


Virtualization

8 March 2007

Håvard Bjerke




- Hosted
 - Microsoft Virtualization Server
 - VMWare
- Non-hosted
 - Xen



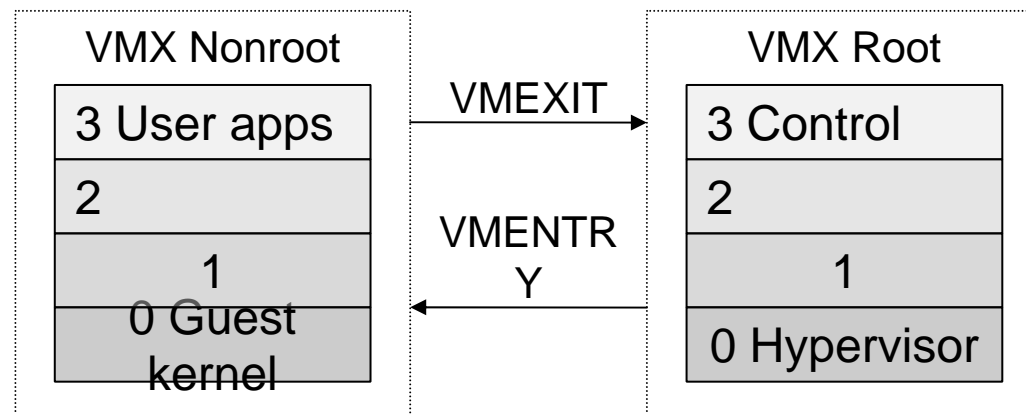
- General application:
 - Server consolidation
- GRID/HPC specific:
 - Secure isolation between VMs
 - Isolation of users and malicious software
 - Software flexibility
 - Let each user manage their own OS
 - And satisfy their own software dependencies
 - Utilisation of SMP and multi-core resources
 - Live Migration of VMs across nodes
 - Checkpointing
 - Utilisation of public computing resources

- Paravirtualization
 - -> Close to native performance
 - Requires kernel modification
- Also supports full virtualization with Intel VT
- Supports IA32, EM64T and IA64 architectures
- Mature and rich feature set
 - Live Migration
 - Flexible CPU/core allocation



- New player, but quite stable
- Takes benefit from hardware extensions for virtualization (VT-x)
- A monolithic hypervisor
 - The hypervisor is compiled as a kernel module and loaded into kernel-space at runtime
- Supports live-migration
- Para-virtualization under development
- IA64 port under development
- No in-VM SMP support

- Virtualization hardware extensions alleviate fundamental performance problems
- EM64T – VTx
- IA64 – VTi
- VTd will give better I/O and MMU virtualization



