

gLite on Grid'5000 : towards a real-size testbed for production grids

Sébastien Badia and Lucas Nussbaum

Partially funded by **Simglite project**
Appel Interfaces Recherche en grilles – Grilles de production
Institut des Grilles du CNRS — Action Aladdin INRIA



Nancy-Université




Goal

- ▶ **Use Grid'5000 as a testbed for gLite**
- ▶ Use cases : developers of gLite components, and of applications interacting with the gLite middleware
 - ▶ Be able to **run experiments in a stable environment** (no variation between experiments) \rightsquigarrow compare results
 - ▶ Be able to **create experimental conditions required by an experiment**, possibly hard to meet in a production environment (e.g service crash)
 - ▶ Be able to **replace components of the infrastructure** \rightsquigarrow test new versions, test interoperability
 - ▶ **Avoid overloading or influencing the production infrastructure** with test jobs

Grid'5000


- ▶ **Experimental platform for research on distributed systems** and high performance parallel computing
- ▶ 1700 nodes (7000 cores), 10 sites in France
- ▶ **Reconfigurable by users** : operating system on nodes can be replaced using Kadeploy, network isolation with KaVLAN



Deployed gLite infrastructure

- ▶ One **VO** and its **VOMS** (*Virtual Organization Membership Service*), users directory
- ▶ Several sites, composed of :
 - ① One **BDII** (*Berkeley Database Information Index*), directory of resources available on each site
 - ② One **CE** (*Computing Element*), task submission service for a given computing site
 - ③ **Worker nodes** and a **batch scheduler** to access them. **Torque/Maui** was used
 - ④ One **UI** (*User Interface*), used by users to access the resources

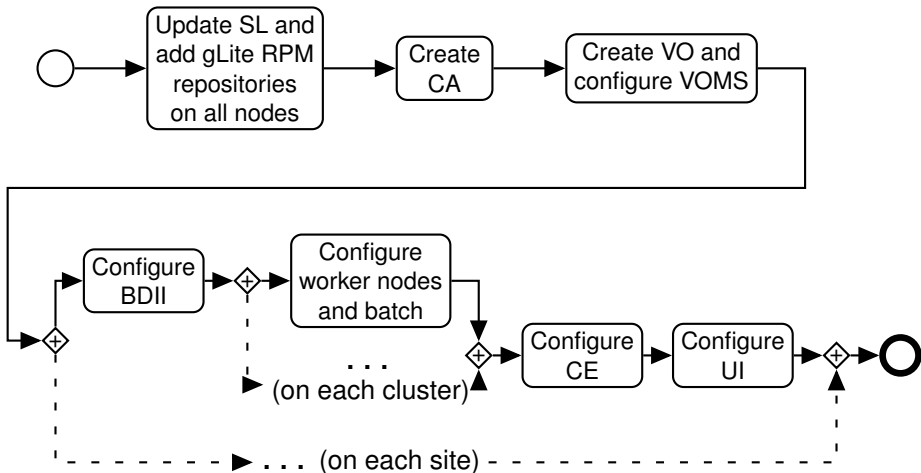
Tools developed



- ▶ **Scientific Linux 5.5** image, minimal and generic (working on all Grid'5000 clusters) for the Kadeploy deployment tool
- ▶ **Ruby scripts** enabling an automated installation of gLite from RPM repositories
 - ▶ Description of the platform to deploy (VO, sites, clusters) in a configuration file
 - ▶ Creation of a certification authority to generate and automatically sign users and machines certificates
 - ▶ Pre-filling of the RPM cache on nodes using Kadeploy to accelerate deployment

<https://github.com/sbadia/gdeploy/>

Deployment process



Results

Use of Grid'5000 to deploy the gLite middleware

- ▶ Deployment up to **926 nodes (17 clusters, 9 sites)**
- ▶ Installation of machines with Scientific Linux 5.5 using Kadeploy :
10 minutes
- ▶ Configuration of gLite with one VO on 597 nodes (6 sites,
10 clusters) : **170 minutes**

Future work

- ▶ **Improvements to the deployment script**
 - ▶ Deployment of several VO
 - ▶ Deployment of other gLite services : storage, monitoring

- ▶ **Collaborations**
 - ▶ **Experiments on evolution of gLite components**
 - ▶ **Experiments on tools interacting with the gLite middleware : workflow engines, pilot jobs managers, etc.**
 - ▶ Simulation of services crash
 - ▶ Load injection
 - ▶ Submission of a large number of fake tasks