



Installed Capacities

WLCG Information System Task Force

Maria Alandes
IT/SDC





Goals of this discussion

- The goal of this discussion is to clarify whether installed capacities are needed or used in WLCG
 - Use cases for installed capacities should be identified
 - If there are use cases for installed capacities
 - Agree how installed capacities are going to be provided
 - What will be the source of information? (i.e. BDII or manual input)
 - What will be the update frequency
 - Will they be published in REBUS? Or somewhere else?



How are installed capacities published today?

INSTALLED CAPACITIES IN REBUS





Two types of installed capacities

- There are two types of installed capacities currently published in REBUS
 - *T1 and T2 Installed Capacities*
 - *Site Installed Capacities*

(Random names have been chosen, any better suggestion?)

- More details in the next slides...



T1 and T2 Installed Capacities

- Calculated on a monthly basis and used by WLCG Management to generate Accounting reports (since January 2006!)
- Installed CPU, disk and tape
- Per VO: grid and non-grid CPU and Wall Time, disk allocated, disk usage and tape usage
- T1 installed capacities are pre-filled automatically from Accounting Portal but can be manually modified
- T2 installed capacities are automatically taken from Accounting Portal
- Details:
https://twiki.cern.ch/twiki/bin/view/EGEE/AllAboutREBUS#Accounting_Reports

✓ Report submitted for CH-CERN!

You can no longer make any changes. Please e-mail any subsequent modifications to lcg.office@cern.ch

April 2014

installed cpu	installed disk	installed tape
356000	29100	67400

VO	Grid CPU Usage	Grid Wall Time	Non-Grid CPU Usage	Non-Grid Wall Time	Disk Allocated	Disk Usage	Tape Used
ALICE	1520047	1934335	40512	393637	8434	6888	8635
ATLAS	883247	1018011	534023	1745745	10166	6377	33000
CMS	805282	951343	1445960	3006399	8080	5977	28673
LHCb	491149	532166	44560	192521	4558	4124	8965

Site Installed Capacities



- Site Capacities taken from the BDII/MyOSG every hour
- Based on https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf
- Physical CPU, Logical CPU and HS06
 - Taken from GlueSubClusterPhysicalCPU, GlueSubClusterLogicalCPUS, GlueHostProcessorOtherDescription (For OSG HS06 is taken from MyOSG portal)
 - For CREAM -> No complex algorithm or script calculating this! It goes from a configuration file to the BDII, so it's as if sys admins would edit this manually in REBUS
 - For ARC CE: Values are taken from arc.conf file which is manually configured by the sysadmin
 - For OSG: Resource administrators enter the information manually through the OIM interface. MyOSG reads from the OIM DB to serve the information to REBUS.
- Total online and nearline storage
 - Depending on the storage, it relies on algorithms or scripts calculating these numbers (DPM or StoRM), or configuration files maintained by the sys admins (EOS, CASTOR or dCache)
- It seems most of these values are defined manually and just transported through BDII to REBUS
 - No complex algorithms behind in most cases!
 - Suggestions to make them more reliable are to enter them manually in REBUS, but would this make a real difference apart from putting more burden on the sys admin?

WLCG Resource, Balance & Usage

REBUS: Site Capacities

VO: ALL Year: 2015 Month: 7

Note: Sorting by multiple columns at the same time can be activated by Shift-clicking on the column headers which they want to add to the sort. Hovering mouse over the column headers to get descriptions of table columns.

Site Name	Physical CPU	Logical CPU	HEPSPEC06	Total Online Storage (GB)	Total Nearline Storage (GB)
AGLT2	710	5,724	68,200	3,722,400	0
Australia-ATLAS	122	920	10,465	1,014,939	0
BEgrid-ULB-VUB	209	2,545	19,138	2,228,911	0
BEIJING-LCG2	122	1,020	11,147	940,368	0
BelGrid-UCL	220	1,320	12,144	984,000	0
BNL-ATLAS	688	11,008	110,000	11,000,000	15,000,000
BU_ATLAS_Tier2	482	3,388	28,888	2,450,000	0
BUDAPEST	159	764	8,954	0	0
CA-MCGILL-CLUMEQ-T2	5,458	21,992	193,039	958,612	0
CA-SCINET-T2	212	848	10,600	846,538	0
CA-VICTORIA-WESTGRID-T2	240	1,440	20,679	884,364	0
CBPF	64	392	3,277	127,981	0
CERN-PROD	63,220	65,952	531,596	60,135,326	101,256,861
CIEMAT-LCG2	328	1,698	22,974	1,298,206	0
CIT_CMS_T2	848	6,216	14,085	550,000	0
Crane	0	0	0	0	0



Site Installed Capacities

- Identified use cases
 - Resource utilisation calculation
 - For both ATLAS and CMS
 - It uses Logical CPU + HS06
- Any other use cases?
 - CMS also expressed interested in using information in REBUS if proven to be reliable, but which information exactly?



What is Installed Capacities?

DEFINITIONS





Two definitions?

- There are two types of Installed Capacities in REBUS as already discussed
 - T1 and T2 Installed Capacities
 - Installed Capacities for WLCG Management
 - This is a clear use case for WLCG management that probably doesn't need to be discussed in this TF
 - It will be in any case properly documented in https://twiki.cern.ch/twiki/bin/view/EGEE/AllAboutREBUS#Accounting_Reports since the existing documentation is a bit old! <https://espace2013.cern.ch/WLCG-document-repository/Accounting/archive/miscellaneous%20documents%20and%20reports/Reporting%20of%20Accounting%20Data%20during%20SC4.doc>
 - Site Installed Capacities
 - Installed Capacities for experiments and other services
 - This is what this TF is interested in understanding!
 - Attempt to define Installed Capacities in Glue 1.3 in 2009
 - https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf
 - Are these definitions still valid today?
 - Some feedback sent in the ML
 - https://twiki.cern.ch/twiki/bin/view/EGEE/WLCGISEvolution#Use_Cases