

# Grid Interoperability Centre

CERN openlab II Quarterly Review  
9 October 2007

Sverre Jarp  
CERN openlab CTO

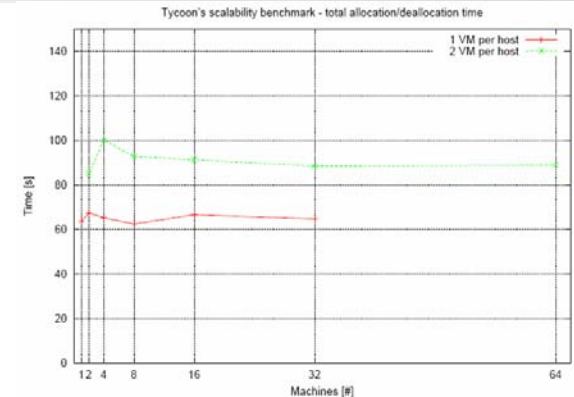
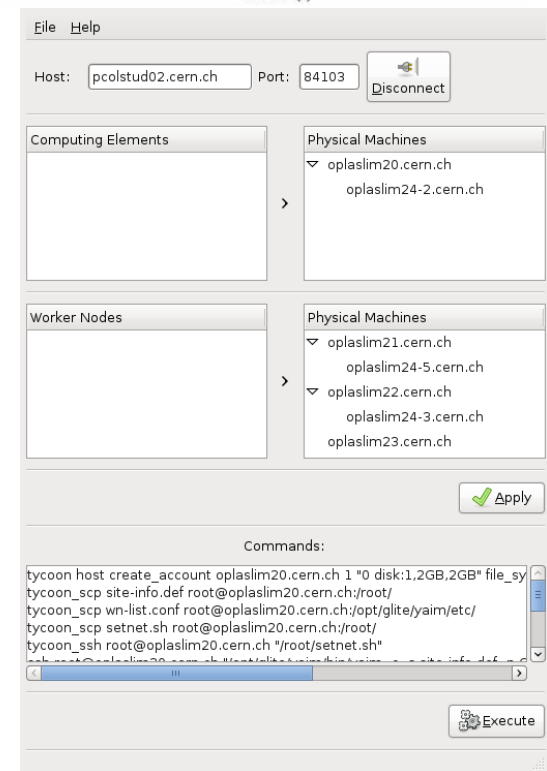
Main contributors:  
H.Bjerke, J.Dana, X.Gréhant





- Tycoon
- Grid scheduling
- Virtualization
- Conclusions

- Scalability tests done
  - with support from summer student, Andrea Sottoriva
  
- Software ported to Scientific Linux CERN 4 and Xen 3.0.3
  
- Integration with EGEE was completed
  
- Presentation given at EGEE Conference 2007 in Budapest
  - within the Business Track

File Help

Host: pcolstud02.cern.ch Port: 84103 Disconnect

Computing Elements

Physical Machines

- o plaslim20.cern.ch
  - o plaslim24-2.cern.ch

Worker Nodes

Physical Machines

- o plaslim21.cern.ch
  - o plaslim24-5.cern.ch
- o plaslim22.cern.ch
  - o plaslim24-3.cern.ch
  - o plaslim23.cern.ch

