



NGS

NGS BDII Service

NGS Surgery

25th Aug 2010



Connecting Infrastructure

Connecting Research



NGS

Summary

- What is the BDII Service ?
- How does it interact with my site ?
- What was wrong with the old service ?
- New BDII service and “washing”
- How / where is the data used.
- Questions

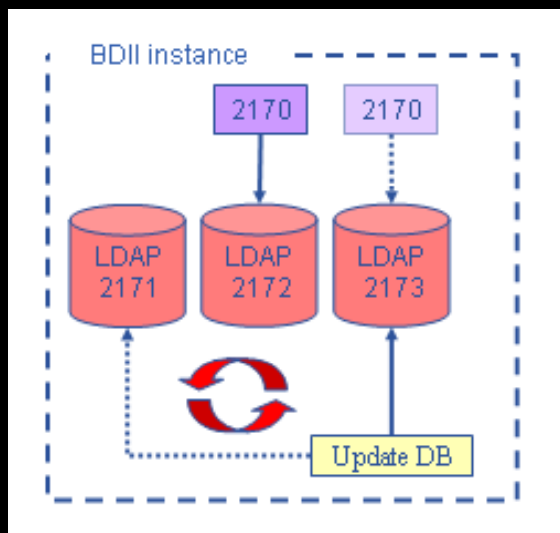


NGS

What is the BDII ?

- Berkeley Database Information Index
- Central resource for machine GRIS/GIIS info:
 - Email addresses, Physical location, ...
 - OS, CPUs, Mem, LSF/PBS, #running, #waiting Jobs, ...
 - VO's , URI's, Policies(MaxWallClock,...)
- LDAP database(s)
 - std SL/RHEL slapd daemon
 - 'Glue' schema hooked into a "Top level" schema
- **Glue** (Grid Laboratory for a Uniform Environment)
 - v1.3 (2007) EGEE, NGS but v2.0 supported by latest BDII v5.
 - Authors inc. Jens Jenson, Matt Viljoen, Steven Burke, CERN, RAL, INFN
 - Now an OGF Project

Berkeley Database Information Index \leq V4

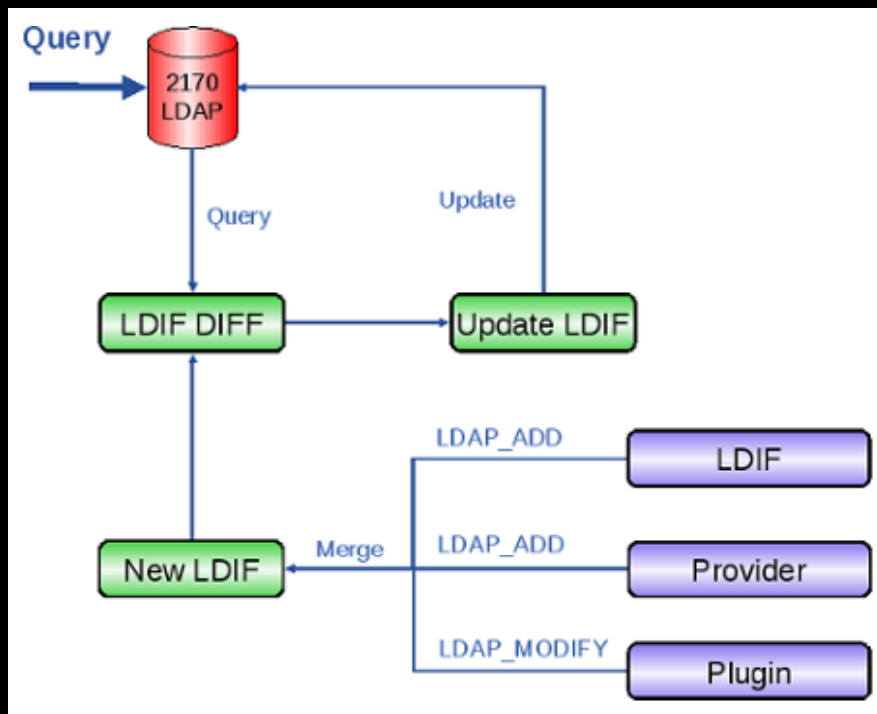


- Eg ngsinfo.grid-support.ac.uk
- Three or more DB's
- Port Forwarding to latest DB
- Inactive one(s) rebuilt with new data.
- ngsinfo.. on old port 2135



