

EGI 1st User Virtualisation Workshop

12-13 May 2011 – Amsterdam

<http://go.egi.eu/uvw1>

EGI brings together over 300,000 processors and more than 100 PB of storage space located at 350 sites in 50 countries, supporting over 12,000 researchers across almost every discipline. The current EGI services are designed to satisfy the requirements of these already established user communities, however, there are many other research groups across a variety of domains that have a strong demand for using e-Infrastructures, but with a different set of services.

In order to meet this demand and increase the diversity of the users, EGI needs to improve the flexibility and efficiency of the infrastructure, scaling out the support model, ultimately, empowering virtual research communities with direct control on the service environments they offer to their end-users.

Virtualisation provides the technology needed to achieve this vision and cloud computing as the operational model for delivering it to the virtual research community. EGI has, therefore, been investigating virtualisation and cloud computing, which demonstrate how new technologies can enable dynamic execution environments or on-demand elastic service deployment with new, clear cost measurements and business models [1].

Under this context, EGI.eu organised a dedicated workshop on User Virtualisation on 12-13 May 2011 in Amsterdam [2][3] to bring together representatives of three critical groups within the European production infrastructure – End-users, Resource Providers and Technology Providers.

Focus

The workshop focused on discussing a model for the integration of virtualisation within EGI through the adoption of an EGI Cloud Integration Profile [4] that would provide virtual research communities with mechanisms for deploying software environments to meet the needs of their end-users. Specific issues were covered through a series of dedicated presentations and breakouts in the key areas of

monitoring, accounting, virtual machine management, and information services.

Discussions were tailored to understanding if and how EGI should move towards providing an Infrastructure as a Service (IaaS) model to support data intensive research communities and if this service could be provided by federations of resource providers from the EGI community. Other key topics included potential uses of such capabilities by virtual research communities and what are the major issues (e.g. technical, policy, governance) that need to be resolved for this use to happen.

Expected Outcomes

Due to recent community experiments, the workshop was able to converge rapidly towards a shared understanding as to how EGI should move towards providing virtualised resources as the next generation service dedicated to expanding its user base.

These discussions allowed best practices, available or required standards, issues around software maturity and availability, development priorities, technology gaps, issues and solutions to be identified. Experts were also identified within the community who could further support each activity.

As a result, a set of critical observations were recorded that will be developed into a roadmap that will meet the growing need for virtualised resources from the European research communities.



Workshop Participants

Identifying and Validating Use Case Scenarios

EGI.eu has drafted an EGI Cloud Integration Profile [4] in order to help define a technical roadmap for the interoperable integration of virtualised resources from different resource providers to provide an integrated federated virtualised resources infrastructure for exploitation by EGI's user communities. The profile defines a minimal set of usage scenarios that when supported will provide key functionality for the virtual research communities wishing to utilise the 'cloud' interfaces provided as part of EGI to deploy software environments to meet the needs of their communities.

The workshop offered an opportunity not only to present these scenarios, but also to obtain feedback from the community as well as technology providers on the applicability of each scenario. These will feed back into the EGI Technology Coordination Board (TCB) in order to coordinate the NGIs and technology providers interested in these developments.

6 Key Usage Scenarios

- Running a pre-defined VM image.
- Running my VM image (with my data).
- Deciding which virtualised resource to use.
- Accounting across resource providers.
- Reliability/availability of the resource.
- State change notification from the VM manager.

StratusLab compared these scenarios with their own use cases being able to map all, either matching or extending each. They were also able to add two potential new scenarios covering credit-based resource utilisation and transparent access to federation resources. Both were identified as low priority in the short-term.

Other use cases were also provided by the CSC-ELIXIR, a model for Grid and Cloud integration presented at ISGC, and the WLCG, including an additional talk on Trusted Virtual Images.

Topic Specific Breakouts

During the first afternoon, four breakout sessions were organised around key areas to address specific questions in key areas: monitoring, accounting, virtual machine management, and information services.

Each session discussed the identification of technology gaps and operational issues, what work would be needed to remove them, specific standards that could be used, current best practices to adopt and the availability and maturity of possible software

solutions. The common guideline from all breakout sessions was to start from the existing standards/services and then extending them to cover the needs of the virtualisation technologies.

In general, the breakout sessions focused on Infrastructure Providers in terms of how virtual machines should be monitored, managed and accounted for; while the services deployed within the virtual machines are monitored and accounted for with the tools specific to the needs of that virtual research community.

The slides and notes from each breakout session are available [2] along with a consolidated set of notes recording the whole meeting [3].

Issues

Some of the main questions and issues from the breakouts were around monitoring and the necessity to support the monitoring of what is running inside the machine, establishing when a VM has started (or not) and who looks after the systems that have been set up. Accounting and billing was a reoccurring subject in terms of how to bill and who to bill. It was agreed that further discussions were needed around the separate issues of what to account for and how to bill for its use.

A number of issues pointed to policy matters, mainly covering trust around endorsements of VM images and authentication on image access.

Standards are always important for interoperable systems. A number of specifications were identified as potentially useful to be considered for the EGI Cloud Profile [4]. A list of standards supporting implementations was recorded even though support, conformance and suitability for real-life scenario needs to be evaluated.

Common points included the need for a common taxonomy of information, measurement and accounting and differentiating between small volume, single location and large distributed data.



Workshop Chair and EGI.eu Director Steven Newhouse

Summary

The consensus emerging from the workshop was to introduce virtualised resources alongside the current grid services to increase the flexibility of infrastructure while retaining the current federated model used within EGI.

This approach would require no 'big bang' migration, but rather a gradual change transparent to the end user. The current production services could continue with no change.

In order to accomplish this objective, a 'testbed' environment will be established leveraging existing expertise and activity in the community. The planning should also consider the current commitment that experts have in current projects and evaluate the possibility of changing priorities, if allowed by the various projects in which experts are engaged.

Next Steps

A number of Technology Providers agreed to investigate and test the functionalities and capabilities of some identified technologies in which several of them were working. Individuals were identified and allocated a technology to evaluate, which comprises: Azure, Eucalyptus (Ubuntu EC), Hyper-V, KVM, OpenNebula, OpenStack, Platform ISF, StratusLab, and VMware. EGI.eu will set up a wiki page in order to host the information and track its progress.

A number of resource providers also informally committed a part of their infrastructure for virtualisation efforts. Formal communication channels are being established in order to take this forward.

A final result of the workshop was for EGI.eu to produce a 1-page mission statement outlining the EGI Virtualisation Vision for wider circulation within the EGI community.

References

- [1] Integration of Clouds and Virtualisation into EGI
<http://go.egi.eu/258>
- [2] User Virtualisation Workshop Agenda
<http://go.egi.eu/uvw1>
- [3] User Virtualisation Workshop Minutes
<http://go.egi.eu/535>
- [4] EGI Cloud Integration Profile
<http://go.egi.eu/435>

Join the mailing list to continue discussions and provide feedback on activities.
cloud-discuss@mailman.egi.eu