

# Virtualization in the gLite Grid Middleware software process


Use cases, technologies and future plans

*Lorenzo Dini*

*Tomasz Wolak – Alberto Resco Perez – Andrew Elwell*

*HEPIX Spring 2010 - 22 April 2010*


- gLite and GT
- Cloud-like virtualization
  - vNode
- Batch-like virtualization
  - VMLoader
- Future Plans






- ~2 M lines of code
- 258 RPMs produced
- ~ 70 external dependencies
- 17 programming languages
- 15 platforms/architectures (4 supported)
- 21 nightly builds
- ~3 hours to build on VM
  - 2 GB RAM
  - 1 of 8 cores 2.33Ghz Intel XEON
  - Heavy IO


- Development
  - DPM, LFC, FTS, BDII, Grid Monitoring
- Build
  - ETICS Build and Test system
- Integration
- Automatic Testing
  - RPM deployment and YAIM configuration tests
- Certification and Test-Beds
  - Manual certification
  - Grid Test-Bed infrastructure




- Cloud-like virtualization
  - Development support
  - Manual certification
  - Grid Services in Test-Beds
- Batch-like virtualization
  - Build
  - Installation test
  - YAIM Configuration test
  - Single- and Multi- Node functional test

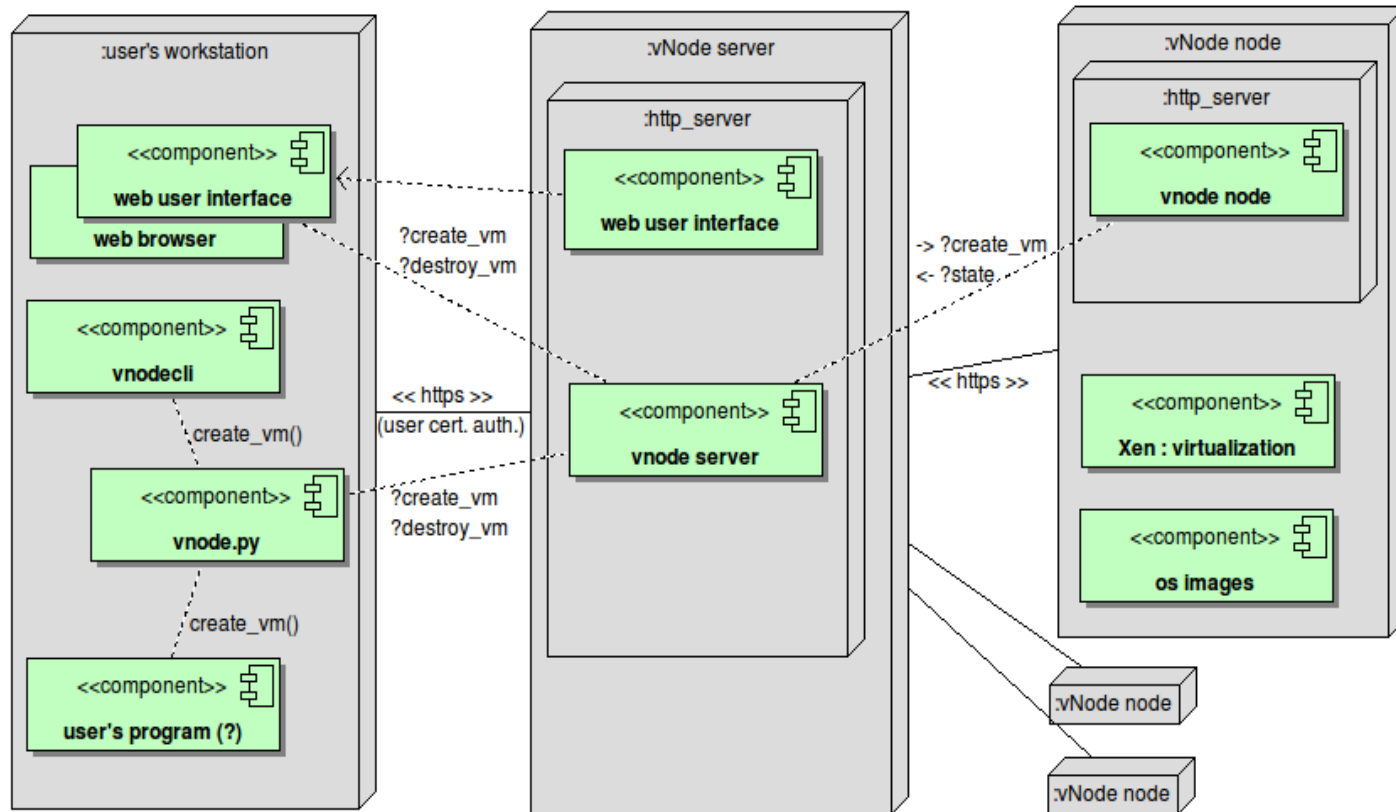


- Usage
  - Interactive **root** access via SSH
  - Many VM deployed
  - **Low** memory / processor / IO
- Requirements
  - **Fast** creation (~ 5 min) of VMs for daily use
  - Variety of platforms: **SLC, SL, Debian**
  - **CLI** and **APIs** for scripting
  - Host certificates in VMs
  - Monitoring and statistics
  - Revert to snapshot



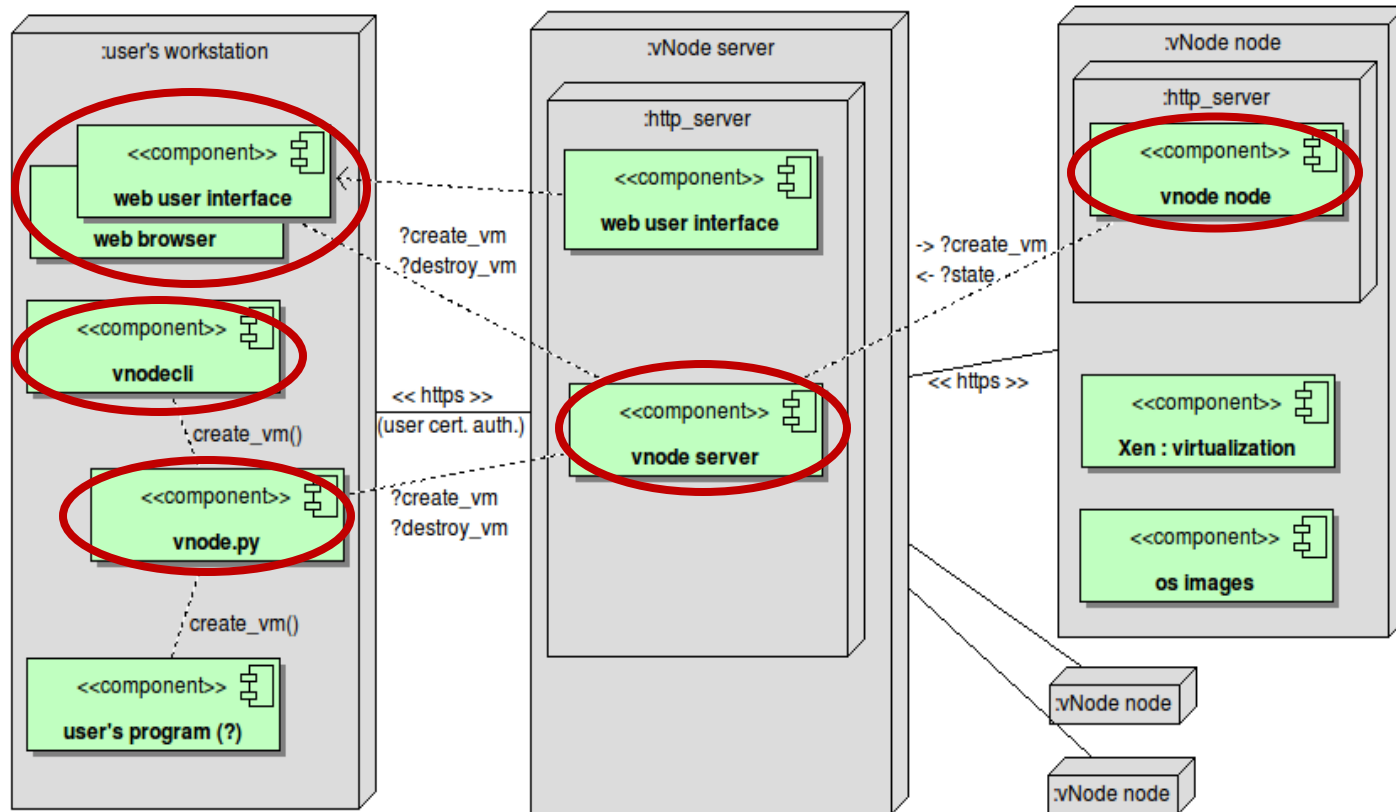
- A simple system to manage XEN virtual machines. It can deploy, terminate, check and show state of virtual machines
- Main developer: Ricardo Mendes
- API
  - AJAX Web Interface
  - Python API
  - CLI interface
  - Server and Deployment node JSON WebServices
- Security
  - Login via SSH keys to VM
  - HTTPS and X509 to access the service