NGS

Berkeley Database Information Index V5

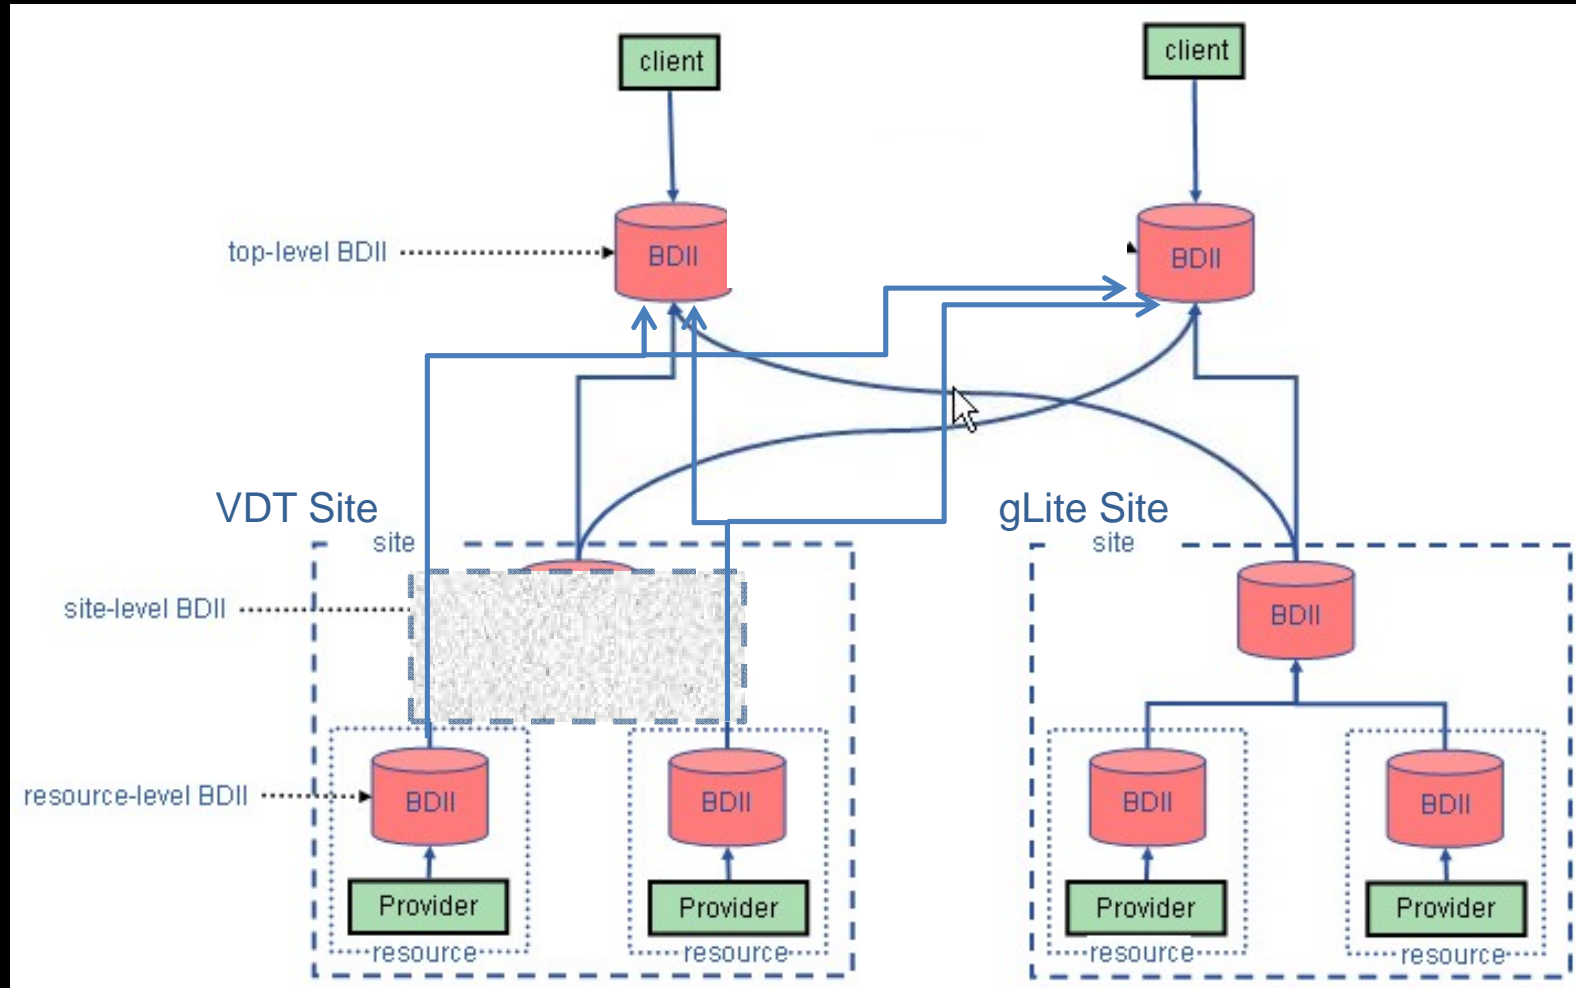


- Eg. bdii.ngs.ac.uk
- One DB
- Updates add, mod, delete
- No port forwarding



NGS

How does it interact with my site ?



