



Enabling Grids for E-scienceE

# Using GStat 2.0 for Information Validation

*Joanna Huang (ASGC, Taiwan)*

*EGEE'09 Conference,  
Barcelona, 24<sup>th</sup> September 2009*

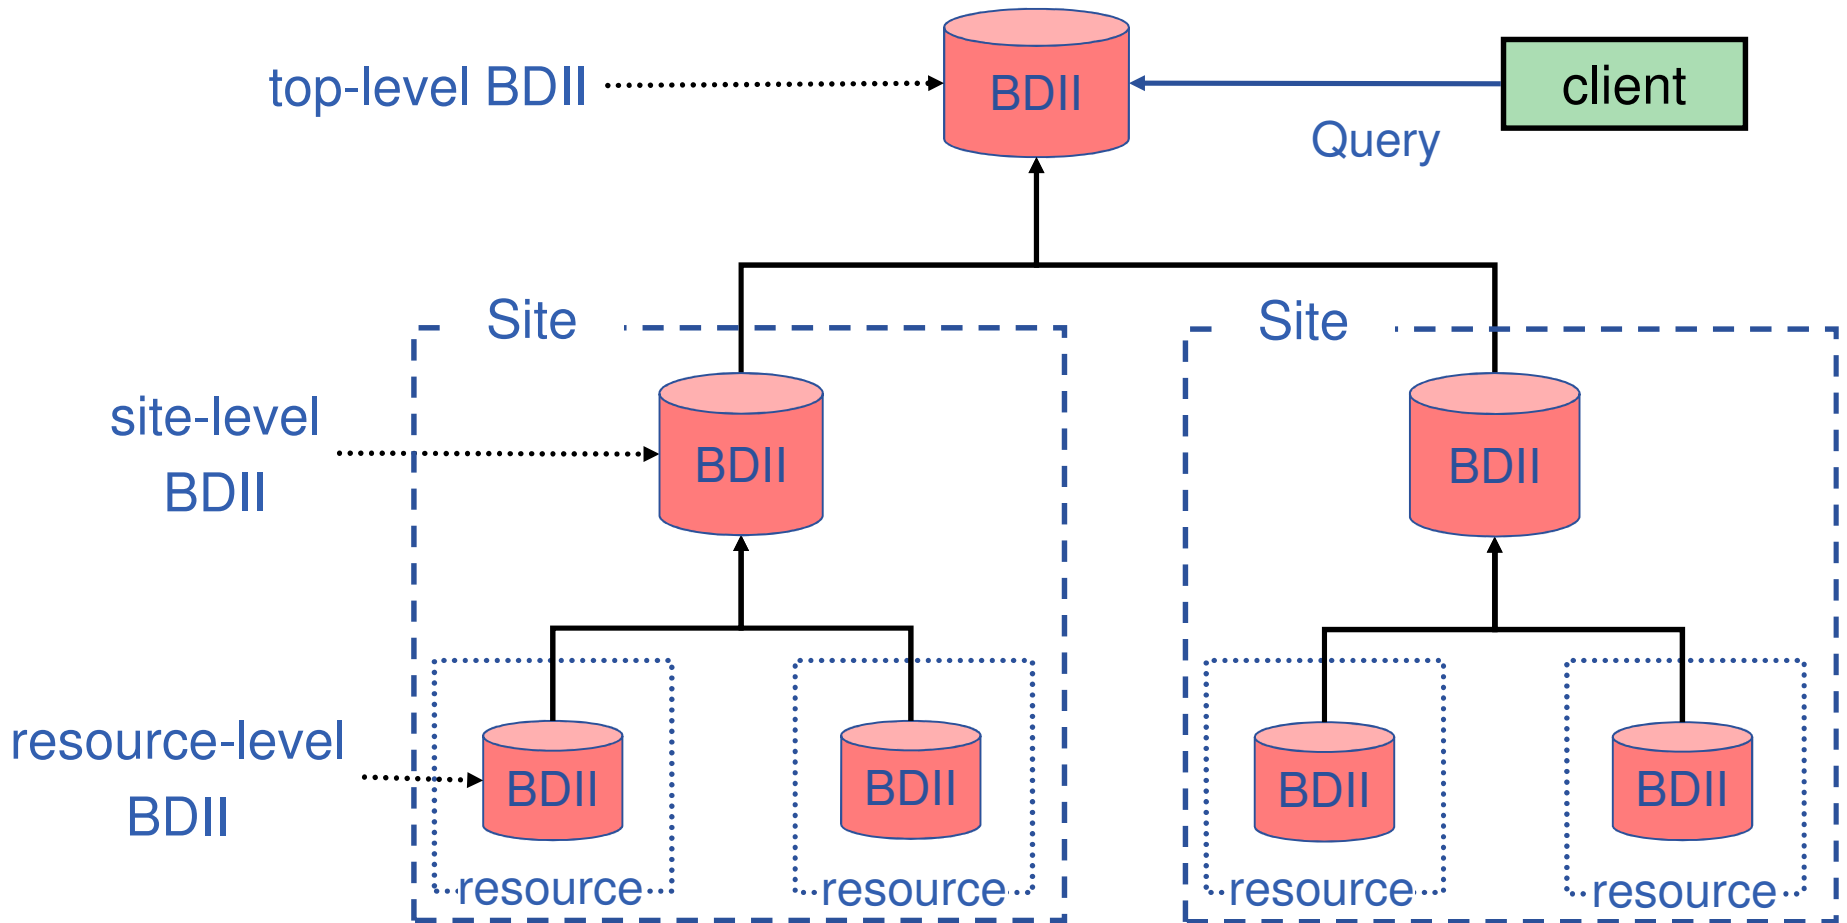
[www.eu-egee.org](http://www.eu-egee.org)



