


# WLCG Information System Status

Maria Alandes Pradillo, CERN

CERN IT Department, Support for Distributed Computing Group

GDB 9<sup>th</sup> April 2014



# Contents

- BDII release highlights
- BDII and GLUE deployment
- GLUE 2 validation status
- Storage Information Providers
- Cloud resources in the BDII
- ginfo, AGIS and GSR



# BDII release highlights

- Releases since the last GDB presentation (October 2013)  
<http://gridinfo.web.cern.ch/sys-admins/bdii-releases>
  - Decommission of FCR (Freedom of Choice of resources) mechanism in top BDII
  - Default top BDII cache validity from 12h to 4 days
  - Glue-validator bug fixes and new features
- No BDII releases scheduled
  - NOTE that BDII is IPv6 compliant since EMI 2
- Only glue-validator bug fixes and improvements are expected
- No GGUS tickets reporting any issue since last October

# BDII deployment status

- <https://wlcg-mon/dashboard/request.py/siteview#currentView=BDII+deployment>  
– All GOCDDB + OSG sites

Site Name	bdi 5.2.10	bdi 5.2.12	bdi 5.2.13	bdi 5.2.17	bdi 5.2.20	bdi 5.2.21	bdi 5.2.22	BDII Endpoints
bdi_site	1	34	0	1	3	7	282	328
bdi_top	2	4	0	3	6	2	71	88

EMI 2

EMI 3

# Cleaning SW tags

- [https://twiki.cern.ch/twiki/bin/view/EGEE/GLUEMonitoring#Cleaning SW Tags](https://twiki.cern.ch/twiki/bin/view/EGEE/GLUEMonitoring#Cleaning_SW_Tags)
  - One year ago: 120 MB of SW tags!
  - SW tags removed (except VO-lhcb-pilot) for LHCb
  - Ongoing cleaning for ATLAS
    - Most of them deleted, still 1771 SW tags published
  - To be done for CMS (still in use)
    - 24275 SW tags published in GLUE2

# GLUE deployment status

- <https://wlcg-mon.cern.ch/dashboard/request.py/siteview#currentView=Glue+Deployment>
  - All GOCDB + OSG BDIIs

Site Name	GLUE 1 DNs	GLUE 1 Data Size	GLUE 2 DNs	GLUE 2 Data Size	GLUE 1 Sites	GLUE 2 Sites	GLUE 1 Endpoints	GLUE 2 Endpoints	GLUE 1 Endpoint Types	GLUE2 Endpoint Types	GLUE 2 Services	GLUE 2 Service Types	Non obsolete GLUE 2 SW tags
lcg-bdii	60578	61MB	123975	114MB	379	329	3960	4607	30	52	1924	33	32219

