



# **WLCG Information System Use Cases Review**

## **WLCG Operations Coordination Meeting 18<sup>th</sup> June 2015**

Maria Alandes  
IT/SDC





# Scope of this presentation

- These slides summarise the current use of the Information System by the following identified Information System users or consumers:
  - ALICE
  - ATLAS
  - CMS
  - LHCb
  - SAM
  - REBUS
  - Miscellanea: GFAL, WLCG Google Earth Dashboard, CERN-IT management C5 report
- The list may not be complete!
  - I haven't gathered feedback from sites
  - It's difficult to identify all possible use cases !
    - Some interactions of the BDII may be hidden and only known to SW architects of a particular SW tool or framework
  - So please speak up during this meeting if you have an important use case to report!



# ALICE use cases

- ALICE relies on the following information published in the resource BDII of CREAM CEs and also the site BDII for ARC CEs:
  - Status of the CEs (Production vs anything else)
  - Number of waiting jobs in the VOView
  - Number of running jobs in the VOView



# ATLAS use cases

- ATLAS runs a BDII collector on AGIS to gather the following information:
  - List of CEs, submission queues and associated parameters
  - List of SEs, protocols, storage areas and paths
  - Site latitude and longitude
  - Batch system type and version
- In many cases the information taken from the BDII is reviewed and updated by ATLAS

## ATLAS Grid Information System

ATLASSite DDMEndpoint PANDA Queue Service Central Services DDM Groups **Services** Docs TWiki OLD JSON

Show 500 entries First Previous 1 Next Last

give me url of this page

Name	Service Type	Site	Endpoint	State	Status
	CE			ACTIVE	
AGLT2-CE-gate02.grid.umich.edu	CE/GLOBUS	AGLT2	gate02.grid.umich.edu:2119	ACTIVE	production
AGLT2-CE-gate03.aglt2.org	CE/GLOBUS	AGLT2	gate03.aglt2.org	ACTIVE	
AGLT2-CE-gate04.aglt2.org	CE/GLOBUS	AGLT2	gate04.aglt2.org:2119	ACTIVE	production
AM-04-YERPHI-CE-ce.yerphi-cluster.grid.am	CE/CREAM-CE	AM-04-YERPHI	ce.yerphi-cluster.grid.am:8443	ACTIVE	production
ARNES-CE-jost.arnes.si	CE/ARC-CE	ARNES	gsiftp://jost.arnes.si:2811	ACTIVE	production
Australia-ATLAS-CE-agcream1.atlas.unimelb.edu.au	CE/CREAM-CE	Australia-ATLAS	agcream1.atlas.unimelb.edu.au:8443	ACTIVE	production
Australia-ATLAS-CE-agcream2.atlas.unimelb.edu.au	CE/CREAM-CE	Australia-ATLAS	agcream2.atlas.unimelb.edu.au:8443	ACTIVE	DISABLED
REFRING-LCG2-CE-arc01.iben.ac.cn	CE/ARC-CE	REFRING-LCG2	gsiftp://arc01.iben.ac.cn:2811	ACTIVE	



# CMS use cases

- CMS relies on the following information published in the Infosys:
  - List of CEs
    - This is used to generate
      - The Vofeed for the SAM tests
      - To bootstrap the glideinWMS Factory XML file
    - HTCondor CEs are not taken from the BDII
      - The list of HTCondor CEs is taken using OSG collectors



# LHCb use cases

- LHCb uses the information published in the Infosys to feed DIRAC:
  - List of CEs
    - CEs are discovered whenever they are deployed
    - The Max CPU Time and the CPU Scaling Reference attributes are also used
- SW Tags published in the Infosys are also used
  - In particular the VO-lhcb-pilot SW Tag

## Information on Queues provided to LHCb for Distributed Data Processing

Legend: LHCb Site Name — WLCG Site Name — Tier Level — CE Type — Submission Style — Queue Name — Max CPU Time (Minutes) — SI2K Specint 2000 Power — Max Amount of Running Jobs — Max Amount of Waiting Jobs — Running to Waiting Ratio — Queue Length in HepSpec06 seconds (MaxCPU\*SI2K\*60/250) — Longest Queue on the Site (if several with the same length will only print one)? — Site Associated to another SE

Queue Length Colorcode: infinite > black > 10M > green > 1.4M > orange > 700k > red > = 0 (HepSpec06.seconds)

	Site	WLCG	Tier ▲	Type	Subm	Queue Name	Max CPU	SI2K	Max Run	Max Wait	Ratio	Queue length
1	DIRAC.Jenkins.ch		Test	Test		jenkins.cern.ch:8443/jenkins-queue_not_important	200	2,400	5	10		115,200
2	DIRAC.JobDebugger.ch			Test		ce.debug.ch:8443/default	1,000,000					0
3	DIRAC.ONLINE.ch			CREAM		OnlineCE.ch:8443/OnlineQueue	400,000					0
4	DIRAC.OSC.us		OSC	SSH	Direct	oakley.osc.edu:8443/serial	2,000	2,400	500	100	1	1,152,000
5	DIRAC.RAL.uk		RAL-LCG2			vac01.gridpp.rl.ac.uk:8443/default	24,000,000					0
6	DIRAC.Syracuse.us		Siracuse	SSHCondor	Direct	phy.syr.us:8443/batch	6,666	2,400	500	10		3,839,616
7	DIRAC.Test.ch		Test	Test		ce.test.ch:8443/batch	6,666	2,400	500	10		3,839,616
8	DIRAC.Test.ch		Test	Cloud		vcycle-lhcb-cloud.cern.ch:8443/default	100					0
9	DIRAC.VANBEK.be		VanBek	SSH	Direct	vac01.gridpp.rl.ac.uk:8443/default	24,000,000					0



