


Testing and Benchmarking Grid Infrastructures using the g-Eclipse Framework

Nicholas Loulloudes


On behalf of the g-Eclipse Project

Overview



- Grid Reality – The Problem
- g-Eclipse Framework
- The Idea
- Technical Overview
- Tests Framework
- Benchmark Framework
- Future Work

Grid Reality – The Problem



Grid Infrastructures: a huge collection of computational and storage resources.


The EGEE* infrastructure in Europe alone has available 24 hours, 7 days a week:

- 267 Sites (in 54 countries)
- ~114,000 CPUs
- 20 Petabytes of disk space.
- Supporting ~ 15 application domains


• Enabling Grids for E-science: www.eu-egee.org



Grid Reality – The Problem II




- Heterogeneity among Grid Sites:
 - Hardware
 - Middleware
 - Application Software
- Failures
 - Network (Congestion, Faulty links, Attacks)
 - Machines (Hardware failures)
 - Other (Power Disruptions, Nature Reasons)



Degradation of the expected Quality of Service (QoS)

Users need to evaluate a priori the resource offerings of different infrastructures and Virtual Organizations (VO).

The g-Eclipse Project




Funded by the European Community

- Duration: July 2006 – December 2008
- Funding: 2 Million Euro
- Consortium members:




An Open Source Framework

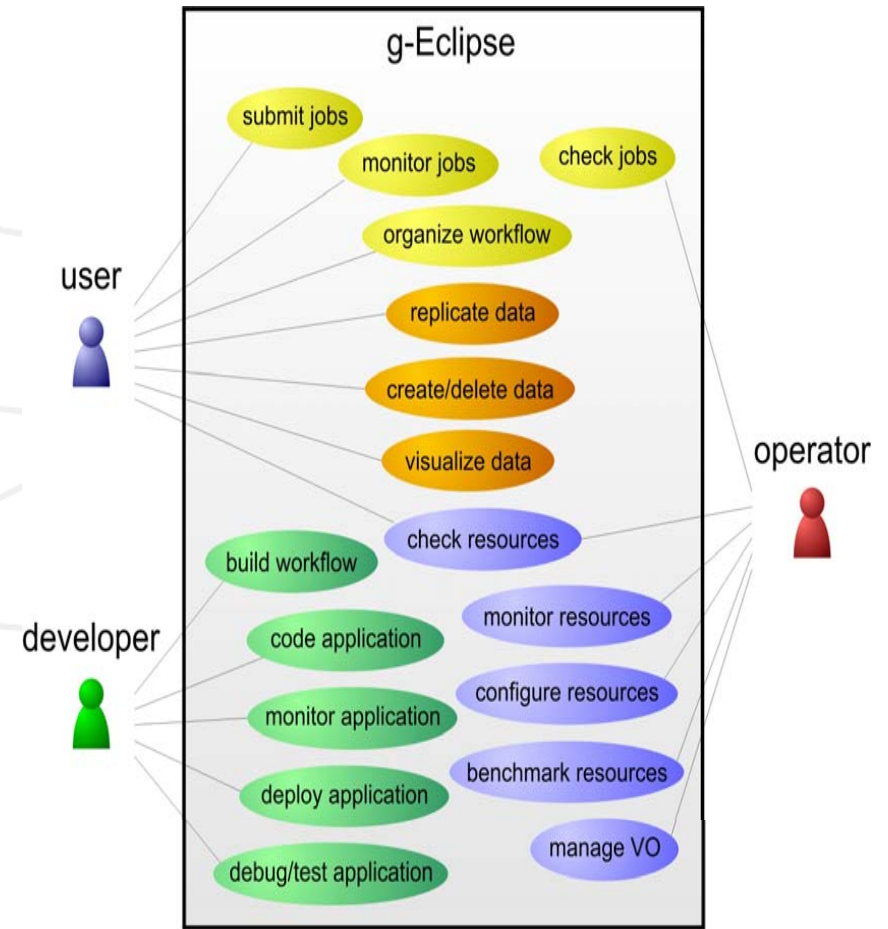


- g-Eclipse is an official Eclipse Technology Project
 - Around 20 developers among the 8 partners
 - 14 members with committer status at Eclipse.org
 - Community established and started to grow.
- Source code released under Eclipse Public License
 - Version 1.0 released in January 2009
 - Available for free download
 - Supported platforms: Linux, Windows, Mac OS X
 - Roughly 70 plug-ins, 4000 classes and interfaces aka 350.000 LOC