- **Grid Information Systems**
- **Typical Operational Problems**
- **The Approach Taken for Improving Quality**
- **Deployment Scenarios**
- **GStat/Nagios Screen Shots**
- **Summary**

- **A fundamental aspect of a Grid infrastructure**
  - Enable users, applications and services to
    - **discover** which services exist
    - further information about their structure and state
- **Key Concepts**
  - Information is described to conform with an **information model**
    - Ensures agreement between the service and the consumer
  - The Grid service itself is the **information source**
    - Information obtained by interrogating the backend system
    - Exposed via parallel information service
  - Information **aggregated** at various level
    - To improve performance and scalability

- **Berkley Database Information Index (BDII)**
  - LDAP based information aggregator
  - Uses the **Glue Schema** Information model
  - 1200+ @resource-level, 280+ @ site-level, 70+ @ top-level



- **Information System Infrastructure**
  - Is the information system healthy?
  - Are the BDIs functioning correctly?
- **Information System Content**
  - Is the information content correct?

- **How do we know that the infrastructure is functioning?**
  - Need to check each individual component (BDII).
- **How do we know which BDII exist in the infrastructure?**
  - Need to query the information system
  - And double check with other sources eg GOC DB.
- **How do we know the BDII is functioning?**
  - Do some checks based on the functions of BDII
    - Can we do a simple LDAP query?
    - Is the response time with acceptable parameters?
    - Is the information recent?
    - Are the number of entries published stable over time?

- **How to obtain the information content?**
  - By querying a BDII
- **How to ensure the information integrity and quality?**
  - Do different checks based on the different entries
    - Does the information agree with the **schema**?
    - Are the entities **we expect** to see published?
    - Are there things published that **we don't expect**?
    - Is the information logically **self consistent**?
      - *Is the number of free CPUs greater than the TotalCPUs?*
    - Is there agreement with **external information sources**?
      - *Is the host registered in DNS?*
    - Is there conformance with **extra project constraints**?
      - *Valid gLite version?*
- ***...How should the above be implemented?***

- **What is Nagios?**
  - An open source application designed as a extensible framework for **monitoring**, **scheduling** and **alerting**
- **Why we use Nagios?**
  - **Manage** the execution of tests and results
  - Probes can be **re-used** by other OAT applications
  - Probes can also be run on the **command line**
- **How we use Nagios?**
  - Nagios config script
    - **Automatically configures Nagios** to monitor BDIIs found in the infrastructure
  - Nagios probes
    - **BDII service monitoring** tests
    - Information **content validation** tests
    - Useful by many different actors
      - *System Administrators, Grid Operations, Other Tools etc.*