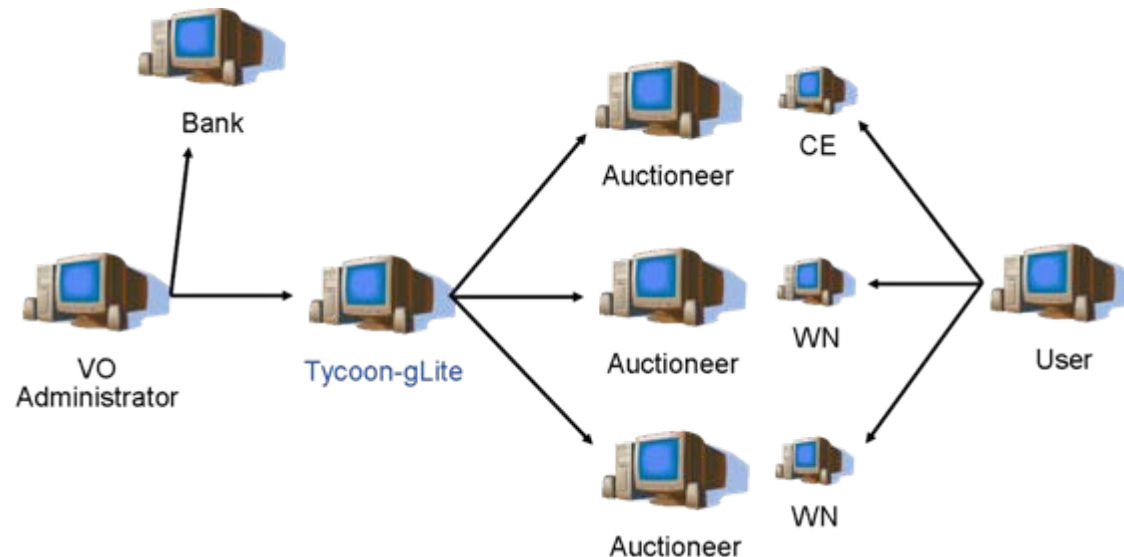
Apply

Commands:

```
tycoon host create_account opaslim20.cern.ch 1 "0 disk:1,2GB,2GB" file_sy
tycoon_scp site-info.def root@opaslim20.cern.ch:/root/
tycoon_scp wn-list.conf root@opaslim20.cern.ch:/opt/glite/yaim/etc/
tycoon_scp setnet.sh root@opaslim20.cern.ch:/root/
tycoon_ssh root@opaslim20.cern.ch "/root/setnet.sh"
```

Execute

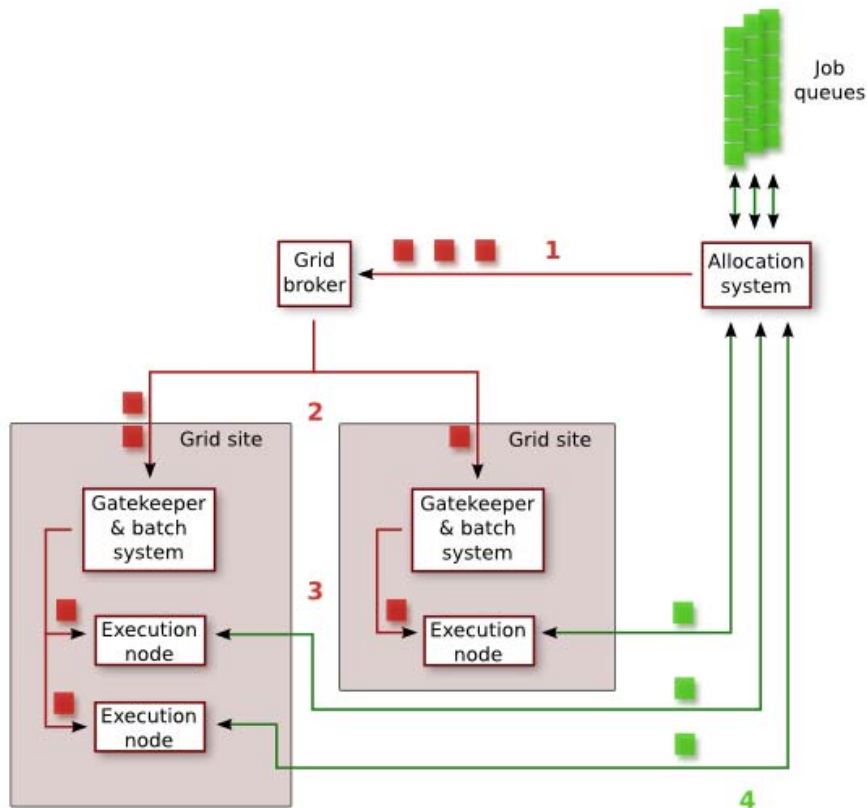
- Our goal: A dynamic Grid infrastructure using a market-driven approach
  
