Usage of virtualization in gLite certification

Andreas Unterkircher
CERN Grid Deployment

Information Society and Media

Contents

1. gLite certification
2. Virtualization with the CERN gLite certification team.
3. Image generation
4. SmartDomains: Xen management with SmartFrog
5. Our portal: vGrid
6. Screenshots
7. Final remarks

gLite certification

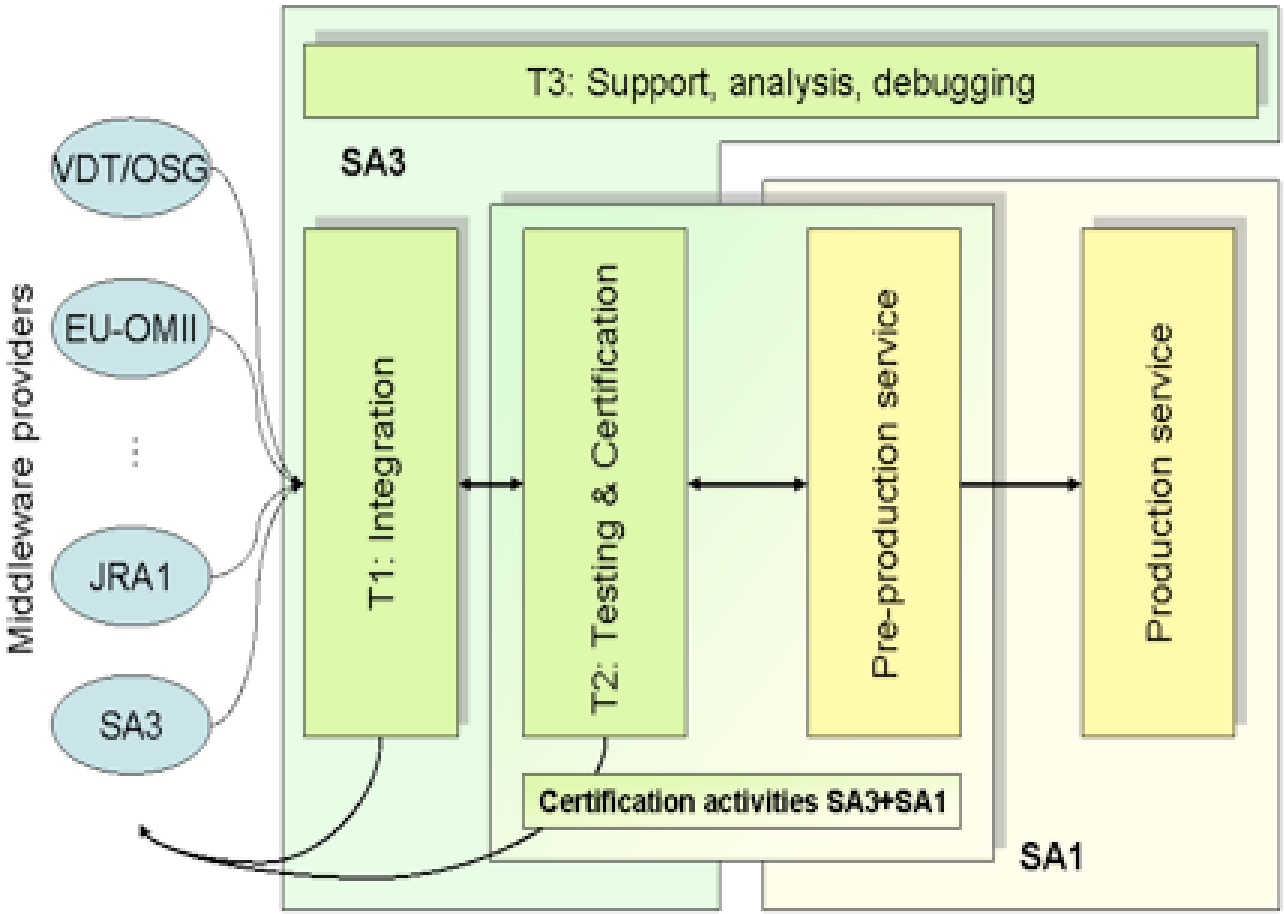

Certification testbed

- Approx. 60 machines at CERN plus several external sites.
- All gLite services are present.
- Daily regression tests.
- Installation (rpm) and configuration of patches.
- Special tests for patches.

