



Towards a Global Service Registry for the World-Wide LHC Computing Grid

Maria ALANDES, Laurence FIELD,
Alessandro DI GIROLAMO
CERN IT Department

CHEP 2013

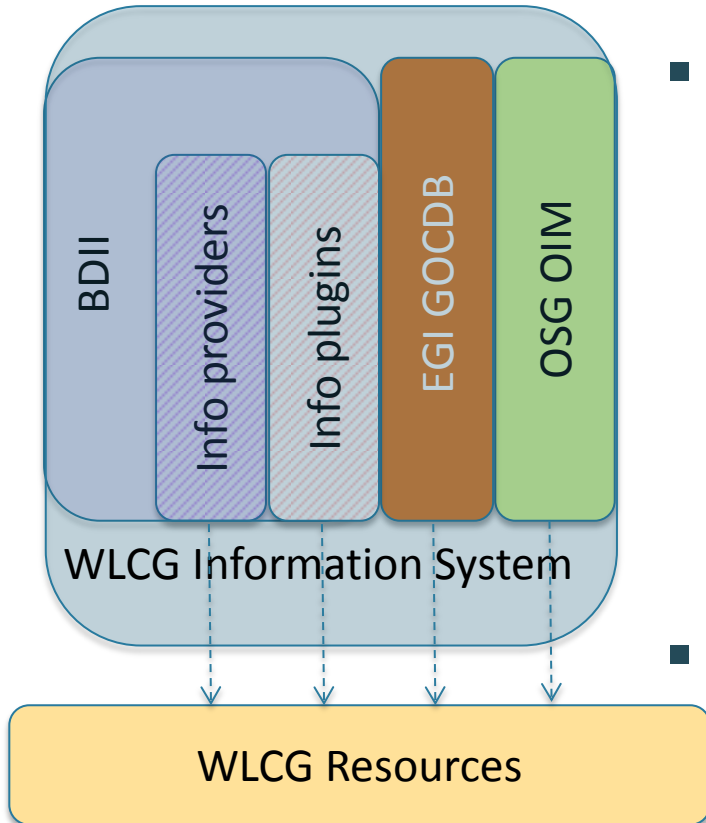


Contents

- WLCG Information System
- Existing Problems
- WLCG Global Service Registry
- Advantages for the LHC VOs
- AGIS use case
- Conclusions



WLCG Information System




- Heterogeneous and distributed
- Comprised of several independent components
 - OSG OIM
 - EGI GOCDDB
 - Resource, site and top BDIIs
 - Information providers: to collect the information from each service
 - Information plugins: to update dynamic information (i.e. Number of running jobs)
- Each component offers a different view of the information
 - All of them are needed, adding complexity to the whole Information System

WLCG Information System

	Information Type	GLUE schema	Deployment model	Installed vs Registered
BDII	Static and Dynamic	GLUE 1 and 2	Distributed	Installed (collected automatically)
Providers/plugins	Dynamic	GLUE 1 and 2	Distributed	-
GOCDDB	Static	GLUE 1	Central	Registered (collected manually)
OIM	Static	GLUE 1	Central	Registered (collected manually)

