CernVM

SoFTware Appliance

CernVM - a virtual software appliance for LHC applications

C. Aguado-Sanchez [1)] , P. Buncic [1)] , L. Franco [1)] , S. Klemer [1)] , P. Mato [1)]

[1)] CERN, Geneva, Switzerland

Predrag Buncic (CERN/PH-SFT)

CernVM

SoFTware Appliance

- Talk Outline
  - Introduction
  - CernVM Project
    - Building blocks
    - Scalability and performance
    - User Interface and API
    - Release status
  - Future plans & directions
  - Conclusions

CernVM

SoFTware Appliance

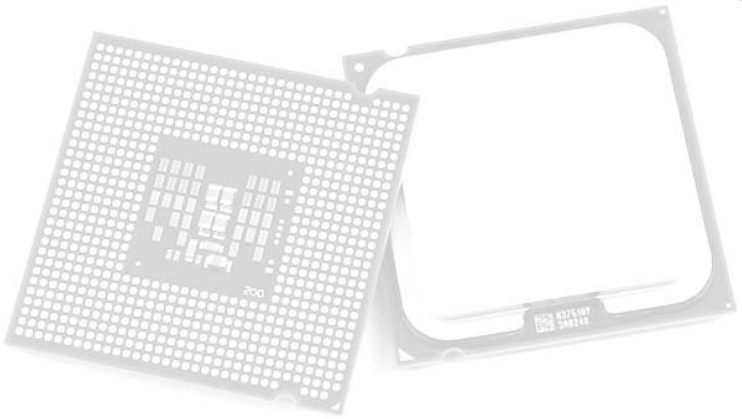

Introduction

CernVM
SoFTware Appliance

Enjoying “Frequency scaling Era”

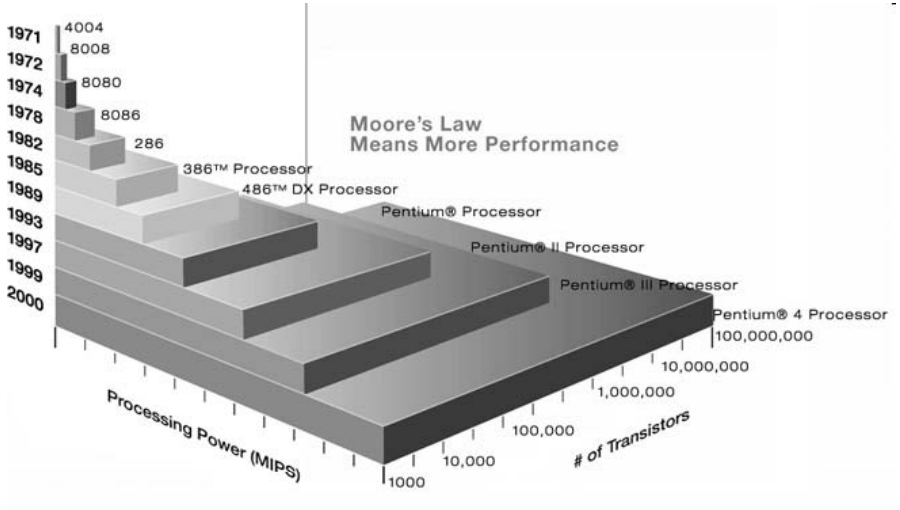

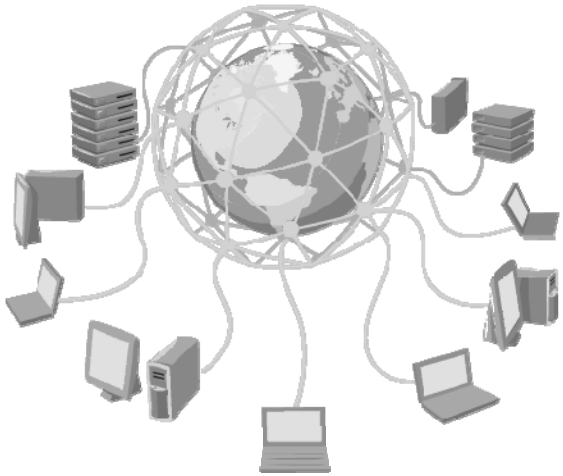

Cluster of Clusters (GRID)

Workstation & PC Clusters

Mainframes

IBM-VM 360, CERNVM,1988

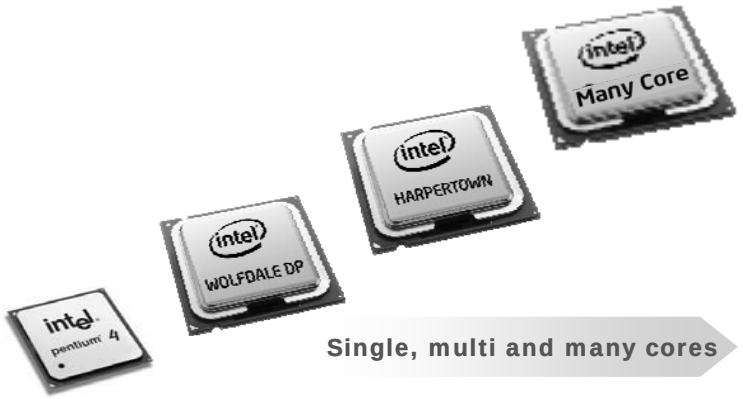

Single, multi and many cores

CernVM
SoFTware Appliance

Recent Trends (h/w)

- Multi & many cores
  - Software benefits from multicore architectures where code can be executed in parallel
    - Under most common operating systems this requires code to execute in separate threads or processes.

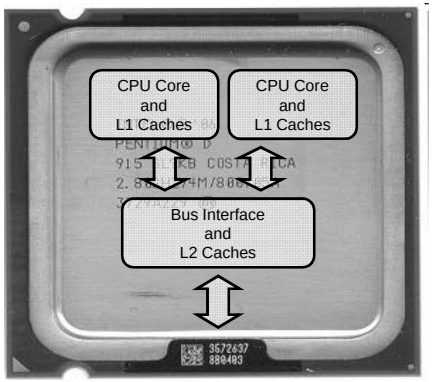

  - Unfortunatelly, HEP/LHC applications were developed during period when it looked like any performance issue can be easily solved by simply waiting 2 more years
- Support for hardware assisted virtualization
  - VMM can now efficiently virtualize the entire x86 instruction set
    - Intel VT and AMD-V implementations
    - VMware, Xen 3.x (including derivatives like Virtual Iron), Linux KVM and Microsoft Hyper-V
  - Running Virtual Machine will benefit from adoption of multiple core architectures since each virtual machine runs independently of others and can be executed in parallel
    - In some cases (KVM,VMware) running concurrent VM can save memory

CernVM
SoFTware Appliance

Recent Trends (s/w)

- Building on this new CPU capabilities, the software industry giants quickly came up with Cloud computing paradigm that relies on
  - virtualization as enabling technology
  - densly packed and energy efficient multi core systems as a basis fo next generation data centers

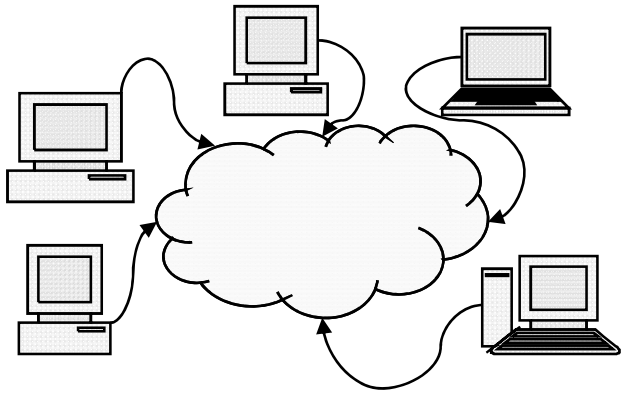

- Utility Computing on demand
  - Amazon Elastic Computing Cloud
  - Blue Cloud (IBM)
    - 13 dedicated data centres

- Software as a Service (SaaS)
  - Google App Engine
    - Delivering an application through the browser to thousands of customers
    - Investing ~$2 billion per year in new generation data centres
  - Microsoft Azzuro
    - Hosted office applications & data
    - Plans to match Google in terms of number of data centres, $600 million per site

