reporting CPU - Current

- Current - WLCG Report - by WLCG topology
 - Uses <http://gridops.cern.ch/mou> and gridmap as source.

Federation	Site	CPU Count	Reli-ability	Avail-ability	Reliability History				
					Feb-09	Mar-09	Apr-09		
AT-HEPHY-VIENNA-UIBK (Austria, Austrian Tier-2 Federation)									
	HEPHY-UIBK	39	100 %	100 %	99 %	97 %	100 %		
	Hephy-Vienna	408	94 %	92 %	95 %	92 %	88 %		

- Current - EGEE Report - by region.
 - Uses GOCDDB and gridmap as source.

Region	Site	Phy. CPU	Log. CPU	KSI2K	Avail-ability	Reli-ability	Availability History		
							Mar-09	Apr-09	May-09
AsiaPacific (Asia Pacific)									
	TOKYO-LCG2	480	480	1,080	100 %	100 %	99 %	100 %	91 %
	TW-FTT	249	1,904	3,821	84 %	93 %	0 %	81 %	90 %



From next month WLCG report will have KSI2K.⁶

Reporting CPU - next steps.

- Aim of Report - installed capacity vs pledges.
- All data can now be published.
- Data is available from:
 - MOU topology DB. <http://gridops.cern.ch/mou/>
 - Raw BDII cpu numbers
 - GridMap gives summary CPU data.
 - GStat2 gives summary CPU data now, later history.
 - OSG - OIM.
- Three possibilities - various stages of development.
 - CESGA portal produces a report along side accounting.
 - A gstat2 application
 - A similar iReport to availability report could be generated.



Storage Resources

- Storage Providers
 - DPM information providers are well deployed for a while.
 - Castor providers are released and deployed at RAL only.
 - Storm providers are available and deployed.
- Monitoring
 - GStat 2 SE probes were completed some time ago.
 - Today's Nagios release means these can now be released.
 - GStat 2 is generating used/free space charts calculating on capacity document.
 - A gstat2 preview of CE/SE tests is available.
 - <https://gstat-dev.cern.ch/nagios/cgi-bin/status.cgi?host=all>
- Visualization
 - A gridmap is needed.
- Reporting.
 - gstat2 is providing storage numbers.
 - Use these - create a gstat2 application.



Actions for Sites.

- Run HepSPEC Benchmarks.
- Upgrade CE info providers and yaim
 - Publish HEPSpec Benchmark
- gstat2 probes released end of week.
 - These will give you information now on compliance.
- Consult EGEE CPU report:
 - <https://edms.cern.ch/document/963325/>
- Consult GridMap in strict mode.
 - <http://gridmap.cern.ch>:
 - Tick SI2K, “more”, strictly to #Logical CPUs.

