



Contribution ID: 93

Type: Oral

## Operating a Grid Site in the Cloud

*Wednesday 4 March 2009 11:25 (25 minutes)*

Cloud technologies have matured quickly over the last couple of years and now provide an interesting platform on which to host grid services. The dynamic nature of these resources could ease life-cycle management for system administrators and could provide customized resources for users. However, questions remain about how these resources can meet the grid's security and operational policies. This presentation explains the challenges raised by using cloud resources for a EGEE grid site.

### Impact

For system administrators, having a pool of virtualized resources may ease the management of a grid site. Specifically the upgrade process can be more efficient because upgraded services can be deployed in tandem with existing services and tested in place. Switching to the new service can be done after the service has been verified. This means less downtime when upgrading, but also provides a more secure fallback solution when something (inevitably) goes wrong with the first installation of the upgraded service.

The lowered downtime, increased reliability, and extensibility are clear benefits for users as well. The virtualization of the cloud resources also permits the execution environment to be customized. This would allow user communities to provide standard images with their software pre-installed. Heterogeneous software environments are one of the leading causes of job failures, and to date, the grid offers no comprehensive solution to this problem.

### URL for further information

<http://stratuslab.org/>

### Conclusions and Future Work

Running grid services within the Amazon cloud is feasible as shown by the two month trial described in this presentation. In the future, site administrators may want to virtualize their own computer centers using open source cloud implementations. For them, bridging external and internal cloud resources may provide an interesting alternative to purchasing hardware. Users may want to take advantage of additional data transfer protocols (http, bittorrent, etc.) offered by the cloud resources.

### Keywords

cloud, grid, AWS, amazon web services

### Detailed analysis

A typical (minimal) grid site provides computing and storage to supported Virtual Organizations (VOs) and runs a few services to make those resources visible on the grid. Amazon Web Services (AWS), the most mature of the available platforms, offers "bare metal" interfaces to virtual machines and to persistent disk images, meaning that standard EGEE tools for machine configuration and management work with little or no changes. Specifically, we take advantage of the Elastic Computing Cloud (EC2), the Elastic Block Store (EBS), and Elastic IP services for the grid site.

The machine-like interfaces mean there are few technical barriers to running grid resources on those services. The full environment where the machines are distant, in Amazon's IP space, and behind firewalls, poses challenges. We describe how various issues such as obtaining grid certificates, keeping logs, etc. were solved. We also describe the operational issues we encountered during the two month trial.

**Authors:** Dr LOOMIS, Charles (CNRS/LAL); Dr LLORENTE, Ignacio (UCM); BÉGIN, Marc-Élian (SixSq Srl); Dr MONTERO, Rubén (UCM); FLOROS, Vangelis (GRNET)

**Presenter:** Dr LOOMIS, Charles (CNRS/LAL)

**Session Classification:** Novel Services and Technologies

**Track Classification:** Emerging Technologies within the EGEE infrastructure