

Sky computing




A grid of clouds
Sander Klous, Nikhef
29-06-2009

<http://indico.cern.ch/conferenceDisplay.py?confId=56353>

Content

- Use Cases
- Classification of Virtual Machines
- Security issues
- Virtual Machine Management
- Batch scheduling and Virtual Machines
- Conclusions

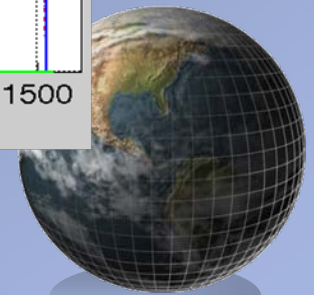
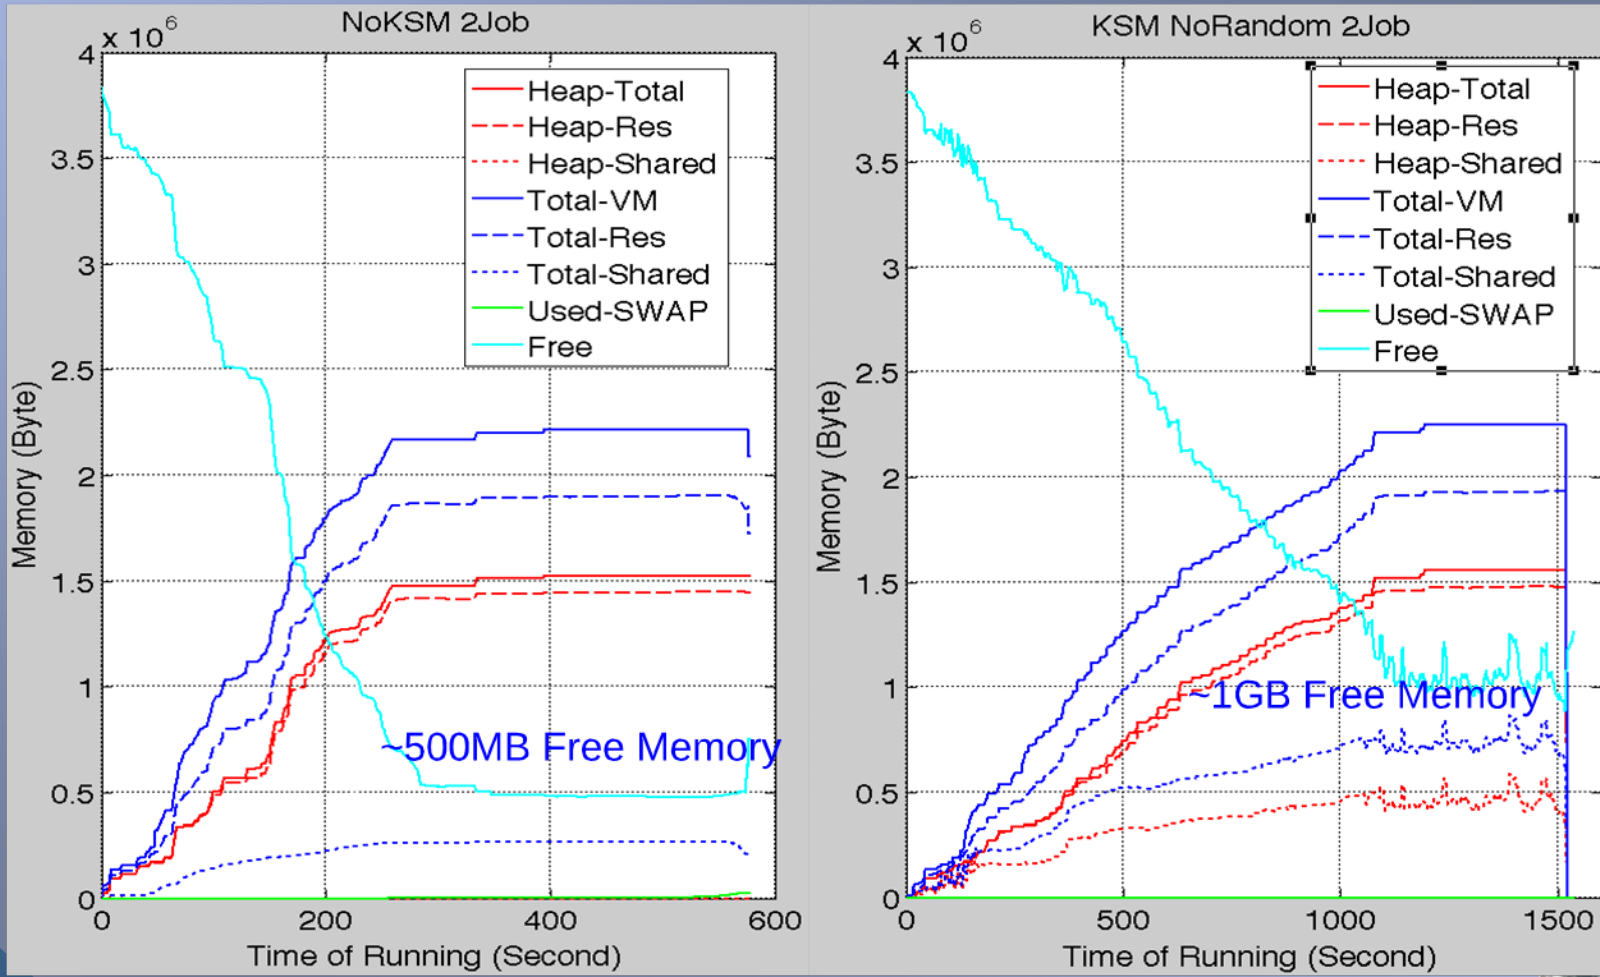