CernVM
SoFTware Appliance

SaaS example: eyeOS

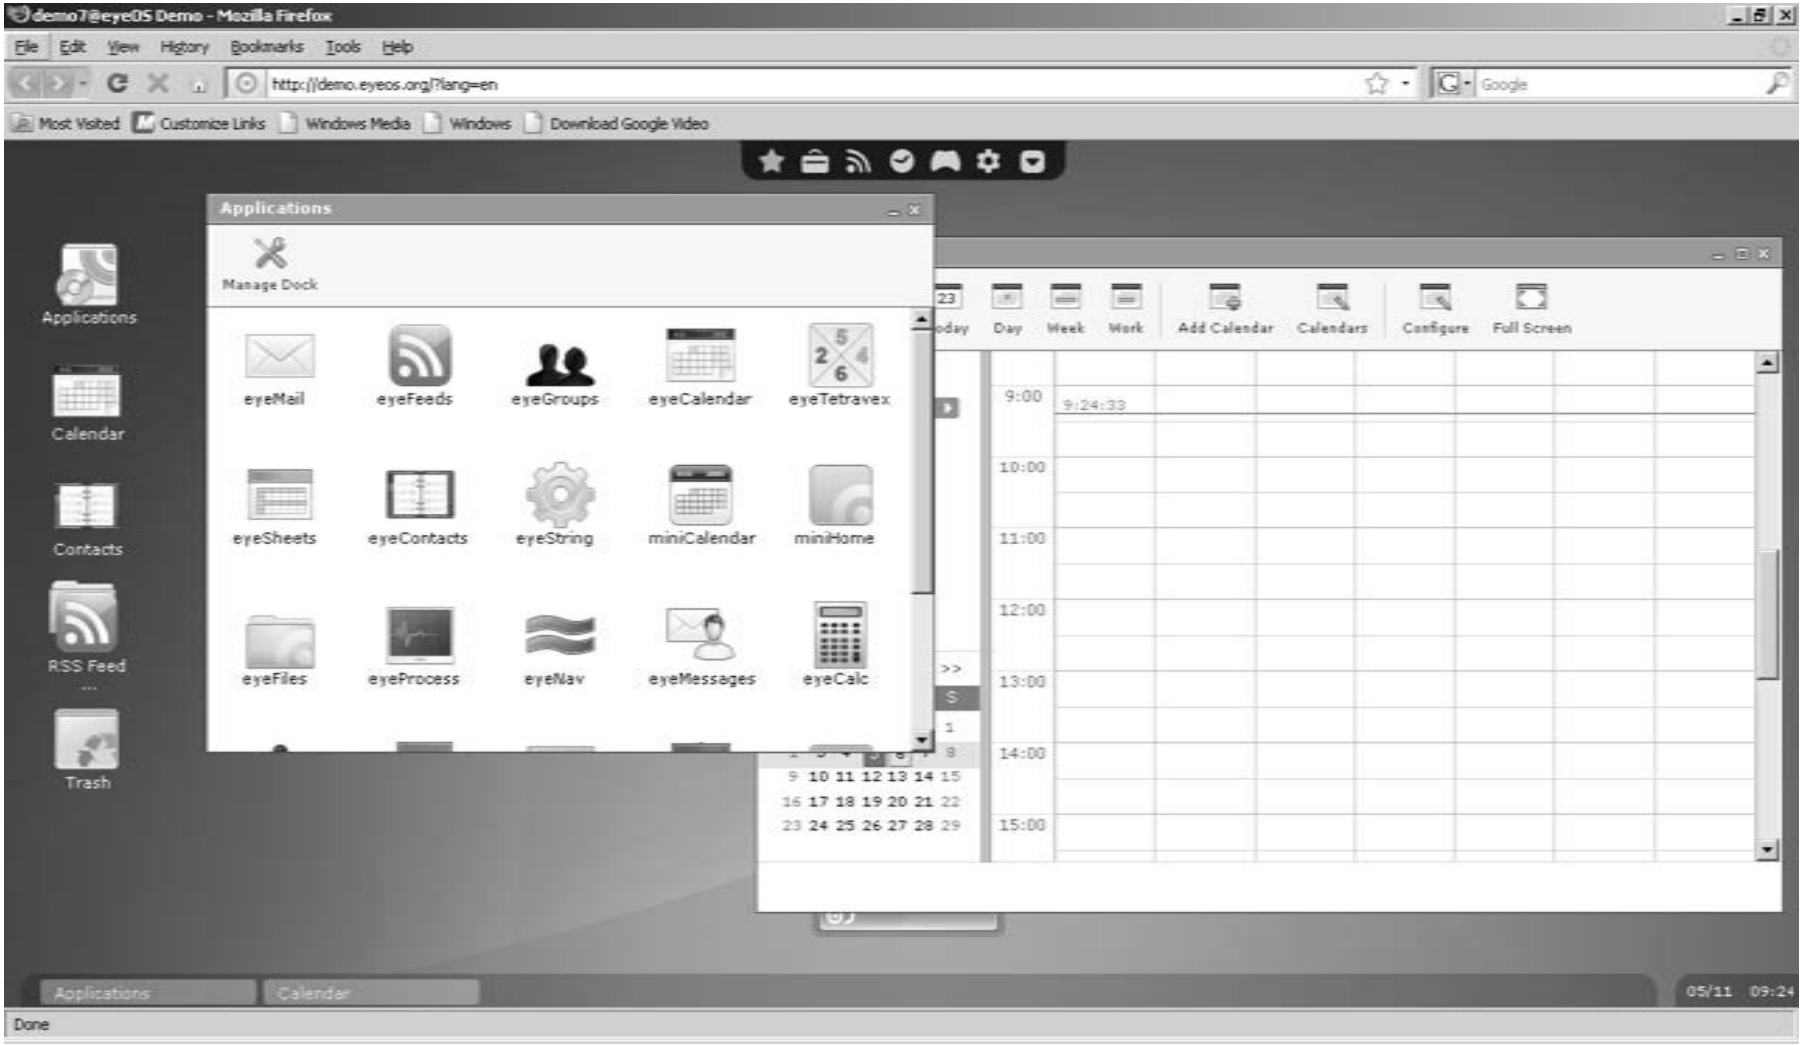

http://eyeos.org

CernVM
SoFTware Appliance

Grid vs Cloud

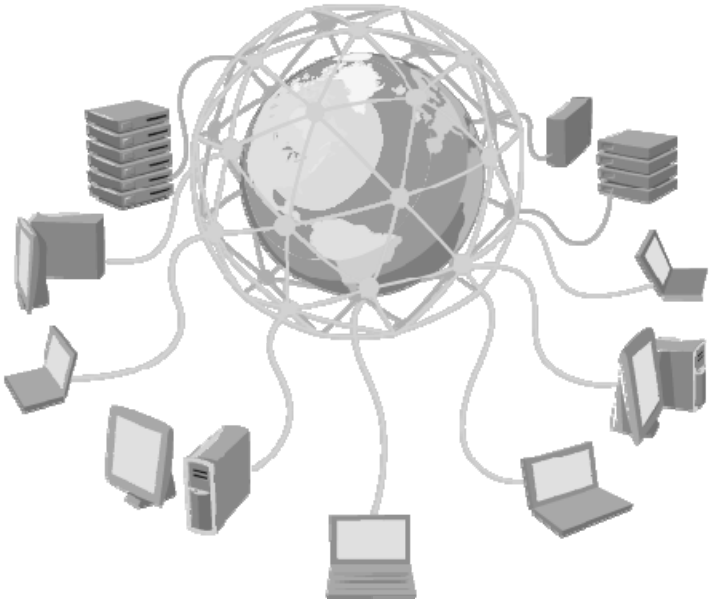

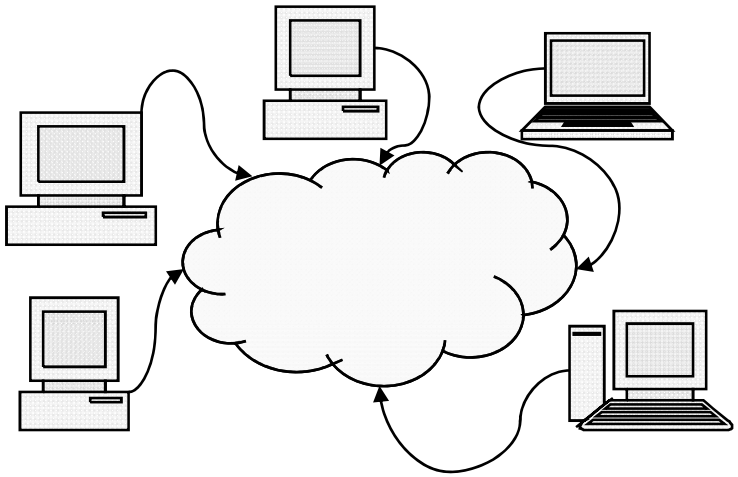

- Localised data (Storage Element)
  - Complicated scheduling
- Grid middleware in charge of abstacting complexity and differences of heterogenious resources
  - Requires complex middleware

- Distributed storage
  - Simplified scheduling
- Virtualization used as a tool to create a homogenious overlay for end user(s) and reduce scale of the problem
  - Simplifies middleware

CernVM

SoFTware Appliance

WIKIPEDIA

*"Cloud Computing is a paradigm in which information is permanently stored in servers on the Internet and **cached** temporarily on clients that include desktops, entertainment centers, table computers, notebooks, wall computers, handhelds, sensors, monitors, etc."*

CernVM

SoFTware Appliance

CernVM Project

CernVM
SoFTware Appliance

CernVM Project

- Portable Analysis Environment using Virtualization Technology (WP9)
- Project goal:
  - Provide a complete, portable and easy to configure user environment for developing and running LHC data analysis locally and on the Grid independent of physical software and hardware platform (Linux, Windows, MacOS)
    - Decouple application lifecycle from evolution of system infrastructure
    - Reduce effort to install, maintain and keep up to date the experiment software
    - Lower the cost of software development by reducing the number of compiler-platform combinations
- Approved in 2007 (2+2 years) as R&D activity, started January 2008

CernVM
SoFTware Appliance

How do we want to do that?

CernVM
SoFTware Appliance

Key Building Blocks

- rPath Linux 1 (www.rpath.org)
  - Slim Linux OS binary compatible with RH/SLC4
- rAA - rPath Linux Appliance Agent
  - Web user interface
  - XMLRPC API
- rBulder
  - A tool to build VM images for various virtualization platforms
- CVMFS - CernVM file system
  - Read only file system optimized for software distribution
    - Aggressive caching
  - Operational in offline mode
    - For as long as you stay within the cache

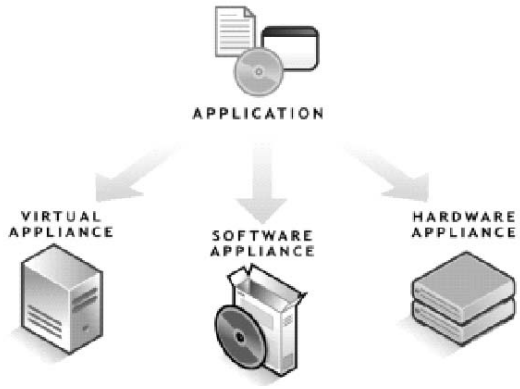

Build types

- Installable CD/DVD
- Stub Image
- Raw Filesystem Image
- Netboot Image
- Compressed Tar File
- Demo CD/DVD (Live CD/DVD)
- Raw Hard Disk Image
- Vmware ® Virtual Appliance
- Vmware ® ESX Server Virtual Appliance
- Microsoft ® VHD Virtual Apliance
- Xen Enterprise Virtual Appliance
- Virtual Iron Virtual Appliance
- Parallels Virtual Appliance
- Amazon Machine Image
- Update CD/DVD
- Appliance Installable ISO

CernVM
SoFTware Appliance

Easy packaging with Conary

```
class Root(CPackageRecipe):
    name='root'
    version='5.19.02'

    buildRequires = ['libpng:devel',
                     'libpng:devellib','krb5:devel',
                     'libstdc++:devel','libxml2:devel',
                     'openssl:devel','python:devel',
                     'xorg-x11:devel', 'zlib:devel',
                     'perl:devel', 'perl:runtime']

    def setup(r):
        r.addArchive('ftp://root.cern.ch/root/%(name)s v%(version)s.source.tar.gz')
        r.Environment('ROOTSYS',%(builddir)s')
        r.ManualConfigure('--prefix=/opt/root ')
        r.Make()
        r.MakeInstall()
```