- XEN 3.1
- vNode Service:
  - apache 2.2.3 with mod\_ssl
  - python-2.4, with httplib, ssl and simplejson





- Reserve / Create / Pause / Destroy VMs
  - Options: Memory / Disk / Image
- List / Status / Uptime
- Availability



vNode v2 server - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Address bar: <https://lxbf2706.cern.ch/vnode-2.0/>

Most Visited Getting Started gLite CERN TestBed Latest Headlines CERN market: M...

Navigation: CERN-IT-... vNo... Network... soappy g... Web Ser... 12.3. Firs... Python ... python s... The Pyth... How to d... 16.7. get...

## vNode

virtual nodes on demand

Virtual Machines Virtual Grids Admin OS Farm

### State of Xen Virtual Machines

State of all Xen Virtual Machines

Deployed

<< first < prev 1 2 3 4 5 6 7 8 9 10 next > last >>

VirtualHost	StateMessage	Message	StateColor
vtb-generic-87.cern.ch	deployed	Could verify connection to VM	Green
lxb7607v1.cern.ch	deployed	Could verify connection to VM	Green
lxb7604v5.cern.ch	deployed	Could verify connection to VM	Green
lxb7608v3.cern.ch	deployed	Could verify connection to VM	Green
vtb-generic-26.cern.ch	deployed	Could verify connection to VM	Green
lxb7604v4.cern.ch	deployed	Could verify connection to VM	Green
vtb-generic-59.cern.ch	deployed	Could verify connection to VM	Green
vtb-generic-25.cern.ch	deployed	Could verify connection to VM	Green
vtb-generic-91.cern.ch	deployed	Could verify connection to VM	Green
vtb-generic-107.cern.ch	deployed	Could verify connection to VM	Green

<< first < prev 1 2 3 4 5 6 7 8 9 10 next > last >>

#### Uptime of Xen Virtual Machines

Uptime of all deployed Xen Virtual Machines

<< first < prev 1 2 3 4 5 6 7 8 9 10 next > last >>

Owner	PhysicalHost	VirtualHost	UpTime
malandes	lxb7604.cern.ch	lxb7604v1.cern.ch	111 days, 8 hours, 8 minutes, 8 seconds
malandes	lxb7604.cern.ch	lxb7604v3.cern.ch	280 days, 10 hours, 12 minutes, 5 seconds
malandes	lxb7604.cern.ch	lxb7604v5.cern.ch	280 days, 4 hours, 11 minutes, 18 seconds
malandes	lxb7604.cern.ch	lxb7604v4.cern.ch	261 days, 8 hours, 58 minutes, 17 seconds
malandes	lxb7604.cern.ch	lxb7604v2.cern.ch	269 days, 6 hours, 16 minutes, 58 seconds
pguerrer	lxb7605.cern.ch	lxb7605v4.cern.ch	3 days, 6 hours, 51 minutes, 33 seconds
pguerrer	lxb7605.cern.ch	lxb7605v2.cern.ch	3 days, 6 hours, 51 minutes, 33 seconds
pguerrer	lxb7605.cern.ch	lxb7605v1.cern.ch	3 days, 6 hours, 51 minutes, 33 seconds
pguerrer	lxb7605.cern.ch	lxb7605v3.cern.ch	3 days, 6 hours, 51 minutes, 33 seconds
pguerrer	lxb7605.cern.ch	lxb7605v5.cern.ch	3 days, 6 hours, 50 minutes, 51 seconds

<< first < prev 1 2 3 4 5 6 7 8 9 10 next > last >>

Refresh

#### Deploy a Xen Virtual Machine

Configure your Xen Virtual Machine

Physical Hosts: lxbf2707.cern.ch Virtual Hostnames: vtb-generic-84.cern.ch Expiry Time(days): 0 Memory (MB): 1000 Partition (GB): 10 OS Image: SL4-32-DH Service Type: Select...

