Virtualizing services in gLite: Increasing productivity without loosing performance

Enabling Grids for E-sciencE

Introduction

A totally virtualized EGEE infrastructure has been implemented at CESGA during the last years. The new infrastructure allows us to easily start gLite services on demand. This infrastructure is based on Xen, sharing all computer elements and storage systems. It is possible to deploy a complete and scalable infrastructure to house new testbeds with greater flexibility. Extensive benchmarking has been done in our EGEE site showing similar performance between physical machines and modern paravirtualized servers using Xen. In the case of hardware virtualized machines (HVM) the performance is degraded and they are currently not recommended.

Virtualized Architecture

Two years ago we started the migration of our Worker Nodes (WNs) to Virtual Machines (VM) due to limited hardware support in Scientific Linux 3 (SL3). All grid services at CESGA are running under a totally virtualized infrastructure which allows to support new projects on demand. When a new physical server arrives it is quickly configured and installed using a specific kickstart script which installs Xen dom0 software from a local repository using Preboot eXecution Environment (PXE) in an automated way. Each new deployed dom0 could run several VM, like SL3 or SL4, depending on the service requirements. When it is needed to configure a new service, the administrator use a golden-copy from our local VM repository to deploy it into any available dom0 and then it is based on gLite, it is configured using Yaim in a few minutes, otherwise additional configuration is created.

Virtualization Advantages

- Better resource utilization. It can be a solution for sites and users who increasingly demand more services. New services can be installed fast and there is no need to spend money on new hardware.
- Easy replication. To deploy a VM from an existing template or golden copy is done in just a few minutes.
- Load balancing. If a VM requires more physical resources, it can be given more resources or migrated to another physical asset.
- Fault-tolerance: Roll-back possibility (snap-shots) combined with LVM. Ability to take snap shots of existing states. Since a VM is really just a set of files it can take snap shots and revert the VM to previous states within seconds.
- Flexibility: Possibility of using old OS versions, like Scientific Linux 3 with modern hardware.