CernVM
SoFTware Appliance

CernVM Components

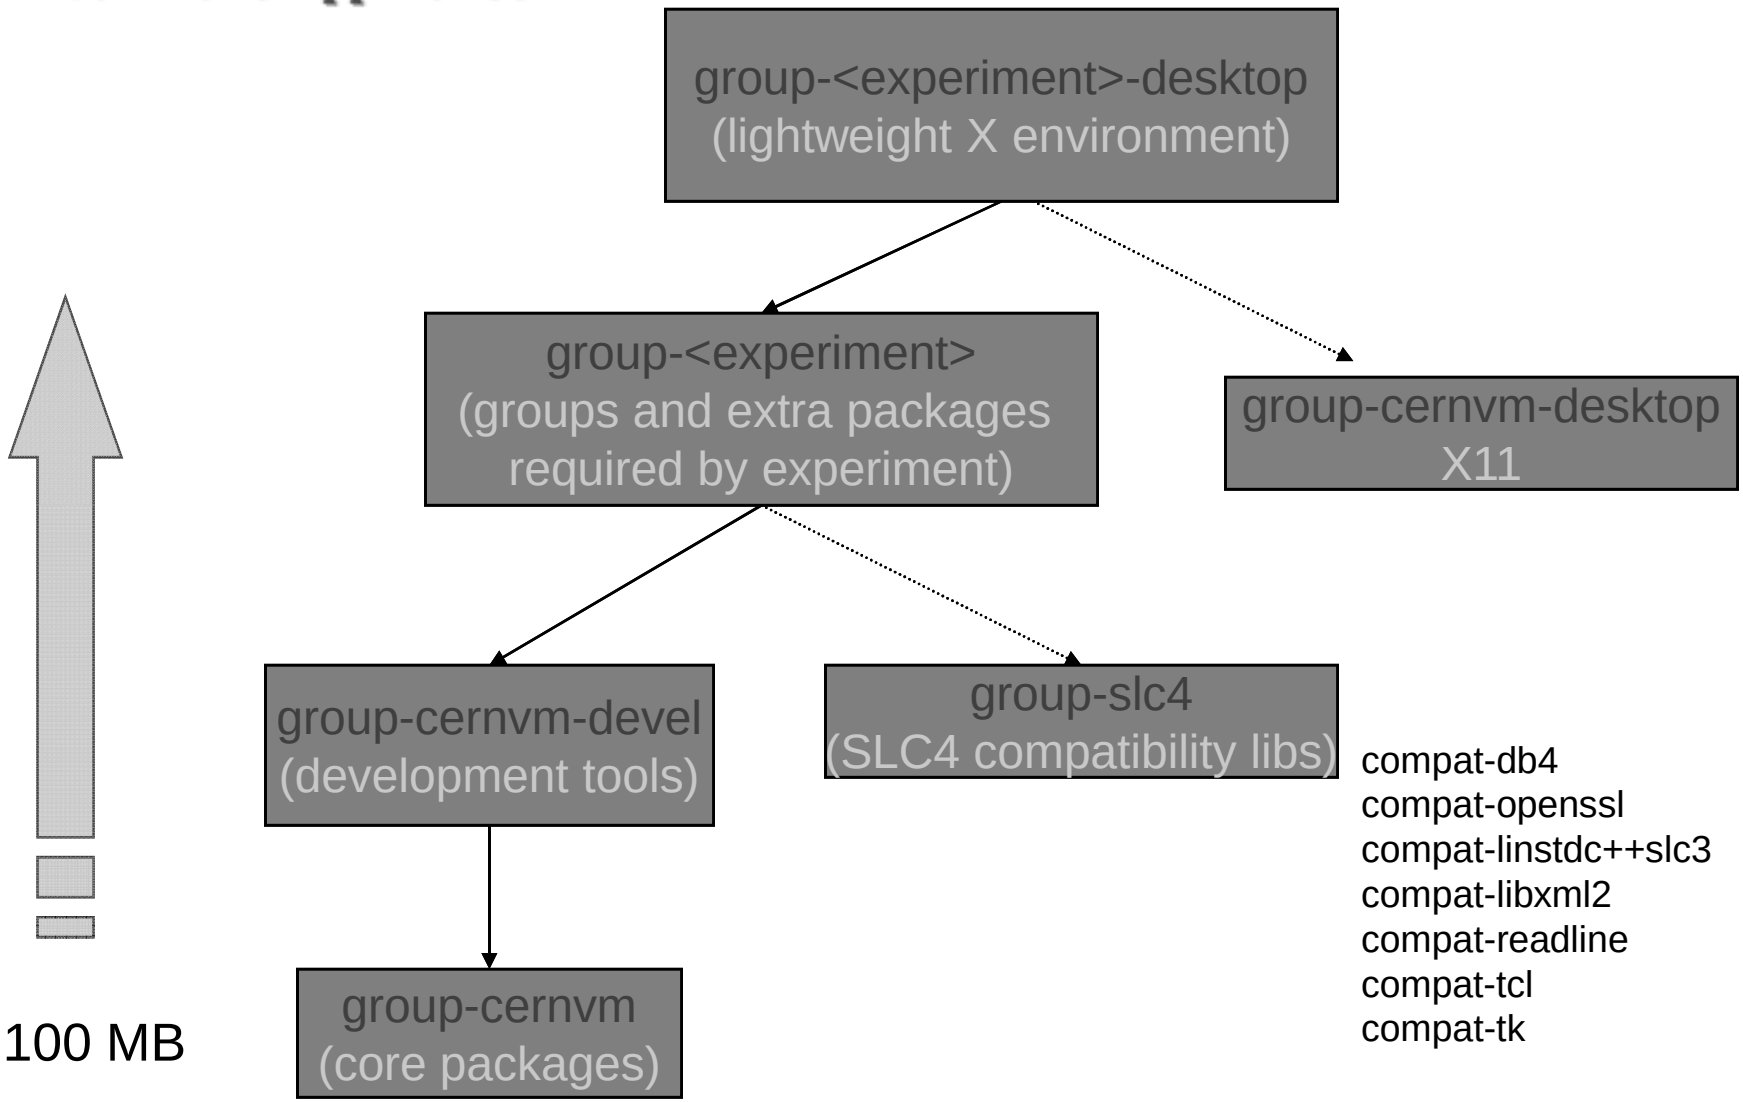

CernVM
SoFTware Appliance

CVMFS

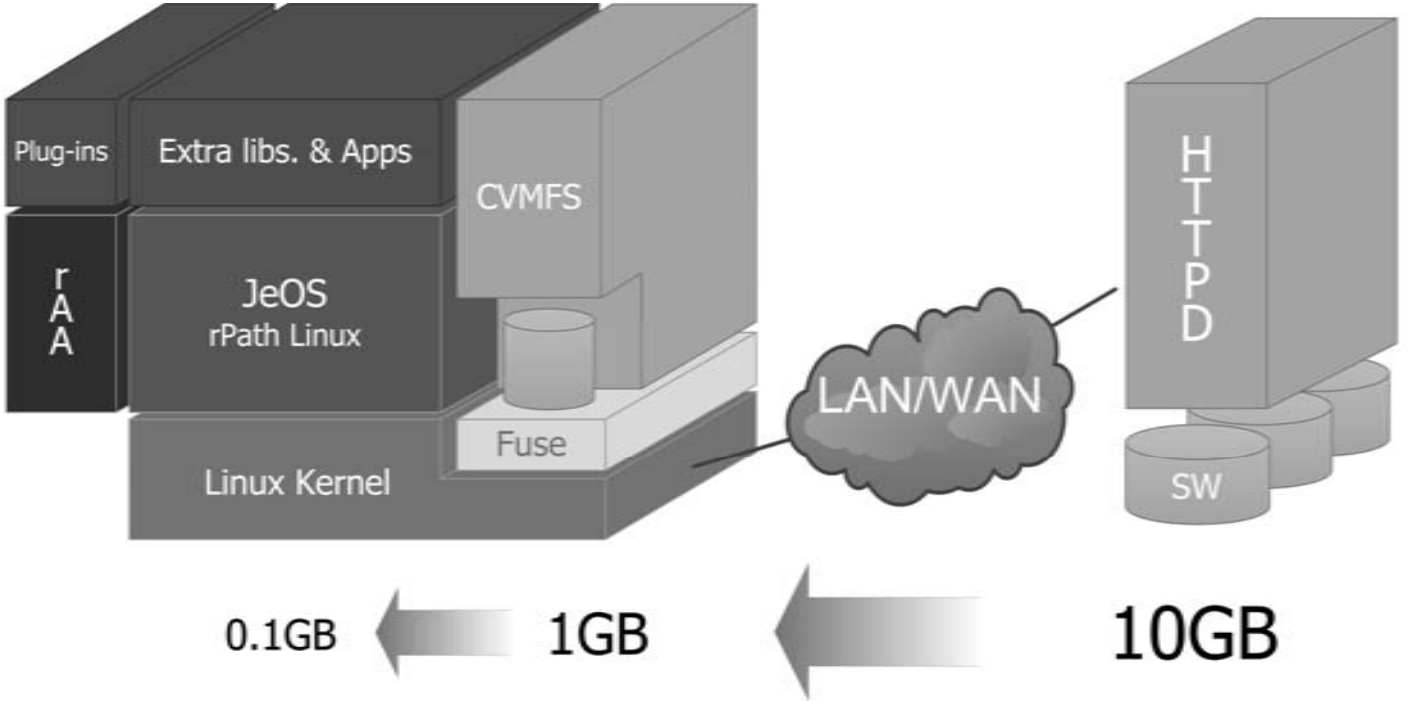

- **CernVM File System (CVMFS) is derived from Parrot (http://www.cctools.org) and its GROW-FS code base and adapted to run as a FUSE kernel module adding extra features like:**
  - **possibility to use multiple file catalogues on the server side**
  - **transparent file compression under given size threshold**
  - **dynamical expansion of environment variables embedded in symbolic links**

