A Symbiotic Deployment of a Service Grid and an Opportunistic Grid over the same e-Infrastructure: the EELA-2 Experience

Francisco Brasileiro

Universidade Federal de Campina Grande (UFCG), Brazil

JRA1 Manager for EELA-2 project

fubica@dsc.ufcg.edu.br

Motivation

- **As the grid technology matures, more and more infrastructure enter in production**
- **Following the sharing spirit of grids, a lot of attention has recently been placed in the integration and interoperation of production grid infrastructures**
- **Integration/interoperation is normally achieved at the grid middleware level**
  - Several approaches have already been proposed
  - Choosing the right approach depends not only on the middleware that are used, but also on the motivations for the integration/interoperation

The “Why” question

- Before we present **how** we approach integration/interoperation of a Service Grid with an Opportunistic grid in the EELA-2 project, let us try to understand **why** one would like to integrate/interoperate two PGIs

- As in any partnership, a good understanding of the **value** that each partner gets and yields is crucial for the long-term **sustainability** of the partnership

The “Why” question

- **Let us take a look at the main stake holders involved**
  - Resource owners
    - These, ultimately, are the ones that decide for or against the “operation”
    - Possibly driven by technical arguments from application users and infrastructure administrators
  - Users
    - Want more/better resources (that they find in the other PGI)
      - *Ideally without giving up the resources they already have access*
    - Want more/better functionalities (that they find in the other PGI)
      - *Ideally without having to change their application*
  - Administrators
    - “Is there a change? I am against!” ☺

The "Why" question in the EELA-2 context

- **Very dissimilar application users and resource providers**
  - Few large/medium labs with strong links with large labs and their associated projects
    - They normally place non-trivial processing, storage, communication, and, most importantly, coordination requirements on the grid middleware
    - Although a non-negligible amount of jobs are in fact embarrassingly parallel sub-tasks of the same application (ie. Bag-of-Tasks – BoT)
    - They need and can cope with the difficulties of installing and, most importantly, maintain operational a sophisticated service grid
  - Many small labs, mostly working in isolation or starting to develop collaborations with other labs
    - Nevertheless, access to non-trivial amounts of computing resources enable these labs to improve their methodological approach
      - *Almost all their jobs are BoT*
    - They do not need, nor are able to, cope with the difficulties of installing and maintain operational a sophisticated service grid
      - *Nevertheless, they want to be able to take part in larger cooperations with larger research labs*

The “Why” question in the EELA-2 context

- **In summary:**
  - Large/Medium labs need a service grid based on a rich grid middleware
    - They have chosen gLite as such middleware
    - The first phase of the EELA project has successfully deployed such a PGI
    - But it would be nice if the resources of the small labs (that could not use gLite) could also be integrated in the grid
  - Small labs would rather build a grid that is supported by a simpler middleware target to efficiently execute BoT applications
    - Opportunistic grid middleware are the correct choice for them
    - They have chosen OurGrid as such middleware
      - *More details about OurGrid later*
    - But they want to cooperate with the larger labs and for that, they are required to provide their share of contribution to the whole system

Promoting symbiosis

- **By "integrating" the gLite-based Service Grid with the OurGrid-based Opportunistic grid we aim at:**
  - Increase the size of the OG, by exploiting idle resources in the SG (more and better resources for the OG)
  - Leverage the virtualization infrastructure of OurGrid to provide gLite worker nodes on the shared machines of the OG and on dedicated machines running OurGrid (more resources to the SG)
    - OurGrid's incentive mechanism/scheduling policy guarantees fair sharing of resources
  - Provide a specialized service for the execution of some of the BoT applications that currently execute on the SG (better functionality for SG users)
    - This allows improved performance for these applications
    - Has the side effect of improving the performance of the other applications that run on the SG, due to the smaller workload it will need to deal with

The “How” question

- **Now that we know why, let us discuss how we integrated a gLite-based SG with an OurGrid-based OG**
- **Several integration approaches have already been proposed**
  - Gateways/Bridges
    - “Transforms” the grid resources of one grid in resources of the other grid by implementing a suitable bridge interface
      - *The EDGeS project is working on a generic bridging technology*
  - Gliding-in
    - “Install” part of the grid middleware of one grid in the resources of the other grid on the fly (through pilot jobs)
- **We propose a conviviality approach that is based on having the two middleware (or parts of them) installed in the resources that wish to provide this feature**
  - It achieves similar effect of the gliding-in approach

If only everyone were trustworthy

- **Security issues is arguably the most complicated matter when interoperating two PGIs**
- **OurGrid is meant to operate in a very promiscuous environment**
  - Peers do not trust each other
  - Yet, they need to cooperate
- **There are different security aspects to take into account**
  - How to protect resources from malicious users?
  - How to protect the applications from malicious resources?
  - How to prevent free-riding?