Frequent problems

- **Certification of several patches at the same time can cause conflicts but we have to certify patches in parallel.**
- **A non functional patch may spoil the whole testbed.**
- **Patches often fail already at an early stage.**
  - Rpm installation fails.
  - Configuration fails.
- **Failed patches can pollute a machine, a complete reinstallation might be necessary.**
- **For testing one might want to switch between different versions of a service quickly.**

Virtualization as a solution

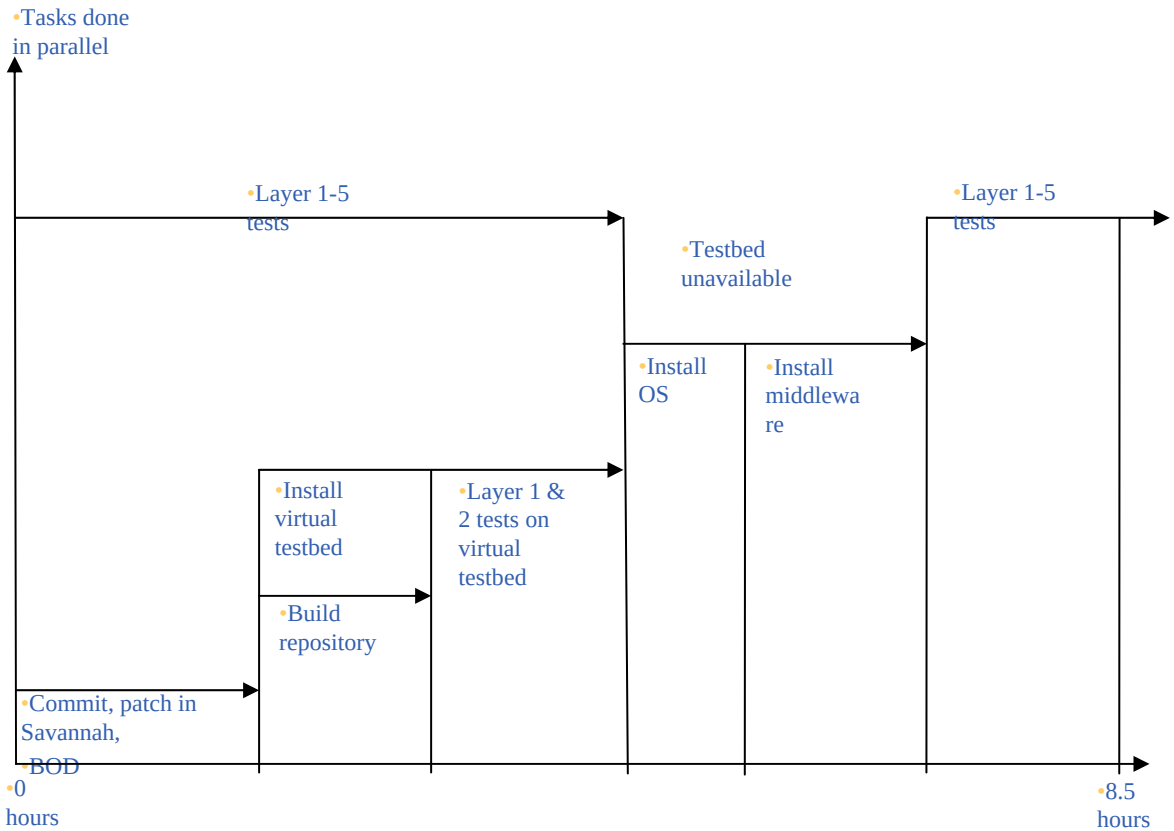

Usage of virtualization to prevent patches failing at installation or configuration step to stop the whole certification process

Virtualization at CERN CTB needs

- OS: SLC3, SLC4 others to come according to needs (Debian).
- Virtualization tool: Xen.
- Need to generate images easily.
- Management of VMs: start/stop several machines (whole grid site) with one click.
- Management of images: save/retrieve images for later use.

Tools developed at CERN (ongoing)

- Custom image generation for SLC3, SLC4 and Debian on SL machines.
- Management of Xen virtual machines with SmartFrog (SmartDomains project).
- vGrid portal to manage the whole chain (image generation and storage, VMs).
- IT Linux support team provides Xen rpms.
- Done in collaboration with CERN openlab.

Image generation problems

- How to make sure that some image boots under Xen ?
- Xen domU kernel must match the image.
- Image might start services that don't work in the guest environment.
- Networking might not work in the guest environment. For gLite services we need to know the hostnames in advance.
- Image generation should be possible without (re)installing a physical machine…

Image creation scripts

- **Bash scripts that generate a complete SLC3, SLC4 or Debian distribution under some directory and produce a .tar.gz or .img.**
- **Support for yum/apt groups make it possible to produce special images easily (e.g. glite-UI on SLC4).**
- **Further customization:**
  - Also install individual rpms.
  - Place files into the image (e.g. ssh keys).
  - Set a root password.
- **Network is configured to take parameters from Xen.**
- **Produces all the necessary files to ensure that the image boots (/dev,/etc/fstab,…).**

Image creation scripts

- Can be launched by the user on his/her machine (but you need to be root).
- Driven by command line options.
- Can be easily integrated into some portal.
- Can be run in a VM.

VM management with SmartFrog

- Smart Framework for Object Groups
- Developed by HP Labs Bristol.
- Java, Open Source.
- Describes distributed software systems as collections of cooperating components and allows to activate and manage them.