### Current Network Status

Last Updated: Wed Sep 16 18:08:49 CEST 2009  
 Updated every 90 seconds  
 Nagios® Core™ 3.2.0 - [www.nagios.org](http://www.nagios.org)  
 Logged in as /C=TW/O=AS/OU=GRID/CN=Hui-Tzu  
 Huang 155651

- [View History For all hosts](#)
- [View Notifications For All Hosts](#)
- [View Host Status Detail For All Hosts](#)

### Host Status Totals

Up	Down	Unreachable	Pending
249	75	0	0
All Problems		All Types	
75		324	

### Service Status Totals

Ok	Warning	Unknown	Critical	Pending
1057	4	0	918	0
All Problems		All Types		
922		1979		

### Service Status Details For All Hosts

Host ↑↓	Service ↑↓	Status ↑↓	Last Check ↑↓	Duration ↑↓	Attempt ↑↓	Status Information
<a href="#">agh1.atlas.unimelb.edu.au</a>	<a href="#">check-bdii-freshness</a>	OK	09-16-2009 17:48:59	2d 9h 19m 51s	1/2	OK: freshness=93s, entries=1
	<a href="#">check-bdii-sites</a>	OK	09-16-2009 18:04:11	3d 20h 34m 37s	1/2	OK: time=2s, entries=310
<a href="#">agh6.atlas.unimelb.edu.au</a>	<a href="#">check-bdii-freshness</a>	OK	09-16-2009 17:55:08	2d 8h 43m 42s	1/2	OK: freshness=10s, entries=1
	<a href="#">check-bdii-services</a>	OK	09-16-2009 17:53:04	2d 9h 15m 44s	1/2	OK: time=1s, entries=6
	<a href="#">check-ce</a>	CRITICAL	09-16-2009 17:52:22	2d 10h 21m 26s	2/2	CRITICAL - errors 21, warnings 33, info 0
	<a href="#">check-sanity</a>	OK	09-16-2009 18:06:33	91d 1h 42m 17s	1/2	OK - errors 0, warnings 0, info 0
	<a href="#">check-se</a>	CRITICAL	09-16-2009 17:45:59	2d 9h 27m 49s	2/2	CRITICAL - errors 13, warnings 22, info 0
	<a href="#">check-service</a>	CRITICAL	09-16-2009 17:55:18	2d 8h 43m 30s	2/2	CRITICAL - errors 1, warnings 0, info 0
	<a href="#">check-site</a>	CRITICAL	09-16-2009 17:44:48	2d 9h 29m 0s	2/2	CRITICAL - errors 3, warnings 0, info 0

## Service State Information

Current Status: **CRITICAL** (for 2d 11h 14m 29s)

Status Information: CRITICAL - errors 21, warnings 33, info 0

ERROR: agh2.atlas.unimelb.edu.au:2119/jobmanager-icgpbs-dteam, GlueCEPolicyAssignedJobSlots has negative or null value.

ERROR: agh2.atlas.unimelb.edu.au:2119/jobmanager-icgpbs-belle, GlueCEPolicyAssignedJobSlots has negative or null value,

ERROR: agh2.atlas.unimelb.edu.au:2119/jobmanager-icgpbs-atlas, GlueCEPolicyAssignedJobSlots has negative or null value,

ERROR: agh2.atlas.unimelb.edu.au, Manditory attribute does not exist in the GlueSubCluster entry., Glue attribute GlueHostProcessorOtherDescription does not exist.