Connecting Infrastructure

Connecting Research



NGS

BDII DB Structure (old)

LDAP Browser/Editor v2.8.2 - [ldap://ngsinfo.grid-support.ac.uk:2135/o=grid]

File Edit View LDIF Help

o=grid

- Mds-Vo-name=ngsinfo
 - Mds-Vo-name=O8C
 - Mds-Vo-name=HECToR
 - Mds-Vo-name=RAL-LCG2
 - Mds-Vo-name=ngs-york
 - Mds-Vo-name=ral-ngs2
 - GlueSiteUniqueID=ngs.rl.ac.uk
 - GlueClusterUniqueID=ngs.rl.ac.uk
 - GlueCEUniqueID=ngs.rl.ac.uk:2119/jobmanager-lsf-ngs
 - GlueVOViewLocalID=ops
 - GlueVOViewLocalID=dteam
 - GlueVOViewLocalID=mott2.org
 - GlueVOViewLocalID=ngs.ac.uk
 - GlueVOViewLocalID=gin.ggf.org
 - GlueVOViewLocalID=skatelescope.org
 - GlueCEUniqueID=ngs.rl.ac.uk:2119/jobmanager-lsf-mott
 - GlueCESEIRingGroupCEUniqueID=ngs.rl.ac.uk:2119/jobmanager-lsf-mott
 - Mds-Vo-name=keele-ngs
 - Mds-Vo-name=ral-scarf
 - Mds-Vo-name=leeds-ngs2
 - Mds-Vo-name=NGS-Cardiff
 - Mds-Vo-name=bristol-ngs
 - Mds-Vo-name=compservgl
 - Mds-Vo-name=oxford-ngs2
 - Mds-Vo-name=reading-ngs
 - Mds-Vo-name=UKI-LT2-QMUL
 - Mds-Vo-name=UKI-LT2-RHUL
 - Mds-Vo-name=edinburgh-ngs
 - Mds-Vo-name=lancs1-nwgrid
 - Mds-Vo-name=UKI-LT2-Brunel
 - Mds-Vo-name=uki-lt2-ic-hep
 - Mds-Vo-name=manchester-ngs2
 - Mds-Vo-name=westminster-ngs2
 - Mds-Vo-name=UKI-SCOTGRID-DURHAM
 - Mds-Vo-name=uki-scotgrid-glasgow
 - Mds-Vo-name=uki-southgrid-ox-hep
 - Mds-Vo-name=UniversityOfSheffield
 - Mds-Vo-name=uki-northgrid-liv-hep
 - Mds-Vo-name=uki-northgrid-man-hep

Top Level Schema

Glue Schema

Attribute	Value
objectClass	GlueCE
objectClass	GlueCEAccessControlBase
objectClass	GlueCEInfo
objectClass	GlueCEPolicy
objectClass	GlueCEState
objectClass	GlueCETop
objectClass	GlueInformationService
objectClass	GlueKey
objectClass	GlueSchemaVersion
GlueSchemaVersionMinor	3
GlueSchemaVersionMajor	1
GlueInformationServiceURL	ldap://ngsinfo.grid-support.ac.uk:2135/mds-vo-name=ral-ngs2,mds-vo-name=ngsinfo,o=grid
GlueOriginKey	GlueClusterUniqueID=ngs.rl.ac.uk
GlueOriginKey	ngs.rl.ac.uk:2119/jobmanager-lsf-ngs
GlueCEStateWorstResponseTime	0
GlueCEStateWaitingJobs	43
GlueCEStateTotalJobs	468
GlueCEStateStatus	Production
GlueCEStateRunningJobs	425
GlueCEStateFreeJobSlots	4
GlueCEStateFreeCPUs	4
GlueCEStateEstimatedResponseTime	0
GlueCEPolicyPriority	1
GlueCEPolicyPreemption	0
GlueCEPolicyMaxWallClockTime	6000
GlueCEPolicyMaxWaitingJobs	0
GlueCEPolicyMaxTotalJobs	0
GlueCEPolicyMaxSlotsPerJob	0
GlueCEPolicyMaxRunningJobs	0
GlueCEPolicyMaxObtainableWallClockTime	0
GlueCEPolicyMaxObtainableCPUTime	0
GlueCEPolicyMaxCPUTime	6000
GlueCEPolicyAssignedJobSlots	429
GlueCEName	ngs
GlueCEInfoTotalCPUs	429
GlueCEInfoLRMSVersion	not defined
GlueCEInfoLRMSType	lsf
GlueCEInfoJobManager	lsf
GlueCEInfoHostName	ngs.rl.ac.uk
GlueCEInfoGatekeeperPort	2119
GlueCEInfoDefaultSE	ngs.rl.ac.uk
GlueCEInfoDataDir	UNAVAILABLE
GlueCEInfoContactString	ngs.rl.ac.uk:2119/jobmanager-lsf-ngs
GlueCEInfoApplicationDir	UNAVAILABLE
GlueCEImplementationVersion	UNDEFINED
GlueCEImplementationName	UNDEFINED
GlueCEHostingCluster	ngs.rl.ac.uk
GlueCECapability	UNDEFINED
GlueCEAccessControlBaseRule	VO:dteam
GlueCEAccessControlBaseRule	VO:gin.ggf.org
GlueCEAccessControlBaseRule	VO:mott2.org
GlueCEAccessControlBaseRule	VO:ngs.ac.uk
GlueCEAccessControlBaseRule	VO:ops
GlueCEAccessControlBaseRule	VO:skatelescope.org