CernVM
SoFTware Appliance

Scalability: Cache Hierarchy

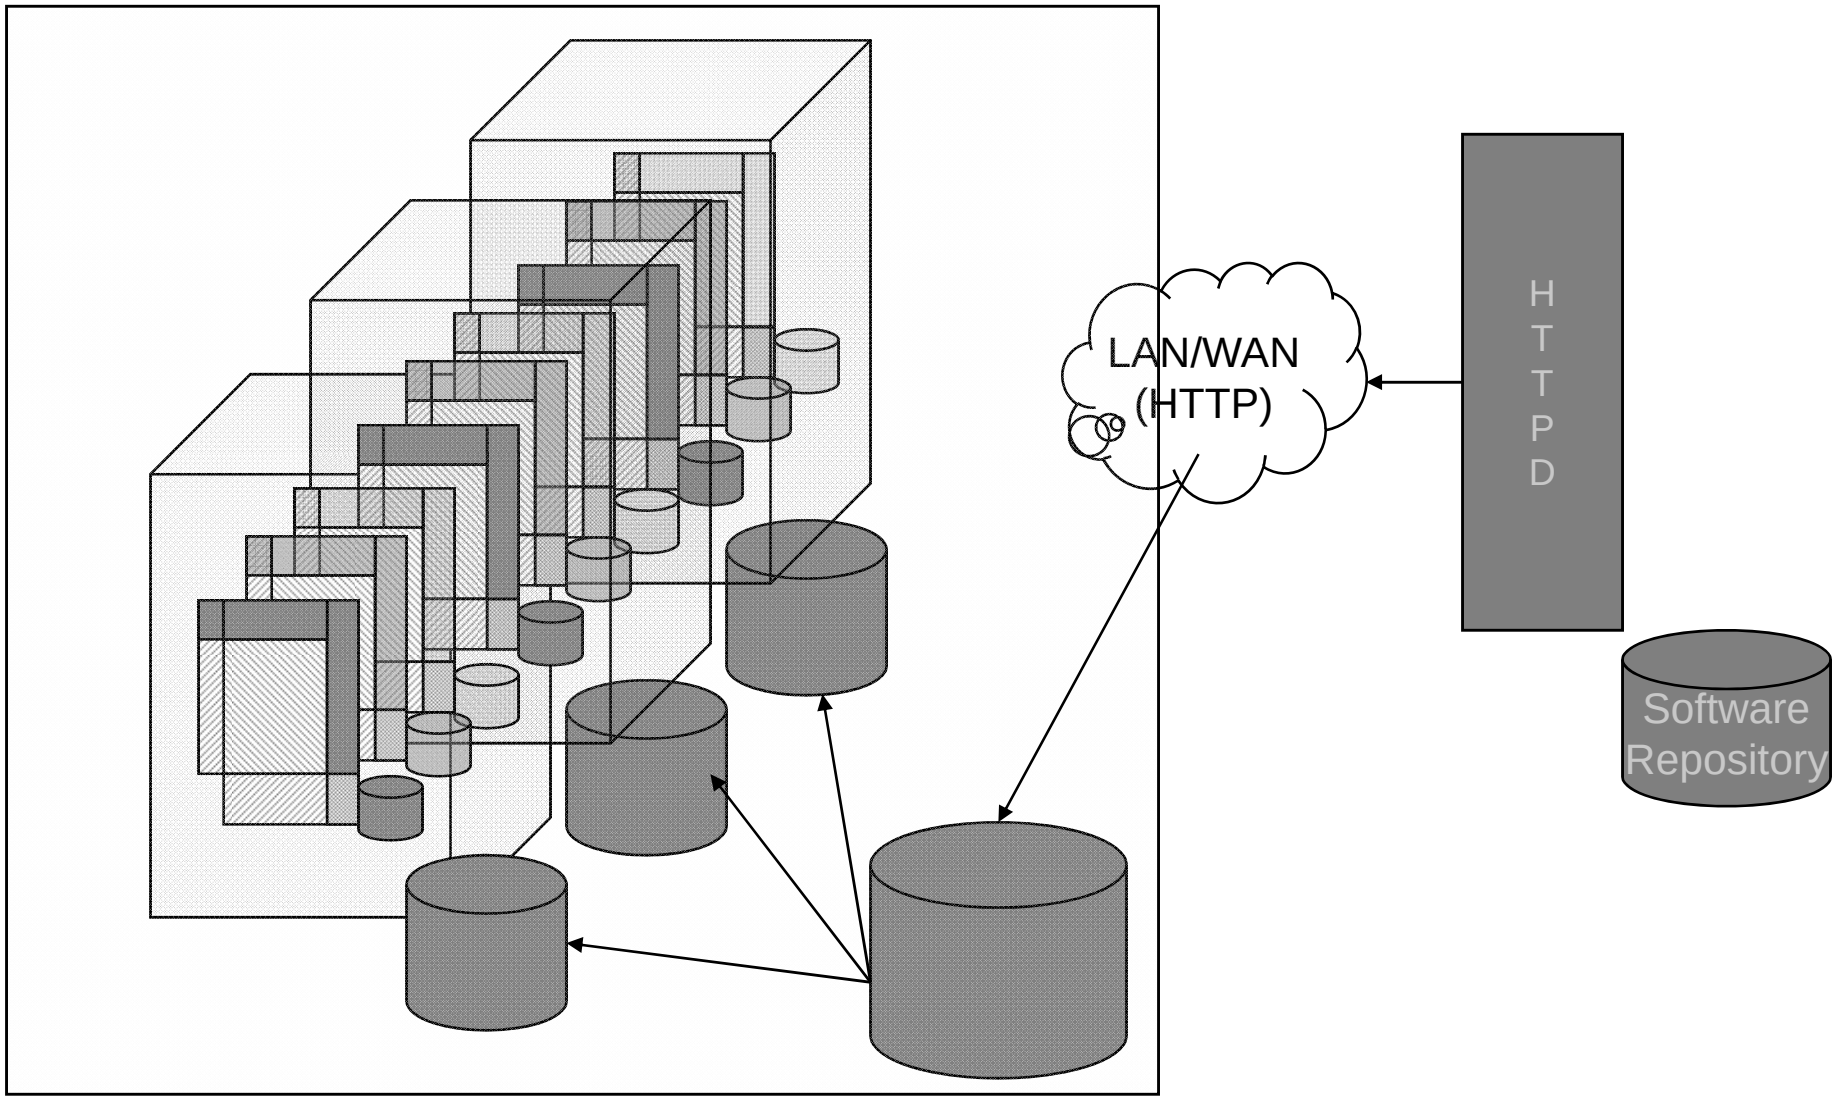

CernVM
SoFTware Appliance

Scalable Infrastructure

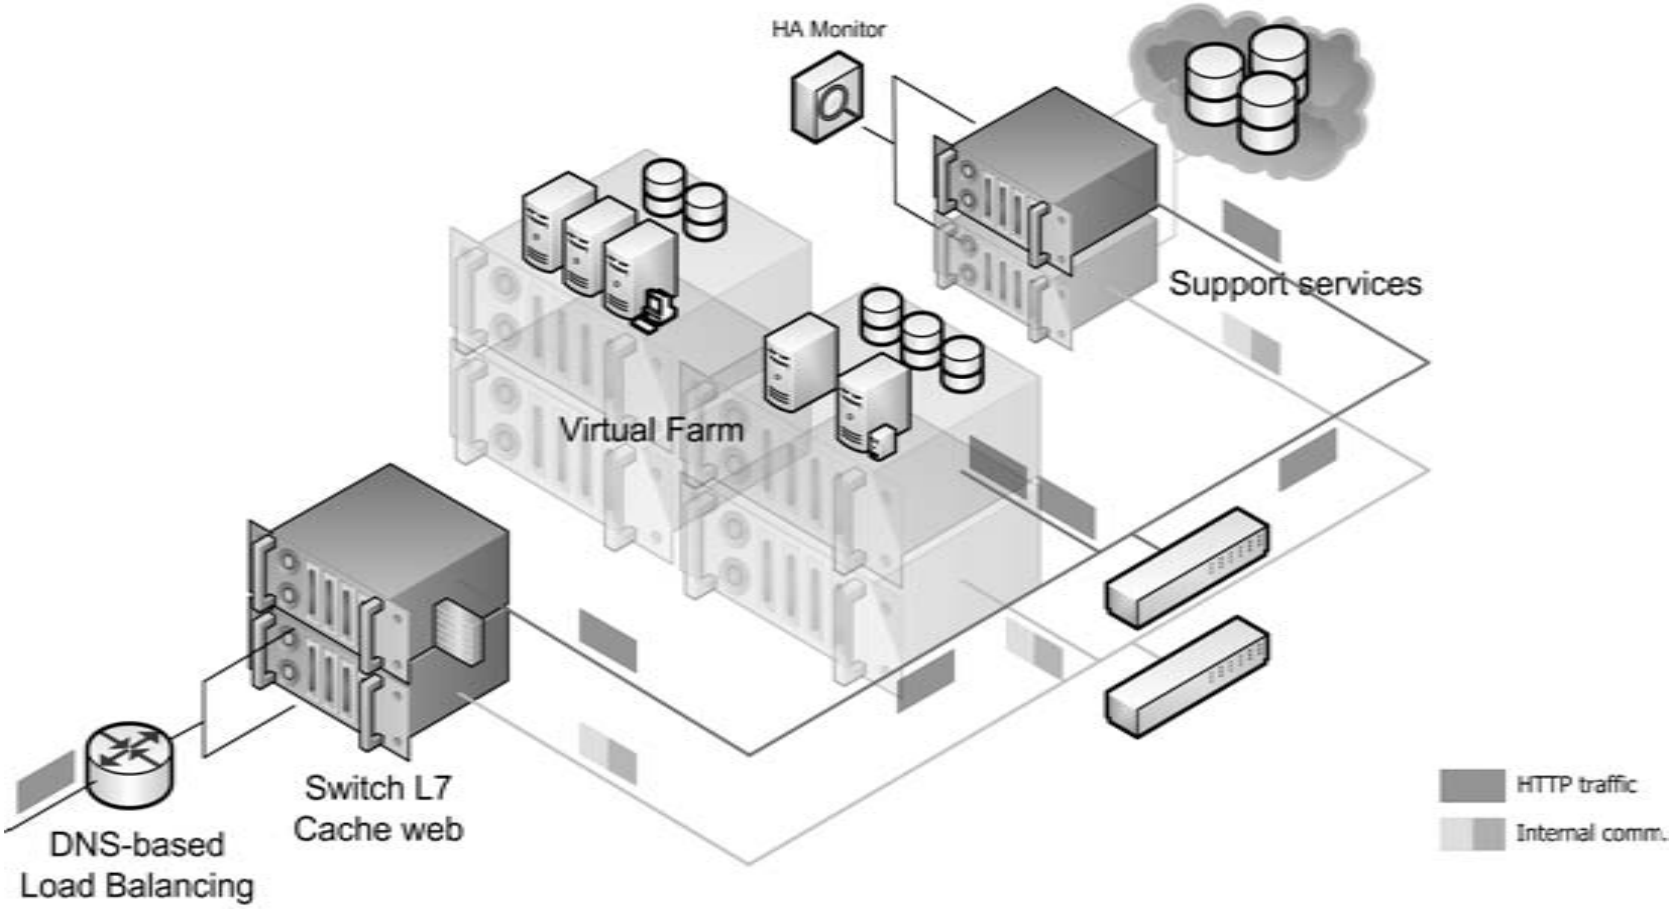

CernVM

SoFT

Publishing Releases

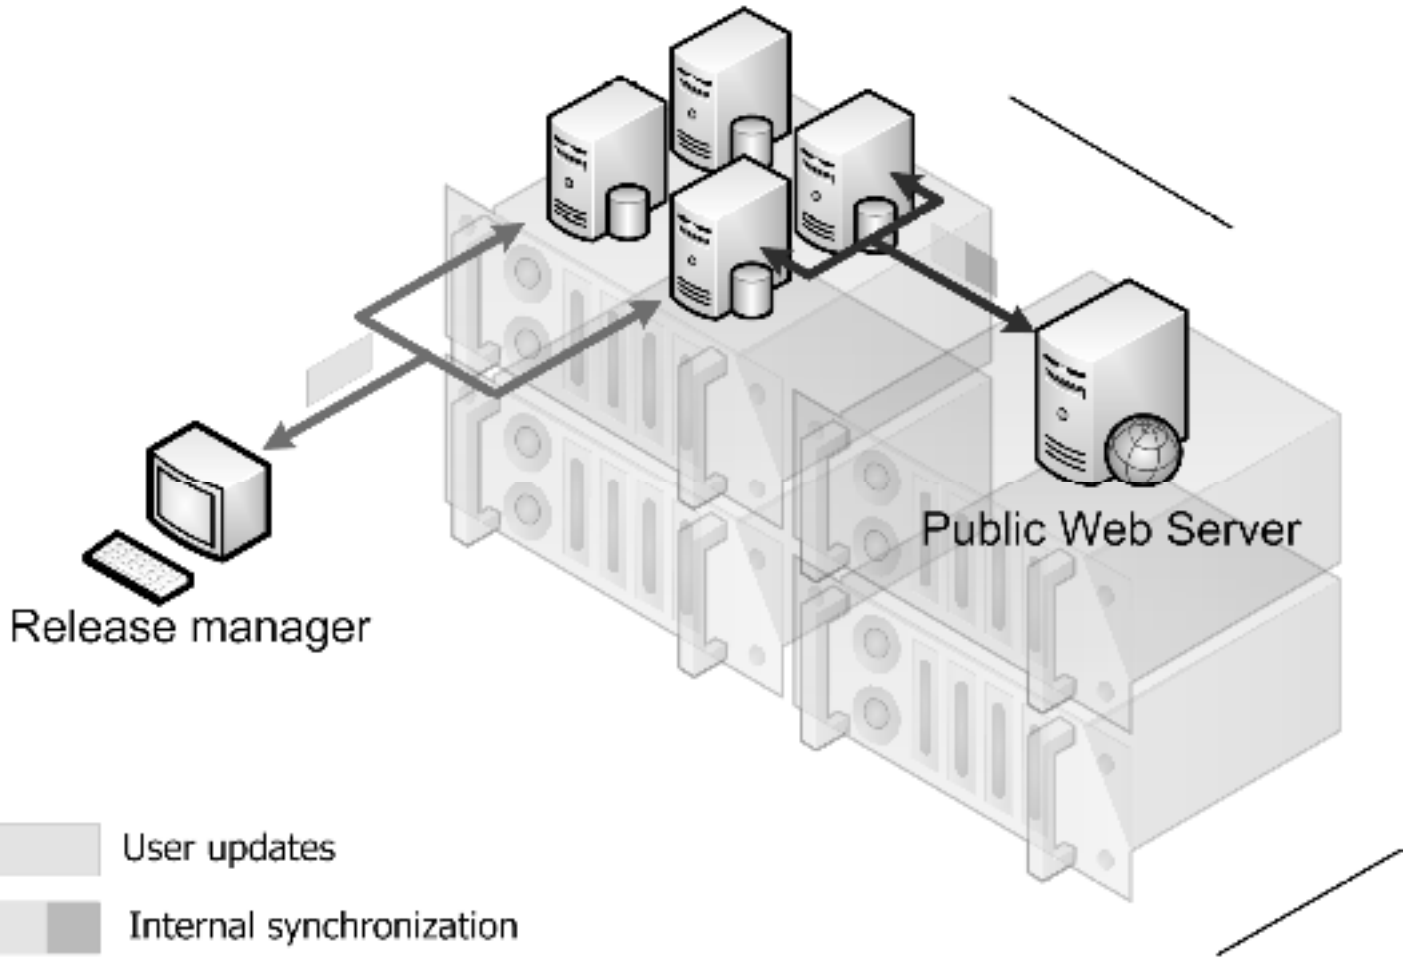

CernVM
SoFTware Appliance

CernVM v.s. LXPLUS

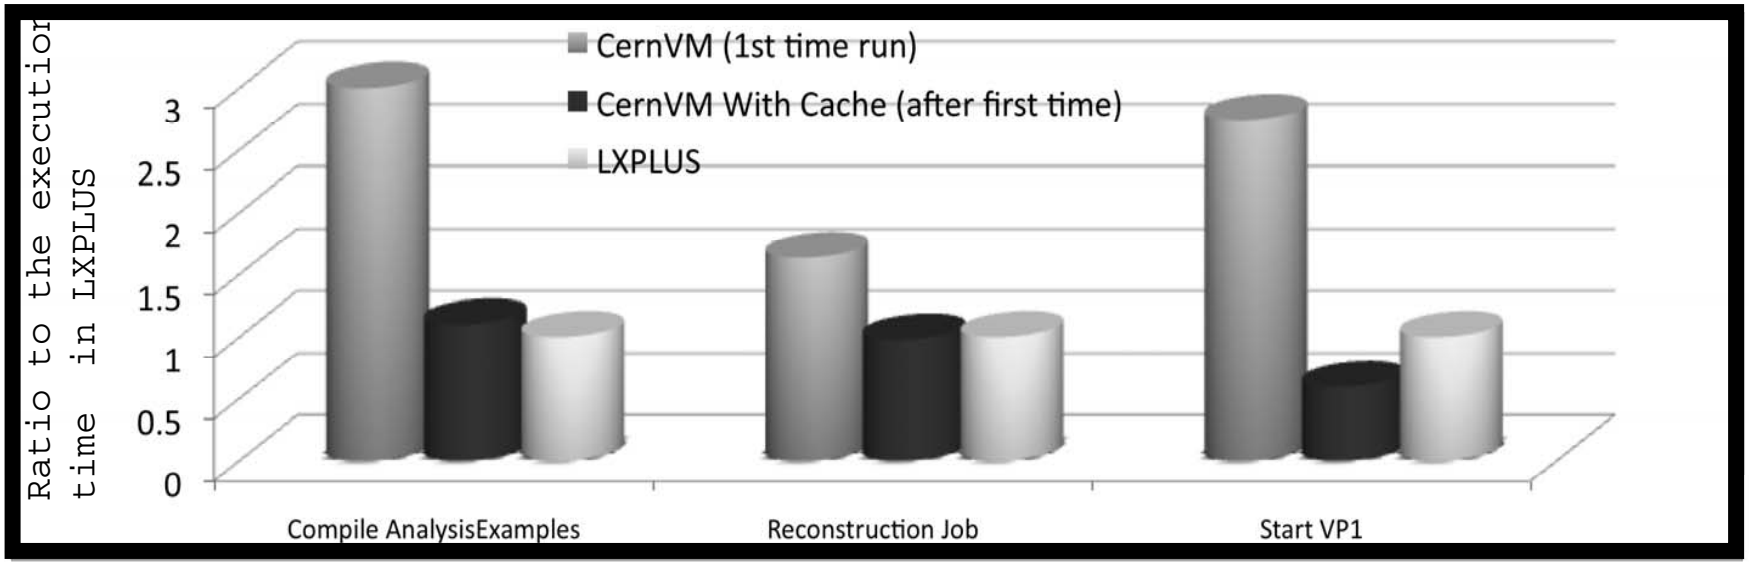

- CernVM for ATLAS (benchmarks by Yushu Yao)
  - For the first time running, CernVM is 2-3 times slower, depending on the network speed, since it needs to cache files
  - Once the files are cached, the speed on CernVM is roughly the same as on LXPLUS.
  - CernVM is faster when reading large files (since they are local)
    Note: LXPLUS has a higher CPU rate than test computer (3GHz vs. 2.66GHz)

CernVM

SoFTware Appliance

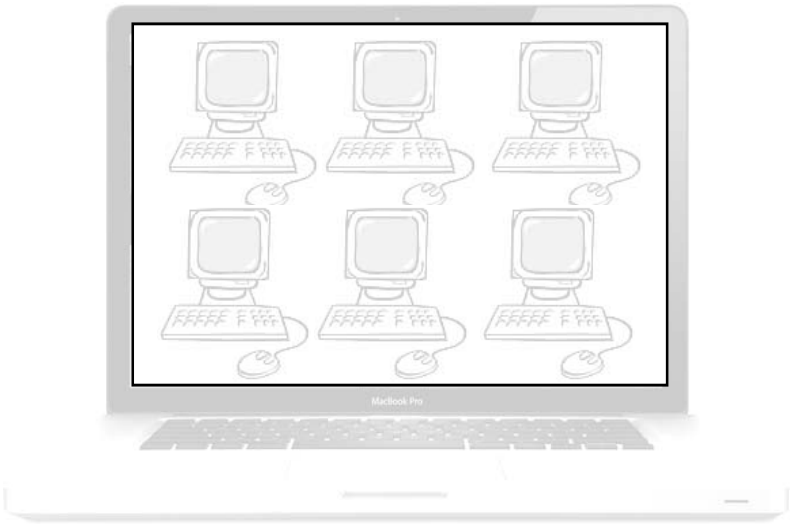

User Interface

rAA Web UI


rPath
CernVM
Software Appliance
Schedule
Updates

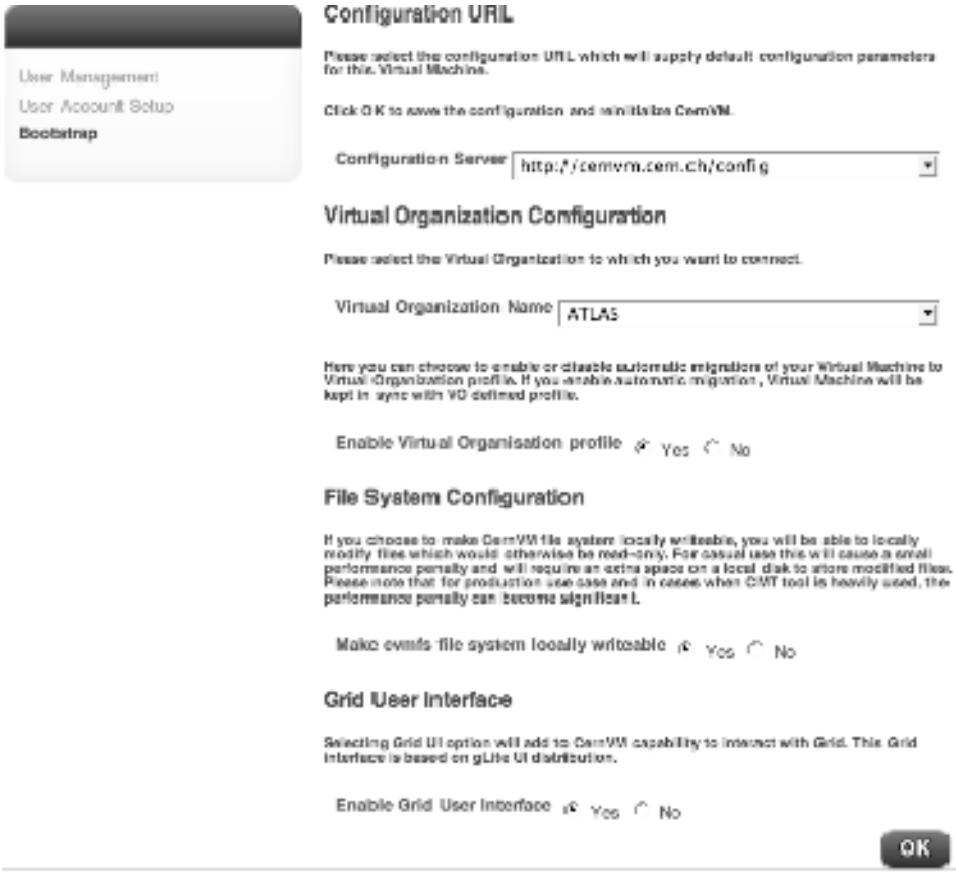
User Management
User Account Setup
Bootstrap
Configuration URL
Please select the configuration URL which will supply default configuration parameters for this Virtual Machine.
Click O K to save the configuration and reinitialize CernVM.
Configuration Server
http://cernvm.cern.ch/config
Virtual Organization Configuration