ERROR:  
GlueSubClusterUniqueID=agh2.atlas.unimelb.edu.au,GlueClusterUniqueID=agh2.atlas.unimelb.edu.au,MdsVo-name=Australia-ATLAS,o=grid, A value is incorrrect, GlueSubClusterPhysicalCPUs must be a valid integer

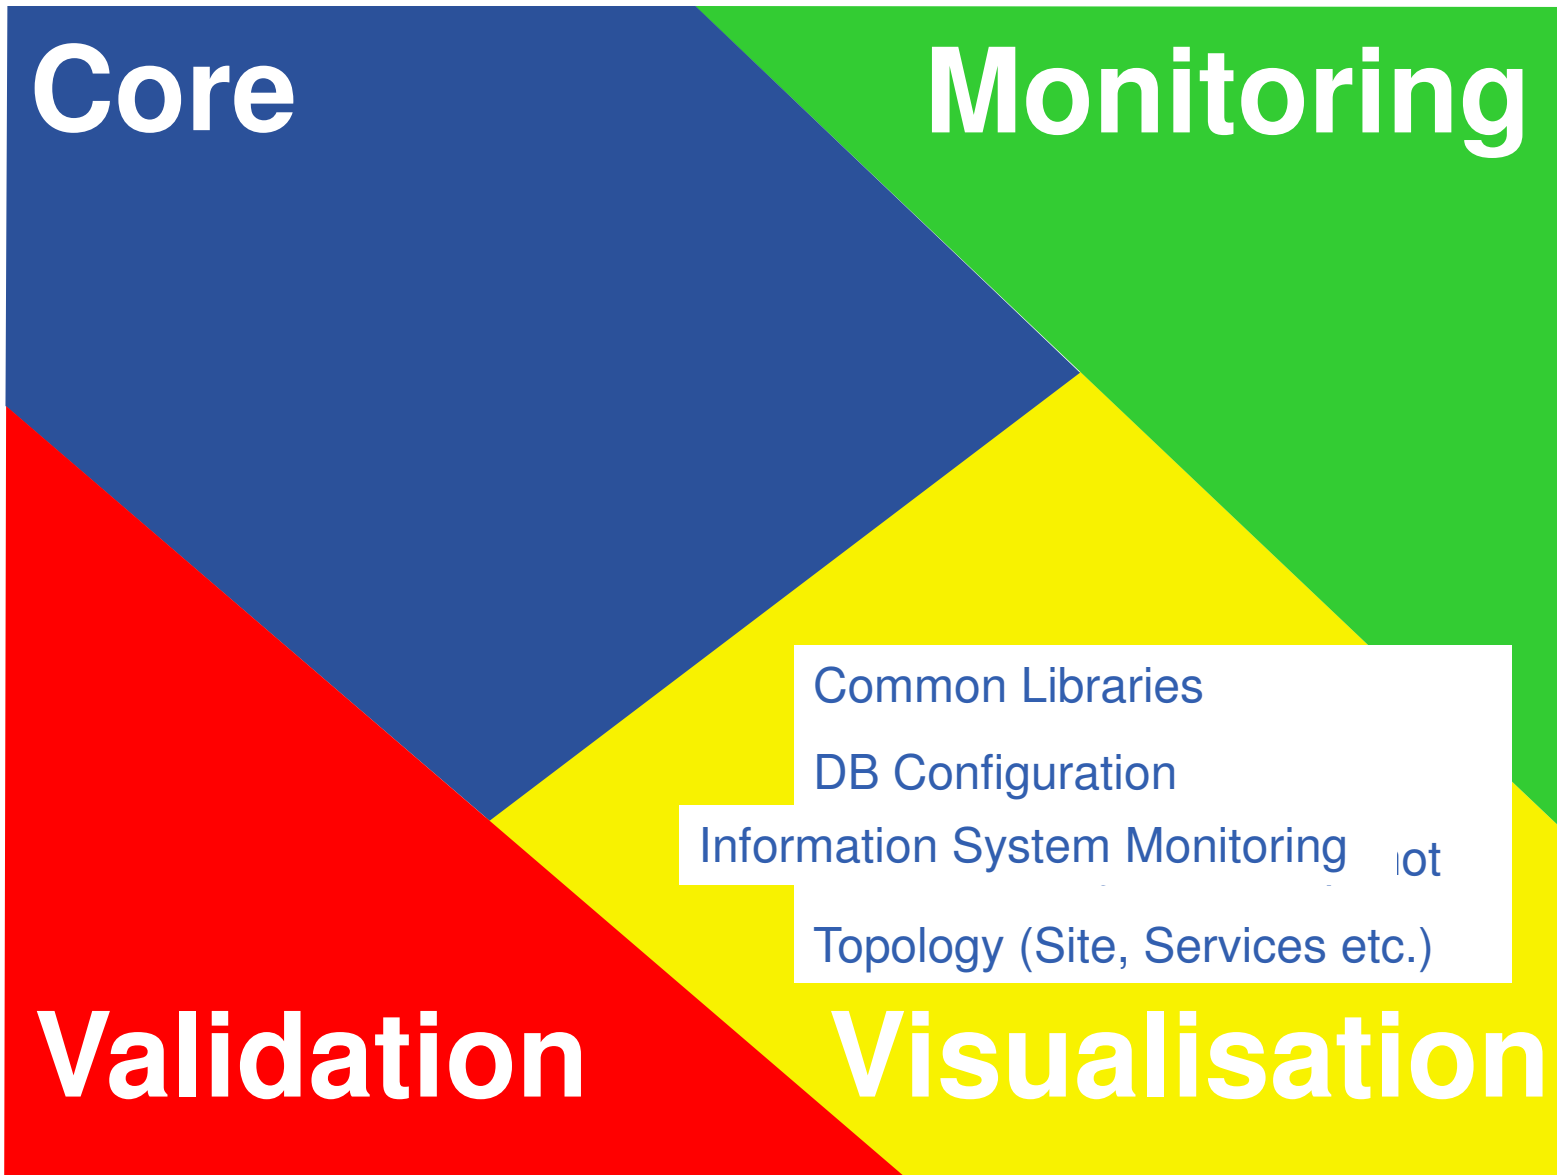
ERROR:  
GlueSubClusterUniqueID=agh2.atlas.unimelb.edu.au,GlueClusterUniqueID=agh2.atlas.unimelb.edu.au,MdsVo-name=Australia-ATLAS,o=grid, A value is incorrrect, GlueSubClusterLogicalCPUs must be a valid integer