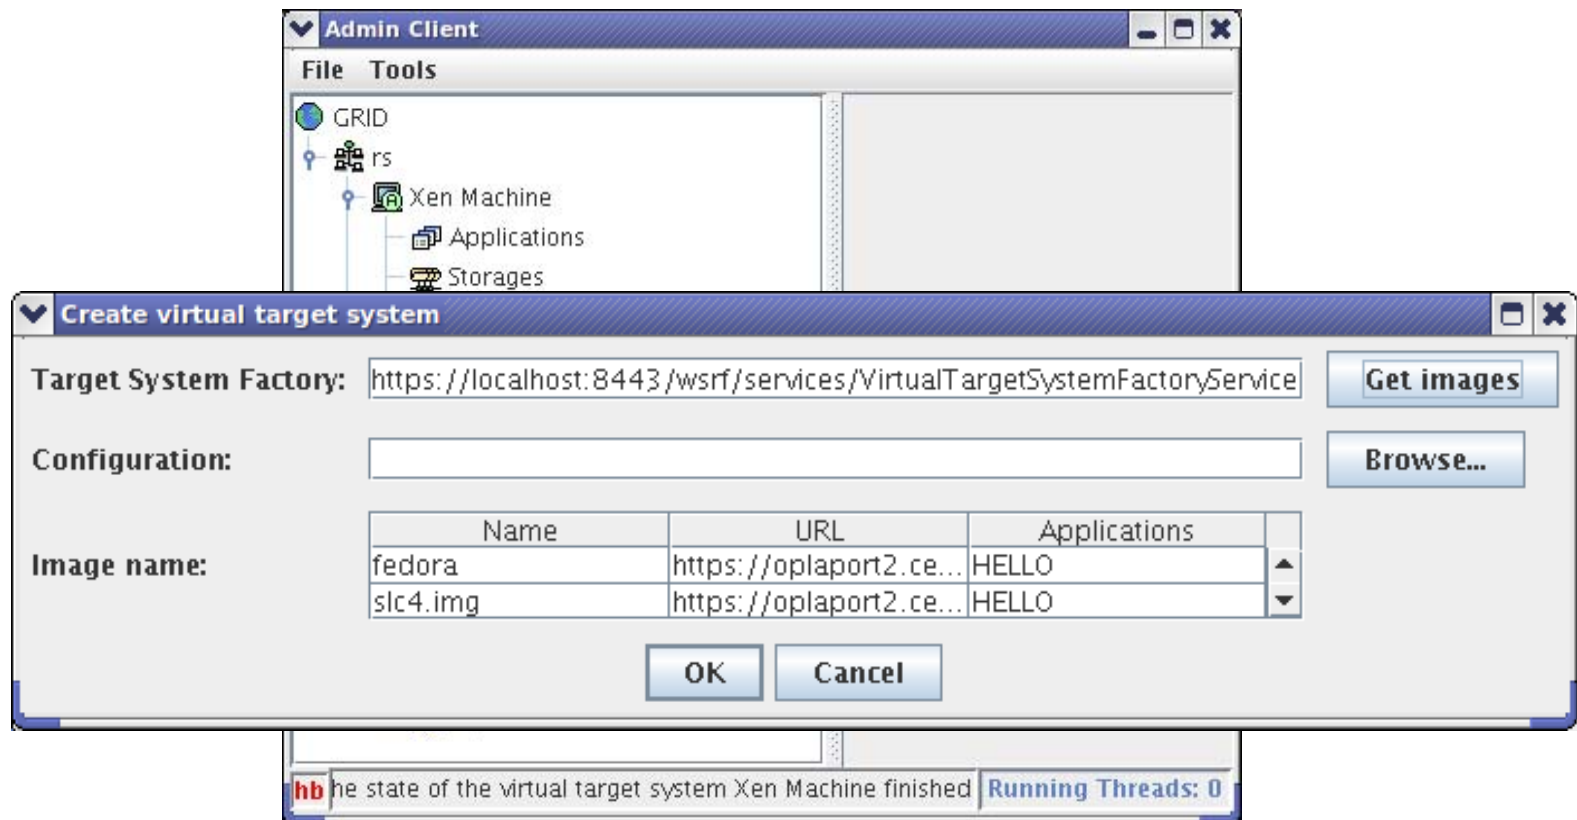


Grid Programming Environment

- API to develop and deploy Grid application beans
- Service-oriented Architecture
 - Target systems (computational resource)
 - Job management
 - Storage management
 - File transfer
- Integrates with Virtual Workspaces for dynamic resource provisioning
- The only Grid middleware to offer full platform virtualization support

- Standards
 - Job Submission Description Language (jobs)
 - Open Grid Forum's standard for specifying requirements for jobs
 - Business Process Execution Language (workflows)
 - Orchestrating interaction between webservice
 - Backed by IBM, SAP, Microsoft
 - Common Information Model (resources)
 - Model description of resources
 - Backed by DMTF

- Can use Virtual Workspaces' built in scheduler to deploy VMs



- OS Farm
- Execution Environment Centric

OS Farm

[Repository](#) [About](#) [Log](#) [Status](#) [Simple request](#) [gLite request](#) [Advanced request](#)

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

If you do not enter a "Name", the image will be named after the md5 checksum of the image configuration parameters. If an image with the exact same parameters exists in the repository, it will not be recreated and can be downloaded immediately.

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

Please allow a few minutes for the image to be created.