Ready.



NGS

What was wrong with the old service ?

- Old: Pre V4, No updates, hand 'patched'
- Dynamic updates slow for WMS (v4 vs v5)
- But mainly because:
- EGEE SAM tests (WMS submit) fail on NGS.
 - Why ? NGS UI/WMS uses O.O.T.B. gLite WMS
 - Because ...
EGEE top level BDII's reject VDT site BDII data
 - Because ...
VDT "Top Level" schema different to gLite ...



NGS

VDT GRIS

LDAP Browser/Editor v2.8.2 - [ldap://ngs.rl.ac.uk:2135/Mds-Vo-name=local,o=Grid]

File Edit View LDIF Help

Attribute Value
objectClass GlobusStub

Mds-Vo-name=local,o=Grid

- GlueCEUniqueID=ngs.rl.ac.uk:2119Vjobmanager-lsf
- GlueVOViewLocalID=ngs.ac.uk
- GlueVOViewLocalID=ops
- GlueVOViewLocalID=dteam
- GlueVOViewLocalID=gin.ggf.org
- GlueVOViewLocalID=mott2.org
- GlueVOViewLocalID=skatelescope.org
- GlueCEUniqueID=ngs.rl.ac.uk:2119Vjobmanager-lsf
- GlueCESEBindGroupCEUniqueID=ngs.rl.ac.uk:2119
- GlueClusterUniqueID=ngs.rl.ac.uk
- GlueSiteUniqueID=ngs.rl.ac.uk

Ready.

Connecting Infrastructure

Connecting Research



NGS

gLite GRIS

LDAP Browser/Editor v2.8.2 - [ldap://lcgce05.gridpp.rl.ac.uk:2170/o=grid]

File Edit View LDIF Help

o=grid

- Mds-Vo-name=local
- Mds-Vo-name=resource**
 - GlueClusterUniqueID=lcgce05.gridpp.rl.ac.uk
 - GlueCESEBindGroupCEUniqueID=lcgce05.gridpp.rl.ac.uk
 - GlueServiceUniqueID=lcgce05.gridpp.rl.ac.uk_org.glit
 - GlueCESEBindGroupCEUniqueID=lcgce05.gridpp.rl.ac.uk
 - GlueServiceUniqueID=lcgce05.gridpp.rl.ac.uk_org.glit
 - GlueCESEBindGroupCEUniqueID=lcgce05.gridpp.rl.ac.uk
 - GlueCESEBindGroupCEUniqueID=lcgce05.gridpp.rl.ac.uk
 - GlueCESEBindGroupCEUniqueID=lcgce05.gridpp.rl.ac.uk
 - GlueCESEBindGroupCEUniqueID=lcgce05.gridpp.rl.ac.uk
 - GlueClusterUniqueID=sl5-700.gridpp.rl.ac.uk
 - GlueClusterUniqueID=sl5-500.gridpp.rl.ac.uk
 - GlueClusterUniqueID=sl5-2000.gridpp.rl.ac.uk
 - GlueClusterUniqueID=sl5-3000.gridpp.rl.ac.uk
 - GlueClusterUniqueID=sl5-1000.gridpp.rl.ac.uk
 - GlueClusterUniqueID=sl5-4000.gridpp.rl.ac.uk
 - GlueCESEBindGroupCEUniqueID=lcgce05.gridpp.rl.ac.uk
 - GlueCEUniqueID=lcgce05.gridpp.rl.ac.uk:8443Vcream
 - GlueCEUniqueID=lcgce05.gridpp.rl.ac.uk:8443Vcream

Attribute	Value
objectClass	Mds
Mds-Vo-name	resource

Ready. A

Connecting Infrastructure

Connecting Research



NGS

- gLite BDII has schema objectClass definitions for:
 - ‘organization’ & ‘mds’
- No definitions for VDT classes:
 - ‘GlueTop’ & ‘GlobusStub’
- ObjectClasses define LDAP schema data types.
 - E.g. Attribute names and datatypes; String, Integer
- Schema checking rejects unknown objectClasses.
- Adding objectClasses to NGS BDII propagates problem.
- The ‘washing’ service fixes VDT GRIS objectClasses.