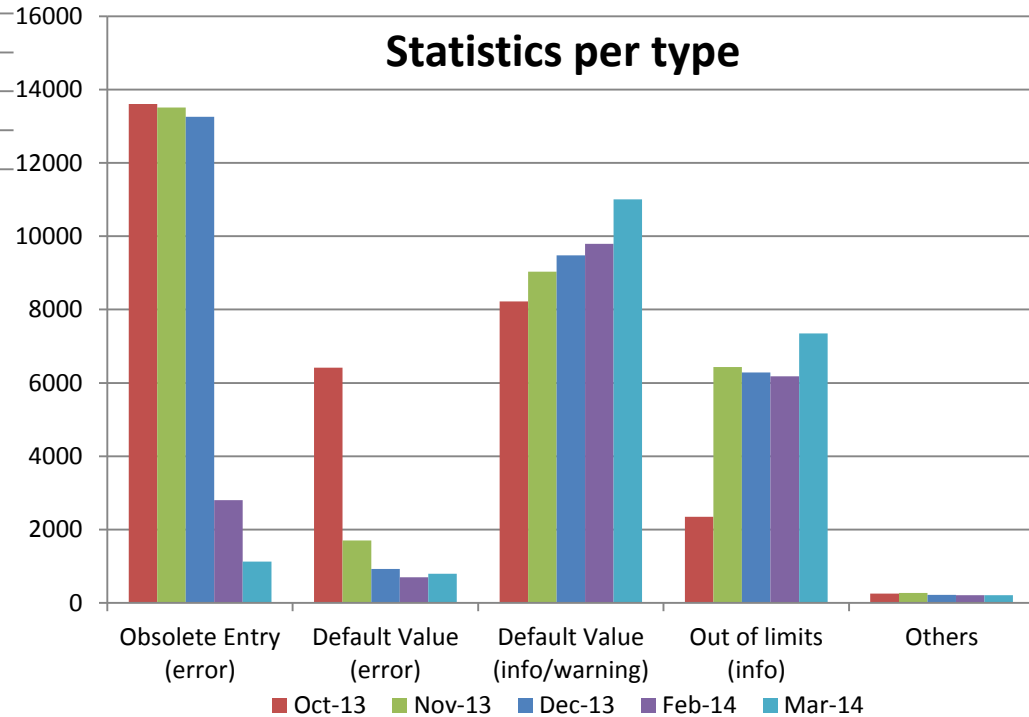
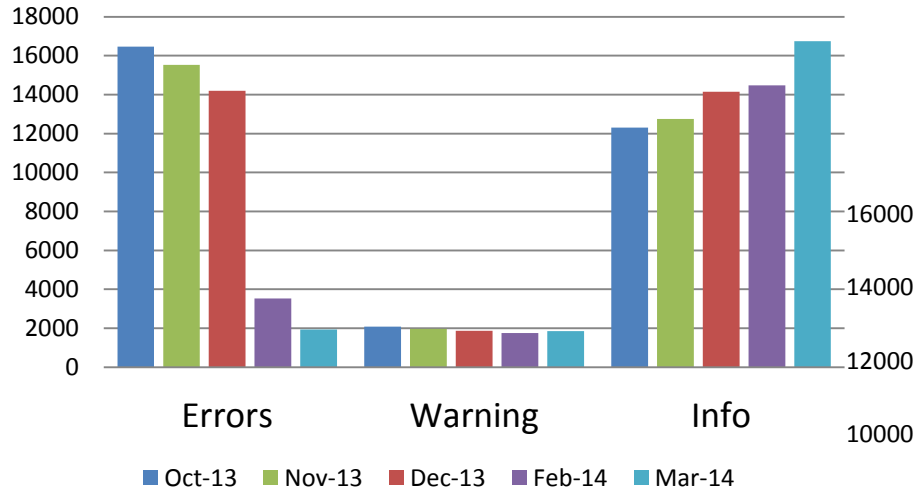
- GLUE deployment per REBUS site:
- <https://wlcg-mon/dashboard/request.py/siteview#currentView=Glue+Deployment+per+site>

# GLUE 2 validation

- GLUE 2 validation for WLCG has been automated since February using the Dashboard and GGUS automatic submission
  - Main focus on major sources of errors at sites that publish GLUE 2:
    - GLUE 2 obsolete entries (Due to a bug in bdii < 5.2.20 old entries are not deleted in GLUE 2. Solved in EMI 3)
      - 76 tickets opened, 73 tickets solved
    - 444444 Waiting jobs (It also affects GLUE 1!)
      - 37 tickets opened, 35 tickets solved
      - This is an old problem but now there is an automatic way of monitoring it!
  - <https://wlcg-mon.cern.ch/dashboard/request.py/siteview#currentView=Glue+Validator+WLCG>
  - This has been in place until glue-validator has been deployed in Nagios as a production probe
    - Available since 18.03.2014
    - 101 tickets opened, 60 tickets solved
  - [https://midmon.egi.eu/nagios/cgi-bin/status.cgi?servicegroup=SERVICE\\_Site-BDII&style=detail](https://midmon.egi.eu/nagios/cgi-bin/status.cgi?servicegroup=SERVICE_Site-BDII&style=detail)
  - NOTE: There will be a new version deployed in the next day to fix a bug in the test that checks the Total Number of Jobs as reported in GGUS:102853.*

# GLUE 2 validation statistics

## GLUE 2.0 Site Validation (exclude-known-issues)





# Specific LHCb campaigns (I)

- BDII vs SRM in T1s (in good shape!)
  - <https://wlcg-mon.cern.ch/dashboard/request.py/siteview#currentView=BDII+vs+SRM+LHCb+Storage>

