



Enabling Grids for E-science

Usage of virtualization in gLite certification

Andreas Unterkircher
CERN Grid Deployment

www.eu-egee.org



EGEE and gLite are registered trademarks

- **Overview of the gLite certification process**
- **Benefits for using virtualization in the process**
- **Virtualization tools**

- **GLite uses a continuous release process. Services are updated individually on top of a baseline release.**
- **Updates are added via a patch**
 - We use Savannah for bug and patch tracking
 - Has one or more bugs attached
 - Can also be used to introduce new features or services
 - Has all relevant information: OS, architecture, affected services, baseline release, configuration changes, rpm lists etc.
- **Patch is being certified**
 - Update and configure affected services
 - Run basic and regression tests
 - Verify if attached bugs are fixed and write regression tests.
 - Put patch into “certified” or “rejected”. For the former the patch is ready for release to pre production and later for production.

- **We operate a certification testbed**
 - Contains all gLite service types
 - Provides SLC3 and SLC4 (x86 and x86_64)
 - Is self contained with own BDII, CA and VOMS server
 - Around 60 nodes at CERN
 - 6 external sites provided by SA3 partners

- **Certification of several patches at the same time can cause conflicts.**
- **A non functional patch may spoil the whole testbed**
- **Patch certification often fails already at an early stage (rpm installation, configuration)**
- **A failed patch can pollute a machine. A complete reinstallation is necessary.**
- **Many scenarios and interactions have to be considered: gLite 3.0/gLite 3.1, SLC3/SLC4, x86/x86_64**

- **Quickly start and stop VMs via web interface**
- **VMs with different Linux flavours and architectures: SLC*, Debian,... x86 and x86_64**
- **Usage of predefined hostnames for VMs**
- **Start/stop several VMs at the same time**
- **Well integrated with CERN computing center**
- **Start a grid site with one click**
 - Must pass information about the service to the VM
 - A script is placed into the image. Running this script will enable the desired yum repository, produce a site-info.def for yaim and finally run the installation and configuration.

Activity started in summer 2006. Within 2 months we already had a running system (thanks to collaboration with openlab and Netops)

- **Xen VM management system VNODE**
- **Image generation tool libfsimage**
- **OSFarm web interface to libfsimage**
- **Quattor profiles for SLC4 Xen hosts running VNODE**

We collaborate with various groups in CERN IT

- Openlab: OSFarm
- FIO: Quattor profiles, SLC Xen support
- Netops: hostnames for VMs
- IS: MS Virtual Server was considered

- Python library for generating Linux file systems and populating them
- Produces a tar.gz file that can be used as an image to boot with Xen
- Supported distributions: SL(C)3/4/5, Debian, Ubuntu, CentOS, Fedora on x86 and x86_64
- Command line interface available, web interface OSFarm
- OSFarm allows layered generation of images in order to reuse already existing (parts of) images
- Available in the xenvirt module in CERN's CVS

File Edit View Go Bookmarks Tools Help

http://oplaslim26.cern.ch/advanced

Zeitungen SA3 < EGEE < TWiki WebHome < LCG < T... WebHome < EGEE < ...

OS Farm

Repository About Log Status Simple request Virtual Appliances request Advanced request

OS Farm dynamically generates OS images, and "virtual appliances" for use with Xen VMs. To create an image, enter a name for the image and select a "Class" and software packages if needed. Click "Create image...", and the image will be created and put in the [repository](#). If you check the "Synchronous" checkbox, the image will be downloaded when the image creation is finished.

If you want to use wget, then here is an example url:
`"http://www.cern.ch/osfarm/create?name=&transfer=http&class=SLC4&arch=i386&filetype=.tar&group=core&group=base&package=glite-BDII"`

Please report [bugs](#) and problems on the SourceForge [project page](#).

Name

Synchronous

Class

Architecture

Software packages

Parse repository

Software packages

Select	Yum	Groups
<input type="checkbox"/>	core	
<input type="checkbox"/>	base	
<input type="checkbox"/>	printing	
<input type="checkbox"/>	base-x	
<input type="checkbox"/>	dialup	
<input type="checkbox"/>	gnome-desktop	
<input type="checkbox"/>	kde-desktop	
<input type="checkbox"/>	graphical-internet	
<input type="checkbox"/>	documentation-guides	
<input type="checkbox"/>	text-internet	