If only everyone were trustworthy

- **Different security aspects are treated differently**
  - How to protect resources?
    - Very conservative approach based on “the administrator denial-of-service” approach ☺
      - *Remote jobs run inside a virtual machine with limited disk access and no access to the network (remember it only supports BoT)*
  - How to protect applications?
    - Very liberal approach – I mean, the user is left on its own
      - *Tasks may indicate a user-defined checking mechanism to be applied to the output of jobs (eg. to verify the existence of a watermark in an appropriate place)*
      - *We have implemented credibility-based scheduling (a la BOINC), but it is not in the supported software*
  - How to prevent free-riding
    - The Network-of-Favors incentive mechanism
    - Asymmetric cryptography used to prevent impersonation

OurGrid in a nutshell (1/2)

```
job:
   label: my_rendering_example
   requirements: (os=linux)
task:
   init:        store render $STORAGE
                put input-1 $PLAYPEN
   remote:      render < input-1 > output-1
   final:       get output-1 output-1
   check:       ./my_check output-1
task:
   init:        store render $STORAGE
                put input-2 $PLAYPEN
   remote:      render < input-2 > output-2
   final:       get output-2 output-2
   check:       ./my_check output-1
...
```

OurGrid in a nutshell (2/2)

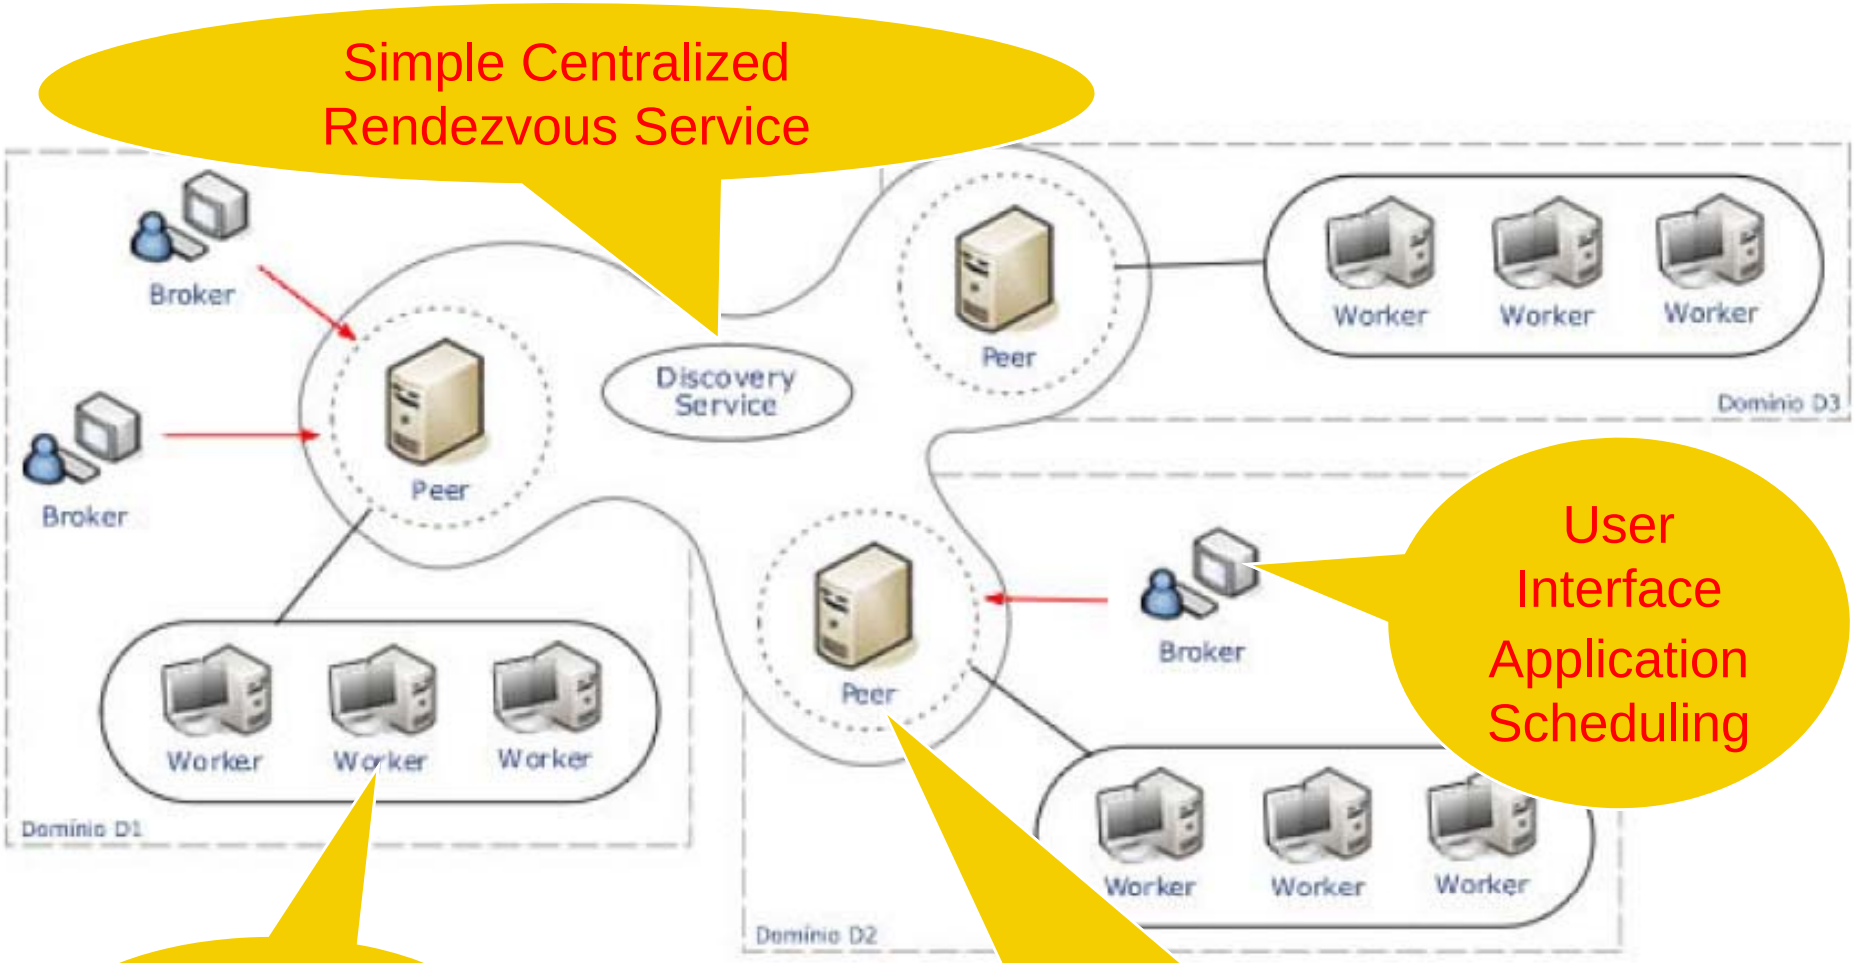

A closer look at the gLite CE

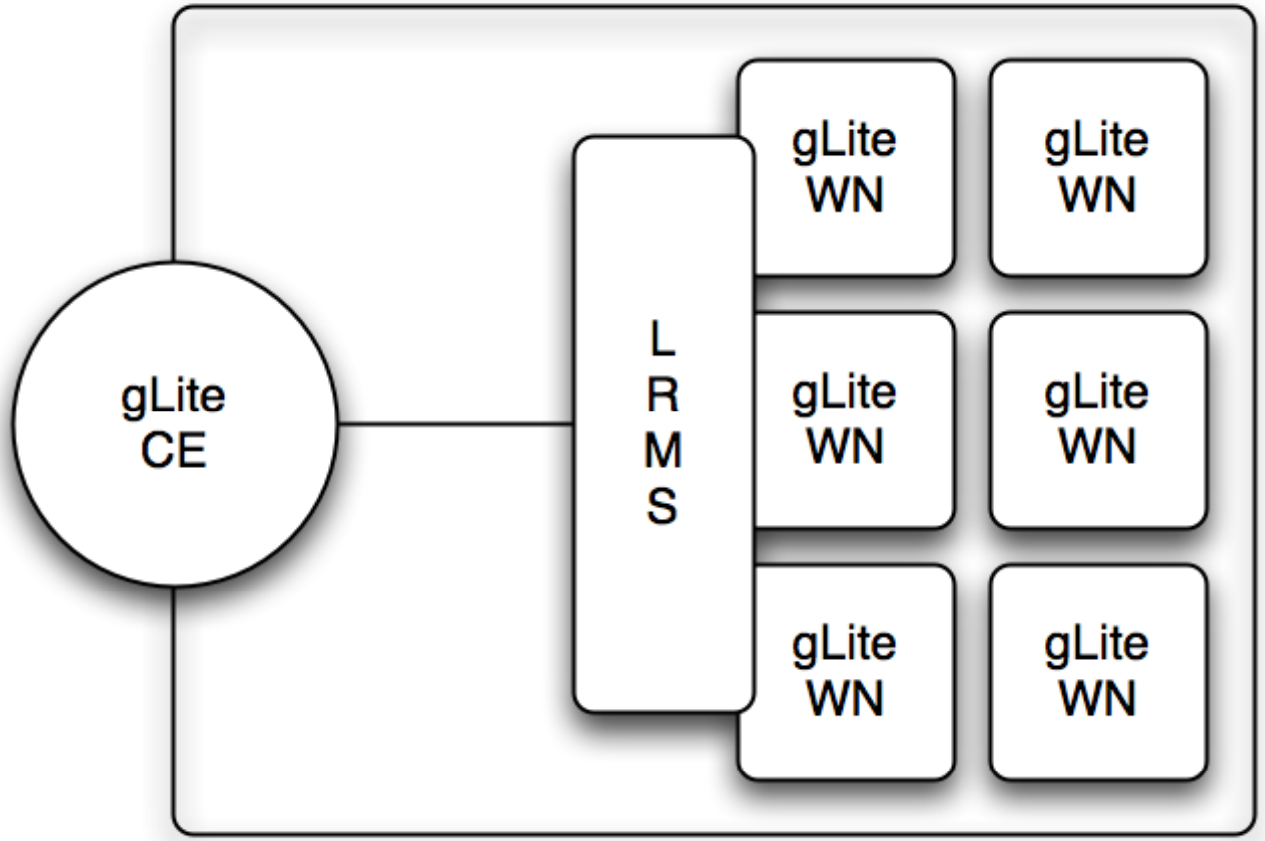

Road map for the middleware conviviality

- **Allowing idle resources in an EELA-2 gLite resource centre to be exposed as OurGrid resources**