Please select the Virtual Organization to which you want to connect.
Virtual Organization Name
ATLAS
Here you can choose to enable or disable automatic migrations of your Virtual Machine to Virtual Organization profile. If you enable automatic migration, Virtual Machine will be kept in sync with VO defined profile.
Enable Virtual Organisation profile
Yes
No
File System Configuration
If you choose to make CernVM file system locally writeable, you will be able to locally modify files which would otherwise be read-only. For casual use this will cause a small performance penalty and will require an extra space on a local disk to store modified files. Please note that for production use case and in cases when CVMT tool is heavily used, the performance penalty can become significant.
Make cvmfs file system locally writeable
Yes
No
Grid User Interface
Selecting Grid UI option will add to CernVM capability to interact with Grid. This Grid interface is based on gLite UI distribution.
Enable Grid User Interface
Yes
No
OK

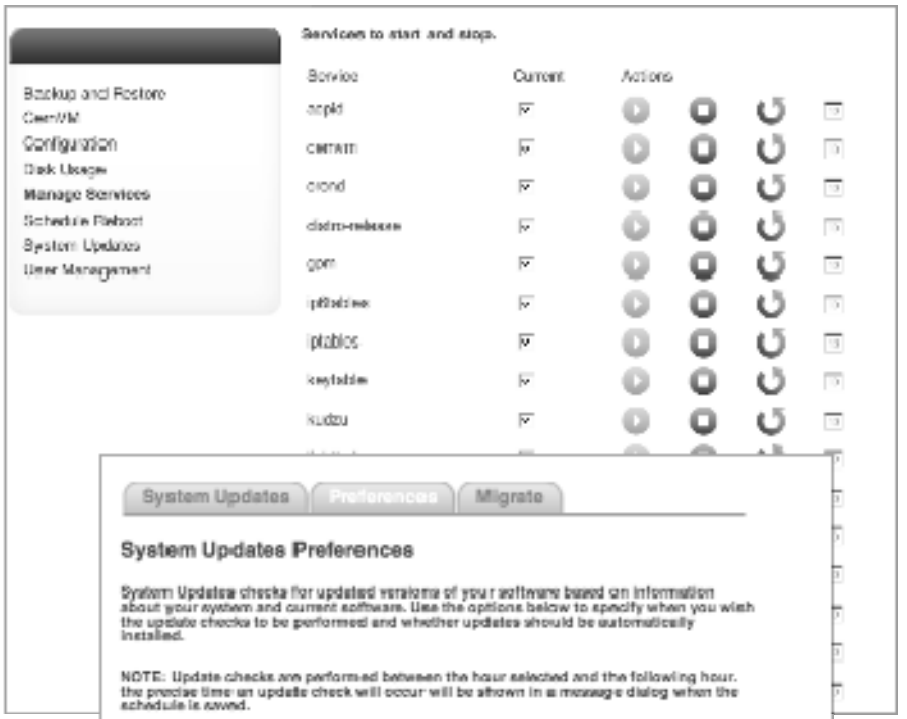
Backup and Restore
CernVM
Configuration
Disk Usage
Manage Services
Schedule Reboot
System Updates
User Management
Services to start and stop.
Service
Current
Actions
crond
gpm
kudzu