- Joint development with HP Labs where:
  - HP Labs did the porting of Tycoon to SLC4 and Xen 3.0.3
  - CERN openlab developed the integration with EGEE in order to deploy CEs and WNs on demand (both running on SLC3)




- Involve more institutions in our project and increase the number of machines running Tycoon-gLite
- Extend Tycoon-gLite functionality deploying Storage Elements (SEs) on demand
- Support for CEs and WNs running on SLC4
  - as soon as gLite developers certify for production the CE!
- Port Tycoon to Xen 3.1?
  - By HP Labs

## ■ Scheduling per Virtual Organization

- New trend in major VOs
  - Atlas: *Cronus*, CDF: *GlideCAF*
  - Alice: *AliEn2*, LHCb: *DIRAC*
  - Application framework: *DIANE*
- We collaborate with Atlas, Alice
  - Virtual environments (custom, controlled)
  - Allocation performance
- Allocation frame
  - Central system submits monitors to the grid
  - From execution nodes, monitors contact central system back
  - Thousands of nodes controlled by VO
  - Fine and dynamic resource allocation now possible

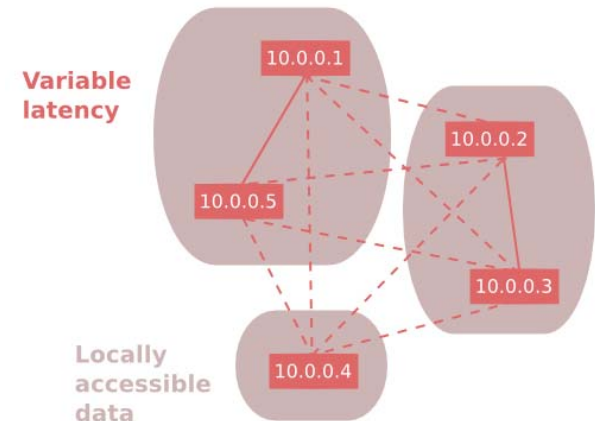
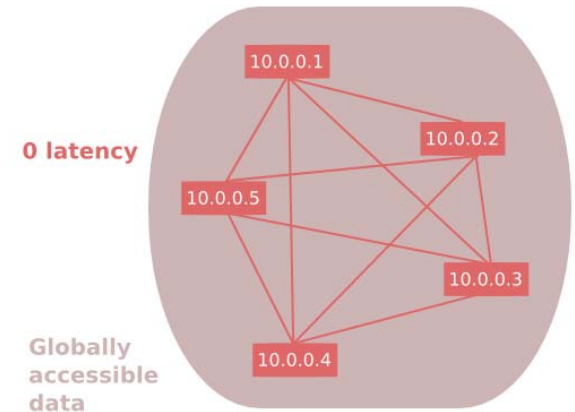


- Goal: can we save “years” of data analysis?
  - By doing precise, dynamic allocation
    - Analyse new constraints
    - Choose allocation algorithm
- Simulating this environment
  - Level-lab:
    - Simulates the environment a VO gets on the grid
    - Evaluates performance of allocation algorithms
    - Development done jointly with Serena Cameirano (Summer student)
  - Status
    - 3000 lines of code, 5000 of unit tests
    - Simple model working. Successive refinements to come (job and resource profiles accuracy)
    - Of course, 3D visualization to complete ;-)



OSG and EGEE sites on Level-Lab  
visualization module

- Elaborating a model
  - What's new compared to cluster scheduling?
  - How does it affect the allocation decisions?
  - How does it affect the decision process?
- Level-Lab implements the model to simulate the environment
  - We know when a bunch of jobs fits well on a node
  - Challenging question: when is the allocation good globally?
    - Try your algorithms on Level-Lab and compare
    - Then choose the best to really deploy