SmartFrog elements

- **Language**
  - templates / descriptions
- **Engine**
  - interprets descriptions to activate running services
- **Components**
  - make up the running service
  - deployed, configured & activated by the engine

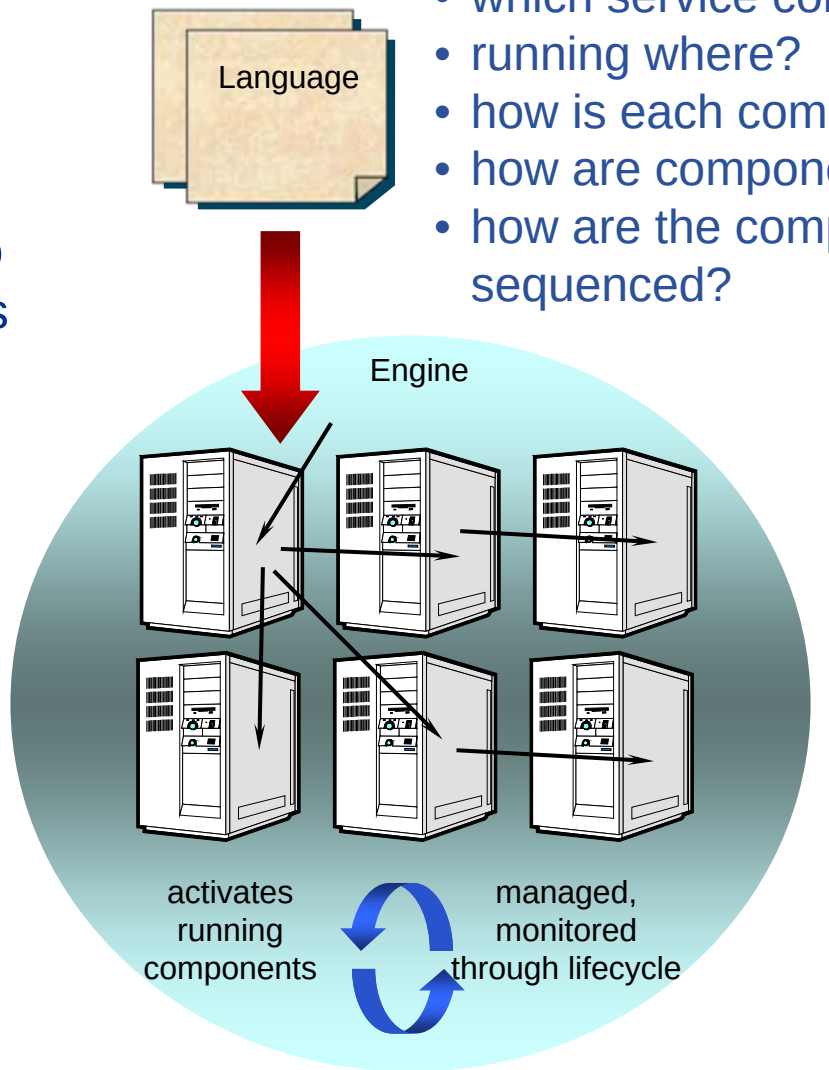

- which service components?
- running where?
- how is each component initialised?
- how are components related?
- how are the component lifecycles sequenced?

Slide taken from SmartFrog Overview Presentation on http://www.smartfrog.org

Managing Xen with SmartDomains

- **SmartFrog components to manage Xen VMs.**
- **Developed at CERN openlab.**
- **Two main components**
  - StorageBackend (LVM or loopback): contains image, creates/deletes images, possibility to save image.
  - XenDomain: manages domU (start, stop).

SmartDomains example

```
sfConfig extends Compound {
 sfProcessHost “host2.cern.ch”;
 myShell extends BashShell;
 lvm extends LVMStorageBackend {
  shell LAZY ATTRIB myShell;
  baseImage “/data/slc3.tar.gz”;
  volumeSize “1G”;
  swapSize “512M”;
  volumeBaseName “xen-domain”;
  usingExistingVolumes false;
  keepVolumes “false”;
  saveImage true;
  saveImagePath “/tmp”;
  saveImageExtension “tgz”;}
```

```
Vm1 extends XenDomain {
 shell LAZY ATTRIB myShell;
 storageBackend LAZY ATTRIB lvm;
 domainName “vm1”;
 hostname “vm1.cern.ch”;
 ip “111.222.333.444”;
 netmask “255.255.0.0”;
 kernel “/boot/vmliuz-xen”;
 ramdisk “/boot/initrd-xen”;
 memory 512;
 extra “fastboot”
 domainLivenessDelay 2000;
 domainLivenessFactor 3;
 domainLivenessCheck true;
 }}
```

```
Commandline: sfStart host1.cern.ch myVM textfile.sf
```

VM management with SmartDomains

- **Manages the lifecycle of one VM session: put image in place, start VM, stop VM, delete/leave/save image.**
- **Error handling: if start of Xen domU fails, the image is also being deleted.**
- **Leverages several SmartFrog features:**
  - Include directives and overriding of attributes allows for easy to read configuration files. Default values can be grouped in common include files.
  - Start/stop several VMs on different physical machines with one click. Other workflows also possible.
  - Management console allows to view and change the status of the deployed VMs.
  - SmartFrog security features.

