

Report on Installed Resource Capacity



Flavia Donno/Steve Traylen
CERN/IT-GS/D

*WLCG GDB,
CERN 11 March 2009*

Outline

- The Installed Capacity document (v. 1.8-7) has been approved by the WLCG Management Board on the 3rd of February 2009.
 - https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf
 - V 2.0 will be published soon with some more operational clarifications and OSG examples
- We are preparing a detailed operational plan that includes EGEE and OSG sites.
 - **To tentatively start deployment during March 2009.**
- In what follows we give a short summary of the deployment plan for both EGEE and OSG sites and some operational details concerning EGEE sites.



Goals

- Publishing information in order to provide the WLCG management with a view of the total installed capacity and resource usage by VOs at sites. Focus on:
 - Static information that produces monthly reports
 - Publishing information to provide the WLCG management with a view of the total installed capacity at the sites
 - Publishing information to provide the WLCG and VO management with a view of the resource assignment per VO at the sites.
 - Dynamic/ongoing information
 - Publishing information to allow VO operators to monitor the VO usage of the resources at a site
 - Allowing clients (consumers of the information published) to select the correct resources and their characteristics.



What are the implications for sites?

■ No need for reconfiguration of resources

- Be aware that validation procedures will be stricter to check for validity of the information provided. Only information providers need more configuration.

■ Deployment of a new version of information software

- For **EGEE** sites using **YAIM** or ncm-yaim, the site administrators need to deploy and configure the new YAIM in order to correctly report CPU resources installed and available at a site.
- For **EGEE** sites using **Quattor**, the new Glue computing attributes described in the document need to be introduced. The documentation available should be sufficient for the job.
- For **OSG** sites, new GIP and configuration scripts for both computing and storage resources are under tests on the Integration TestBed. Detailed instructions on how to configure them will be provided.
- For **storage resources**, very little manual configuration should be needed (mostly for tape resources) or none at all. Automatic GIP are available for most of the storage implementations which include CASTOR, dCache, DPM, StoRM, BestMan and pure xrootd.



Installed Capacity: plan

■ Operational plan in preparation.

- Precise instructions for site admins (OSG and EGEE sites)
- Examples
- gstat test suites for both computing and storage resources. They will probably be translated into Nagios sensors and/or Apel SAM tests
- Coordination for deployment
- Support for sites
- All changes are released in the usual fashion.

■ Client tools that are being modified.

- YAIM and Storage GIPs
- Gstat sanity checks
- GridMAP
- APEL
- APEL/gstat SAM tests
- gfa/lcg-utils will improve resource selection algorithm with next releases.
- OSG validation and publishing tools



Installed Capacity: the plan

■ Support group of experts

- During March 2009, **GGUS** will include under the GLUE support unit a group of experts to help sites with questions concerning the publication of computing and storage capacity.
- Such support group is also connected with the **OSG GOC/ROC** for OSG specific configurations. Experts in USA will answer questions concerning OSG sites.

■ New version of YAIM

- **YAIM** provides support for publishing Computing Installed capacity information starting from glite-yaim-lcg-ce-4.0.5-6.noarch.rpm. The evolution of the release process can be followed in the patch <https://savannah.cern.ch/patch/index.php?2840>
- The new variables to be configured by hands by the site admins are described in the “Installed Capacity” document and in the YAIM administration manual

This new version of YAIM will be probably ready for deployment some time in March 2009



Installed Capacity: the plan

■ New GIP and Configuration scripts for OSG sites

- The development of the new information providers and configuration scripts for OSG sites has been completed.
- The OSG release including these new GIPs and configuration scripts will be available around the beginning of March 2009. It is under test on the Integration TestBed as of now.
- Deployment expected March through to May.



Installed Capacity: the plan

■ Storage Services Information Providers

- The new **Castor information providers** v 1.1.3 are available for installation in production. The distribution can be found here: <http://storage.esc.rl.ac.uk/grid-deployment/cip/cip-1.1.3.tbz>. They have been installed in production at RAL. For support, please contact Jens Jensen j.jensen.ral@googlemail.com.
- The new **dCache information providers** will not be available with dCache v.1.9.2 as expected. Preliminary versions are installed at NDGF and are being tested. Support provided through userforum@dcache.org and the dCache tracking system. support@dcache.org. No time scale is available.
- The new **DPM information providers** are available with DPM v1.7.0-6 and already certified. Support provided through GGUS.
- The new **StoRM information providers** available with StoRM v.1.4.0
- The new xrootd information providers for pure xrootd installations are being developed by Claudiu Schiaua [<Claudiu.Schiaua@cern.ch>](mailto:Claudiu.Schiaua@cern.ch) - NIPNE and Eygene Ryabinkin [<rea@GRID.KIAE.RU>](mailto:rea@GRID.KIAE.RU) - Russian Research Centre "Kurchatov Institute" from the Alice collaboration.