g-Eclipse – The Idea



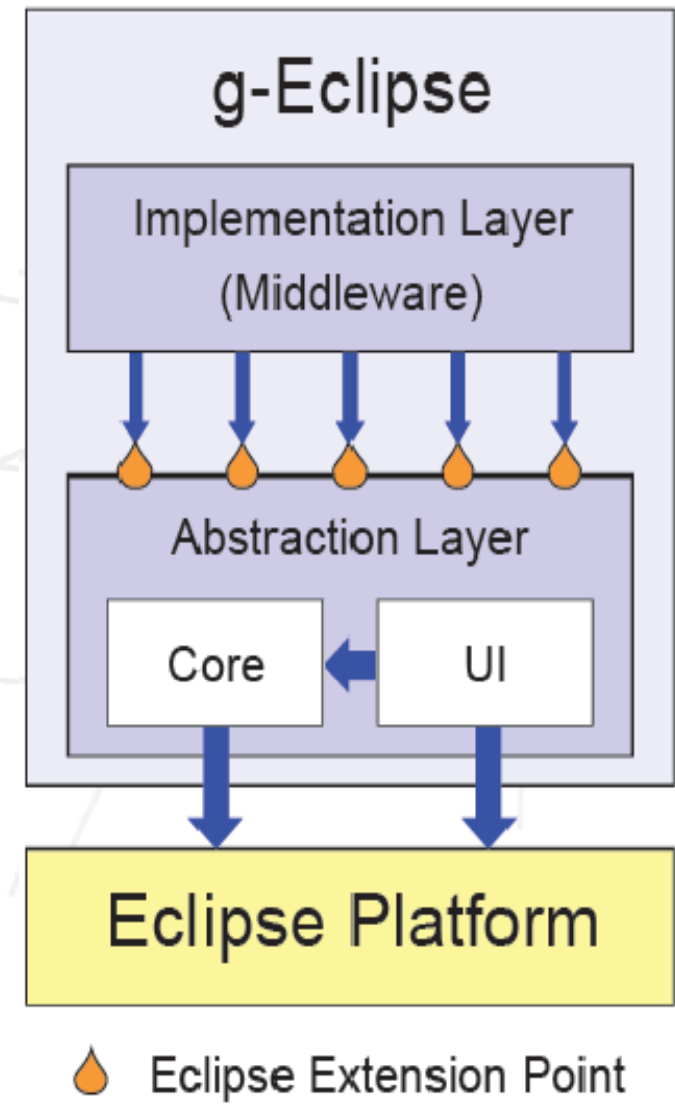
- Provide a friendly UI for accessing Grids.
- Provide an extensible, middleware-independent, framework for accessing Grids.
- Supports the roles of Grid **users**, **developers** and **operators**.
- Provide the necessary tooling to hide the complexity (**wizards**, **editors...**).
- Substitute CLI with GUI.
- Conform to Grid Standards.




g-Eclipse - Technical Overview



- **g-Eclipse is based on the Eclipse platform**
 - It is modular (OSGi) and extensible (Extension point mechanism of Eclipse).
- **Provide a middleware-independent architecture that:**
 - Abstracts common grid concepts.
 - Provides abstract core functionalities.
 - Is extensible by middleware-specific plug-ins.
- **Provide a graphical user interface that:**
 - Is based on the abstract core and therefore.
 - Looks and behaves (at least) the same for any middleware.



g-Eclipse – Middleware support



Currently supported middleware:

- **gLite** - Batch system for the scientific user.