Site Name	LHCb Total Disk	LHCb Used Disk	LHCb Total Tape	LHCb Used Tape	LHCb Total User Space	LHCb Used User Space	GGUS ticket for incoherent storage capacity numbers (LHCb)
CERN	OK=>BDII:65==>SRM:65	OK=>BDII:51==>SRM:51	OK=>BDII:443==>SRM:443	OK=>BDII:429==>SRM:429	OK=>BDII:261==>SRM:261	OK=>BDII:171==>SRM:172	None
CERN-EOS	Missing_LHCb_shares	Missing_LHCb_shares	Missing_LHCb_shares	Missing_LHCb_shares	Missing_LHCb_shares	Missing_LHCb_shares	None
CNAF	OK=>BDII:2284==>SRM:2284	OK=>BDII:1413==>SRM:1413	OK=>BDII:1055==>SRM:1055	OK=>BDII:172==>SRM:172	OK=>BDII:121==>SRM:121	OK=>BDII:67==>SRM:67	None
GRIDKA	OK=>BDII:1280==>SRM:1280	OK=>BDII:1262==>SRM:1262	OK=>BDII:135==>SRM:135	OK=>BDII:2==>SRM:2	OK=>BDII:60==>SRM:60	OK=>BDII:58==>SRM:58	None
IN2P3	OK=>BDII:1022==>SRM:1022	OK=>BDII:1001==>SRM:1001	OK=>BDII:27==>SRM:27	OK=>BDII:0==>SRM:0	OK=>BDII:87==>SRM:87	OK=>BDII:60==>SRM:60	None
PIC	OK=>BDII:730==>SRM:730	OK=>BDII:578==>SRM:578	OK=>BDII:12==>SRM:12	OK=>BDII:0==>SRM:0	OK=>BDII:60==>SRM:60	OK=>BDII:43==>SRM:43	None
RAL	OK=>BDII:199==>SRM:3116	OK=>BDII:149==>SRM:1749	OK=>BDII:3116==>SRM:199	OK=>BDII:1749==>SRM:149	OK=>BDII:199==>SRM:199	OK=>BDII:71==>SRM:71	None
SARA	OK=>BDII:857==>SRM:857	OK=>BDII:845==>SRM:845	OK=>BDII:33==>SRM:33	OK=>BDII:0==>SRM:0	OK=>BDII:62==>SRM:62	OK=>BDII:51==>SRM:51	None

# Specific LHCb campaigns (II)

- Max CPU Time (45 tickets opened, 41 solved)
  - Not clear whether 999999999 means unlimited or means there is a misconfiguration/BDII problem in the site
    - Situation will improve with [https://ggus.eu/index.php?mode=ticket\\_info&ticket\\_id=97721](https://ggus.eu/index.php?mode=ticket_info&ticket_id=97721)
  - <https://wlcg-mon.cern.ch/dashboard/request.py/siteview#currentView=Glue+Validator+LHCb>

Site Name	MaxCPUTime (LHCb)	GGUS ticket for MaxCPUTime (LHCb)
LCG.ARAGRID-CIENCIAS.es	OK	None
LCG.AUVER.fr	Unreachable	None
LCG.BHAM-HEP.uk	OK	None
LCG.BIFI.es	OK	None
LCG.BMEGrid.hu	ERROR	101317
LCG.Barcelona.es	OK	None
LCG.Bologna.it	OK	None
LCG.Bristol.uk	OK	None
LCG.CBPF.br	Unreachable	None
LCG.CERN.ch	OK	None
LCG.CNAF-T2.it	OK	None

# ATLAS xrootd and http endpoints

- <https://wlcg-mon.cern.ch/dashboard/request.py/siteview#currentView=BDII+monitoring+for+ATLAS>
  - Understanding whether it could be useful to monitor this using the BDII
  - Some Storage info providers do not fully publish GLUE properly

Site Name	ATLAS GLUE 1 xroot endpoints	ATLAS GLUE 2 xroot endpoints	ATLAS GLUE 1 http endpoints	ATLAS GLUE 2 http endpoints
AEGIS07-IPB-ATLAS	Unreachable	Unreachable	Unreachable	Unreachable
AGLT2	OK	OK	OK	OK
AM-04-YERPHI	None	None	None	None
AMD64.PSNC.PL	Unreachable	Unreachable	Unreachable	Unreachable
Australia-ATLAS	None	OK	None	None
BEIJING-LCG2	OK	OK	None	None
BELLARMINE-ATLAS-T3	None	OK	None	None
BNL-ATLAS	OK	OK	OK	OK
BU_ATLAS_Tier2	Unreachable	OK	Unreachable	Unreachable
Brandeis-Atlas-T3	Unreachable	OK	Unreachable	Unreachable
CA-MCGILL-CLUMEQ-T2	None	None	None	None
CA-SCINET-T2	None	None	OK	OK
CA-VICTORIA-WESTGRID-T2	None	None	OK	OK
CERN-PROD	OK	None	None	None