NGS

The 'washing' service

- gLite BDII flavours : site , top level
 - site BDII has 'FIX GLUE' config option 😊
 - VDT Sites having a gLite site BDII. Job done ! 😊
 - Site BDII's all resources under 1 Mds-Vo-Name
 - Need 1 Mds-Vo-Name for each 'washed' site.
 - Washing service =
 - hacked site BDII with each resource a separate Mds-Vo-Name
 - NGS Top Level BDII Queries:
 - Site BDII's or the Washing Service



NGS

'Washing Service'

LDAP Browser/Editor v2.8.2 - [ldap://sitebdii.ngs.rl.ac.uk:2170/o=grid]

File Edit View LDIF Help

Attribute	Value
objectClass	organization
o	grid

Ready.

LDAP Browser/Editor v2.8.2 - [ldap://sitebdii.ngs.rl.ac.uk:2170/o=grid]

File Edit View LDIF Help

Attribute	Value
objectClass	Mds
Mds-Vo-name	ral-ngs2

Ready.

Connecting Infrastructure

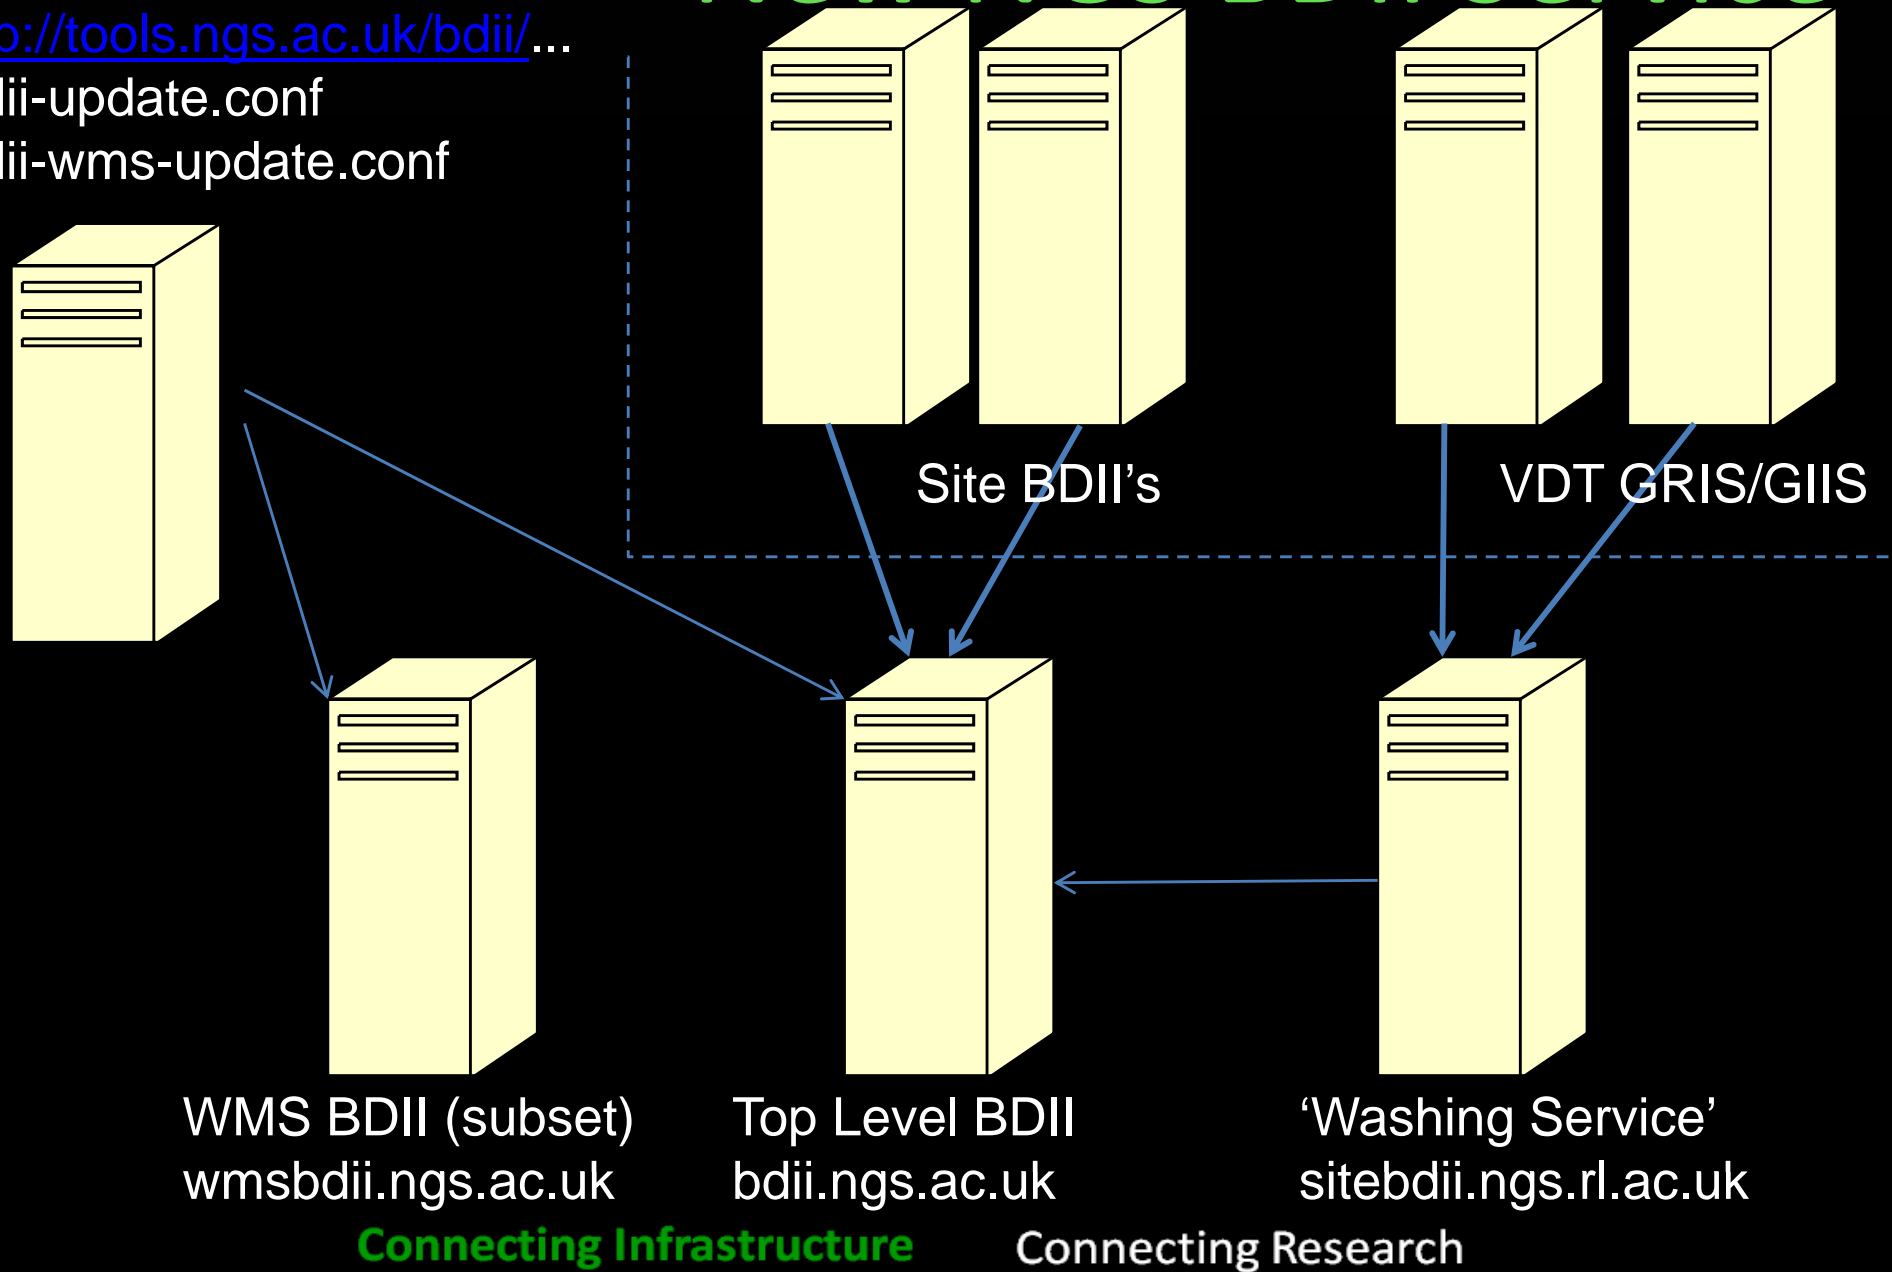
Connecting Research



NGS

[http://tools.ngs.ac.uk/bdii/...](http://tools.ngs.ac.uk/bdii/)
/bdii-update.conf
/bdii-wms-update.conf

New NGS BDII Service





NGS

BDII DB Structure (new)

LDAP Browser/Editor v2.8.2 - [ldap://bdii.ngs.ac.uk:2170/o=grid]