Done

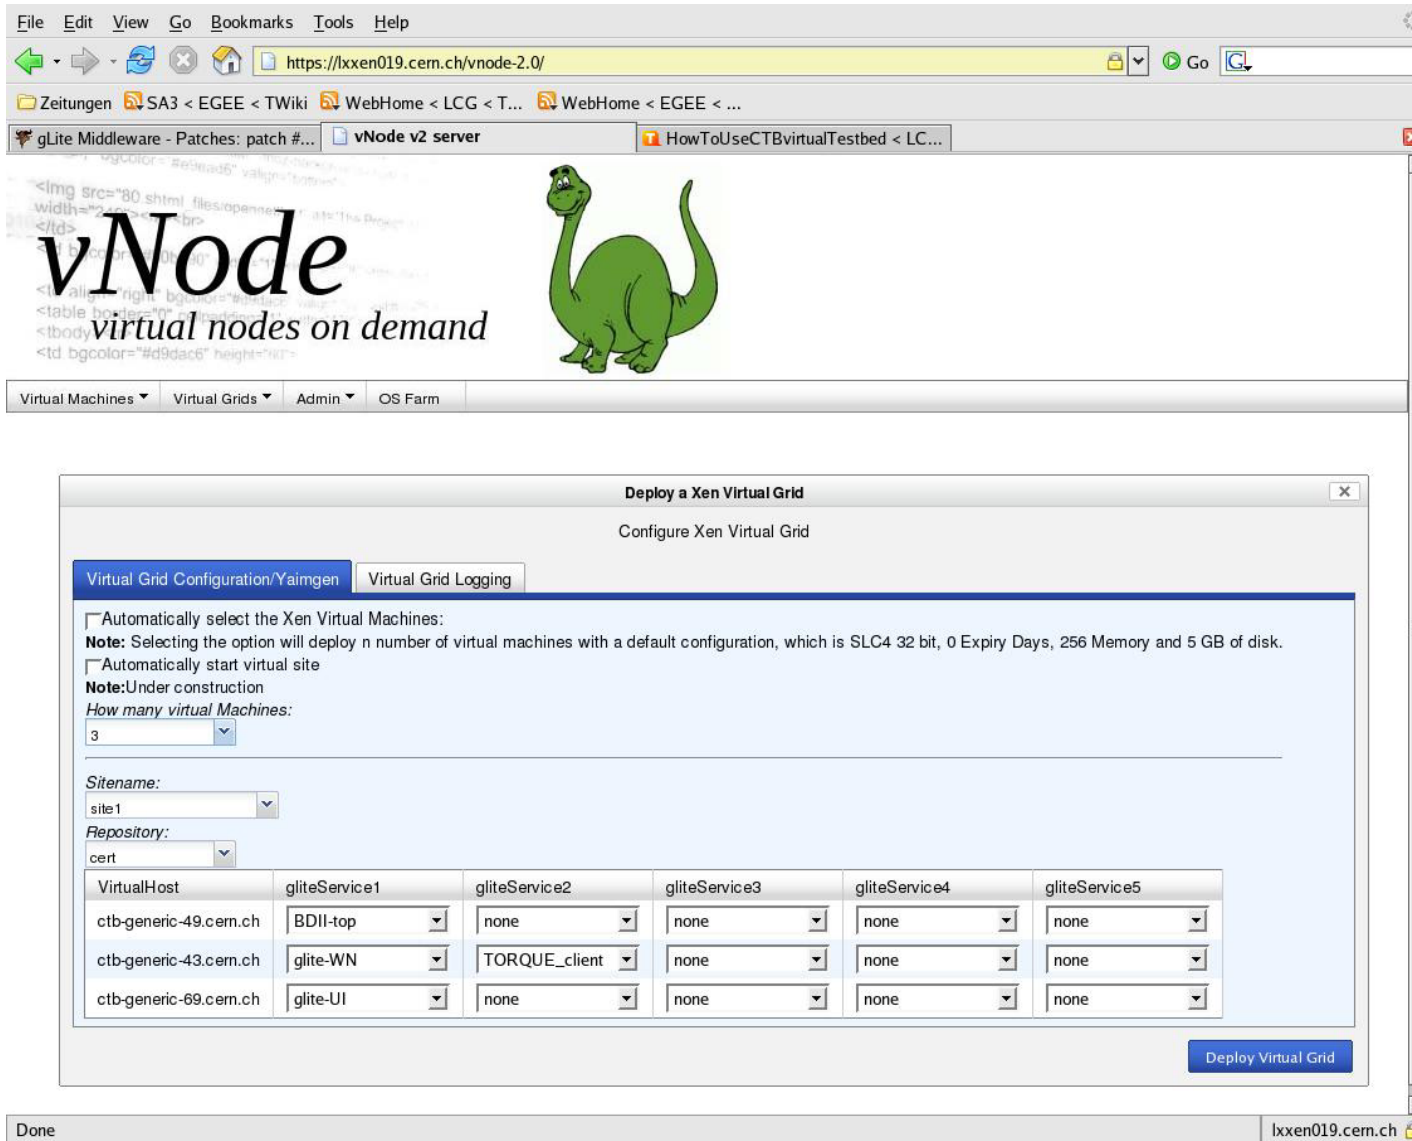
- **Management tool for Xen based VMs**
 - Start VMs with parameters: hostname, RAM, HD size, Linux distribution, architecture, expiry time
 - Start/stop several VMs with one click
 - Start complete grid site with one click
- **Technologies used: Python, Java Script (YUI), https requests to communicate**
- **Information:**
<https://twiki.cern.ch/twiki/bin/view/Virtualization/VirtualNodesOnDemand>