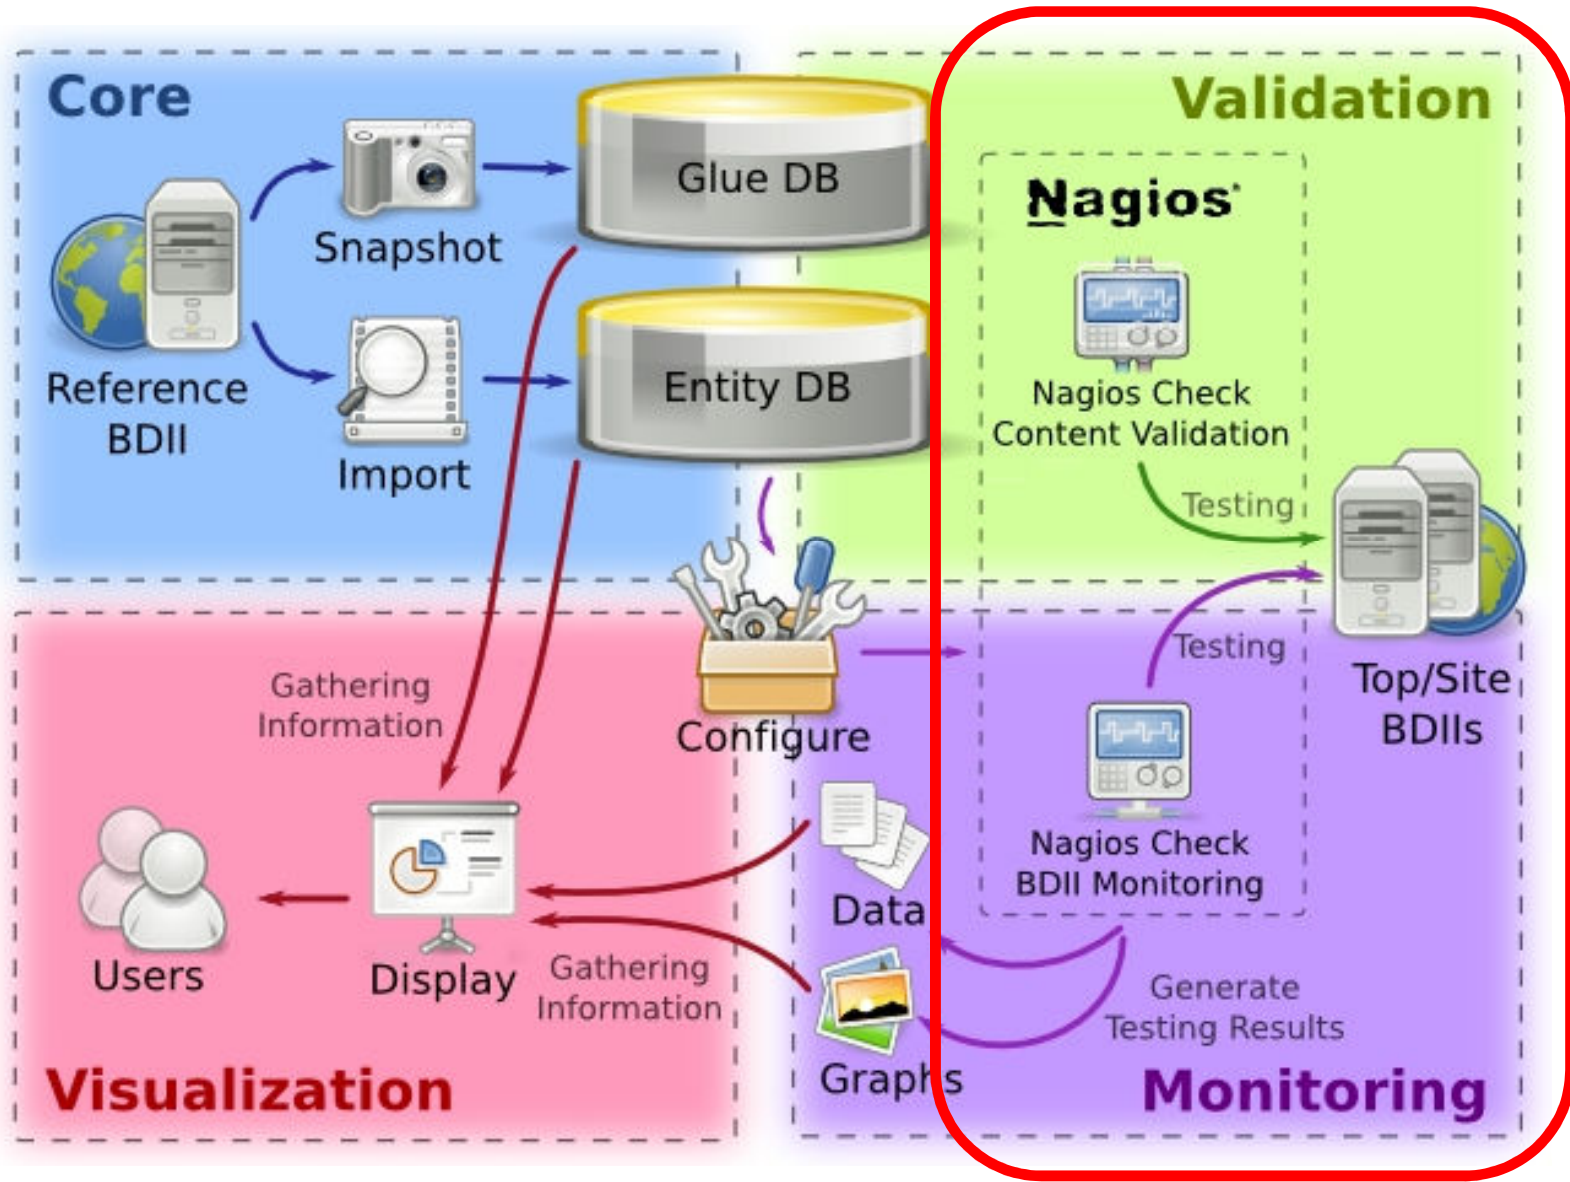
ERROR:  
GlueSubClusterUniqueID=agh2.atlas.unimelb.edu.au,GlueClusterUniqueID=agh2.atlas.unimelb.edu.au,MdsVo-name=Australia-ATLAS,o=grid, A value is incorrrect, GlueHostArchitectureSMPSize must be a valid integer

