

A Mahout Cluster Across France and Luxembourg Using CloudStack

Grid'5000 is a scientific instrument for the study of large scale parallel and distributed systems. It aims at providing a highly reconfigurable, controllable and monitorable experimental platform to its users.

Currently, Grid5000 operates 1195 physical hosts, for a total of 8184 cores across 10 sites.

In 2003, the Grid'5000 project was launched as a project of the French ACI Grid incentive. Grid'5000 has now reached its objective of exploiting the computational power of over 5,000 core processors to form the only research platform of its kind. There are [19 laboratories](#)

involved in France with the objective of providing the community a testbed, allowing experiments in all the software layers between the network protocols up to the applications. Today, the ALADDIN-G5K initiative has been put in place by [INRIA](#) to ensure the development and day to day support of the infrastructure.

Key Solution Benefits

- Level of functionality
- Open, flexible and customizable
- Easy-to-use web interface
- Control
- Simple management

The Challenge: Increase Awareness of Grid '5000 Tools Developed by the Community

An INRIA research engineer, [Alexandra Carpen Amarie](#), developed an amazing tool called [G5k campaign](#). The G5k campaign allows any user of Grid'5000 (G5k) to book nodes on the infrastructure, deploy machines with bare-metal provisioning and then deploy their favorite Cloud IaaS framework (currently [CloudStack](#), [Opennebula](#) and [Nimbus](#)). In this case, Alexandra deployed Apache CloudStack (IaaS) and Apache Mahout (PaaS) on G5k. Apache CloudStack was used to deploy and manage a set of nodes that run Hadoop and Mahout on top providing a high level of functionality.

The Grid'5000 community has developed a wide range of [great tools](#) that can be used to deploy a cloud IaaS framework. However, the awareness of these tools is currently low in the cloud computing community. By using three tools developed by the Grid'5000 community such as [OAR](#), [Kadeploy](#), and [KaVLAN](#), a campaign can be launched and configured with Apache CloudStack and a cloud could be working within twenty minutes.

Technical Situation

The clusters of G5k are operated like regular batch processing clusters. A batch scheduler is used to access the nodes. Tool #1: OAR, a PBS/MOAB like equivalent. Once the nodes are allocated they are provisioned using Tool #2: Kadeploy a [crowbar](#) like equivalent. Tool #3 KaVLAN, a tool to lease the VLANs configured on G5k. The beauty of G5k campaign is that it hides most of the complexity of the provisioning and configuring.

Only a YAML configuration file needs to be written for your deployment. Specify the sites and the number of nodes that you want to run on/at, for example:

```
deployment:
  engine:
    name: CloudStack
    customization_type: multisiteChef
  walltime: 2:00:00
  sites:
    rennes:
      nodes: 10
      subnet: slash_22=1
    nancy:
      nodes: 10
      subnet: slash_22=1
    sophia:
      nodes: 10
      subnet: slash_22=1
  ssh:
    user: username
```

Then, launch the campaign and wait for the nodes to be allocated, provisioned, and then configured with your IaaS. Depending on the number of nodes requested, you could have a cloud working within 20 minutes and can then interact with it.

Solution: Apache CloudStack

CloudStack was deployed across five sites. There is one basic zone per physical site and 97 physical hosts setup, 800 cores, and 100 VMs deployed running Mahout. Within 24 hours, the snapshots are shown in the inbox. It took 30 minutes to deploy the nodes and one hour to configure the hosts in CloudStack. The 100 VMs were deployed in roughly 10 minutes. The 3 physical nodes missing were due to bare metal provisioning problems.

Figure 1 shows the infrastructure/zone view of the CloudStack deployment. Five basic zones were configured at Rennes, Toulouse, Nancy, Sophia and Luxembourg. All cities connected via the [RENATER](#) fiber network.




Figure 1: Infrastructure/Zone View

Figure 2 demonstrates the infrastructure view, showing five zones, 97 hosts, 10 system VMs (console proxy and secondary storage), and 4 virtual routers (One router was not started at the time of the snapshot).