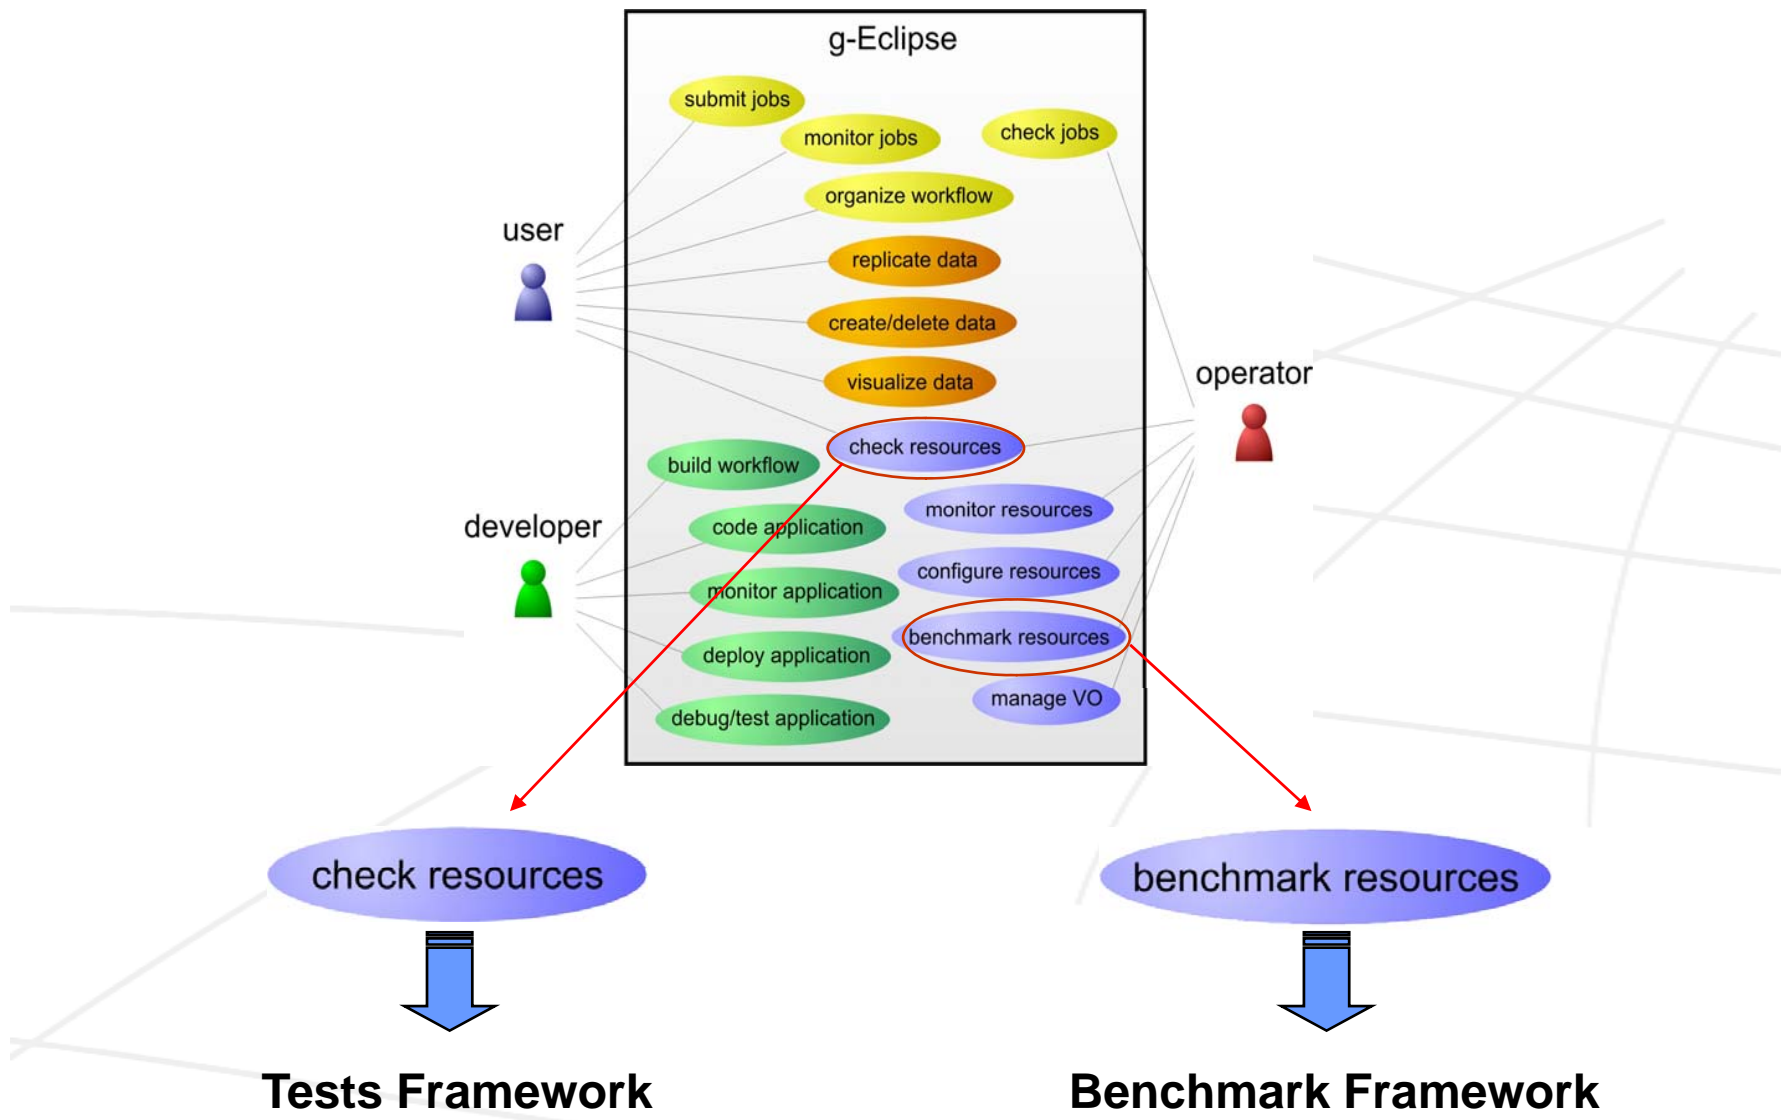

- **GRIA** - Service-oriented infrastructure for industry and commerce.




- **AWS** adapters for cloud computing.



Use Case




Tests Framework




- The framework provides
 - Simple, on-demand tests to determine if a service is currently up.
 - Complex tests to determine if a service complies with the requirements of users.
- Tests can be performed to one or more resources are categorized to:
 - **Simple Tests**
 - **Structured Tests** (Operators Jobs).
- Simple tests *simply* check if a resource is available.
- Exemplary simple tests are provided:
 - Ping Test.
 - Domain Name Lookup Test.
 - Port Scanner.

Tests Framework II



- Structured Tests perform extensive tests that may require the submission of a job.
- Grouped in three categories:
 1. **Direct Tests:** directly tests remote services of a resource using service dependent access protocols
 2. **Submittable Tests:** Must be submitted to an infrastructure. i.e.:
 - Test the configuration of an executing machine
 - Check middleware version
 - Check host certificate expiration
 3. **External Tests:** Accessing an external testing service or database to retrieve results or possibly initiate new tests.
 - SAM service
- Test results and history are **stored locally** in XML format and can be retrieved at any time.