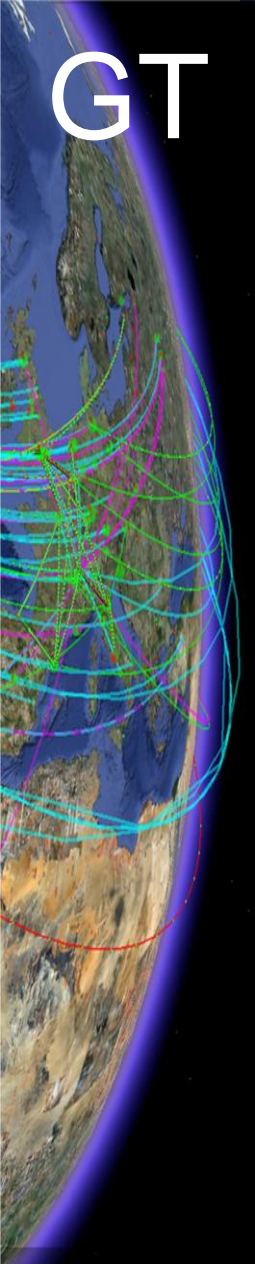
Advanced Options

Image Filename: Virtual Machine Name: W2PsRq2TI00Ce

Xen Virtual Machine(s) that will be deployed or are currently reserved to you							
Name	PhysicalHost	VirtualHost	ExpiryTime	Memory	Partition	OSImage	ExpiryTimeAt
myvm1	lxbf2707.cern.ch	vtb-generic-98.cern.ch	0	1000	10	SL4-32-DH	09:51:11

Deploy Add Delete Refresh

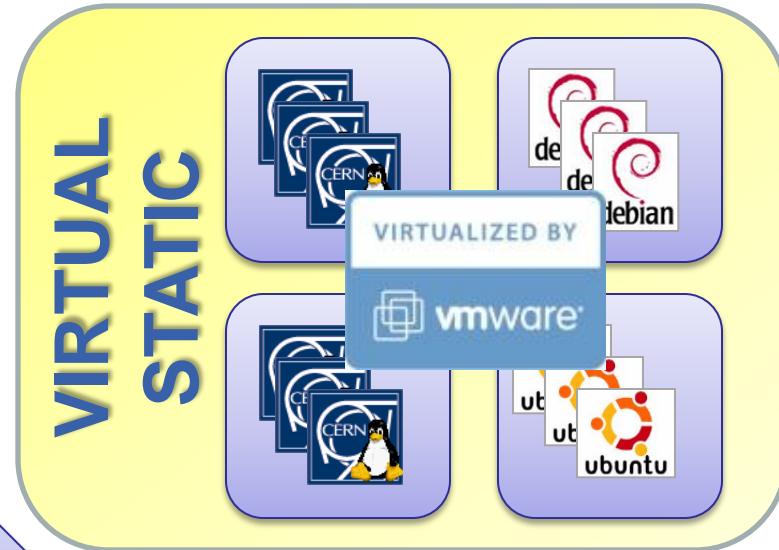
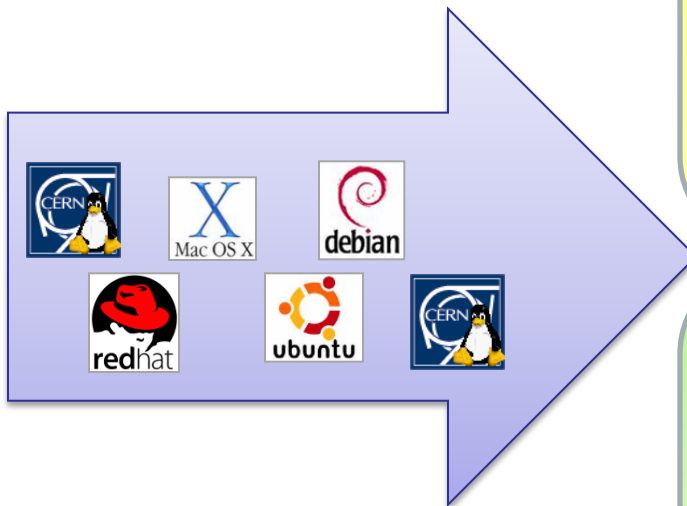
Done