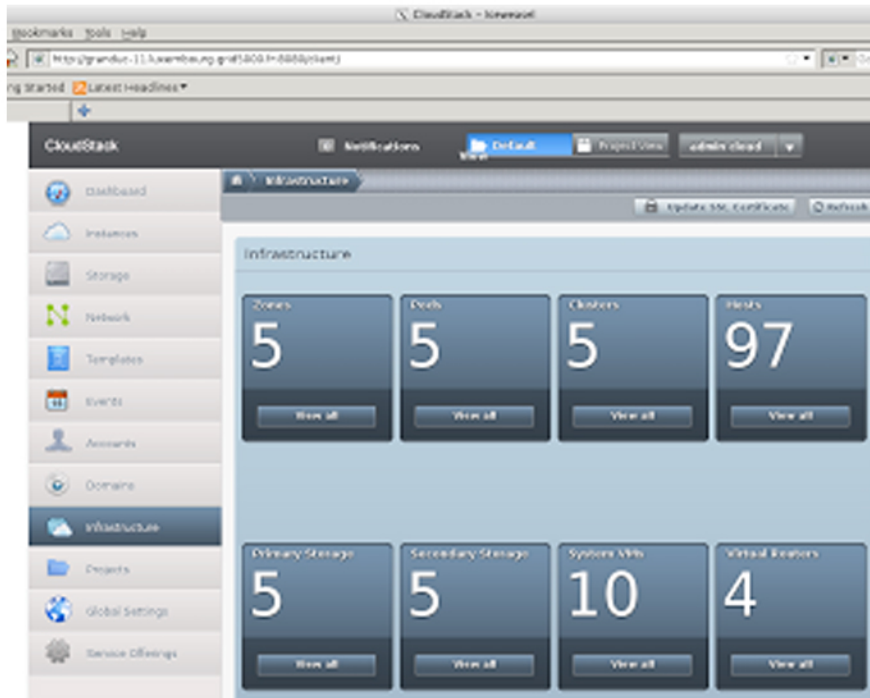


Figure 2: Infrastructure View

By taking a close look at the YAML configuration file, you may notice that this is all based on SSH. Access to G5k is via SSH keys and not through a PKI infrastructure. In G5k the user base is smaller and more trusted. SSH keys are distributed among sites using a basic NFS setup on the private RENATER network. This makes it easy for users to access all sites.

The Benefit: Level of Functionality, Easy-to-Use Web Interface and Control

Apache CloudStack was used to deploy and manage a set of nodes that run Hadoop and Mahout on top providing a high level of functionality. CloudStack allowed the user the flexibility to use tools of their choosing and deploy machines with bare-metal provisioning to build a cloud.

With Apache CloudStack's easy-to-use web interface, it enabled simple management of the cloud infrastructure. The feature-rich user interface implemented on top of the CloudStack API allows the user to see the infrastructure/zone view of the CloudStack deployment and obtain a real-time view of the aggregated storage, IP pools, CPU, memory and other resources in use giving better visibility and control of the cloud.

CloudStack provided orchestration of virtualized resources into one homogeneous environment where it allows the delegation of the creation to users of the virtual machines using the self-service CloudStack portal.

To learn more about CloudStack, please visit www.cloudstack.org.

Services Provided by Other Groups

- Grid'5000 is a scientific instrument for the study of large scale parallel and distributed systems.
- Apache Mahout (PaaS) – allows building of scalable machine learning libraries. Scalable to reasonably large data sets. The core algorithms for clustering, classification and batch based collaborative filtering are implemented on top of Apache Hadoop using the map/reduce paradigm.
- Apache CloudStack - open source software designed to deploy and manage large networks of virtual machines, as a highly available, highly scalable Infrastructure as a Service (IaaS) cloud computing platform.
- OAR - the next generation of the batch management system used on Grid'5000.
- Kadeploy - a scalable, efficient and reliable deployment system (cluster provisioning solution) for cluster and grid computing.
- KaVLAN - used to manage VLAN on Grid'5000 nodes.