File Edit View LDIF Help

o=grid

- Mds-Vo-name=local
- Mds-Vo-name=ral-icg2
- Mds-Vo-name=uki-scotgrid-durham
- Mds-Vo-name=uki-northgrid-liv-hep
- Mds-Vo-name=osc
- Mds-Vo-name=hector
- Mds-Vo-name=ngs-york
- Mds-Vo-name=ral-ngs2
 - GlueSiteUniqueID=ngs.rl.ac.uk
 - GlueClusterUniqueID=ngs.rl.ac.uk
 - GlueCEUniqueID=ngs.rl.ac.uk:2119/jobmanager-Isf-ngs
 - GlueVOViewLocalID=ops
 - GlueVOViewLocalID=dteam
 - GlueVOViewLocalID=mott2.org
 - GlueVOViewLocalID=ngs.ac.uk
 - GlueVOViewLocalID=gin.ggf.org
 - GlueVOViewLocalID=skatelescope.org
 - GlueCEUniqueID=ngs.rl.ac.uk:2119/jobmanager-Isf-mott
 - GlueCESEBindGroupCEUniqueID=ngs.rl.ac.uk:2119/jobmanager-Isf-mott
- Mds-Vo-name=ral-scarf
- Mds-Vo-name=keele-ngs
- Mds-Vo-name=leeds-ngs2
- Mds-Vo-name=bristol-ngs
- Mds-Vo-name=oxford-ngs2
- Mds-Vo-name=compservgla
- Mds-Vo-name=reading-ngs
- Mds-Vo-name=ngs-cardiff
- Mds-Vo-name=lancs1-nwgrid
- Mds-Vo-name=edinburgh-ngs
- Mds-Vo-name=manchester-ngs2
- Mds-Vo-name=westminster-ngs2
- Mds-Vo-name=universityofsheffield
- Mds-Vo-name=uki-It2-ic-hep
- Mds-Vo-name=uki-scotgrid-glasgow

Top Level Schema

Glue Schema

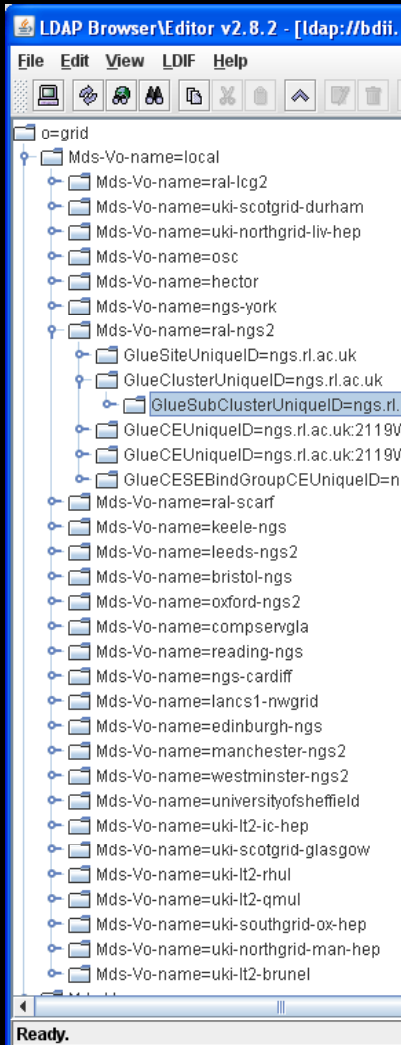
Attribute	Value
objectClass	GlueCEState
objectClass	GlueInformationService
objectClass	GlueCE
objectClass	GlueCETop
objectClass	GlueKey
objectClass	GlueCEAccessControlBase
objectClass	GlueCEInfo
objectClass	GlueCEPolicy
objectClass	GlueSchemaVersion
GlueSchemaVersionMinor	3
GlueSchemaVersionMajor	1
GlueInformationServiceURL	ldap://ngsinfo.grid-support.ac.uk:2135/mds-vo-name=ral-ngs2,mc
GlueForeignKey	GlueClusterUniqueID=ngs.rl.ac.uk
GlueCEUniqueID	ngs.rl.ac.uk:2119/jobmanager-Isf-ngs
GlueCEStateWorstResponseTime	0
GlueCEStateWaitingJobs	0
GlueCEStateTotalJobs	149
GlueCEStateStatus	Production
GlueCEStateRunningJobs	149
GlueCEStateFreeJobSlots	280
GlueCEStateFreeCPUs	280
GlueCEStateEstimatedResponseTime	0
GlueCEPolicyPriority	1
GlueCEPolicyPreemption	0
GlueCEPolicyMaxWallClockTime	6000
GlueCEPolicyMaxWaitingJobs	0
GlueCEPolicyMaxTotalJobs	0
GlueCEPolicyMaxSlotsPerJob	0
GlueCEPolicyMaxRunningJobs	0
GlueCEPolicyMaxObtainableWallClockTime	0
GlueCEPolicyMaxObtainableCPUTime	0
GlueCEPolicyMaxCPUTime	6000
GlueCEPolicyAssignedJobSlots	429
GlueCEName	ngs
GlueCEInfoTotalCPUs	429
GlueCEInfoLRMSVersion	not defined
GlueCEInfoLRMSType	Isf
GlueCEInfoJobManager	Isf
GlueCEInfoHostName	ngs.rl.ac.uk

Ready.