Current setup

- **10 SLC4 machines with Xen 3.0.1, LVM and SmartFrog.**
- **28 hostnames/IP numbers for use with virtual machines. For gLite services we need known hostnames.**
- **vGrid portal to**
  - Generate SmartFrog configuration files.
  - See which hostnames are currently in use.
  - View log files on Xen servers.
  - On demand image generation (under development).
  - Start VMs (under development).

SmartDomains example

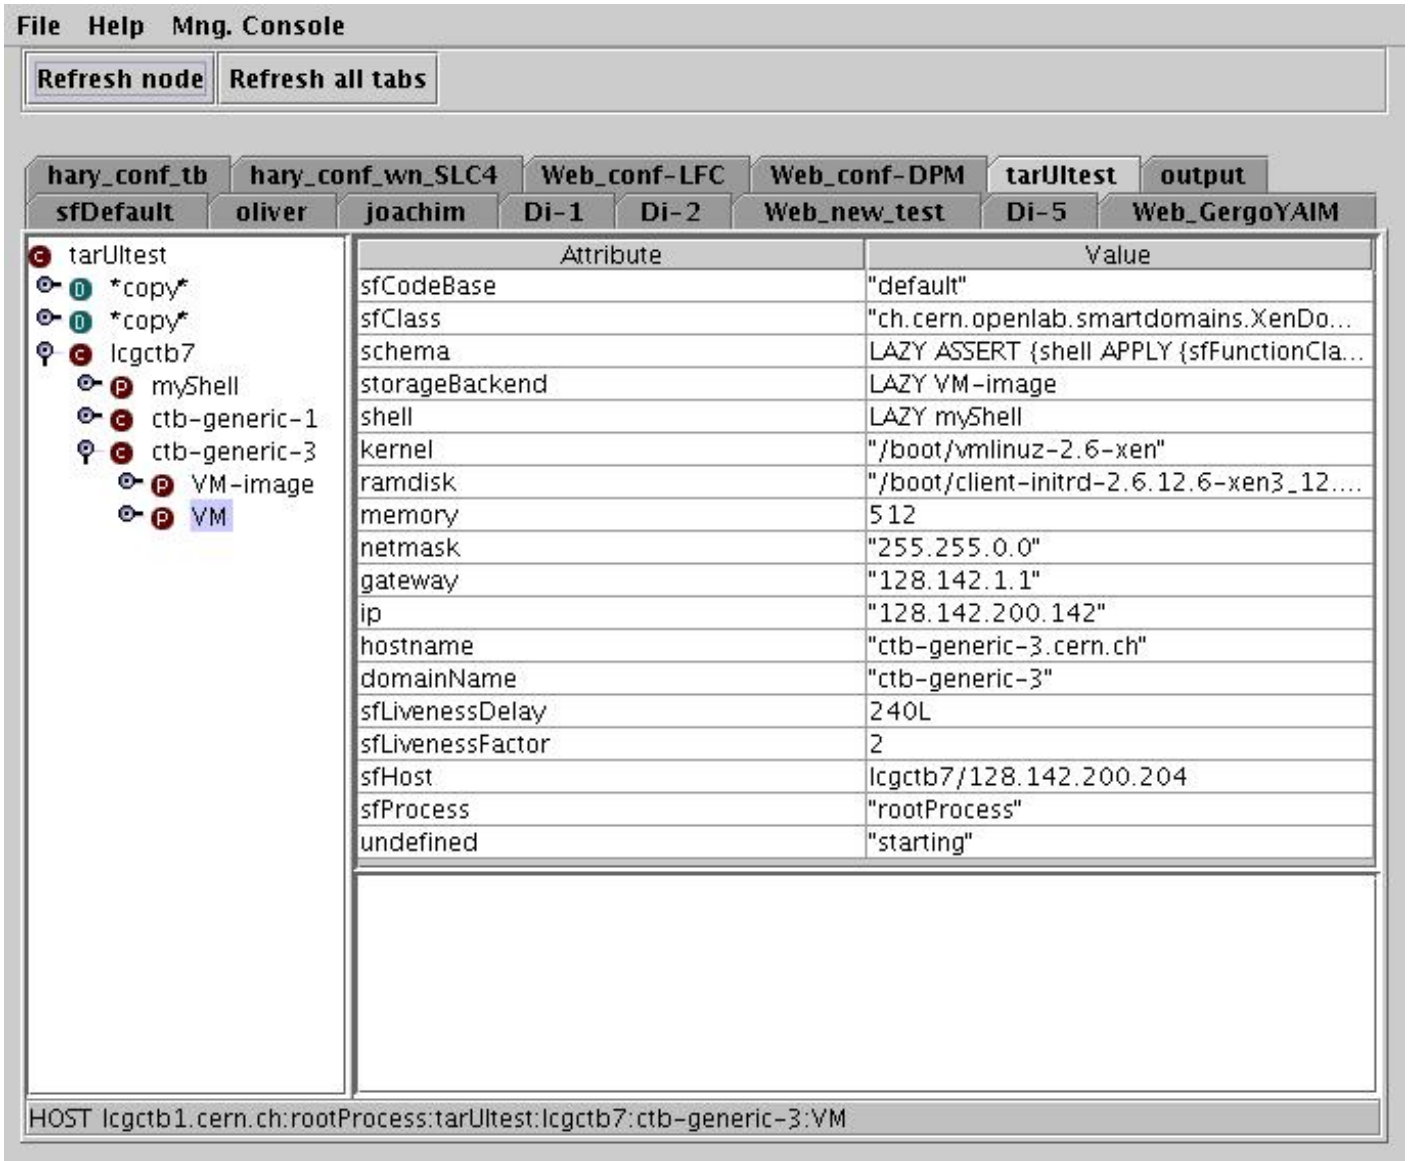

SmartFrog management console

SmartDomains example

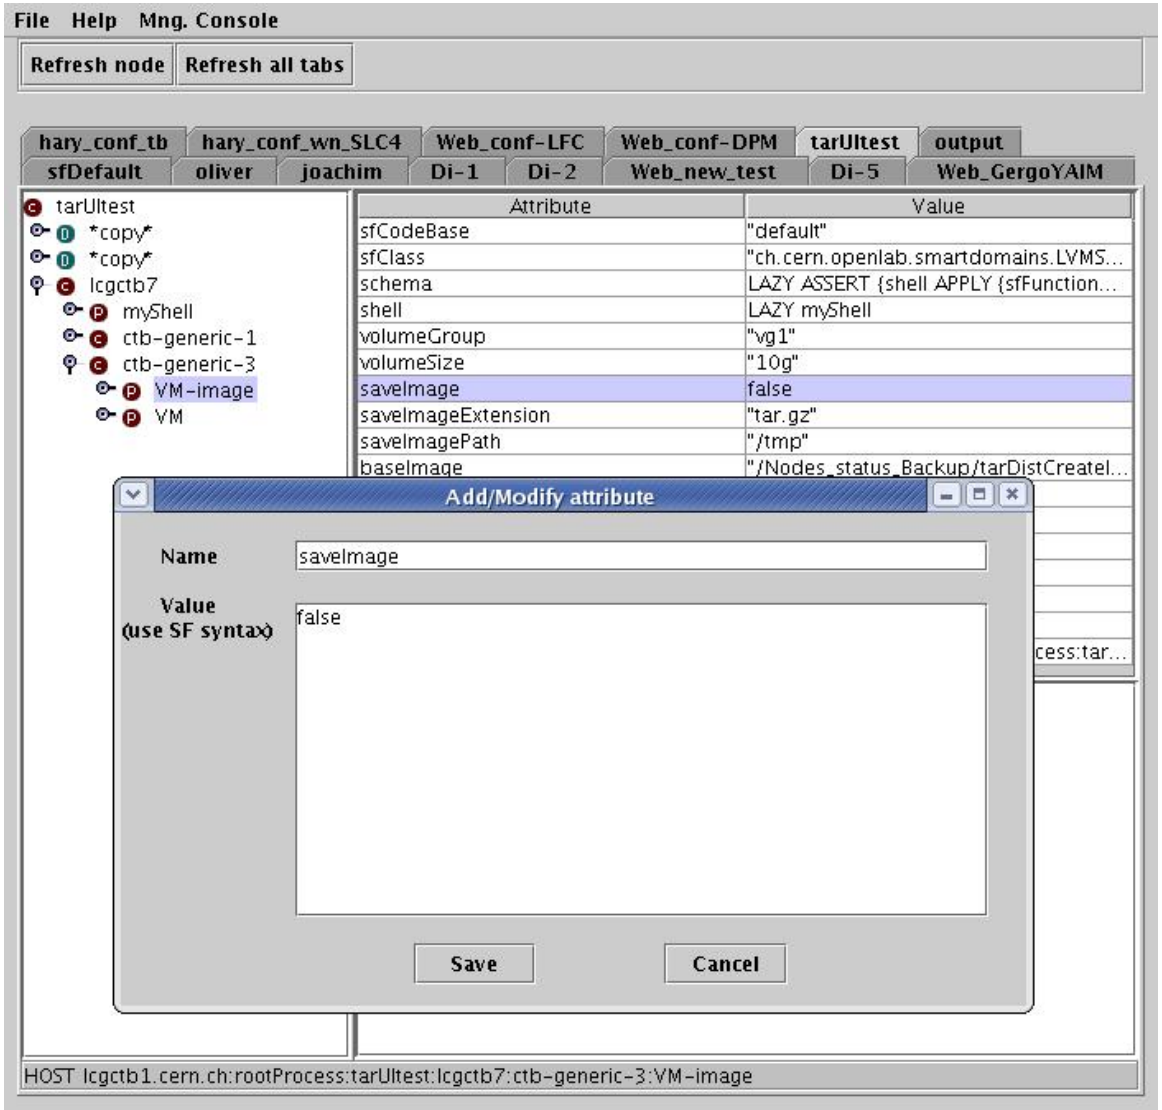

Modification of an attribute while the VM is running.