WLCG Information System



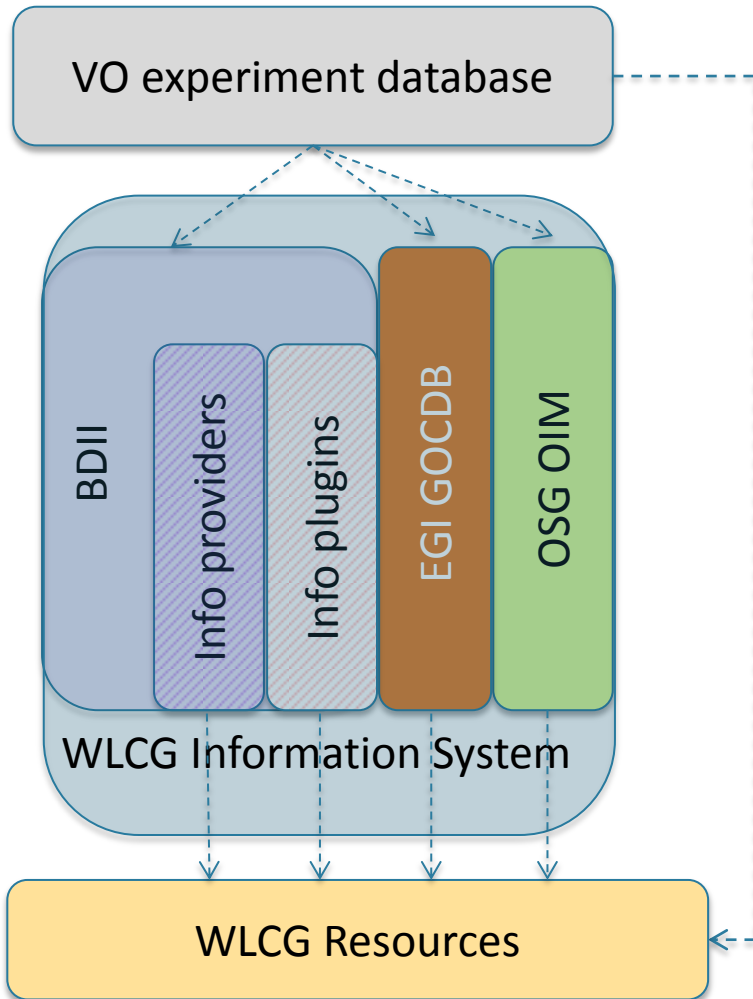
- Interaction with three different services is needed (GOCDDB, OIM and BDII)
 - Because the Information used by VOs is scattered!
 - Downtimes in GOCDDB/OIM and not in BDII
 - GOCDDB/OIM only contains some basic static information for each service, for the rest BDII is needed
 - Because the nature of information is different
 - Registered resources in GOCDDB/OIM may not be aligned with what it is actually installed, so the BDII is needed

Existing problems




- Dealing with different sources of information means
 - Interacting with different development teams
 - Not easy to get fixes in the desired time scale
 - Interacting with different operations teams
 - GOCDB and BDII managed by EGI, OIM managed by OSG
 - Interacting with different versions of the GLUE schema
 - OSG is still publishing GLUE 1 when EGI has already plans to decommission GLUE 1 in 2014
 - Interacting with different sys admins
 - Not easy to issues in the BDII distributed environment
 - There is little control on what it is published
 - A big effort is needed to improve the quality of information

Existing Problems



- VOs have decided to add an extra layer on top of the existing system
 - To add VO specific knowledge
 - To workaround the existing problems
 - With a lot of manual modifications
→ Not sustainable!
 - Having to deal with the three services → A lot of integration effort
 - Sometimes dealing directly with the resources to get extra information
→ Poor information quality in BDII!
- Effort is duplicated since each VO has its own solution!

WLCG Global Service Registry



- Develop an intermediate layer to overcome the existing problems
 - Central and unique place gathering WLCG resource information
 - Focusing on the information requested by LHC VOs
 - More control on the information
 - We can directly improve the information quality to meet VO requirements

WLCG Global Service Registry

Global Service Registry [WLCG Home](#) | [Contact](#) | [GGUS](#) | [Admin](#)

GSR: Service Pledges

Services

[Home](#)

Show 100 entries

JSON CSV

service_id
a2ge.jinr-t1.ru
abaddon.hec.lancs.ac.uk
abisko-ce.hpc2n.umu.se
abitibi.sbgrid.org
acc.grid.umb.sk
accounting.dur.scotgrid.ac.uk
accounting.egi.eu
accounting.ndgf.org

Global Service Registry [WLCG Home](#) | [Contact](#) | [GGUS](#) | [Admin](#)

GSR: Service Records

Services

[Home](#)

Show 100 entries

JSON CSV

Search:

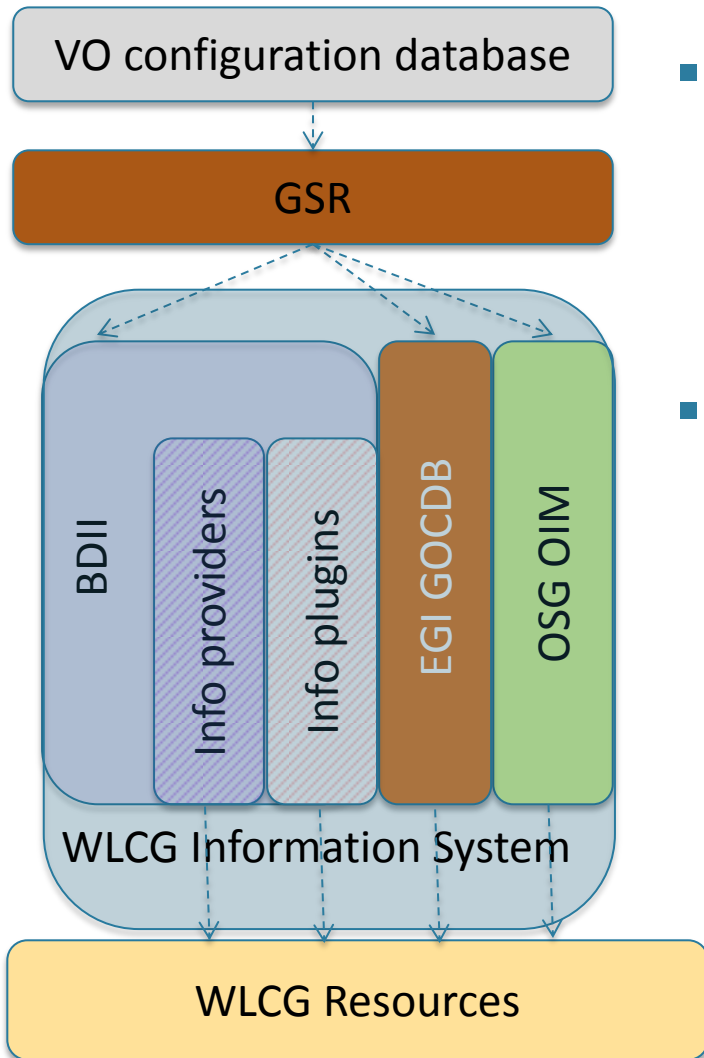
id	admindomainforeignkey	updated
zeus60.cyf-kr.edu.pl_bdii-top_4259157489	CYFRONET-LCG2	2013-10-09 16:01:23
zeus60.cyf-kr.edu.pl_bdii-top_2372158615	CYFRONET-LCG2	2013-08-22 12:02:59
xcreamce.ge.infn.it_ComputingElement	INFN-GENOVA	2013-10-08 00:02:29
wormhole.westgrid.ca_bdii-site_3536672524	SFU-LCG2	2013-09-10 01:02:24
wormhole.westgrid.ca_bdii-site_3171828135	SFU-LCG2	2013-10-09 16:01:25
wmslb101.grid.ucy.ac.cy_wmproxy_333884185	CY-01-KIMON	2013-10-09 16:02:29
wmslb101.grid.ucy.ac.cy_wmproxy_1770776232	CY-01-KIMON	2013-10-09 16:01:14
wmslb101.grid.ucy.ac.cy_lbserver_201164012	CY-01-KIMON	2013-10-09 16:01:20

WLCG Global Service Registry

- Single entry point to get information about WLCG resources
 - Registered resources
 - Actual resources
- Integrates resources from OSG and EGI
 - Independent from GLUE schema version
 - Although GLUE 2.0 is the internal integration model
- Easier for the experiments to gather the information they need in their databases
 - Saves them from dealing with different sources of information
 - Enables the possibility of changing the sources of information in a transparent way
- Single central repository
 - Single place where information is processed
 - Avoids duplication of effort
 - Reduces inconsistencies
 - Improves quality as the processing is done by experts
 - Easier to update and maintain



Advantages for the LHC VOs



- Unique source of information
 - All VO configuration databases interact with the GSR in the same way
 - Reuse of existing solutions is possible in an easy way
- Full control of the information
 - Enforcing consistent information between registered and actual resources
 - Hiding bugs of the underlying components
 - E.g. Middleware provider bugs in the BDII
 - Ensuring quality by fixing the published information when it is wrong

AGIS use case

- AGIS is the ATLAS Grid Information System
 - Implements a set of collectors to gather information from GOCDDB, OIM and BDII
 - A new collector to GSR has been added
 - This experience has helped developed the first GSR prototype
 - Work is still progressing but seems promising so far
 - Easy to integrate
 - Easy to accommodate AGIS requirements to GSR

Conclusions



- The complexity of the Information System relies on its heterogeneity and difficulty to deal with the different sources of information
- The WLCG Global Service Registry hides the current complexity offering an homogeneous layer to LHC VOs
- The first prototype has been integrated within AGIS
 - No showstopper, work progressing