NGS

Using the BDII: (WMS)



```
Requirements = (
# Select any site which has any version of DLPOLY2 installed
  member( "NGS-UEE-DLPOLY2", other.GlueHostApplicationSoftwareRunTimeEnvironment)
);
```

```
Requirements = (
# Select the site ngs.rl.ac.uk, batch queue "ngs", only if it has DLPOLY v2.20
  member( "NGS-UEE-DLPOLY2_2_20", other.GlueHostApplicationSoftwareRunTimeEnvironment)
&& other.GlueCEUniqueID == "ngs.rl.ac.uk:2119/jobmanager-lsf-ngs"
);
```

```
Requirements = (
# Submit to any site with the string "ngs" in the batch queue string
  RegExp( "ngs", other.GlueCEUniqueID)
);
# Rank the sites matching "Requirements" by...
Rank = other.GlueCEStateFreeCPUs;
# Rank = other.GlueCEInfoTotalCPUs;
```

- JDL 'Requirements'
- WMS matching
- WMS Site Status

GlueHostOperatingSystemName	linux-rhel-4
GlueHostApplicationSoftwareRunTimeEnvironment	MPICH
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE_1_0
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-ABAUUS
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-ABAUUS_6_6_1
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-ABAUUS_6_7_1
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-ABAUUS_6_8
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-ABAUUS_6_9
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AMBER
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AMBER_10_0
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AMBER_9_0
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AUTODOCK
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AUTODOCK_4_0_1
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AUTODOCK_4_2_2_1
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AUTOGRID
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AUTOGRID_4_0_0
GlueHostApplicationSoftwareRunTimeEnvironment	NGS-UEE-AUTOGRID 4 2 2 1



NGS

Using the BDII: (CLI)

```
ldapsearch -x -h bdii.ngs.ac.uk -p 2170 -LLL -b "mds-vo-  
name=local,o=grid"
```

```
dn: Mds-Vo-name=local,o=grid
```

```
objectClass: Mds
```

```
Mds-Vo-name: local
```

```
dn: Mds-Vo-name=ral-lcg2,Mds-Vo-name=local,o=grid
```

```
objectClass: Mds
```

```
Mds-Vo-name: ral-lcg2
```

```
dn: Mds-Vo-name=uki-scotgrid-durham,Mds-Vo-name=local,o=grid
```

```
objectClass: Mds
```

```
Mds-Vo-name: uki-scotgrid-durham
```

```
dn: Mds-Vo-name=uki-northgrid-liv-hep,Mds-Vo-name=local,o=grid
```

```
objectClass: Mds
```

```
Mds-Vo-name: uki-northgrid-liv-hep
```

```
dn: GlueSiteUniqueID=RAL-LCG2,Mds-Vo-name=RAL-LCG2,Mds-Vo-name=local,o=grid
```

```
GlueSiteDescription: Rutherford Appleton Laboratory
```

```
GlueSiteSecurityContact: mailto:security@gridpp.rl.ac.uk
```

```
GlueSiteSponsor: none
```

```
objectClass: GlueTop
```

```
objectClass: GlueSite
```



NGS

Using the BDII (web)

The screenshot shows the NGS Scientific Software Applications website. The browser window title is "NGS Scientific Software Applications | NGS - Mozilla Firefox". The address bar shows "http://www.ngs.ac.uk/applications". The page features a green header with the NGS logo and navigation links: Home, About Us, Use the NGS, Learn More, Innovation, Member Sites, Help. Below the header, there's a search bar and a "Use the NGS" button. The main content area is titled "NGS Scientific Software Applications" and includes a list of application categories: AstroPhysics, BioMedical, BioInformatics, Chemistry, Data Analysis, Engineering, and Image Analysis. There are also several small images representing different scientific data. Below this, there's a section for "Application Categories" with links to various categories like AMBER, Autodock, Autogrid, CASTED, etc. A "Sites" section lists various computing centers like Glasgow Scotgrid, Leeds, Manchester, Oxford, etc. On the right side, there's a "Use the NGS" sidebar with sections for Overview, How to Join, Your Certificate, Your Account, User Tools, Applications, Database Service, How do I...?, Support, Communities, and Terms and Conditions. At the bottom of the browser window, there's a status bar showing "18 hosts down, 6 unknown services, 6 service warnings, 121 critical services".

- PHP LDAP Libraries



NGS

Summary

- What is the BDII Service ?
- How does it interact with my site ?
- What was wrong with the old service ?
- New BDII service and “washing”
- How / where is the data used.



NGS

Questions ?

Connecting Infrastructure

Connecting Research