Use Cases (1)

- Resource flexibility
 - E.g. migration from SL4 to SL5
 - Provide both as virtual machine images
 - Balance resources with demand
- Resource management
 - Better scheduling options (preempt, migrate)
 - Difficult to manage many-core machines
 - Best to assign sandboxed resources
 - Virtual machines allow fine grained assignment
 - Simplification for multi-core and MPI jobs?




Multi-core, ATLAS KSM



Use Cases (2)

- User perspective, isolation from environment
 - Run exactly the same on desktop and grid
 - Run a Virtual Machine on the desktop
 - Run the same Virtual Machine on the grid
- Non-batch processing, e.g. pilot jobs overlay or
 - Run some of the workflow tasks locally
 - Replicate Virtual Machine on the grid
 - Run intermediate tasks on the grid
 - Run final tasks on the desktop



Classification of Virtual Machines



Service Level Status overview






Home Search KPIs Tags Admin Documentation Help

CERNVM WebFS Service 24 Jun 2009 Wed 23:13:35


CERNVM WebFS Service

availability: 
(more)
percentage: 100%
status: **available**

this service consists of:

-  CERNVM WebFS ALICE Service
-  CERNVM WebFS ATLAS Service
-  CERNVM WebFS CMS Service
-  CERNVM WebFS LHCb Service
-  CERNVM WebFS Grid UI Service

availability in the last 24 hours (more):




Additional information


full name: **CERNVM WebFS Service**
group: PH-SFT
email: cernvm.administrator@cern.ch
web site: <http://cernvm-webfs...>
manager: **Carlos Aguado Sanchez** 

Availability update

last update: 23:10:28, 24 Jun 2009
(3 minutes ago)
expires after: 15 minutes


 [rss feed with status changes](#)

Part of (subservice of):

 CERNVM Project

Admin

[admin tools](#)


