

Table of Contents

New WLCG IS Roadmap.....	1
Roadmap.....	1
Definitions.....	1
Information Providers.....	2
Site specific recipes.....	2
GLUE 1 - GLUE 2 translator.....	3

New WLCG IS Roadmap

Please, note that this twiki is now obsolete. Please, refer to this twiki for up to date information

Roadmap

The following roadmap will be presented and discussed at the IS TF meeting on the 8th of January.

- Implement the new WLCG IS based on AGIS. The new IS will consume GLUE2 attributes, but only the subset needed by WLCG. This subset will be clearly documented and well defined so that all sys admins and MW developers understand what is required.
- The WLCG IS will query several sources of information. In case the BDII is needed, it will query the resource BDII directly. In this way, WLCG will stop relying on top and site BDIIs.
- New services do not need to deploy a BDII in order to provide resource information. They could provide the needed GLUE2 subset through other means (i.e. OSG will provide GLUE2 JSON)
- REBUS will remain as a WLCG project office tool to collect MoU sites, their pledges and accounting. Installed capacity information will be removed from REBUS and included into the new IS.

Definitions

In order to improve the information published in the new WLCG IS, clear definitions are needed. The following definitions have been proposed by the Information System Task Force. Feedback from site admins is being collected:

Proposal sent to LCG-ROLLOUT:

- `GLUE2ExecutionEnvironmentLogicalCPUs`: the number of processors in one Execution Environment instance which may be allocated to jobs. Typically the number of processors seen by the operating system on one Worker Node (that is the number of "processor :" lines in `/proc/cpuinfo` on Linux), but potentially set to more or less than this for performance reasons.
- `GLUE2BenchmarkValue`: the average HS06 benchmark when a benchmark instance is run for each processor which may be allocated to jobs. Typically the number of processors which may be allocated corresponds to the number seen by the operating system on the worker node (that is the number of "processor :" lines in `/proc/cpuinfo` on Linux), but potentially set to more or less than this for performance reasons.

Another proposal by Andrew after some discussion:

- `GLUE2ExecutionEnvironmentLogicalCPUs`: the number of single-process benchmark instances run when benchmarking the Execution Environment, corresponding to the number of processors which may be allocated to jobs. Typically this is the number of processors seen by the operating system on one Worker Node (that is the number of "processor :" lines in `/proc/cpuinfo` on Linux), but potentially set to more or less than this for performance reasons. This value corresponds to the total number of processors which may be reported to APEL by jobs running in parallel in this Execution Environment, found by adding the values of the "Processor" keys in all of their accounting records.
- `GLUE2BenchmarkValue`: the average benchmark when a single-process benchmark instance is run for each processor which may be allocated to jobs. Typically the number of processors which may be allocated corresponds to the number seen by the operating system on the worker node (that is the number of "processor :" lines in `/proc/cpuinfo` on Linux), but potentially set to more or less than this for performance reasons. This should be equal to the benchmark `ServiceLevel` in the APEL accounting record of a single-processor job, where the APEL "Processors" key will have the value 1.

Another proposal by Brian:

- **GLUE2Execution environment:** The hardware environment allocated by a single resource request.
- **GLUE2ExecutionEnvironmentLogicalCPUs:** the number of single-process benchmark instances run when benchmarking the Execution Environment.
- **GLUE2BenchmarkValue:** the average benchmark result when \$(GLUE2ExecutionEnvironmentLogicalCPUs) single-threaded benchmark instances are run in the execution environment in parallel.
- **GLUE2ExecutionEnvironmentTotalInstances:** The aggregate benchmark results of the computing resource divided by \$(GLUE2BenchmarkValue).

Information Providers

The information providers publish information in an automatic way. How the relevant attributes are calculated by each IP should be aligned with the proposed definitions. The table below collects some information on how services are currently publishing some of these attributes:

Service	Installed Capacities
dCache	It includes all pools (storage nodes) known to the system, whether or not they are enabled
DPM	It includes the pools/disks configured by the site admin. If the admin disables one pool or disk, the installed capacity is updated via de information provider
StoRM	It just reports available space information from the underlying filesystem, there's no concept of resources in downtime visible to StoRM.
CASTOR/EOS (CERN)	it is possible to retrieve information including broken disks or not, it depends on the query performed in CASTOR/EOS
CREAM	Installed capacity is set statically at configuration time in GlueSubClusterUniqueID/GlueSubClusterPhysicalCPUs, GlueHostArchitectureSMPSize and GlueSubClusterLogicalCPUs attributes, while Available capacity is set dynamically in GlueCEUniqueID/GlueCeStateFreeCPUs, GlueCeStateFreeJobSlots, GlueCeInfoTotalCPUs: e.g. Torque infoprovder executes a "pbsnodes -a" and counts all CPUs not in down,offline, unknoww status to set GlueCeInfoTotalCPUs
HTCondorCE	?

Information providers are described in detail under the IS providers twiki.

Site specific recipes

The information that is manually configured by site admins should also be aligned with the proposed definitions. The table below collects some information on how sys admins are currently configuring their sites:

Sites	Service	Scripts	Notes
CERN-PROD	HTCondorCE	htcondorce-cern ↗	There are only 4 values that aren't generated dynamically by calling out to the HTCondor Pool Collector and the Compute Element Schedduler. These are HTCONDORCE_VONames = atlas, cms, lhcb, dteam, alice, ilc (shortend for brevity), HTCONDORCE_SiteName = CERN-PROD, HTCONDORCE_HEPSPEC_INFO = 8.97-HEP-SPEC06, HTCONDORCE_CORES = 16. All our htcondor worker nodes expose a hepspec fact. The averaged hepspec value on the CEs above is taken by a query of all the facts and then averaged.
	CREAM CE	[[[UpdateStaticInfo]]]	It parses the LSF configuration file to extract capacities

GLUE 1 - GLUE 2 translator

In order to start consuming GLUE 2 information, a translator from GLUE 1 to GLUE 2 for the attributes published in the REBUS installed capacities is available in the table below:

GLUE 1 Attribute	Definition	GLUE 2 Attribute
GlueHostProcessorOtherDescription: Benchmark	GLUE2BenchmarkValue	Benchmark Value
GlueSubClusterLogicalCPUs	The effective number of CPUs in the subcluster, including the effect of hyperthreading and the effects of virtualisation due to the queueing system	GLUE2ExecutionEnvironmentLogicalCPUs GLUE2ExecutionEnvironmentTotalLogicalCPUs
GlueSESizeTotal	GLUE2StorageServiceCapacityTotalSize	The total amount of storage of the default type. It is the sum of free, used and reserved)

-- MariaALANDESPRADILLO - 2015-11-24

This topic: EGEE > NewWLCGISRoadmap

Topic revision: r9 - 2016-09-09 - MariaALANDESPRADILLO



Copyright &© by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Ask a support question or Send feedback