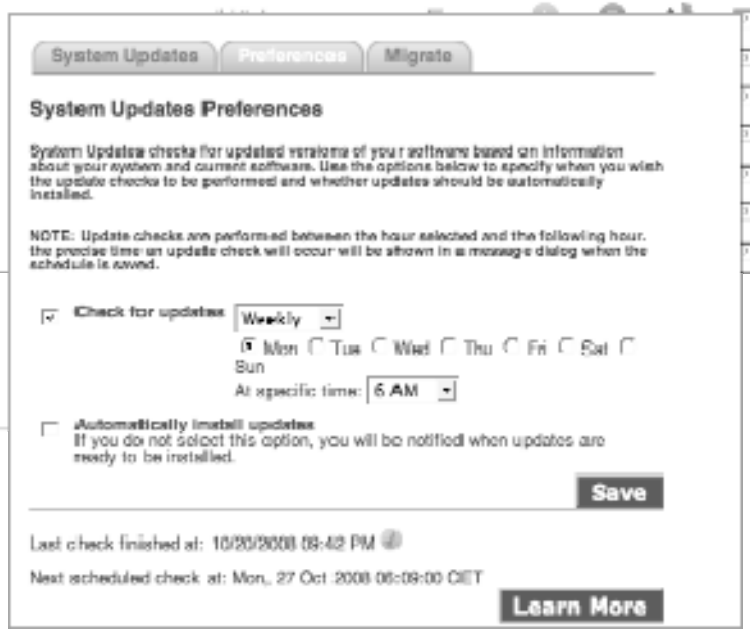
System Updates
Preferences
Migrate
System Updates Preferences
System Updates checks for updated versions of your software based on information about your system and current software. Use the options below to specify when you wish the update checks to be performed and whether updates should be automatically installed.
NOTE: Update checks are performed between the hour selected and the following hour. the precise time an update check will occur will be shown in a message dialog when the schedule is saved.
Check for updates
Weekly
Mon
Tue
Wed
Thu
Fri
Sat
Sun
At specific time:
6 AM
Automatically install updates
If you do not select this option, you will be notified when updates are ready to be installed.
Save
Last check finished at: 10/20/2008 09:42 PM
Next scheduled check at: Mon, 27 Oct 2008 06:09:00 CET
Learn More

CernVM
SoFTware Appliance

rAA: XML-RPC API

```
import xmlrpclib
import os
url = 'http://' + user + ':' + password + '@' + host + ':8004/rAA/xmlrpc'
server = xmlrpclib.ServerProxy(url)
r = server.cernvm.Config.configUpdate('http://cern.ch/cernvm/config',
                    'ALICE',  # (can be None,ATLAS, ALICE, LHCB, CMS)
                    'off',    # (on/off - if on it will make cvmfs locally writeable)
                    'off',    # (on/off - enable/disable grid UI)
                    'off')    # (on/off - enable VO profile)
print r["message"]
####
# Here we create initial user 'test' with the same password as Web admin
####
r = server.cernvm.User.userUpdate('test', password,password)
print r["message"]
####
# Select grid UI version
####
r = server.cernvm.Config.configGridUIVersion(”3.1.22-0")
print r["message"]
```

CernVM

SoFTware Appliance

Developing CernVM

CernVM
SoFTware Appliance

Developing CernVM

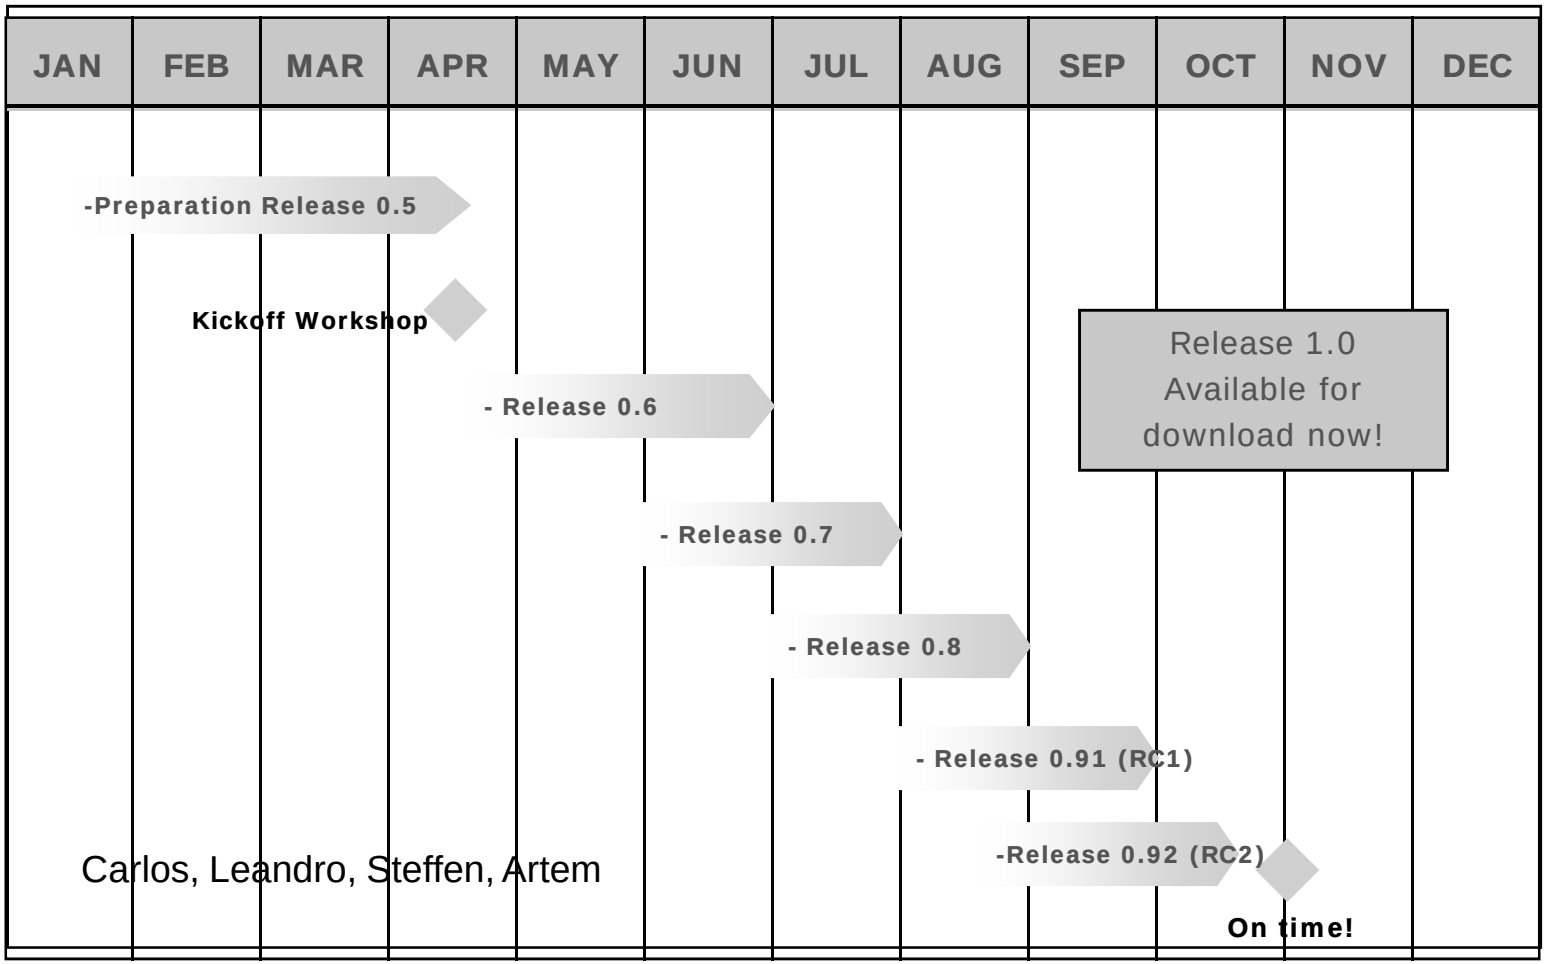

CernVM
SoFTware Appliance

1.0 Release

- Available now for download from
  - http://rbuilder.cern.ch/project/cernvm/releases
- Can be run on
  - Linux (KVM, Xen,VMware Player, VirtualBox)
  - Windows(WMware Player, VirtualBox)
  - Mac (Fusion, Parallels, VirtualBox)
- Release Notes
  http://cernvm.web.cern.ch/cernvm/index.cgi?page=ReleaseNotes
- HowTo
- http://cernvm.web.cern.ch/cernvm/?page=HowTo

- Appliance can be configured and used with ALICE, LHCb, ATLAS (and CMS) software frameworks