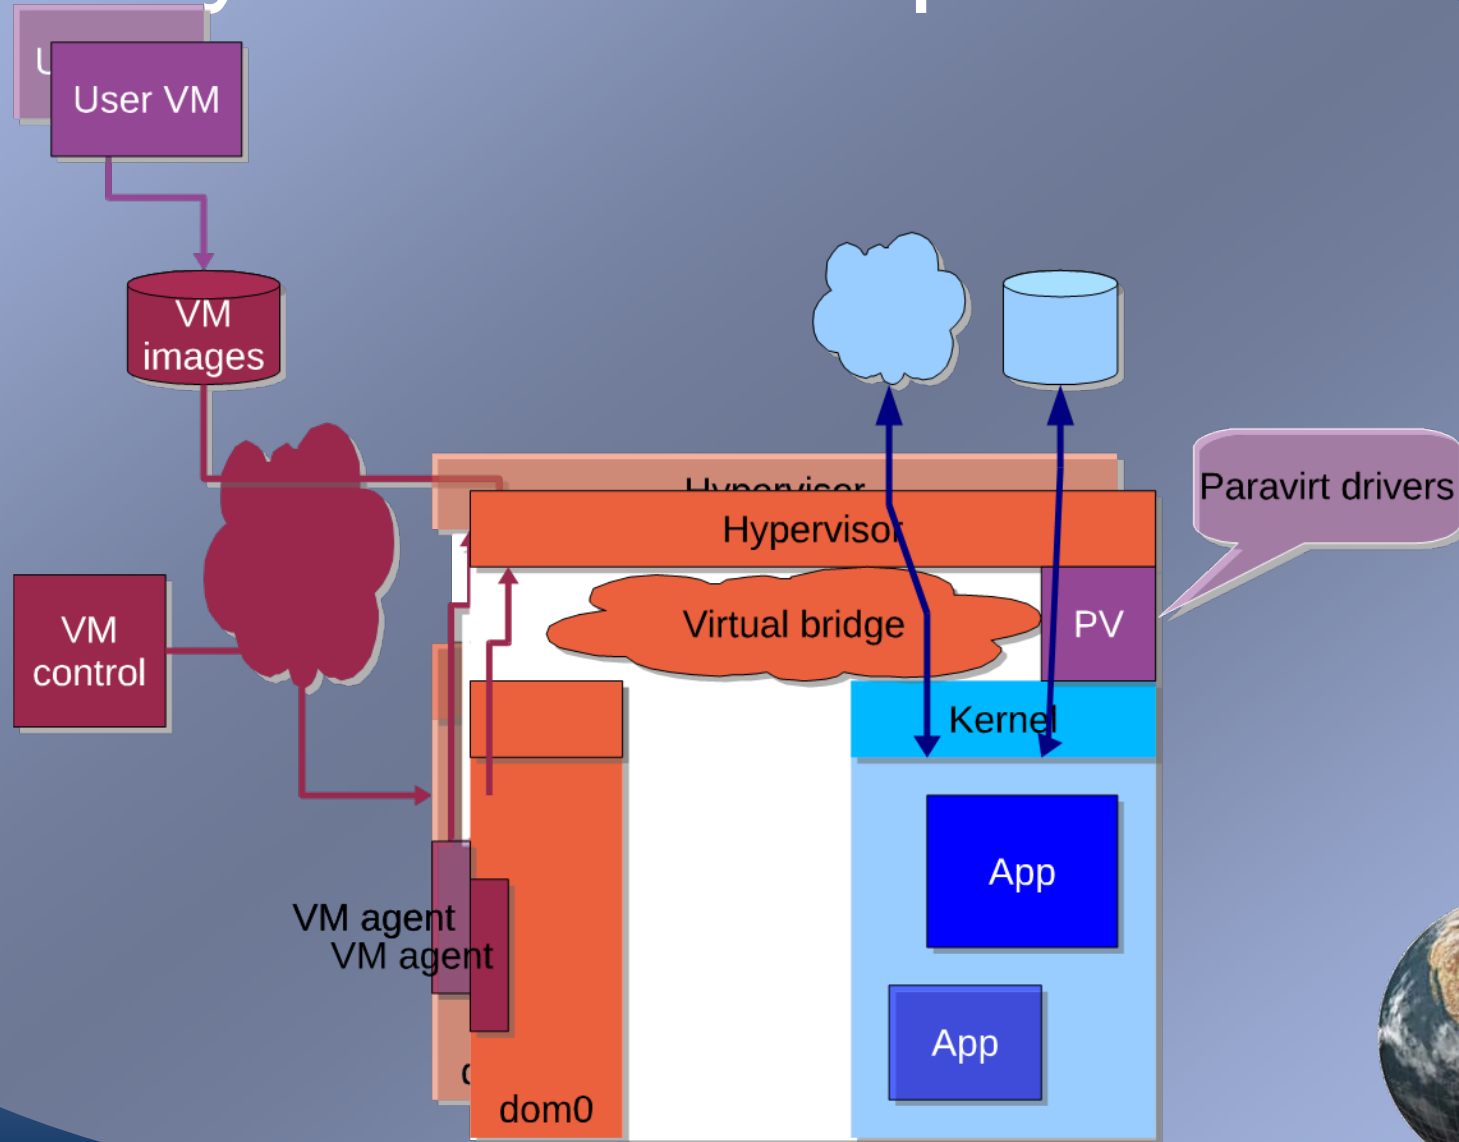


Classification of Virtual Machines

- Different levels of trust
 - Class 1: Site generated Virtual Machines
 - E.g. worker node on SL4 or worker node on SL5
 - Equivalent setup as with existing physical nodes
 - Class 2: VO generated Virtual Machines
 - E.g. CERNVM images with ATLAS software
 - Establish policies between site and VO
 - Inspection and certification
 - Batch system integration and monitoring tools
 - Class 3: User generated Virtual Machines
 - Requires sandboxing and a lot more...