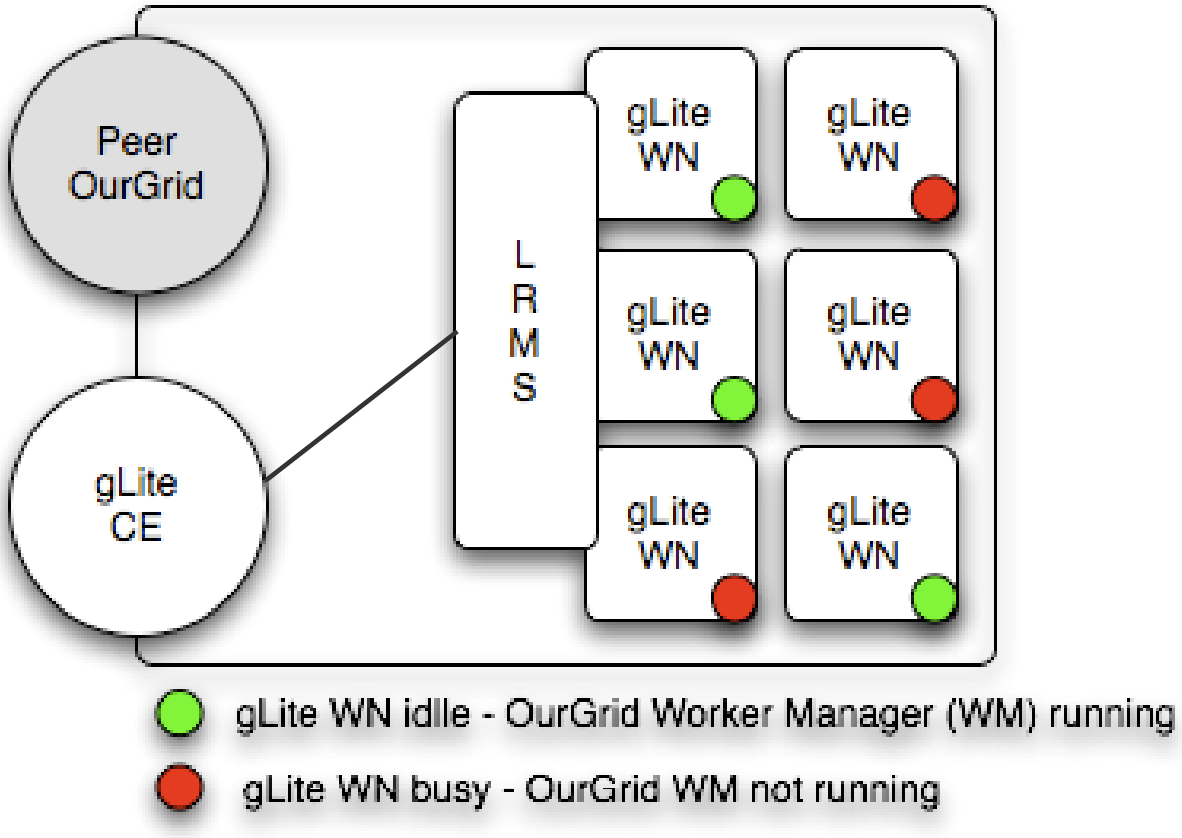

Road map for the middleware conviviality

- **Allowing resources of an OurGrid resource centre to be exposed as virtual gLite resources**

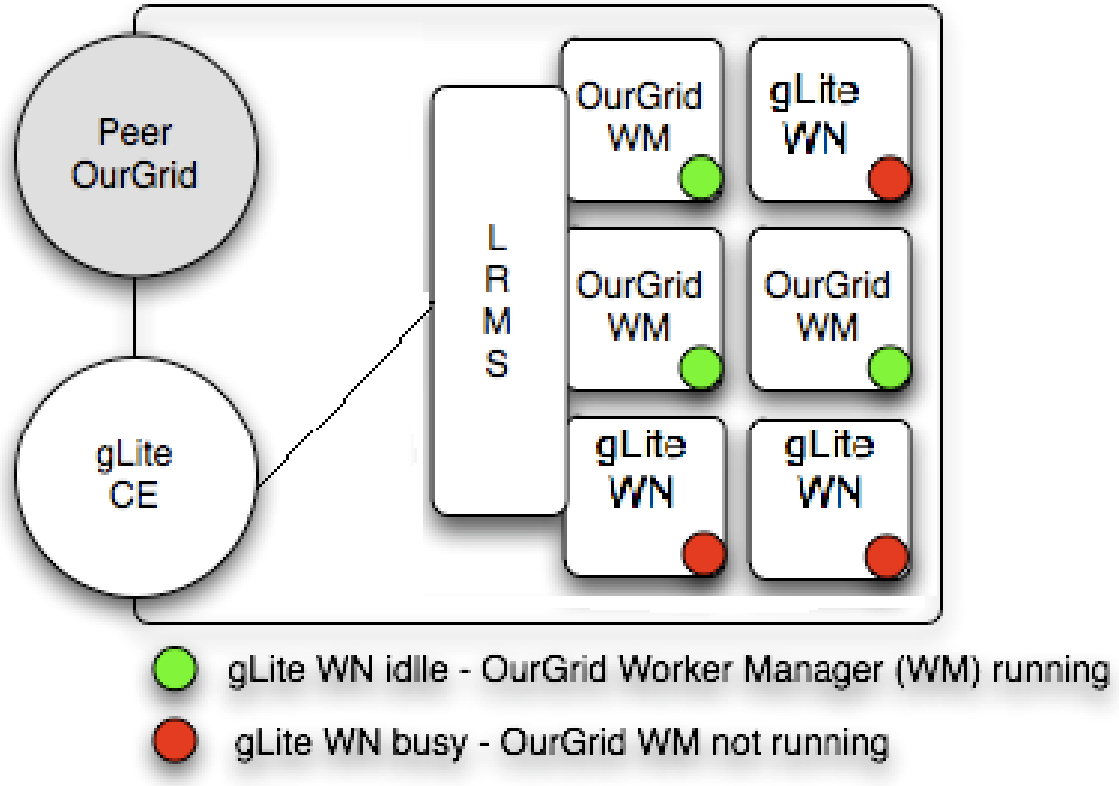

Current status and future work

- **Latest version of OurGrid just released with support for the exploitation of idle cycles in service grids**
  - Available for download at both:
    - *https://forge.eu-eela.eu/projects/ourgrid-eela/*
    - *http://www.ourgrid.org/*
  - Customized installation guide available for sites already running gLite
- **Installation of gLite in OurGrid sites and mapping of gLite BoT jobs in OurGrid jobs to be released within a couple of months**
- **Evaluate the impact of the conviviality in a production environment**

Concluding remarks

- **More information on the EELA-2 JRA1 activities at the project's website:**
  - http://www.eu-eela.eu/
- **For OurGrid specifities, visit:**
  - http://www.ourgrid.org/
- **Contact me at:**
  - fubica@dsc.ufcg.edu.br

Thanks for your attention!