CernVM

SoFTware Appliance

Future plans & directions

CernVM
SoFTware Appliance

Removing Single Point of Failure

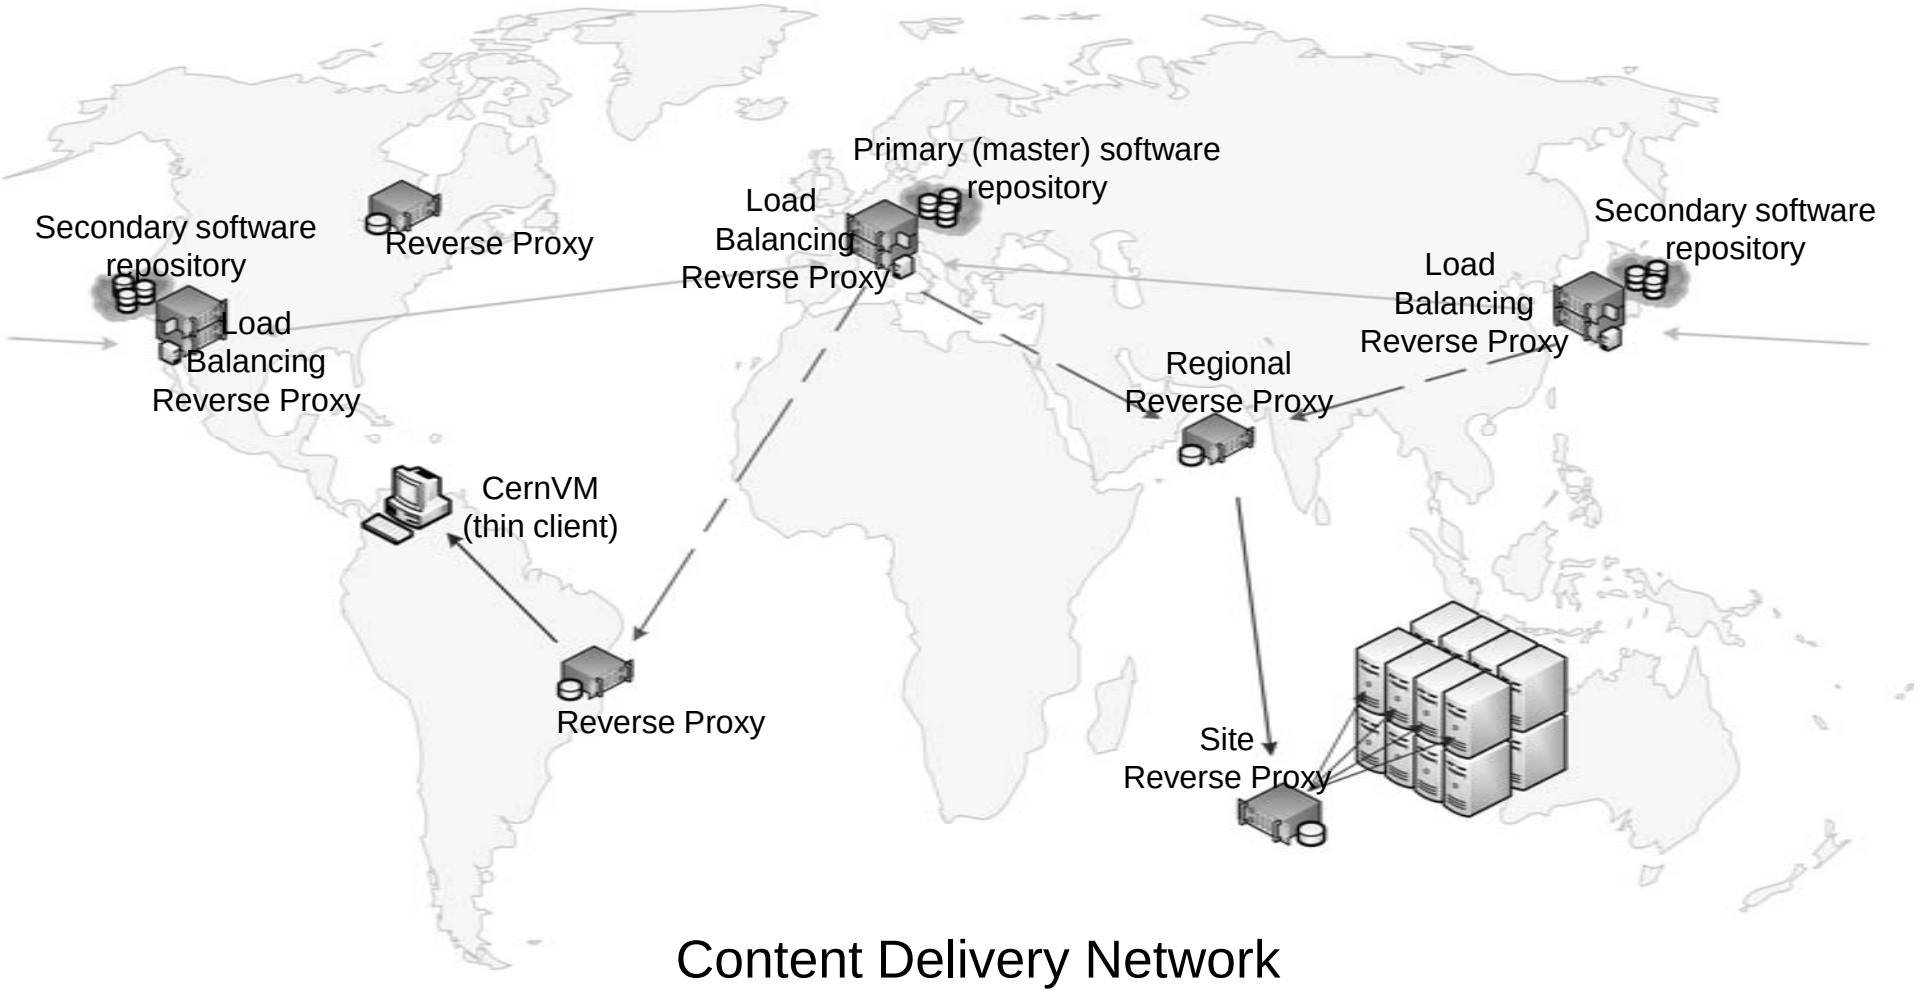

Content Delivery Network

Bridging Grids & Clouds

- BOINC
  - Open-source software for volunteer computing and grid computing
  - http://boinc.berkeley.edu/
  - Summer student project in OpenLab
    - based on LHC@HOME experience and CernVM image, aims to run ATLAS simulation using BOINC infrastructure
  - Successfully run ATLAS PanDA job in CernVM

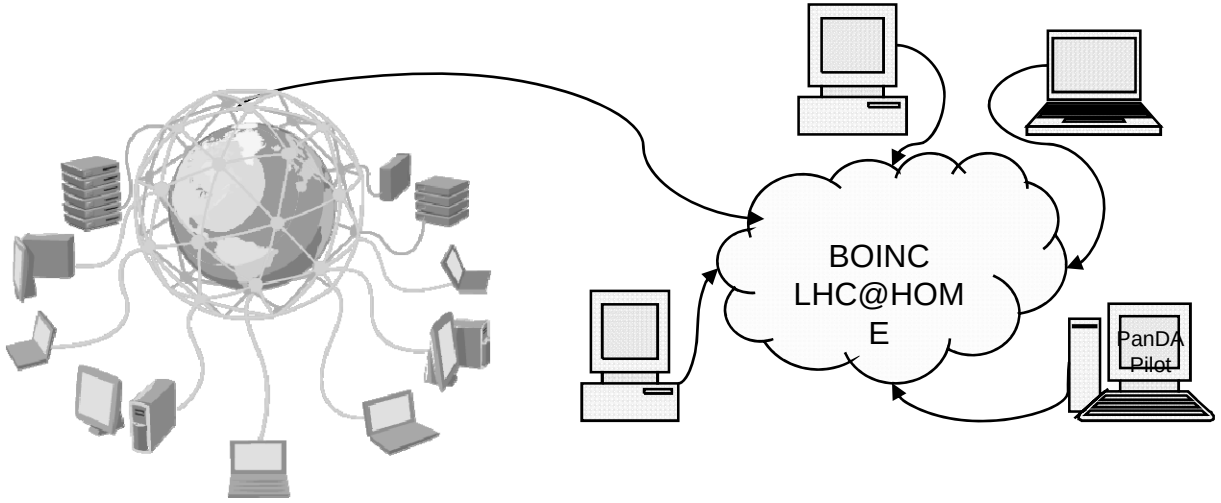

CernVM
SoFTware Appliance

CernVM CoPilot

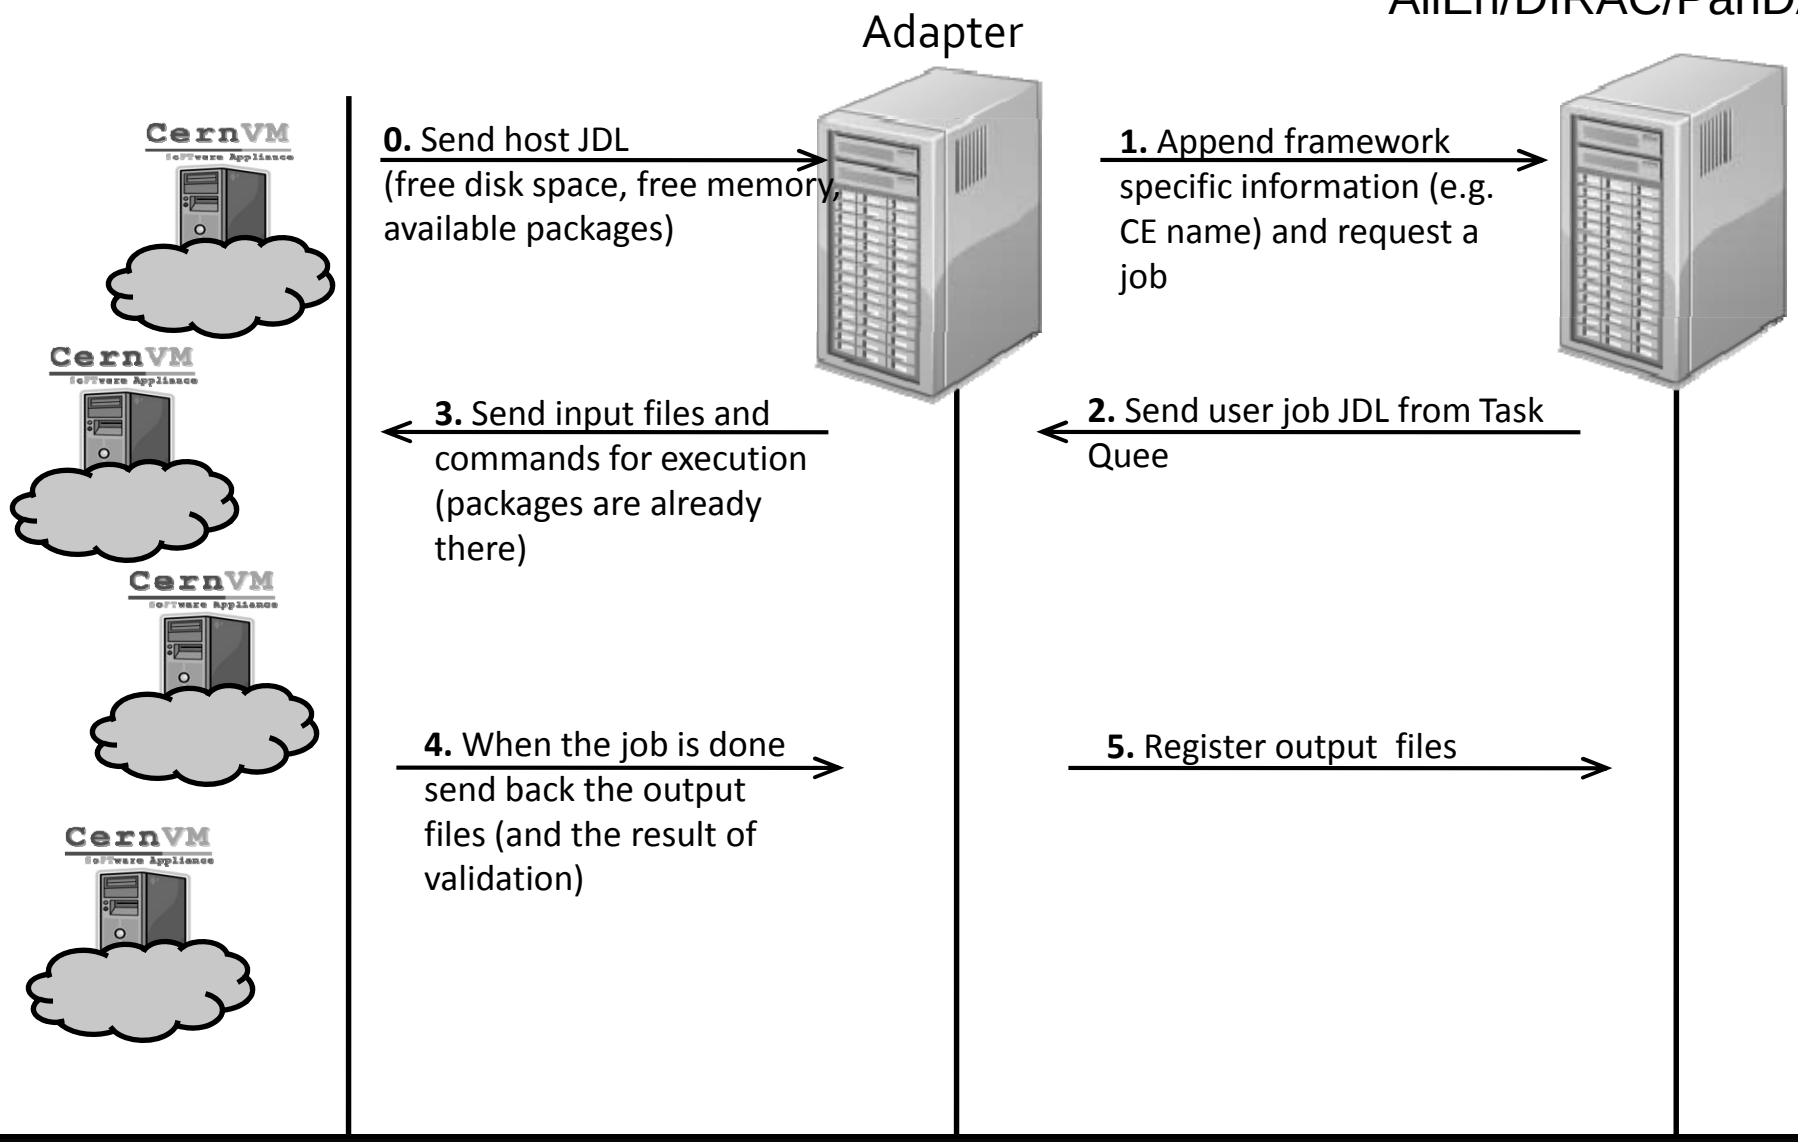

CernVM
SoFTware Appliance

Simplifying Grid deployment

- Nimbus (former Globus Workspace Service)
  - Nimbus is a set of open source tools that together provide an "Infrastructure-as-a-Service" (IaaS) cloud computing solution
    - http://workspace.globus.org/
  - Google Summer School (hosted at ANL) project to deploy a one-click, auto-configuring virtual Grid overlay for Alice/AliEn
    - Successfully created virtual AliEn site for ALICE with one command