Security issues - scope



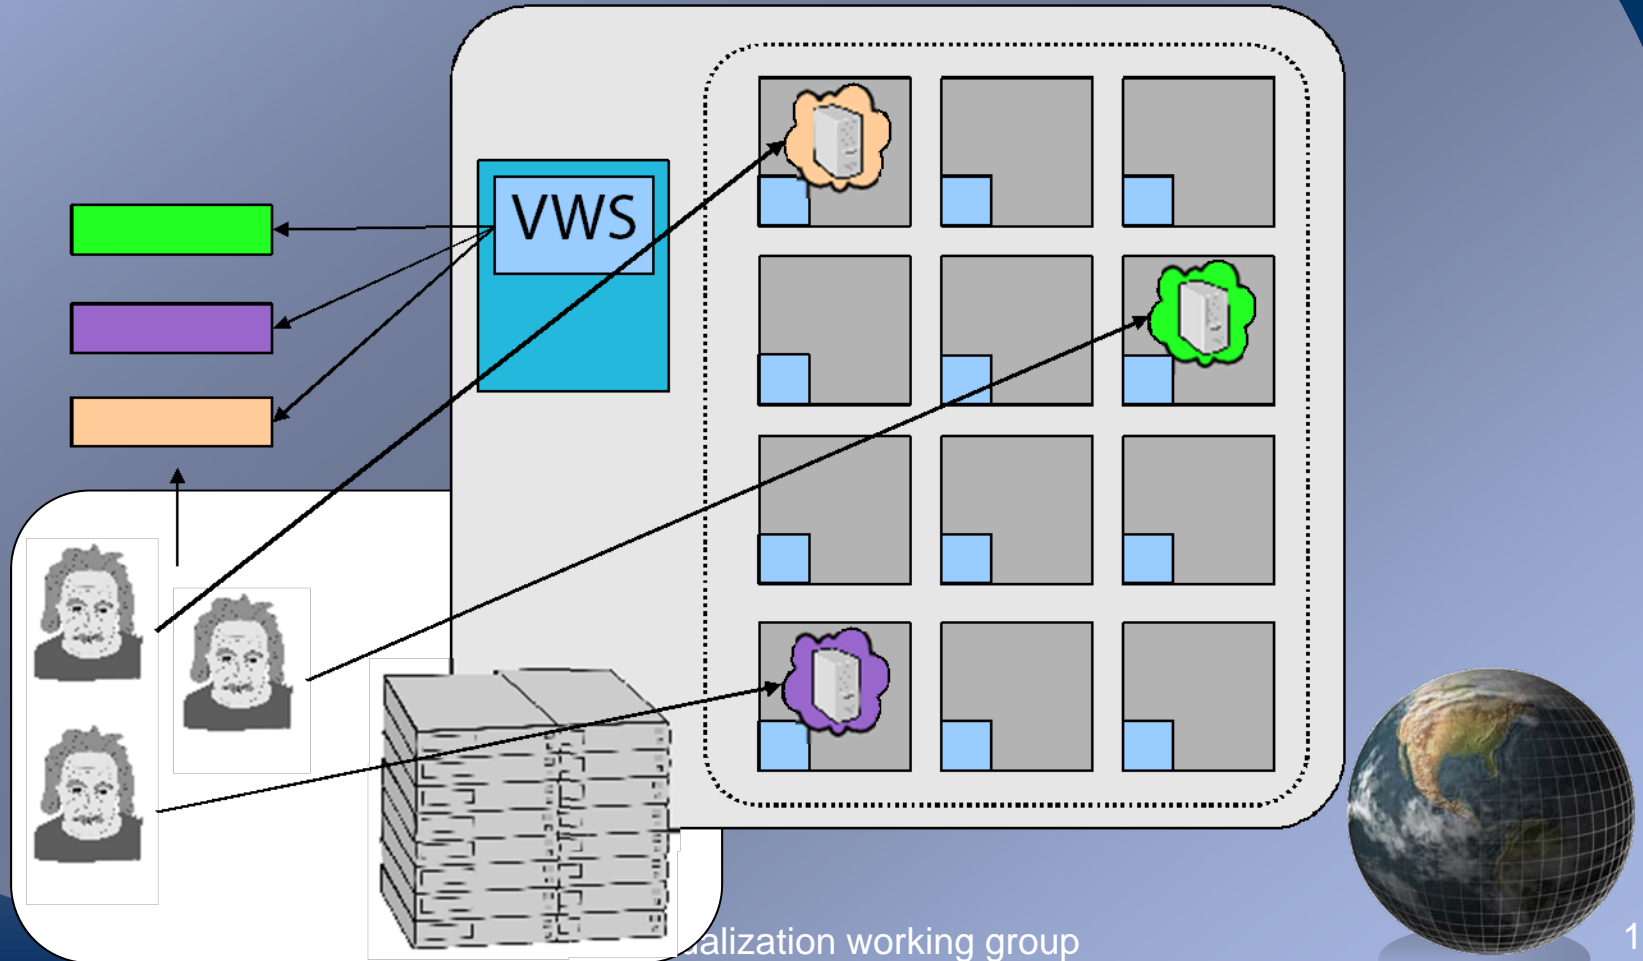
Security issues

- Site generated Virtual Machines
 - User accessible weaknesses in virtualization layer
 - VM break out, VM network, Para virtual drivers
- VO generated Virtual Machines
 - Management and trust infrastructure (suppliers?)
 - IP-based access control (e.g. NFS, firewalls)
 - Responsibility (patches) and reputation (legislation)
- User generated Virtual Machines
 - Transparency, traceability, provenance
 - Job management, credentials