Benchmark Framework



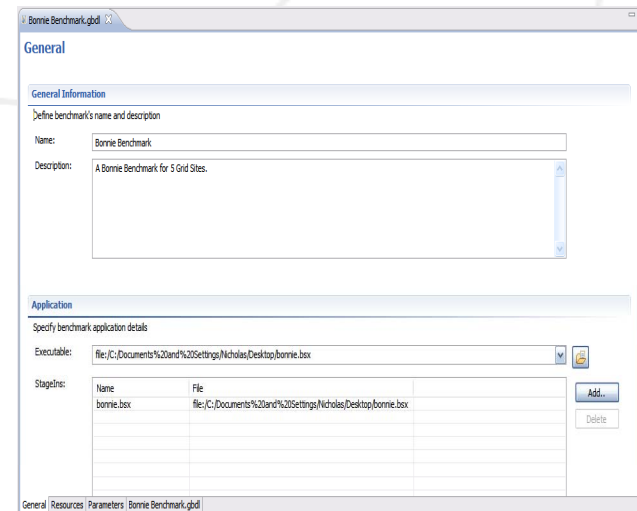
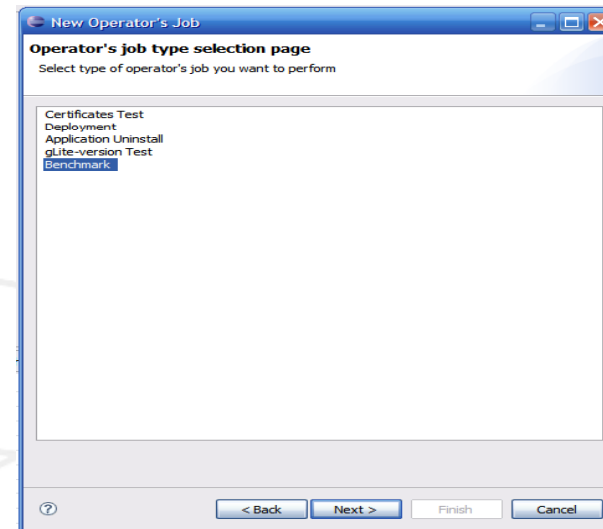
- Port of **GridBench*** to g-Eclipse.
- Allows the Grid User to measure on-demand the performance of selected Grid sites.
 - Performance is measured by applying a known workload (Benchmarks) which might differ among services.
- Allows to **compare** current results with the history of previous results to observe the current “health” of resources.
- Currently 2 exemplary benchmarks are provided (FLOPS, Bonnie) but the framework is **fully extensible** for new Benchmarks.

* <http://grid.uce.ac.cy/gridbench/>


Benchmark Framework II



- Grid Benchmarks require **submission of a job** to a site.
- Benchmarks are wrapped to a single self extractable, self-executable, self-destructible job.
 - Easily created using a guided wizard.
- Benchmark Jobs can be refined through a **Grid Benchmark Definition Language (GBDL)**.
 - A Multi-page form editor is provided.




Benchmark Framework III




- Ability to monitor the Benchmark job status.
- Benchmark jobs output **stored locally** in a relational database (Eclipse Derby).
- Database is queried using SQL to retrieve required data.
- Output can be presented in **tabular** or **graphical** form.
- Local database **import / export** functionality.

Resource	Benchmark
dgrid-ce.fzk.de	ERROR
dgrid-ce.fzk.de	ERROR
egee-ce1.gup.uni-linz.ac.at	OK
ce101.grid.ucy.ac.cy	RUNNING

Resource	Date	input_block...	random_se...	input_char...	output_blo...	random_se...	output_rew...	input_char...	input_block...
ivrce.fzk.de	08/03/2009 16:18:58	100.0	13566.0	99.0	329767.0	90.2	367082.0	96.8	15597.0
ivrce.fzk.de	08/03/2009 16:19:09	100.0	18676.0	93.2	314548.0	70.7	330782.0	96.9	22200.0
egee-ce1.gup.uni-linz.ac.at	08/03/2009 16:21:28	100.0	20359.0	96.6	180453.0	82.5	11832.0	4.8	17081.0
ce101.grid.ucy.ac.cy	08/03/2009 19:48:26	100.0	58282.0	99.6	348331.0	100.0	491431.0	100.3	63906.0




Future Work



- Tests Framework
 - Extension of Simple and Structured tests.
- Benchmark Framework
 - Implementation of additional Benchmarks to reach completeness of GridBench.
 - Ability to rank sites adaptively based on current and previous benchmark results.

Contact the g-Eclipse Team.



Thank you for your attention.

For more information visit:

<http://www.eclipse.org/geclipse>

<http://www.geclipse.eu>

Or contact the:

g-Eclipse User Mail List: [geclipse-user \[at\] eclipse.org](mailto:geclipse-user@eclipse.org)

g-Eclipse Development Mail List: [geclipse-dev \[at\] eclipse.org](mailto:geclipse-dev@eclipse.org)