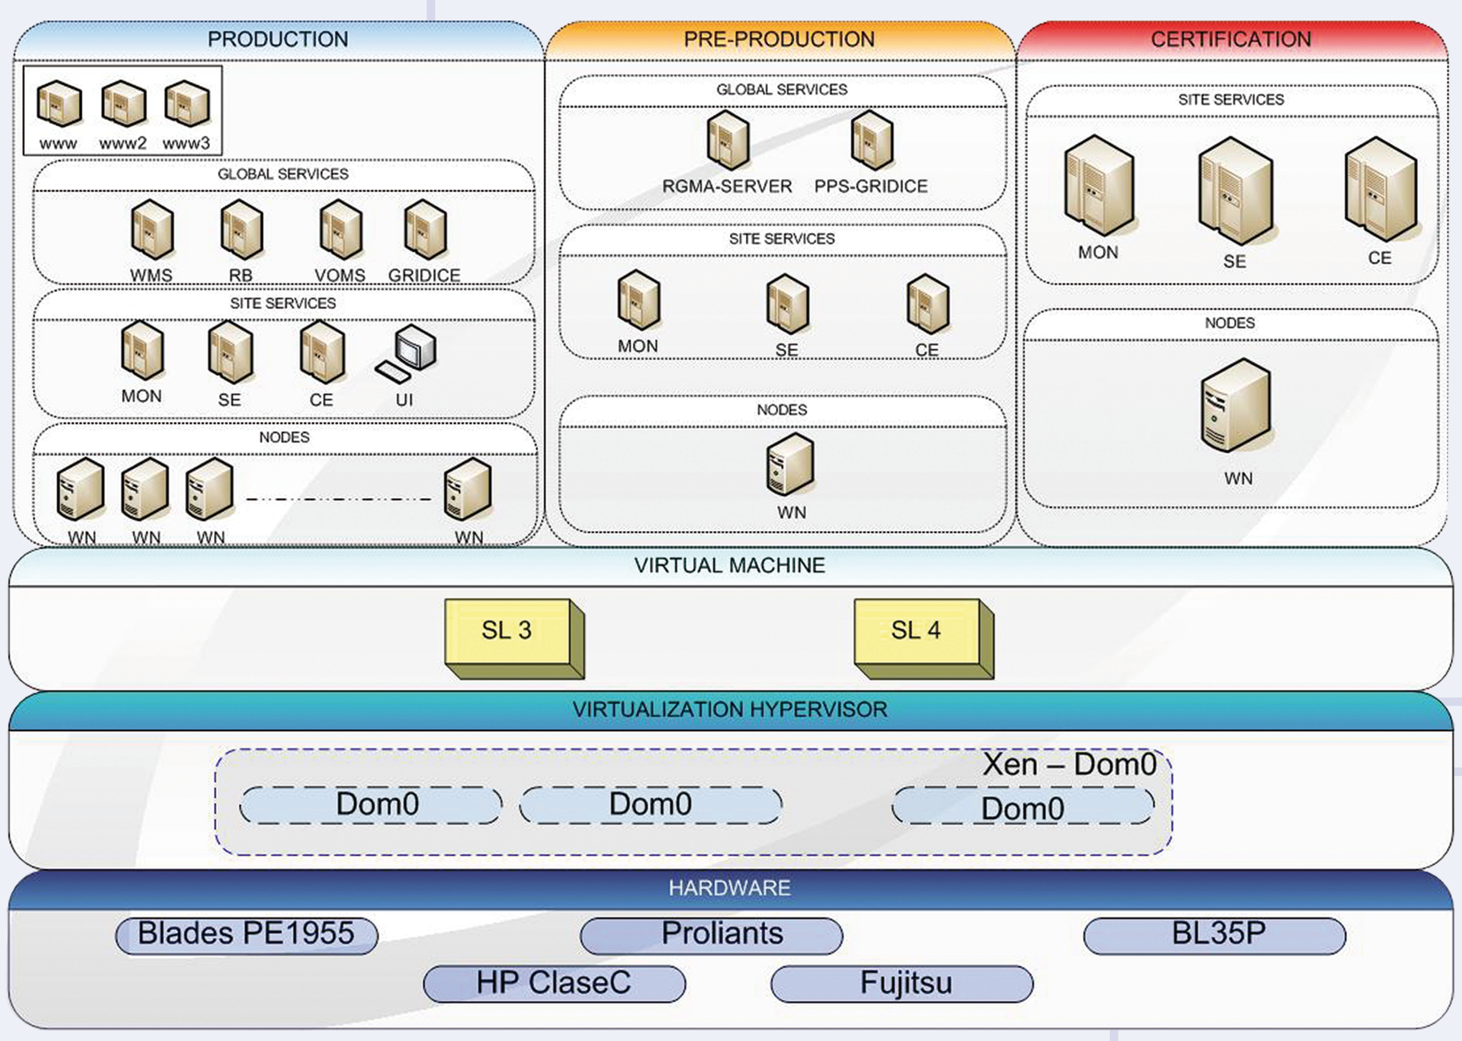

Benchmarks

To study the impact on performance of why virtualization, extensive processor, network and disk benchmarks were executed, both in physical machines and Xen VM inside the EGEE grid environment. The final results show that only a 3% performance loss is incured by virtual machines. In the case of HVM machines the benchmarks showed that the I/O performance was highly degraded so we HVM should only be used in special cases where I/O is not critical.

		Output		Input	
	File Size	Per ch. (K/sec)	Per bl. (K/sec)	Per ch. (K/sec)	Per bl. (K/sec)
Real	1Gb	41327	350687	55085	1848108
ParaVir	1Gb	35378	285954	41540	1316149
HVM	1Gb	10222	8952	15171	14608

Table1.
Bonnie ++ (disk I/O) benchmark results for virtual and real machines

Problem Size	Virtual machine (GFlops)	Real machine (GFlops)
12800	4.374	4.384
6400	4.190	4.207

Table 2.
Linpack (HPL) benchmark results

USE CASES

Formiga Project

FORMIGA and GFLUXO are two regional projects that aim to help universities to take advantage of the computational capabilities of their computer labs when they are idle by making them available to the researchers to perform computational simulations. This objective is accomplished by virtualizing the resources in the computer labs and integrating them in the global virtualized infrastructure of CESGA. These resources are them made available for researchers through the EGEE platform. Researchers can then take advantage of the resources by using a standard EGEE UI or through a visual web portal developed inside GFLUXO project