ERROR: agh5.atlas.unimelb.edu.au:8443/cream-pbs-atlas, GlueCEPolicyAssignedJobSlots has negative or null value,


ERROR: agh5.atlas.unimelb.edu.au:8443/cream-pbs-atlas, GlueCEStatus has unknown value, Value Special is not an allowed value.

- **Use within a Grid site**
  - Troubleshoot BDII configuration issues immediately
  - Validate BDII configuration settings proactively
- **Use within Nagios based Regional Monitoring**
  - Using Nagios for Regional and Site service level monitoring
  - Installing and launching the Nagios probes to monitor BDII
  - Sending out the notifications about the BDII status
- **Use within GStat 2.0**
  - Using Nagios for Information System monitoring and validation
  - GStat 2.0 has a web front end to aggregate overall information
    - Grid infrastructure
    - Resources and services
    - Status
- **Ensuring information system integrity and quality**





Grid Statistics at your finger tips



## GStat 2.0

Geo View
LDAP Browser
Summary Views

---

**Site Information**

Site name: Australia-ATLAS

Description: Australia ATLAS, University of Melbourne

Location: Melbourne, Australia

Web: <http://epp.ph.unimelb.edu.au/>

Contact: [tier2@physics.unimelb.edu.au](mailto:tier2@physics.unimelb.edu.au)

**Monitoring and Validation Status**

Services	All	Overall Status
Top BDII	1	OK
Site BDII	1	OK
CE	2	CRITICAL
SE	1	CRITICAL
All the Services	9	CRITICAL
Site	-	CRITICAL


**GStat Update**

Last update: 17:10:06, 3 Aug 2009  
(8 minutes ago)

**Installed Capacity**

Site Name	Logical CPUs	Physical CPUs	Storage Space	Waiting Jobs
Australia-ATLAS	84	15	<div style="display: flex; align-items: center;"> <div style="width: 33%; height: 10px; background-color: green; border: 1px solid gray;"></div> <span style="margin-left: 5px;">33%</span> </div>	<div style="display: flex; align-items: center;"> <div style="width: 86%; height: 10px; background-color: yellow; border: 1px solid gray;"></div> <span style="margin-left: 5px;">86%</span> </div>

Grid Statistics at your finger tips



## GStat 2.0


Geo View
LDAP Browser
Summary Views

---

Show  entries Search:

Site BDII	Check	Current State	Status Information
agh6.atlas.unimelb.edu.au	check-ce	CRITICAL	CRITICAL - errors 26, warnings 33, info 0

Showing 1 to 1 of 1 entries ⏪ ⏩



- **Addresses typical operations problems with information system**
  - Information system infrastructure monitoring
  - Information content validation
- **Nagios Probes provides an extensible mechanism**
  - Components already reused within the OAT
  - Integrated into the Nagios based regional monitoring
    - Sites will be alerted to information system problems
    - Will be incorporated into the Site availability metrics
  - Sites/Regions can run the checks directly
    - Command line manually
    - Nagios instance
    - GStat 2.0 instance
- **GStat 2.0 is a monitoring framework**
  - Scalable, distributed, extensible, modular approach
  - Improved displays for viewing the results and information
  - Handy tool for debugging operational issues
  - Global, national, regional, site use cases supported



- **GStat 2.0 Beta Release Notes**
  - [http://svnweb.cern.ch/trac/gridinfo/blog/GStat\\_2\\_0\\_beta](http://svnweb.cern.ch/trac/gridinfo/blog/GStat_2_0_beta)
- **Installation and Configuration Instructions**
  - <http://goc.grid.sinica.edu.tw/gocwiki/GSInstallationGuide>
- **GStat Support List**
  - [project-grid-info-support@cern.ch](mailto:project-grid-info-support@cern.ch)
  - any questions or comments or problems
    - please email the GStat support list
- **Feedback on desired features are welcome 😊!**