SmartDomains example

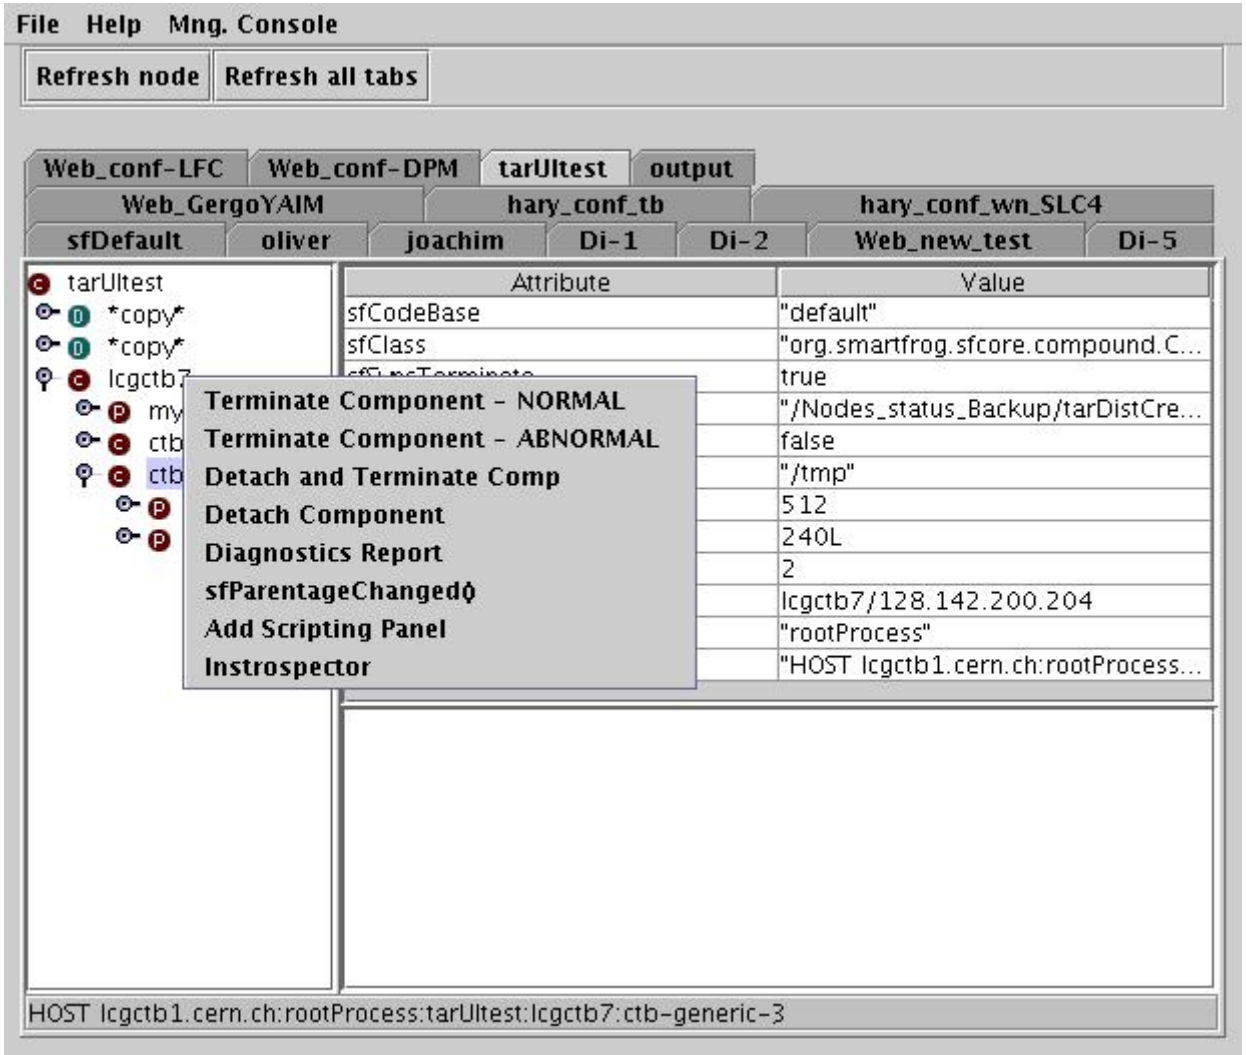

Termination of a VM.