# SAM use cases

- SAM relies on the following information published in the Infosys:
  - Queue name
    - Even if SAM now relies on the experiments Vofeed to get the list of resources, for CREAM CEs, ARC CEs and GRAMs still needs to get the queue name from the BDII
      - This is not provided in the VOfeed

```
<group name="BU_ATLAS_Tier2" type="ATLAS_Federation_US-NE12"/>
<group name="BU_ATLAS_Tier2" type="ATLAS_Site"/>
<group name="Tier-2" type="ATLAS_Tier"/>
<group name="BU_ATLAS_Tier2" type="All sites"/>
<group name="BU_ATLAS_Tier2" type="Tier2s"/>
</atp_site>
- <atp_site name="CA-MCGILL-CLUMEQ-T2" infrast="LCG">
- <service hostname="storm02.clumeq.mcgill.ca" flavour="SRMv2" endpoint="srm://storm02.clumeq.mcgill.ca:8444/srm/managerv2?SFN=">
  <spacetoken name="CA-MCGILL-CLUMEQ-T2_DATADISK" base_path="/atlas/atlasdatadisk"/>
  <spacetoken name="CA-MCGILL-CLUMEQ-T2_LOCALGROUPDISK" base_path="/atlas/atlaslocalgroupdisk"/>
  <spacetoken name="CA-MCGILL-CLUMEQ-T2_PHYS-HIGGS" base_path="/atlas/atlasgroupdisk/phys-higgs"/>
  <spacetoken name="CA-MCGILL-CLUMEQ-T2_PRODDISK" base_path="/atlas/atlasproddisk"/>
  <spacetoken name="CA-MCGILL-CLUMEQ-T2_SCRATCHDISK" base_path="/atlas/atlasscratchdisk"/>
</service>
<service hostname="squids.clumeq.mcgill.ca" flavour="Squid" endpoint="http://squids.clumeq.mcgill.ca:3128"/>
<service hostname="ce02.clumeq.mcgill.ca" flavour="CREAM-CE" endpoint="ce02.clumeq.mcgill.ca:8443"/>
<service hostname="ce03.clumeq.mcgill.ca" flavour="CREAM-CE" endpoint="ce03.clumeq.mcgill.ca:8443"/>
<service hostname="ps-latency.clumeq.mcgill.ca" flavour="Personal" endpoint="ps-latency.clumeq.mcgill.ca"/>
<service hostname="ps-bandwidth.clumeq.mcgill.ca" flavour="PerfSonar" endpoint="ps-bandwidth.clumeq.mcgill.ca"/>
<group name="CA-MCGILL-CLUMEQ-T2" type="ATLAS_Cloud_CA"/>
<group name="CA-MCGILL-CLUMEQ-T2" type="ATLAS_Federation_CA-EAST-T2"/>
<group name="CA-MCGILL-CLUMEQ-T2" type="ATLAS_Site"/>
<group name="Tier-2" type="ATLAS_Tier"/>
<group name="CA-MCGILL-CLUMEQ-T2" type="All sites"/>
```



# REBUS use cases

- Capacities are taken automatically from the information published in the Infosys
  - For OSG HEPSPSPEC is taken from MyOSG
- Pledges are manually reported by sites
  - Pledges are used for the accounting reports
  - It's not so clear who uses the capacity information (?)

WLCG REsource, Balance & Usage | WLCG Home | Contact | GGUS | Admin Login

## REBUS: Site Capacities

Topology | Pledges | **Capacities** | Report | Trends

Capacities > Site Capacities

VO: ALL | Year: 2015 | Month: 6

Note: Sorting by multiple columns at the same time can be achieved by 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.

JSON | CSV | XLSX

Site Name	Logical CPU	HEPSPSPEC06	Total Online Storage (GB)	Total Nearline Storage (GB)	
AGLT2	780	5,856	64,300	3,513,800	0
Australia-ATLAS	122	920	10,465	1,473,939	0
BEgrid-ULB-VUB	208	2,544	19,131	1,997,406	0
BEIJING-LCG2	256	1,088	11,890	940,406	0
BelGrid-UCL	220	1,320	12,144	984,000	0
BNL-ATLAS	688	11,008	110,000	11,000,000	12,000,000
BU_ATLAS_Tier2	482	3,388	28,888	2,450,000	0
BUDAPEST	159	764	8,954	0	0





# Miscellanea

- GFAL 2 has also a dependency on the Infosys:
  - BDII lookups can be disabled in all cases with `DBDII:ENABLE=false`, or modifying the corresponding configuration file (i.e. `/etc/gfal2.d/bdii.conf`)
  - However, if the BDII lookups are enabled and the SURL is not a full SURL, the BDII will be contacted
- Weekly C5 report (used internally in CERN IT for the management) relies fully on the BDII
- WLCG Google Earth Dashboard relies also on the BDII to obtain the site coordinates



# GLUE 1 or GLUE 2?

- All the use cases presented in these slides, query GLUE 1 information



# A future without the Infosys

- OSG is trying to get rid of the dependencies on the Infosys
  - Not all WLCG sites will be published in the BDII anymore
    - Incomplete picture!
- Could WLCG also consider getting rid of the Infosys?
  - Is this the direction we want to go?
  - Could the presented use cases find alternative ways of getting the same information?
    - Maybe we can learn from the OSG experience!
  - One less service to maintain
    - But multi VO sites may still need to maintain the Infosys as it is used by other VOs
  - If we stop the Infosys, what happens with GLUE? how do we ensure interoperability to exchange information?