# To and T1 accounting

- Studying the differences between WLCG accounting report, REBUS and BDII
- <https://wlcg-mon.cern.ch/dashboard/request.py/siteview#currentView=WLCG+Accounting+vs+BDII>

Site Name ▲	WLCG CPU Accounting ◇	WLCG CPU BDII ◇	WLCG CPU REBUS ◇	Total CPU REBUS ◇	WLCG Disk Accounting ◇	WLCG Disk BDII ◇	WLCG Disk REBUS ◇	Total Disk REBUS ◇	WLCG Tape Accounting ◇	WLCG Tape BDII ◇	WLCG Tape REBUS ◇	Total Tape REBUS ◇
ASGC	36165	44508	22254	44508	4275	3071	515	5536	4000	0	0	3906
BNL	74000	48817856	0.0	74000	11300	41729	23777	10742	10000	0	0	8398
CC-IN2P3	67350	161984	56456	161986	7000	20865	5739	8795	11025	212	131	0
CERN	356000	245298	177868	332144	29100	7733	16386	32541	67400	65223	36329	0
CNAF	88050	144067	92253	211864	10252	10887	6342	14296	15800	5774	4142	7388
FNAL	58000	15533952	0	58000	11000	14741	0	9765	22000	16380	0	21484
KIT	106585	150781	8582	0	11000	14809	6789	9279	14400	45164	28070	45384
KR-KISTI-GSDC	14500	28054	0	28055	990	0	0	0	1040	0	0	0
NDGF	30900	24791	10231	342084	5129	9108	1339	6266	5464	5335	2406	5335
NL-T1	47296	99973	31153	128814	5362	8133	2355	7081	5593	0	0	0
RAL	90246	142397	0.0	96435	9667	9661	4567	10194	12122	10602	5592	11413
TRIUMF	70226	102413	0	104506	6420	10788	0	6330	5300	5190	0	5190
pic	31143	39004	18328	46994	3850	8050	0	0	5887	5977	0	0

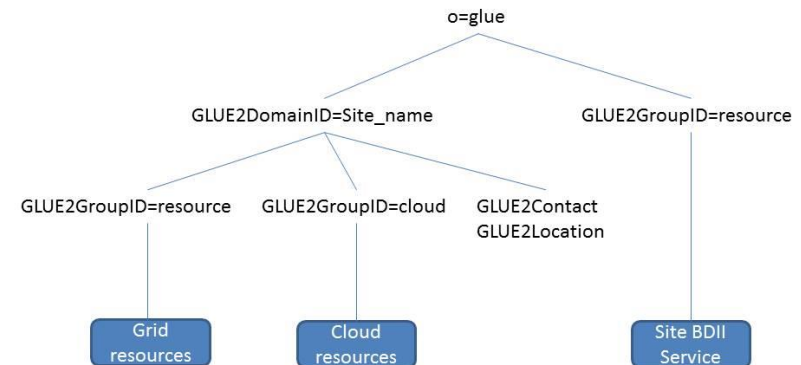
# Storage Information Providers

- Meeting with Storage developers scheduled on 15<sup>th</sup> April
  - <https://indico.cern.ch/event/311528/>
- Meet and discuss common areas of interest related to the Information System
  - Align the way things are published
    - Also agree on issues related to the OGF GLUE WG
  - Improve the quality of the data
    - Storage capacities

# Cloud resources in the BDII

- After discussions with EGI, the LDAP tree of the BDII includes now cloud resources
  - Maintains backwards compatibility
  - Some EGI sites are already being published:
    - 100IT, PRISMA-INFN-BARI, INFN-CATANIA-STACK, INFN-CATANIA-NEBULA
    - They publish cloud resources using the existing GLUE 2.0 schema
  - There are ongoing discussions at the OGF GLUE WG to express cloud resources in GLUE 2
    - i.e. Extending GLUE 2.0 and creating GLUE 2.1

Glue 2.0 DIT in site BDII



# ginfo, AGIS and GSR

- ginfo
  - Reminder: GLUE 2 Replacement of lcg-info and lcg-infosites
  - New version under preparation
    - It is able to combine entries of different GLUE 2 objects
    - Similar approach to lcg-info and lcg-infosites
- AGIS
  - IT-SDC has proposed to CMS an investigation of adopting AGIS as a common information system and a feasibility study is ongoing.
  - The feasibility study and the prototype will be presented within CMS in the upcoming weeks to gather feedback
- GSR
  - It has been put on hold while evaluating AGIS feasibility
    - AGIS has the advantage of allowing VOs to define their own topology of grid and cloud resources