vGrid portal

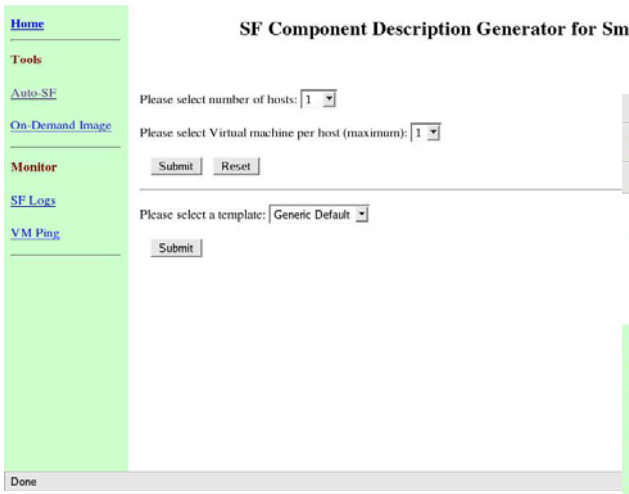

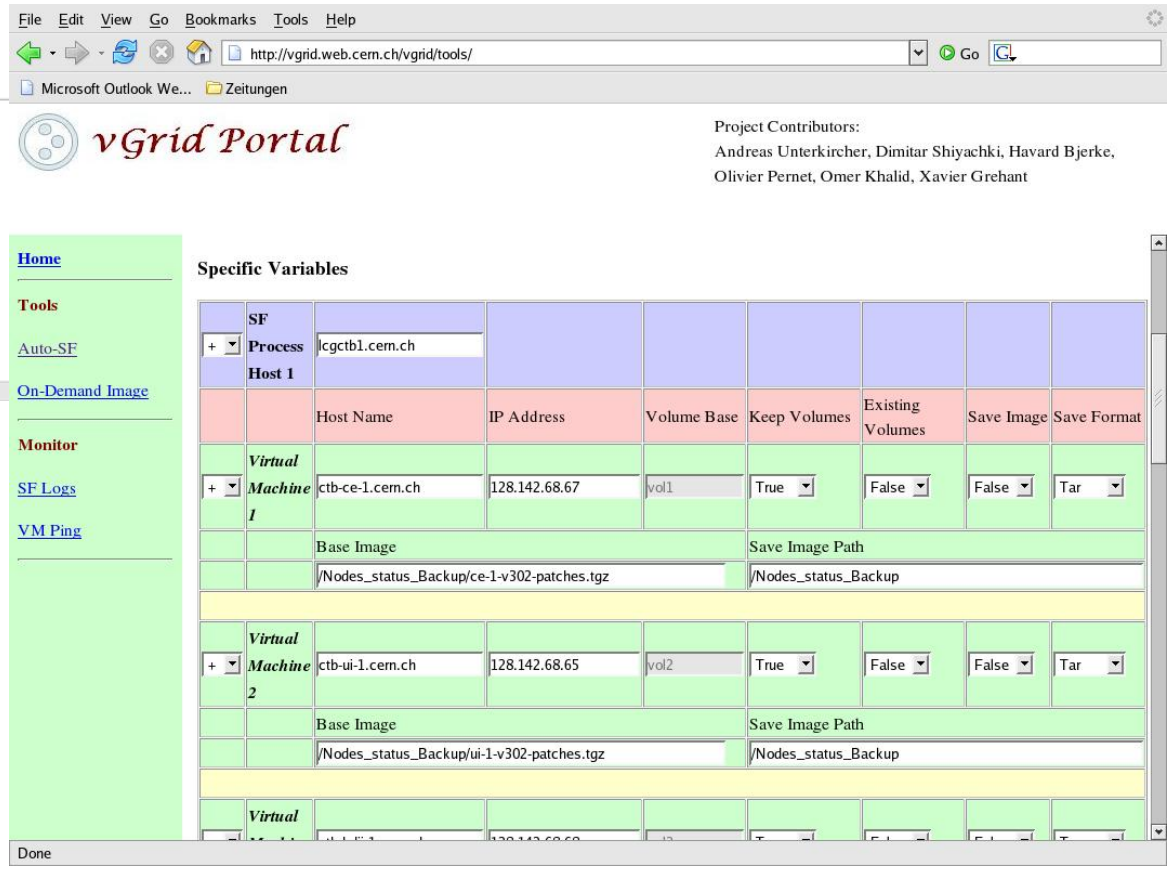

vGrid portal

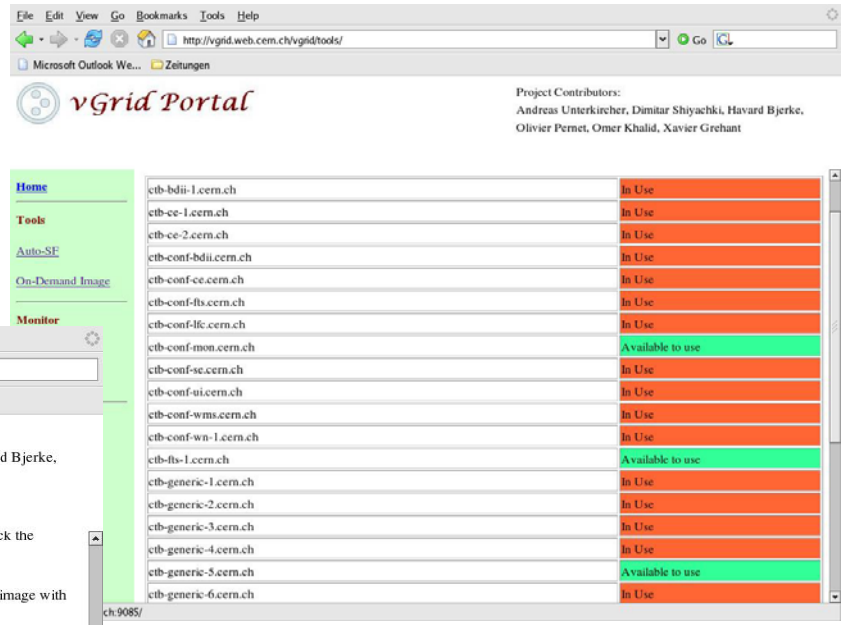

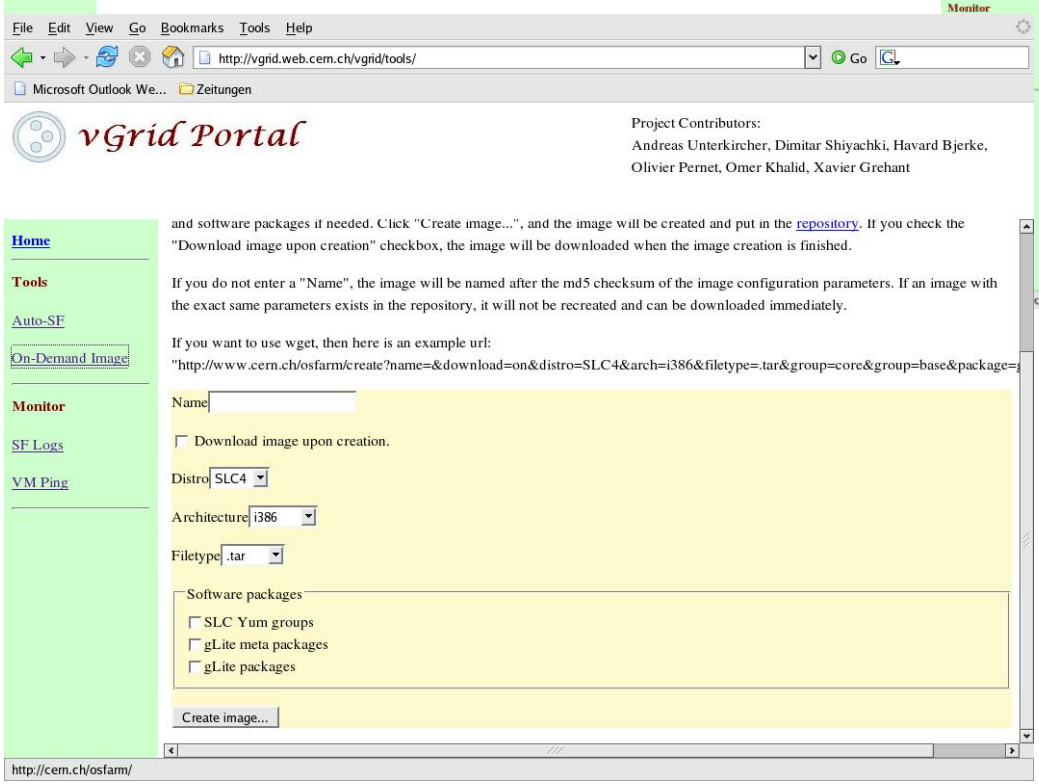

Observations

- **The system is in heavy use since October 2006.**
- **We have basic SLC3/4 images on every Xen server. Users install gLite services on top of them.**
- **About 10 users.**
- **SmartFrog is easy to install and runs on all platforms.**
- **We have no scheduler. Users have to decide which hostname they use on which Xen server.**
- **Some use cases**
  - Creation and testing of tarball UI and WN.
  - Yaim development.
  - A full testbed for WMS patch certification.
  - Testing apt-get dist upgrade.

Next steps

- Upgrade to Xen 3.0.2 using rpms provided by CERN Linux Team. Enables us to have AFS enabled VMs.
- SmartDomains and vGrid are under constant development.
- Image storage.
- On demand image generation.
- Scheduler for VM creation.

Links and credits

SmartFrog: http://www.smartfrog.org/

SmartDomains: http://sourceforge.net/projects/smartdomains/

vGrid: http://vgrid.web.cern.ch

Image creation scripts: http://isscvs.cern.ch/cgi-bin/cvsweb.cgi/vgrid/ch/cern/osfarm/scripts/?cvsroot=vgrid

CERN Linux Support Xen HowTo: https://twiki.cern.ch/twiki/bin/view/LinuxSupport/XenHowTo

CERN openlab: http://cern.ch/openlab

People involved:

Andreas Unterkircher, Omer Khalid, Dimitar Shiyachki, Havard Bjerke, Xavier Grehant, Olivier Pernet, Jarek Polok

Discussion