Virtual Machine management

- The Virtual Workspace Service (Nimbus)
- Separating resource and job management




Batch scheduling and VMs

- More than 1 resource management layer
 - Batch job scheduling
 - Virtual Machine Management
 - (Pilot job frameworks)
- How to make them cooperate
 - LSF/VMO: adapt resources automatically
 - Nimbus/Haizea: Virtual Machine scheduler
 - Desy/INFN: Virtual Machines on demand

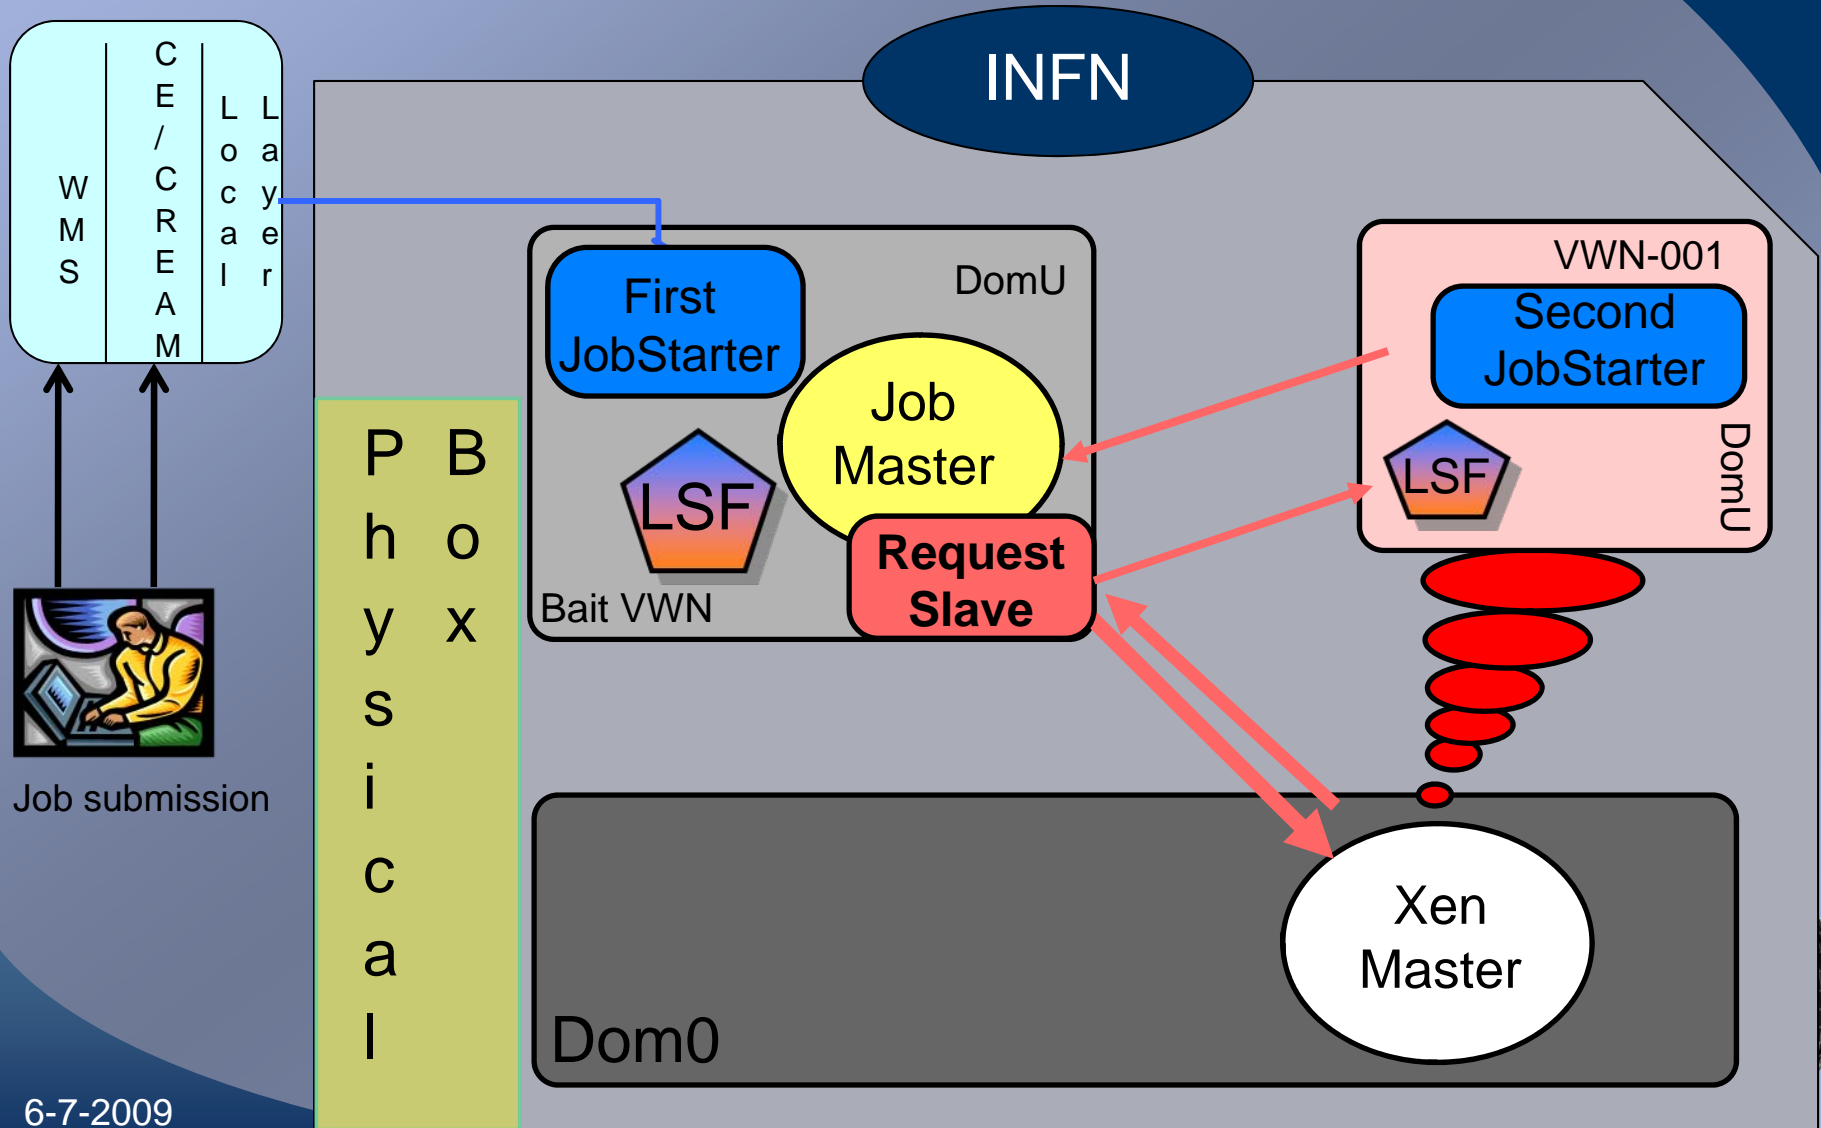


Adapt resources automatically

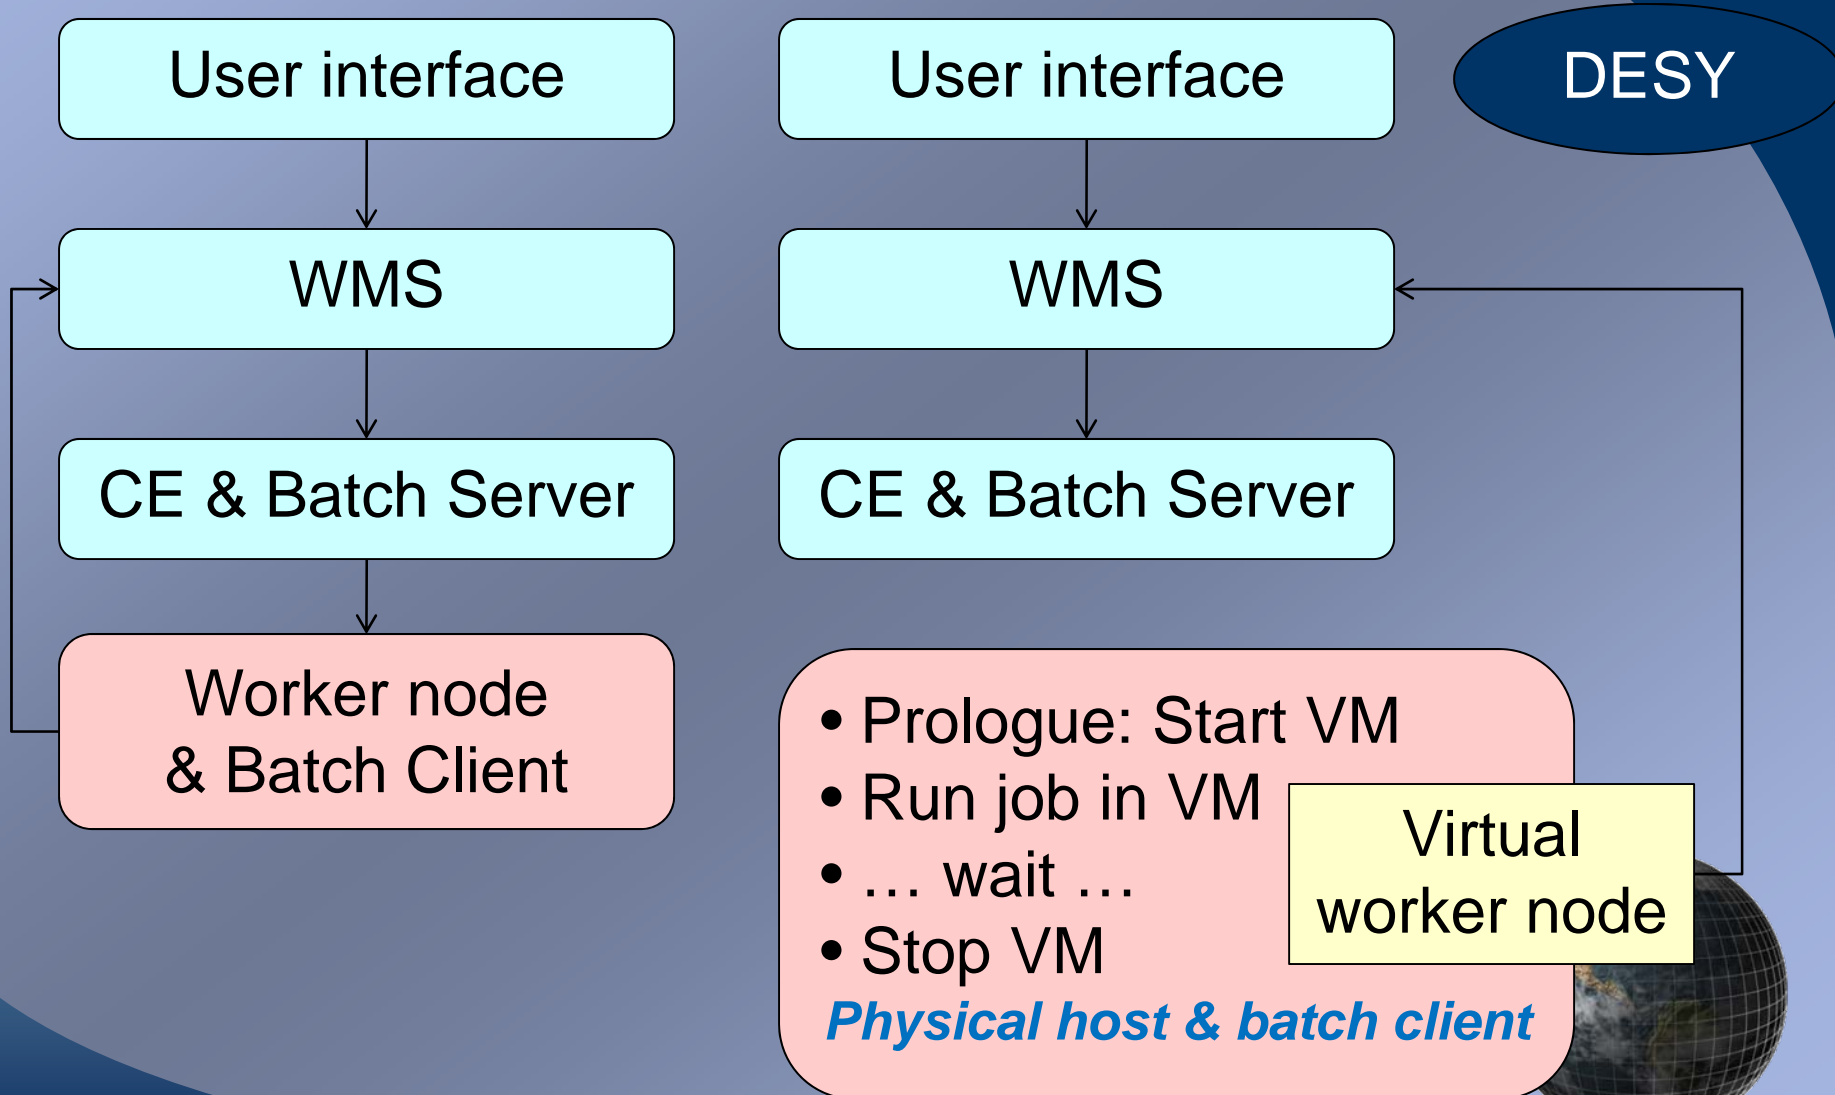
- Virtual Machine Orchestrator



Virtual Machines on Demand (1)


