



# Virtualization in the LHC Grid

Andreas Unterkircher, CERN openlab

1. Overview of virtualization
2. Possible use of virtualization in LCG
3. Work done so far at openlab
4. Virtualization on the Grid and opportunities at openlab

“Virtualization techniques encompass a variety of mechanisms and techniques used to decouple the architecture and user-perceived behavior of hardware and software resources from their physical implementation.”

IEEE Computer special issue on “Virtualization Technologies” (2005)



# Types of virtualization

## Process Virtual Machines (VMs)

Emulators/binary translators  
(Intel IA32-EL on IA64)

High-level-language VMs (Java VM)

## System VMs

Hosted VMs (VMware, same ISA for host & guest)

Para-virtualization (Xen)

Kernel virtualization (VServer)

Whole system VMs (VirtualPC, host & guest ISA differ)

Codesigned VMs (Transmeta Crusoe)

## Higher Level Virtualization

Virtual networks (Violin)

Distributed Virtualization (PlanetLab)

# Virtualization in LCG

At the moment we are mainly interested in system VMs under Linux, i.e. bringing up multiple Linux environments on one host running Linux. This can be achieved with tools like VMware (commercial), VServer or Xen.

# Benefits of virtualization

## Application in LCG:

A minimal LCG site on one node

VO Box

Different Linux distributions for VOs

Catalogs, long jobs

Server consolidation

Security

Customizable environments

Availability (serialization, migration)

# Virtualization tool used

We use the Xen hypervisor:

- + Open source
- + Allows to run different Linux distributions
- + Little overhead
- Not supported on Windows
- Needs a modified Linux kernel.

The situation is expected to improve when virtualization support is added to x86, x86\_64 and IA64 (Intel VT, AMD Pacifica).

# Work done so far

- \* One PC running 3 Linux images hosting a CE, WN and SEclassic runs as a minimal LCG site in the CERN LCG test bed.
- \* Performance tests with HEP software under Xen.
- \* A system has been set up to quickly produce images and test them for use with Xen. An SLC3 image with pre installed LCG software is available now.
- \* On a webpage we describe how to run Xen in the CERN Linux environment (SLC 3/4). This page also contains rpms for easy installation. Thus interested sites should be able to set up LCG nodes with Xen.