Installed Capacity: the plan

■ EGEE/OSG Validation Tools

■ *GridMAP* - <http://gridmap.cern.ch/>

- A **new GridMAP** has/will/may be released **today**.
- Improvements and benefits for all:
 - Reviewers and comments are most welcome.
- Sites are now sized by #LogicalCPUs.
 - Currently if missing #PhysicalCPUs but is highlighted.
- With “tiers” and “SI00” selected = WLCG Installed Capacity.
 - Includes EGEE & OSG (if in SAM and BDII)
- With “regions” = Grid Installed Capacity
 - Includes OSG optionally.
- The WLCG topology data feed (<http://gridops.cern.ch/mou/>)
 - The source for the "tiers" button.
 - Ability to show WLCG sites only.
- VO Shares will come later as the new YAIM starts to allow their publication



Installed Capacity: the plan

■ EGEE/OSG Validation Tools

■ *GSTAT*

- New Gstat sanity checks will be available at the end of March 2009.
- They include consistency checks for both storage and computing resources.
- They have been validated by OSG and are compliant to their deployment. They will be used for the validation of BDII information published by OSG as well.
- They will be provided to ROCs as Nagios tests in the SA1 nagios bundle.
- A process similar to the one used for the APEL/gstat SAM tests will be followed in order to encourage uptake: the tests will first be made critical and after a period of time also influence availability.



Installed Capacity: the plan

■ OSG Validation Tools

- Development of publishing and validation scripts for OSG and WLCG monthly reports has not yet started. OSG will share standalone validation tools with EGEE as far as possible. Availability foreseen for May 2009.



Installed Capacity: the plan

■ APEL reports

- Ongoing development. Scripts to collect storage information need update. Scripts to collect computing information are being written as of now.
- A possible date for producing EGEE first test reports to verify the validity of the system in place is April 2009.
- Dissemination and Validation of published reports for OSG sites will start around July 2009.
- Reports will list both average and peak measurements monthly.
- The reports will be discussed with the WLCG federations before making them public.
- Experience is needed in order to validate the reports. Therefore, we need to start soon!



Computing Resources: changes

■ ***GlueHostBenchMarkSI00 = Average SpecInt2000 rating per LogicalCPU***

- WLCG is currently (feb 2009) in a state of transition with respect to the units of core power measurement. For sites whose MEASUREMENTS are still based on the old SpecInt2000 suite, the measured power must be published only via the GlueHostBenchMarkSI00 attribute. For sites whose core power MEASUREMENTS are done using the new HEP-SPEC06 benchmark, the result of this measurement must be PUBLISHED via the attribute *GlueHostProcessorOtherDescription: Benchmark=<value>-HEP-SPEC06*, where in <value> the CPU power is expressed in the new unit. In addition, this HEP-SPEC06 measurement must be CONVERTED (via the agreed conversion factor) to SpecInt2000 units, and this resulting value must be PUBLISHED via the the GlueHostBenchmarkSI00 attribute.

- ***CECapability: CPUScalingReferenceSI00=<refCPU SI00> the CPU SI00 for which the published GlueCEMaxCPUTimes are valid.***




Computing Resources: an update

$$S(WLCGSubCluster)_{Site} = \begin{cases} \sum_{WLCG\ VO_s} GlueCECapability(Share) / 100 & \text{if WLCG share are published} \\ & \text{through the GlueCECapability attribute (EGEE sites)} \\ \sum_{WLCG\ VO_s} GlueSiteSponsor(Share) / 100 & \text{if WLCG share are published} \\ & \text{through the GlueSiteSponsor attribute (OSG sites)} \\ 1 & \text{if WLCG share are not published} \end{cases}$$

This is one of the quantities collected by the APEL reports at Federation and VO level

$$Total\ Installed\ Computing\ Capacity_{Site}\ (KSI00) =$$

$$\sum_{WLCG\ SubClusters} GlueHostBenchMarkSI00 * GlueSubClusterLogicalCPUs * \frac{S(WLCGSubCluster)_{Site}}{10^3}$$



Storage Resources: an update

■ Changes involve the description of pure xroot installations

- GlueSEImplementationName/Version: It includes xrootd and the version of the protocol.
- GlueControlProtocol/AccessProtocol: Pure xrootd installations will publish both control and access protocol (xroot is both a storage control and file access protocol) to distinguish from “xroot door”-only installations such as dCache, DPM, and CASTOR.
- It is advisable that pure xroot installations will publish only one SA.



$$\text{Installed Capacity} = \left(\sum_{\text{WLCG GlueSA}} \text{GlueSACapability}(\text{InstalledCapacity}) \right)$$

This is one of the quantities to collect for the APEL reports at Federation and VO level



More information

- Further announcements will be given during WLCG daily operations meetings, weekly EGEE operations meetings and OSG weekly production meetings.
- Please, refer to the already mentioned document on “Installed Capacity”
- Further details can also be found in the presentation given at the GDB on December 10th, 2008.