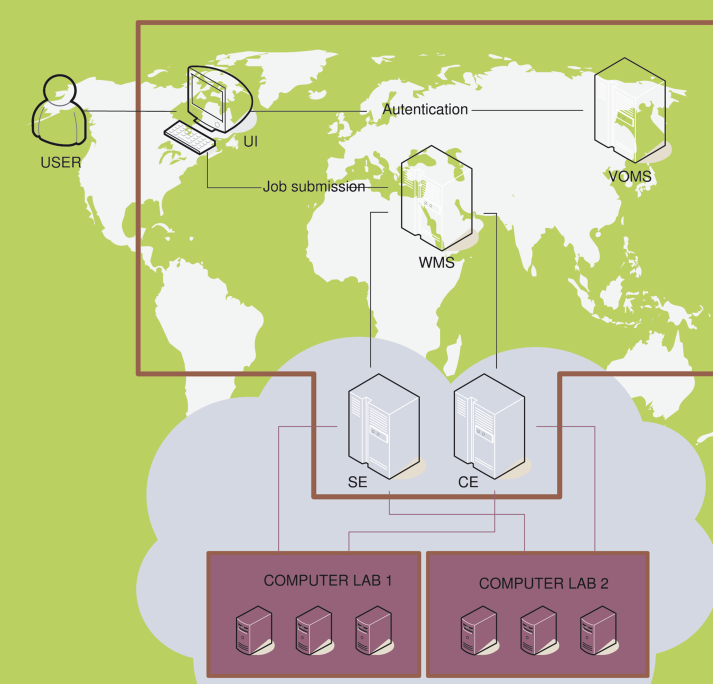

EIMRT

eIMRT is a project which raises new systems of planning for the treatment of the cancer by means of the radiotherapy using distributed calculation and WEB interfaces. In the early stages, eIMRT core application was running in a cluster of computers, submitting jobs through a local scheduler. Soon the jump to the grid was addressed, eIMRT application was adapted to use gLite and Globus Toolkit 4 middleware through the virtual architecture provided by the CESGA.

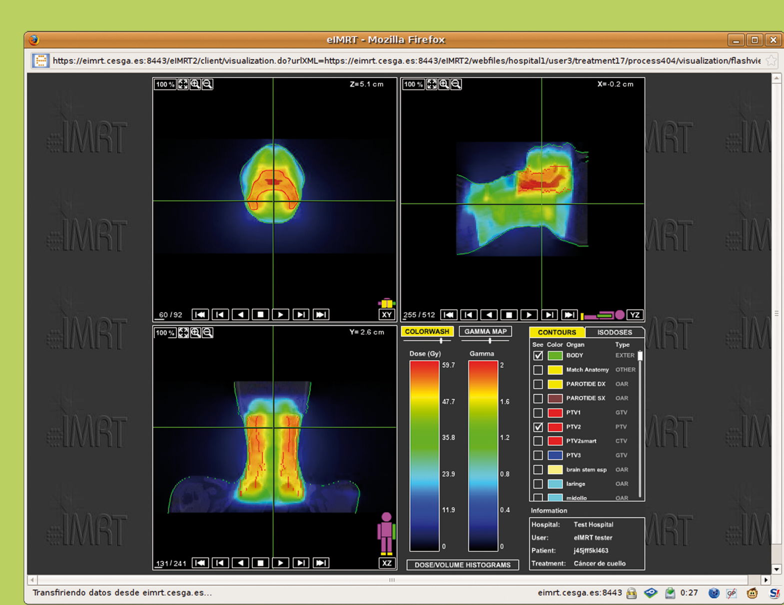

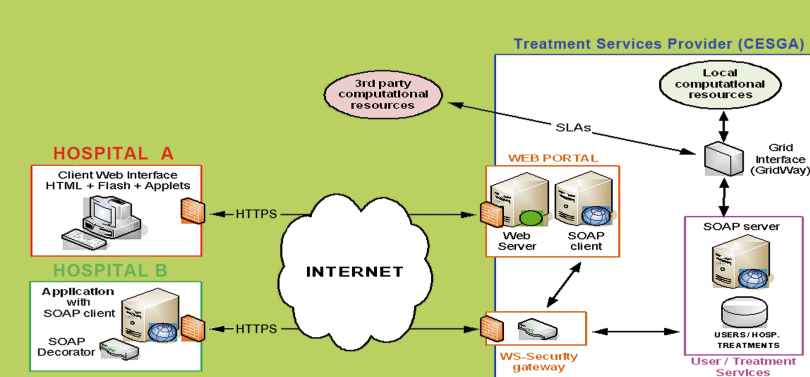

IMATH

IMATH is a national project with the purpose of promoting mathematical presence on the international scene and in the Spanish system of science. CESGA has collaborated with IMATH in the deployment of the infrastructure extending our virtual platform to help different Spanish Mathematics Departments to share their resources. Later, these resources has been integrated in the EGEE project and the IMATH VO has been created. Now, they are able to use them through a standard EGEE UI.

Álvaro Simón García
Carlos Fernández Sánchez
Carlos Fernández Iglesias
David Cordero Placer
Esteban Freire García
Javier López Cacheiro
Pablo Rey Mayo
Sergio Díaz Montes

The EGEE project is building a Grid infrastructure for the scientific community. Grids are networks of computers spread across many sites but able to act together to provide a range of large scale facilities, from incredible processing power and mass storage to a platform for international collaboration.

CESGA http://gfluxo.cesga.es http://formiga.cesga.es

http://eimrt.cesga.es

http://i-math.cesga.es

www.eu-egee.org

EGEE-III is co-funded by the European Commission under contract number INFSO-RI-222667