Name

Synchronous

Distro

Architecture

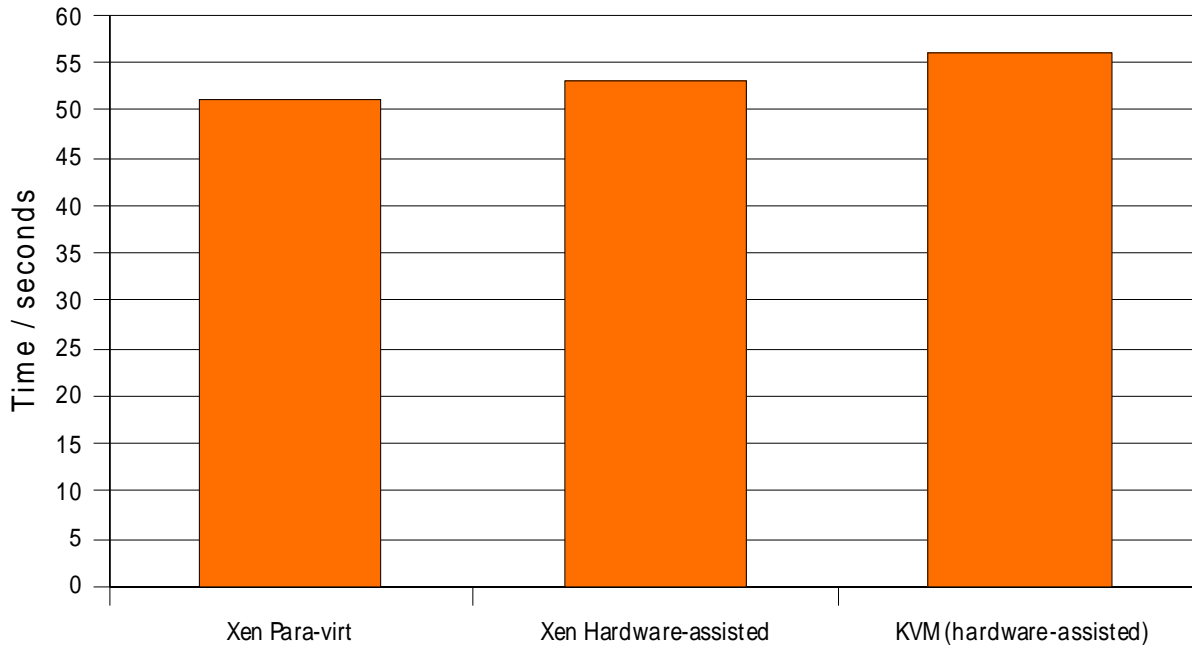
Filetype

Software packages

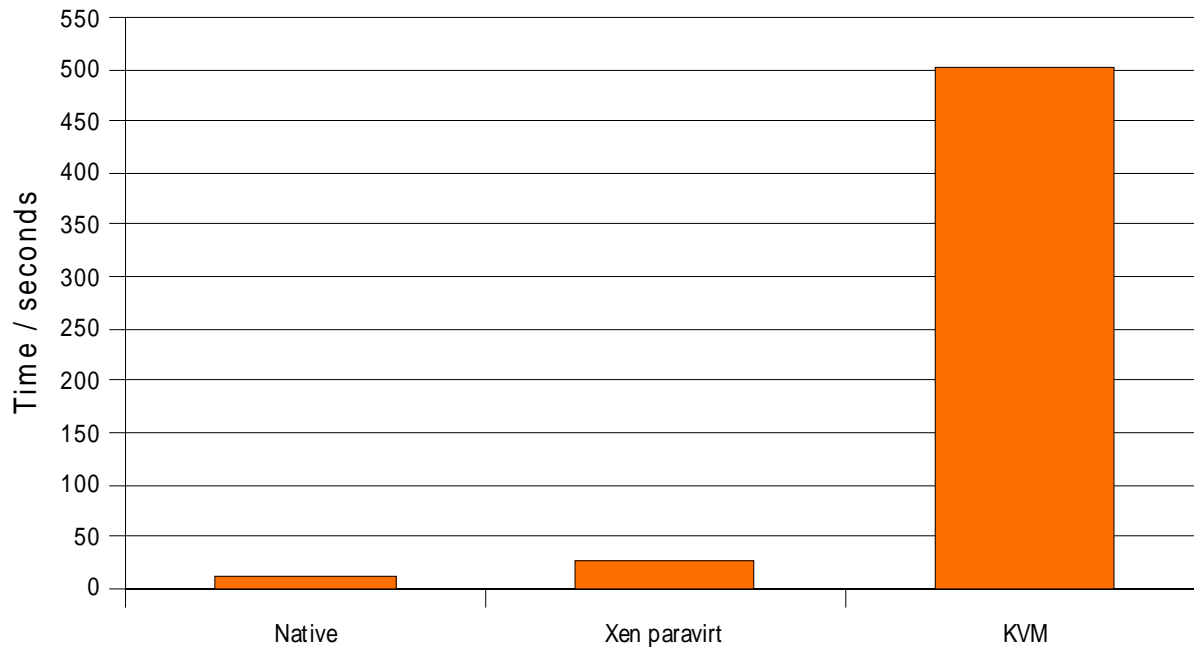
SLC Yum groups

- core
- base
- printing
- base-x
- dialup
- gnome-desktop
- kde-desktop

- CPU-intensive application:
 - Event simulation with Geant4



- I/O-intensive benchmark
 - Write 1 GB file



- MMU microbenchmark
 - 10k page faults