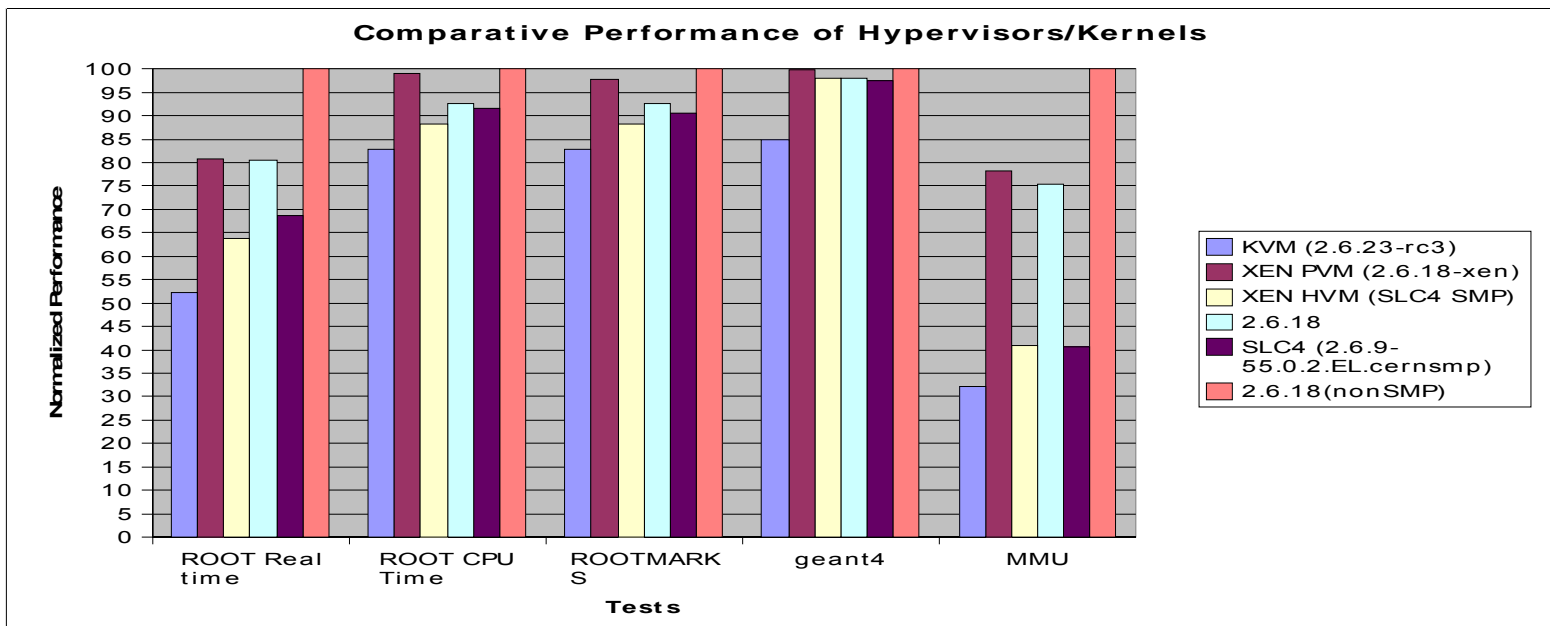




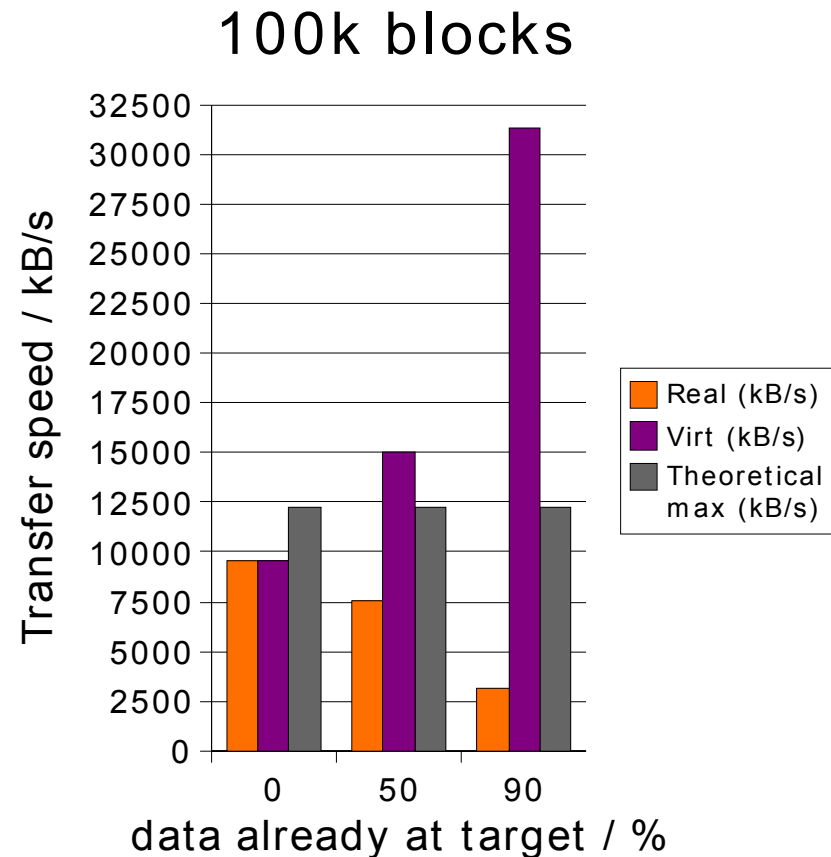
CERN  
openlab

# Virtualization Benchmarks

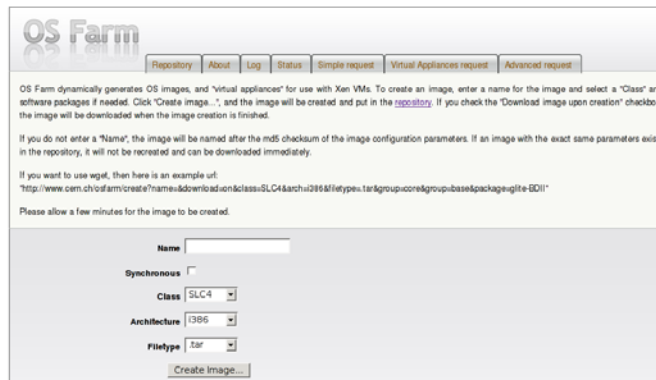
- A set of synthetic and application-benchmarks points to aspects that need improvement
  - Sensitive operations can incur a moderate to significant impact on performance



- VM images are big
  - Long network transfer times
  - Congests network
- But most images are relatively similar
  - Just transfer the delta
- Tool/library developed for efficient transfer of VM images



- Various projects at CERN need quick deployment of clean machines
  - Software building and testing
  - Need a rich set of Linux flavours
  - Complementary tool, OS Farm, provides
    - RedHat (SLC, SL, CentOS) and Debian based Xen VM images
    - gLite and Quattor virtual appliances



**OS Farm**

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

OS Farm dynamically generates OS images, and "virtual appliances" for use with Xen VMs. To create an image, enter a name for the image and select a "Class" 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&groupbase&package=glite-BDII"

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

Name

Synchronous

Class

Architecture

Filetype

- Benchmark next generation hardware assisted virtualization extensions
  - Integrate Content-Based Transfer with OS Farm
  - Investigate Content-Based Transfer tool with perfmon, in order to find hotspots and further optimizations
  - Signing images in OS Farm (DMTF standard)
  - OS Farm web-services interface
-

- A very active period:
  - Good collaboration with partners and EGEE
    - Also: Support letter provided to EGEE-3
  - PhD studies by X.Gréhant have gathered a strong momentum
  - Virtualization made more appealing by providing complementary tools
  
- We expect to maintain the vigour in the coming periods.

# Q & A