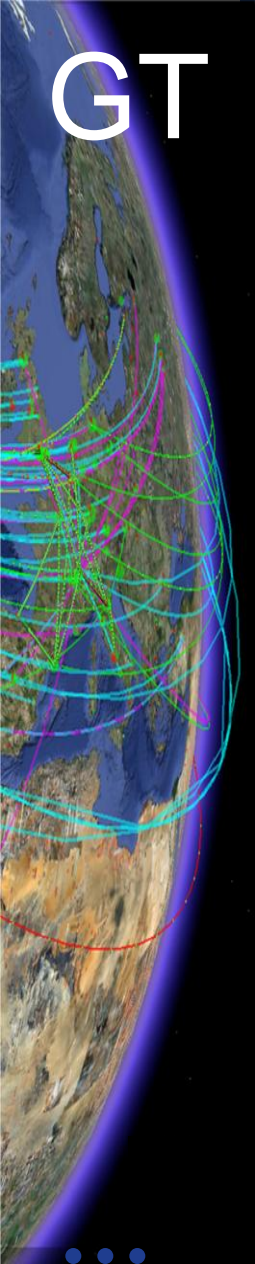
<b>Users:</b>	<b>~ 33 (25 active)</b>
<b>VMs:</b>	<b>&gt; 100 (reaching limit)</b>
<b>VM creation time:</b>	<b>2-3 minutes</b>
<b>VM host-names:</b>	<b>~150 pre-registered</b>
<b>Hypervisors:</b>	<b>9</b>
<b>VM per HV:</b>	<b>15 – 10 – 5</b>
<b>VM Stats:</b>	<b>1 core, 250MB – 1GB RAM</b>
<b>Images:</b>	<b>~10 platforms (SL, SLC, Deb)</b>

- Usage
  - User space (builds) and as **root** (tests)
  - VM starts and ends within a job submission
  - Memory, Processor and IO **intensive**
- Requirements
  - **Fast start (~ 5 min)**: 1 VM start each job
  - **Custom image loading** w/o much validation
  - **Snapshot** available after job for debugging
  - Logging of VM boot and execution
  - Host **certificates** in VM
  - Wide set of images required