# Virtualization and the Grid

## Example: Architecture proposed by the Virtual Workspaces project (Globus Alliance)

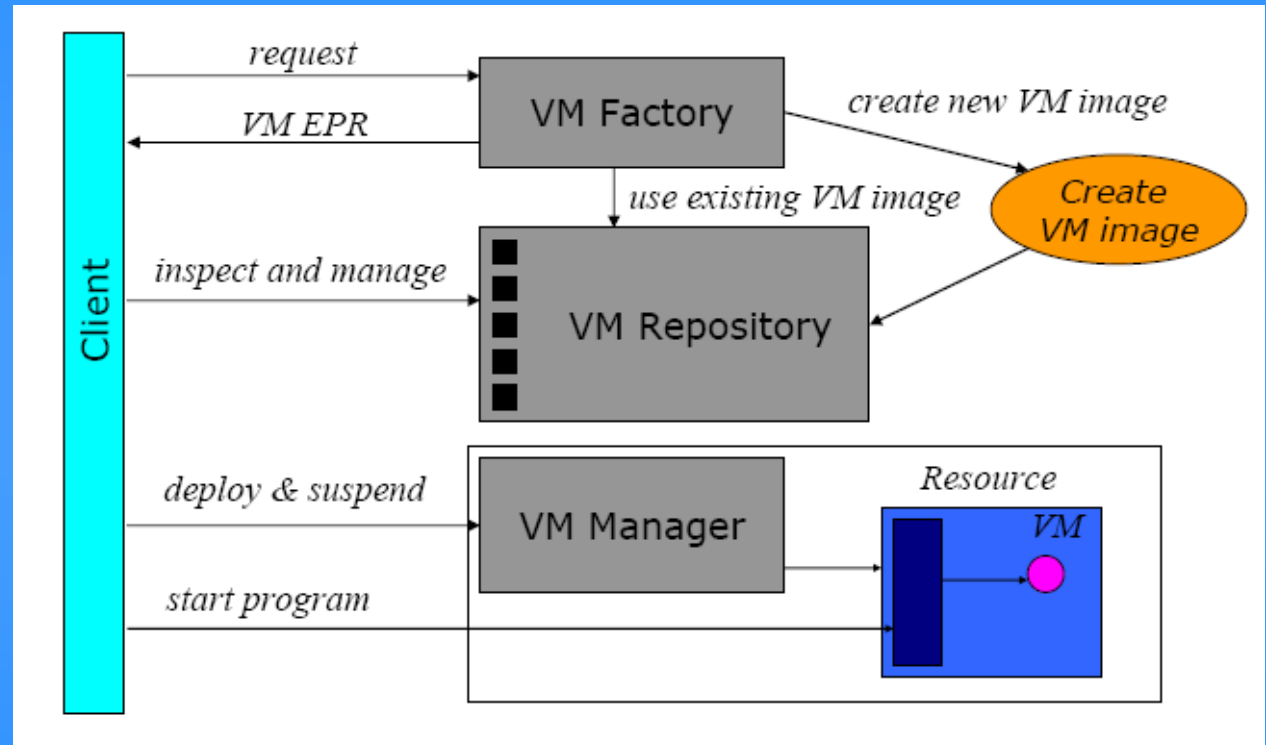
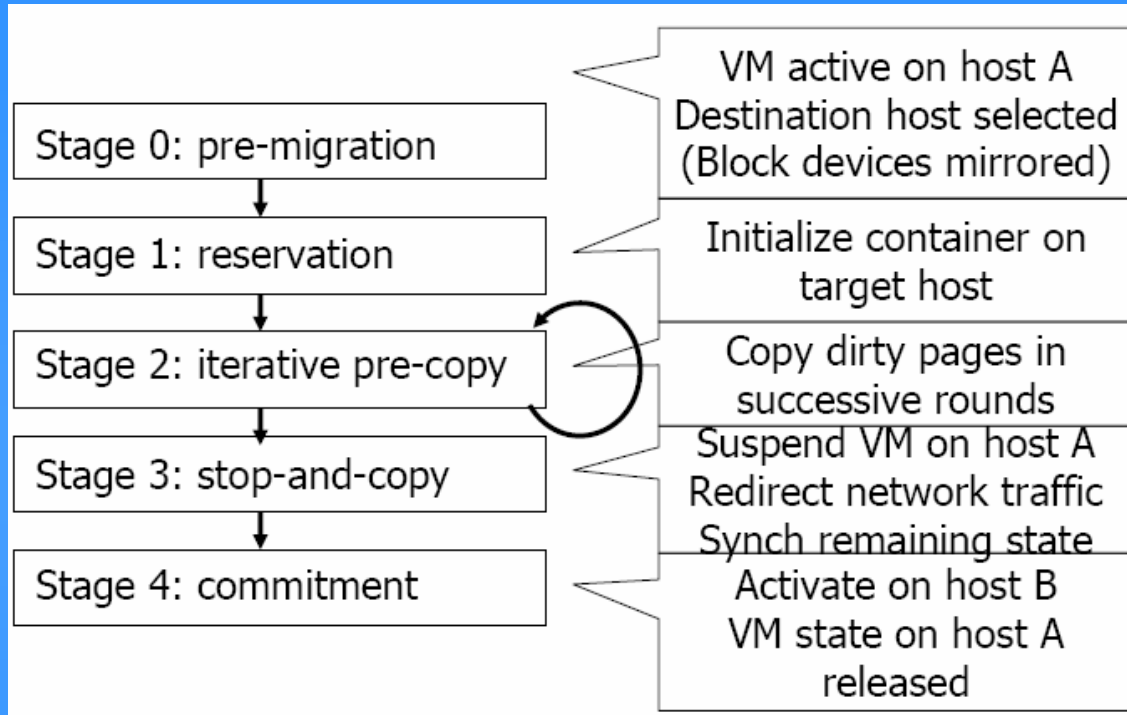


Image taken from "Working Spaces: Virtual Machines in the Grid" by Kate Keahey, Globus World 2005

# VM migration complexity



- \* Lots of state in memory
- \* Soft real-time requirements (web servers, databases)
- \* Minimal downtime

Taken from D. Magenheimer: “Xen and the Art of Open Source Itanium Virtualization”, talk at Oct. 05 Gelato meeting



# Opportunities at openlab

- \* Xen is a fast moving open source project thus continuous tests on different platforms are essential.
- \* Track Intel's hardware role.
- \* Performance tests (HEP software/networking under VMs).
- \* Evaluate the migration feature.
- \* Configuration with SmartFrog.
- \* Automated generation of images on demand.
- \* Integration of virtualization into the LCG infrastructure (how to store images, how to deal with batch systems, consequences for the information system,...).
- \* Talk with the experiments to learn about their demand for virtualization.

Many challenging open questions.

Together with LCG and the HEP community openlab has the possibility to expose virtualization to a large user base. Their feedback will provide us valuable information about possible use cases for virtualization.