The screenshot displays the vNode web interface with the following components:

- Header:** "vNode virtual nodes on demand" with a green dinosaur logo.
- Navigation:** Virtual Machines, Virtual Grids, Admin, OS Farm.
- Deploy a Xen Virtual Machine:**
 - Configuration fields: Physical Hosts (lxxen0x2.cern.ch), Virtual Hostnames (ctb-generic-69.cern.ch), Expiry Time (2 days), Memory (256 MB), Partition (10 GB), OS Image (SLC-4-32).
 - Advanced Options: Image Filename, Virtual Machine Name (Xusr0ognh2wqd).
 - Table of existing VMs:

Name	PhysicalHost	VirtualHost	ExpiryTime	Memory	Partition	OSImage
blag	lxxen007.cern.ch	ctb-generic-49.cern.ch	0	256	5	SLC-4-32
blag	lxxen007.cern.ch	ctb-generic-43.cern.ch	0	256	5	SLC-4-32
blag	lxxen007.cern.ch	ctb-generic-69.cern.ch	0	256	5	SLC-4-32
 - Buttons: Deploy, Add, Delete, Refresh.
- State of Xen Virtual Machines:**
 - Table showing VM status:

VirtualHost	StateMessage	Message	StateColor
ctb-generic-49.cern.ch	available	None	Green
ctb-generic-43.cern.ch	available	None	Green
ctb-generic-47.cern.ch	available	None	Green
ctb-generic-69.cern.ch	available	None	Green
- Terminate a Xen Virtual Machine:**
 - Buttons: Refresh, Terminate.



The screenshot shows a web browser window with the URL `https://lxxen019.cern.ch/vnode-2.0/`. The page content includes the text "vNode virtual nodes on demand" and a green cartoon dinosaur. Below the main content is a navigation menu with items: Virtual Machines, Virtual Grids, Admin, and OS Farm.

Overlaid on the browser is a configuration window titled "Deploy a Xen Virtual Grid". The window has two tabs: "Virtual Grid Configuration/Yaimgen" (selected) and "Virtual Grid Logging".

Configuration options in the window include:

- Automatically select the Xen Virtual Machines:
Note: Selecting the option will deploy n number of virtual machines with a default configuration, which is SLC4 32 bit, 0 Expiry Days, 256 Memory and 5 GB of disk.
- Automatically start virtual site
Note: Under construction
- How many virtual Machines:
- Sitename:
- Repository:

A table lists the configuration for three virtual hosts:

VirtualHost	gliteService1	gliteService2	gliteService3	gliteService4	gliteService5
ctb-generic-49.cern.ch	BDII-top	none	none	none	none
ctb-generic-43.cern.ch	glite-WN	TORQUE_client	none	none	none
ctb-generic-69.cern.ch	glite-UI	none	none	none	none

A "Deploy Virtual Grid" button is located at the bottom right of the configuration window.

The browser status bar at the bottom shows "Done" and the address "lxxen019.cern.ch".

- **12 Xen SLC4 hosts, quattor managed**
 - Elonex 2.8GHz: 2 CPU, 2GB RAM, 160 GB HD
 - On the same switch (for shared VM hostname usage)
- **Base Linux images are stored locally on the Xen hosts and updated regularly**
- **40 hostnames for VMs**
- **Around 15 users**
- **The system is in heavy use**

- Harden the “grid with a click” feature so that it can be an integral part of the certification process
- Shared file system for the base Linux images
- Better handling of hostnames for VMs from the network point of view; probably use SOAP API to landb
- Automatic update of the Linux base images using OSFarm
- Image saving currently is not needed
- Collaborate with other groups at CERN to find a unified virtualization solution
- Information about virtualization as used by FIO:
<https://twiki.cern.ch/twiki/bin/view/FIOgroup/VirtualizationAtFIO>