**run as root:**  
revert to snapshot  
after each job




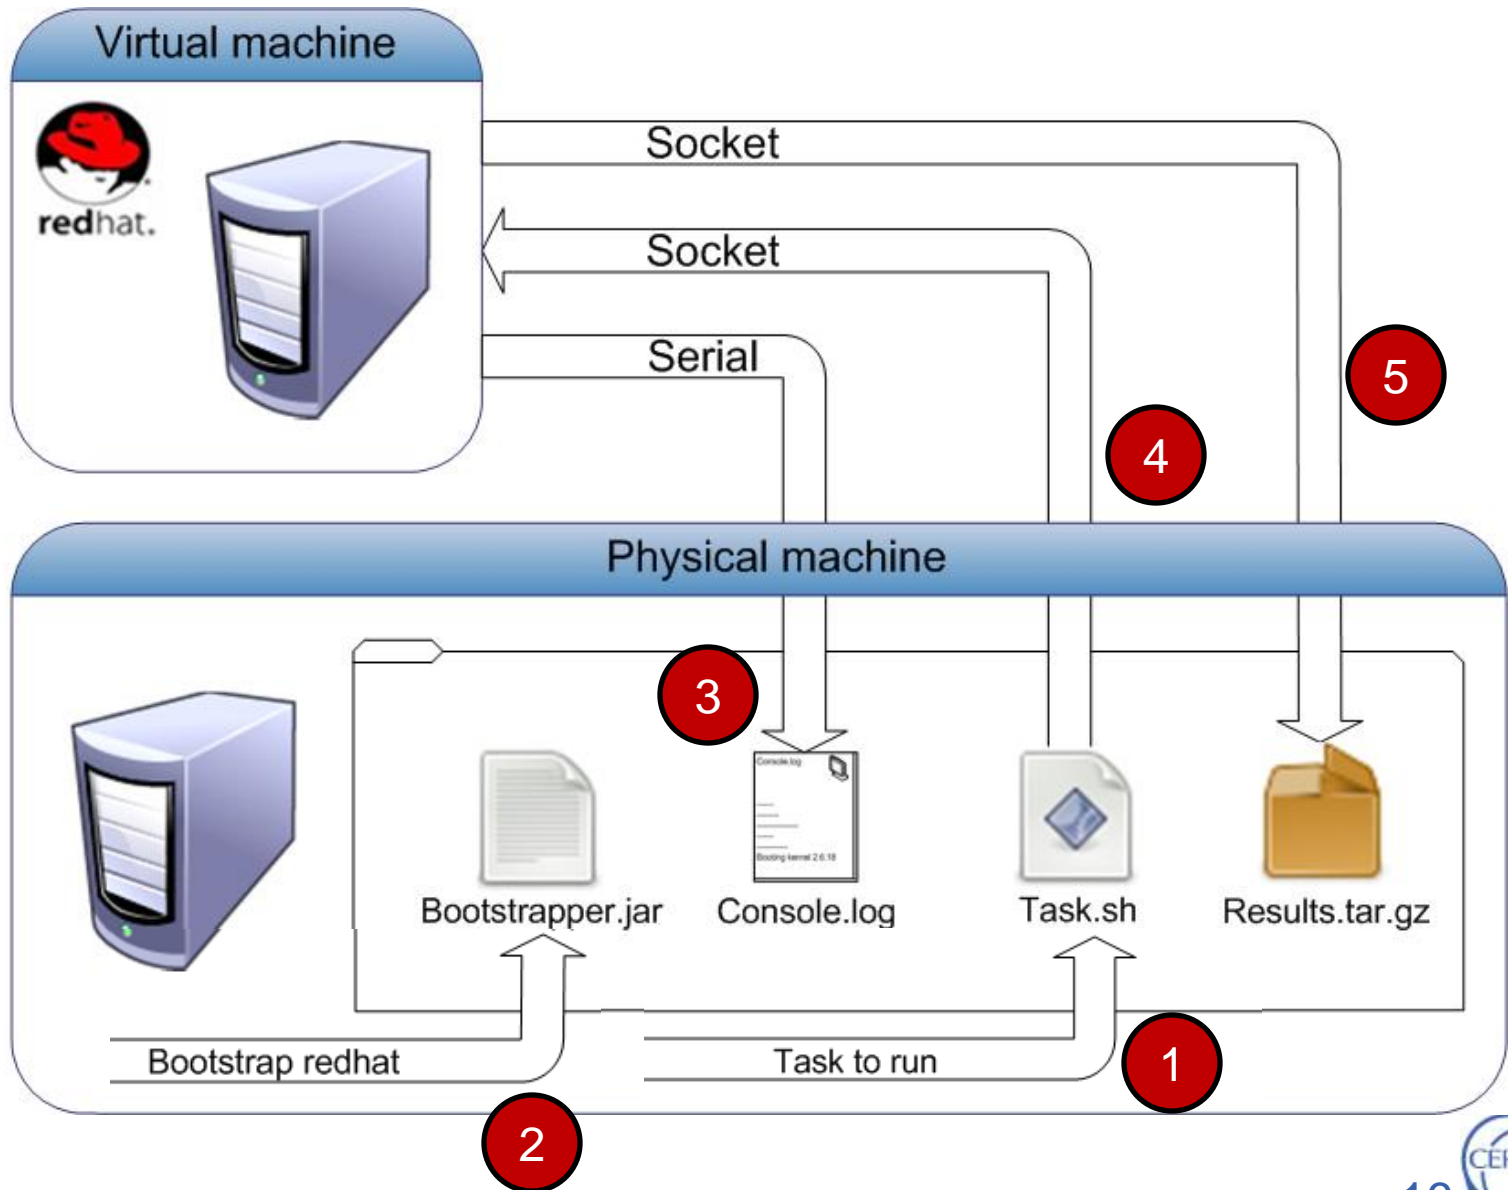

VMWare Server 1.0.4  
Condor 6.9.5 to 7.4.2  
Metronome 2.5.2b



<b>Users:</b>	<b>~ 300 (60 active)</b>
<b>Jobs per month:</b>	<b>~3000</b>
<b>Job duration:</b>	<b>5 minutes to 4-5 hours</b>
<b>Hypervisors:</b>	<b>~20</b>
<b>VM per HV:</b>	<b>3MAX</b>
<b>Slots per VM:</b>	<b>2 (1 for root-enabled)</b>
<b>VM Stats:</b>	<b>1 core, 2GB RAM, 30GB Disk</b>
<b>Images:</b>	<b>~20 platforms + custom</b>

- A Java tool for the execution of jobs in virtual environments
  - **Bootstrapper** in hypervisor and **Agent** in VM
- Prototype stage
- Runs in user space, detects installed hypervisors
  - **VMWare**, XEN, KVM and QEMU
- Images:
  - Downloaded from central repository
  - Locally cached







- Move to the CERN IT central service
  - + Ability to choose hostname
  - + Large availability of resources
  - + SOAP APIs available
  - + Standard guaranteed service
  - + “30 days free” 😊
- ~ Speed of VM creation (up to 20 min)
- ~ Variety of required images (SL, Debian)
- ~ Javascript interface hic-ups with Firefox
- ~ Reuse of hostname takes time due to DNS flush