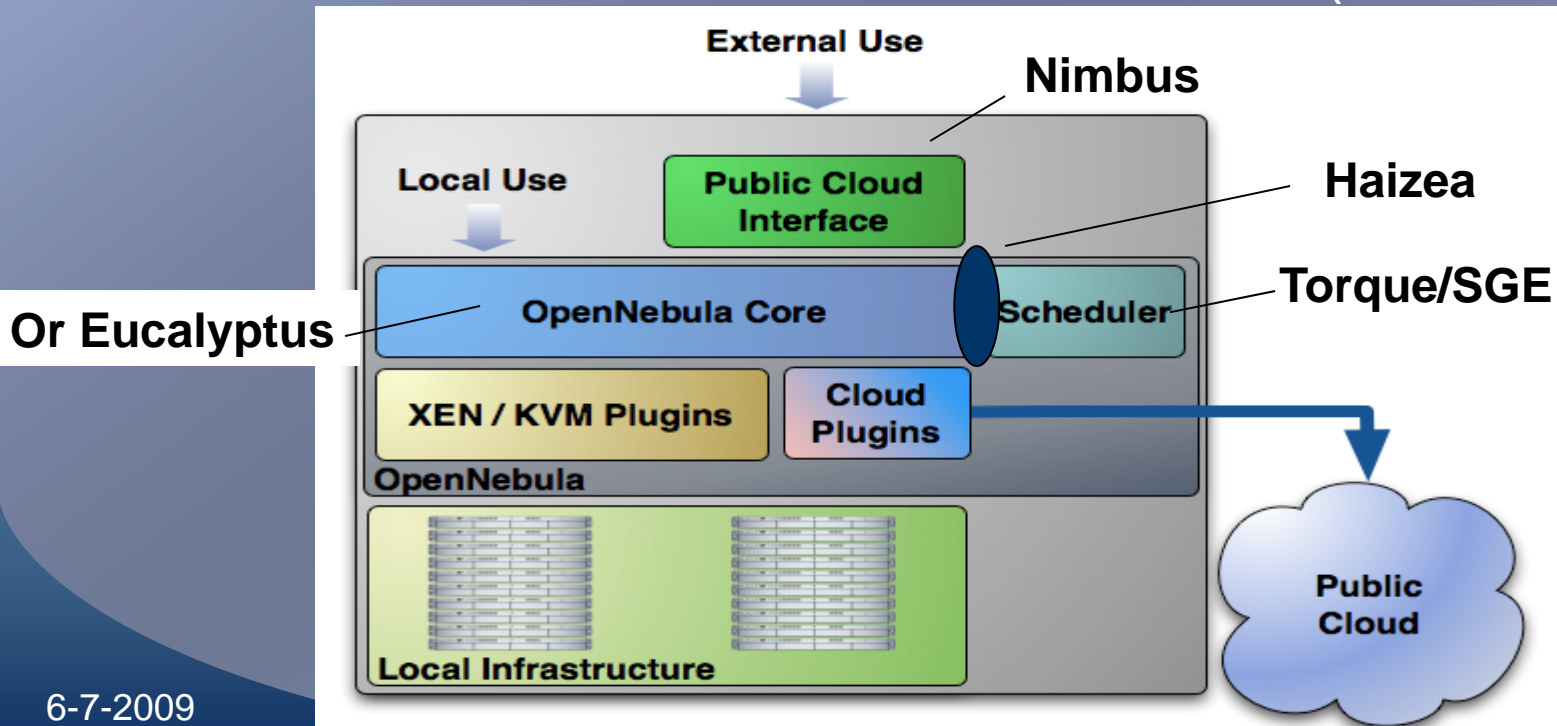


Virtual Machines on Demand (2)



Virtual Machine scheduler

- Haizea
 - Works with OpenNebula (VM management)
 - Backend for torque or Sun Grid Engine
 - Allows combination with Nimbus (cloud model)



Site impact

- Site generated virtual machines
 - Fabric management changes (quattor/jaime, cfengine)
 - IP addressing issues
- VO generated virtual machines
 - Scheduler changes
 - Torque integration with VM management
 - Policy and trust infrastructure development
 - VM/batch integration issues



Conclusions

- Use Cases
 - Infrastructure as a service based
- Security issues
 - Depends heavily on VM classification
 - User generated VMs not (yet) feasible
- Resource management
 - Many efforts (Coordination: EGEE/Reservoir)
 - Site generated: fairly easy and site beneficial
 - VO generated: mostly site solvable
 - User generated: a separate project



Project Scope

- Deploy a system to manage VO generated Virtual Machines at Nikhef.
- Sub efforts
 - Investigate use cases
 - Determine site policy
 - Assess security implications
 - Solve technical issues
 - Generalize (document, disseminate, etc.)