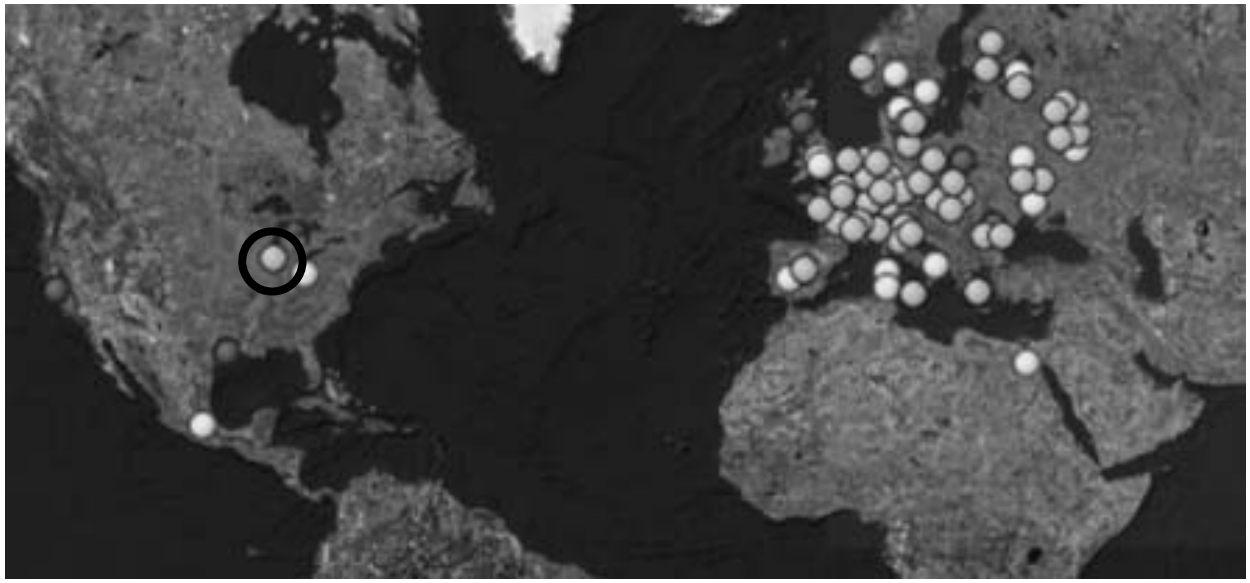

CernVM
SoFTware Appliance

CernVM Links…

- Mailing lists
  - cernvm-talk@cern.ch (Open list for announcements and discussion)
    cernvm.support@cern.ch (End-user support for the CernVM project)
- Savannah Portal
  - Please submit bugs and feature requests to Savannah at
    - http://savannah.cern.ch/projects/cernvm
- CernVM Home Page:
  - http://cernvm.cern.ch
- rBuilder & Download Page:
  - http://rbuilder.cern.ch

CernVM
SoFTware Appliance

Summary

- If you believe that it is time to jump on bandwagon and expolore directioin in which major software industry players are going and if you have one of these problems:
  - You work for LHC experiment and its software is not compatible with your favorite hardware or s/w platform running on your laptop/desktop
  - You do not want to spend time to manually keep software up to date
  - Your want to profit from the latest developments in CPU technology and use your multi/many core CPU to its maximal potential without modifying your application
  - You want to share spare CPU cores/cycles with others
  - You want to run your software on voluntary resources beyond the current Grid

…then CernVM might be what you are looking for.

- Version 1.0 available at http://rbuilder.cern.ch/project/cernvm/releases

CernVM

SoFTware Appliance

The End

CernVM
SoFTware Appliance

Multi core performance

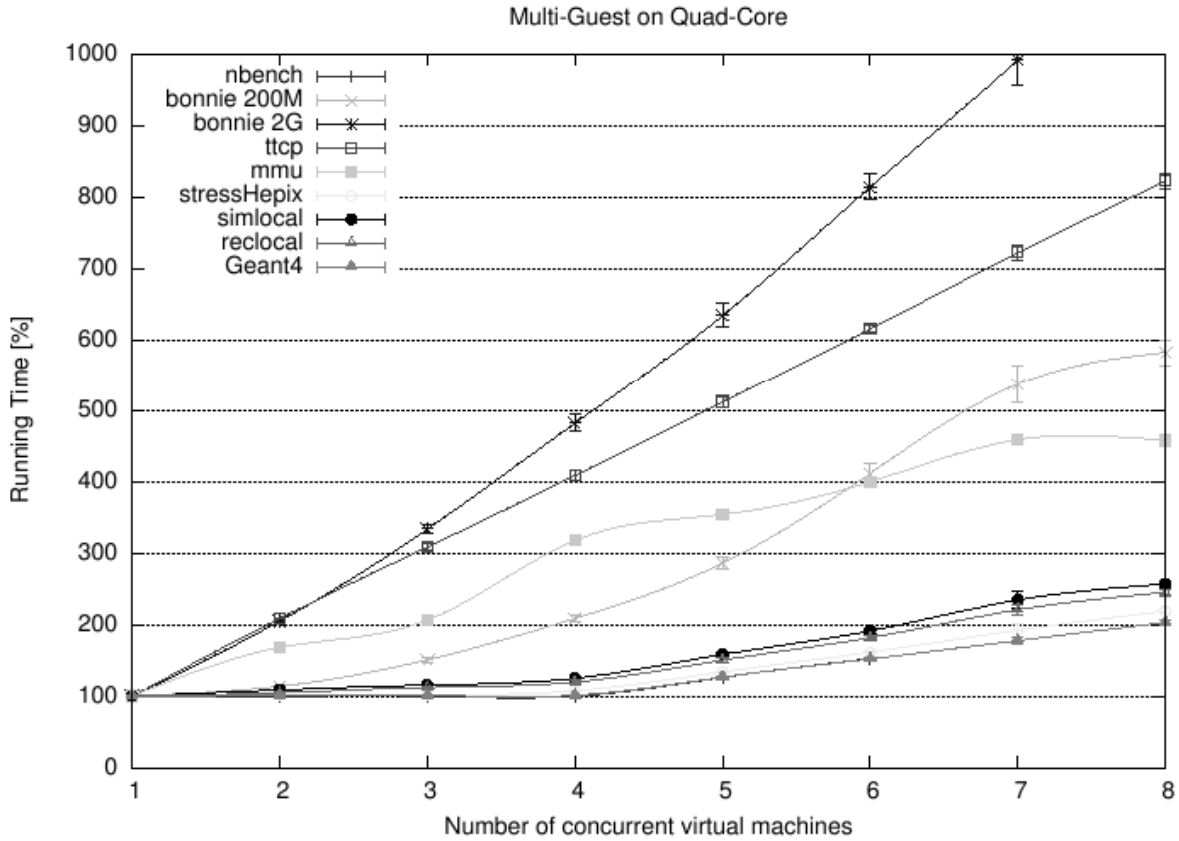

CernVM
SoFTware Appliance

Application benchmarks

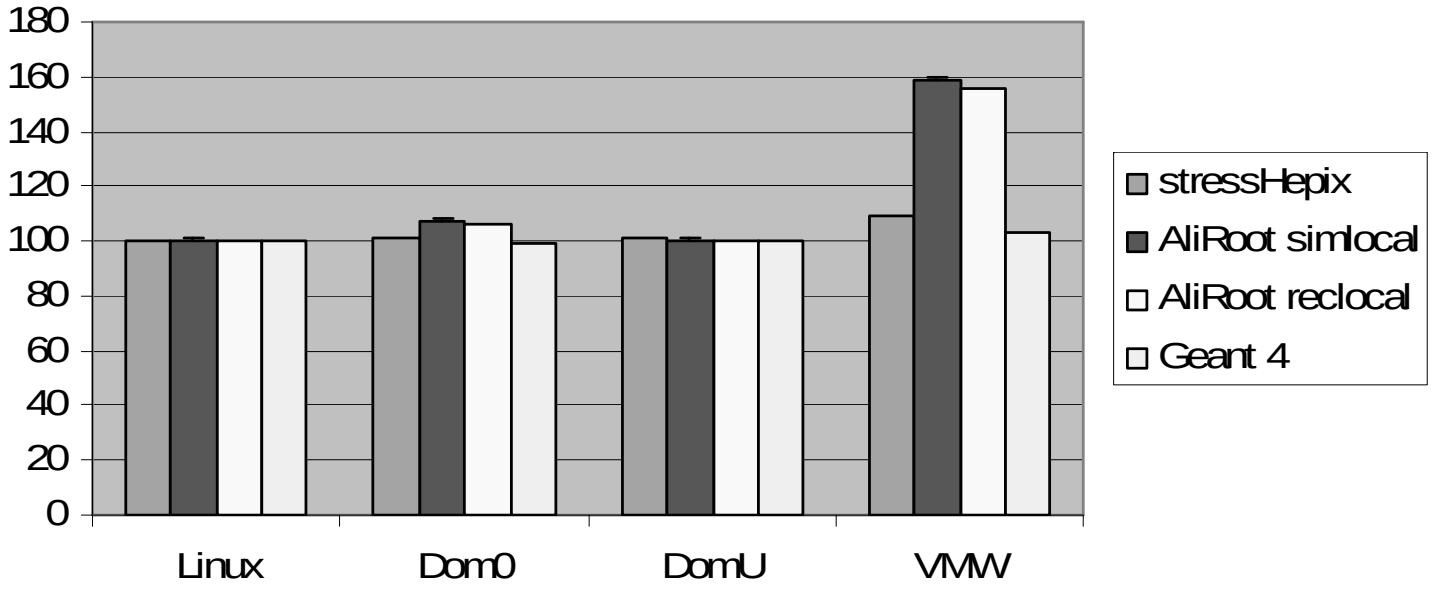

http://cernvm.cern.ch/cernvm/index.cgi?page=BenchmarkResults