

Grid IC

Quarterly Review Meeting
29th January 2007

J. M. Dana,
H. Bjerke,
X. Grehant,
S. Jarp




- Scheduling (X. Grehant)
 - Tycoon (J.M. Dana)
 - Virtualization (H. Bjerke)
-

- Scheduling (X. Grehant)
 - Tycoon (J.M. Dana)
 - Virtualization (H. Bjerke)
-

- Synthesis on Grid Scheduling
 - VO* management, resource access
 - EGEE, OSG, NorduGrid, Naregi, etc.
 - Direct scheduling in a VO*
 - glideCAF, Cronus, GlideInWMS
 - AliEn2, DIRAC, Panda
 - DIANE
- * VO: Virtual Organization, federation of users.
- With the help of several grid doers at CERN
 - Submitted to the Journal of Supercomputing

- Resource supply / consumption is heterogeneous
 - ↳ Benefits of careful allocation and migration?
- Design of a resource model
- Development of Levellab, a simulator
- Example simulation:
 - Optimal dynamic vs random allocation
- Submitted to HPDC



- Resource availability is transient
 - ↳ Resilient service deployment
- Design of a P2P resource election mechanism
 - Decides where to (re-)deploy a service
- Development of SmartCitizens, based on SmartFrog




Figure: SmartCitizens Integration inside a node, and between nodes

- Scheduling (X. Grehant)
 - Tycoon (J.M. Dana)
 - Virtualization (H. Bjerke)
-

- A technical report about Tycoon activities at CERN openlab in 2007 and our future plans for 2008 has been written and sent to HP Labs (Palo Alto)
- Several modifications in the Tycoon-gLite implementation
- Collaboration with BalticGrid has been resumed

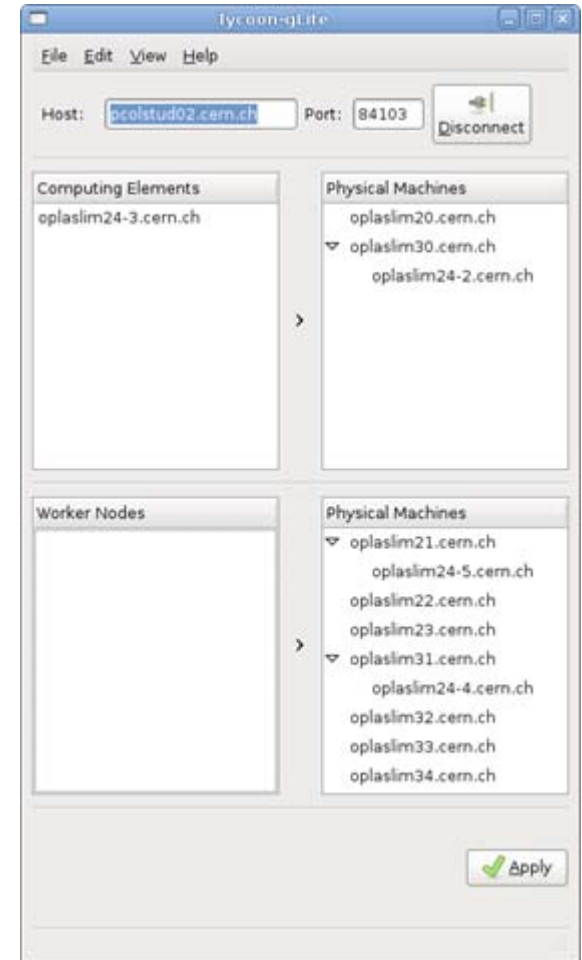
- Worth highlighting:
 - Collaborations (HP Labs, EGEE, BalticGrid)
 - Tycoon-gLite integration
 - Scalability tests
 - Issues concerning security and trust
 - Conferences
 - Future collaboration with Constellation Technologies?

- The implementation has been modified in order to:
 - Deploy different kinds of nodes more easily (i.e. Storage Elements)
 - Allow the modification of the output




```

tycoon host create_account oplasim30.cern.ch 1 "0 dak:1.200.200" file_system=File (/var/lib/tycoon/raid/ra33/Model/Default.ec2)
tycoon_stop [mp]tycoon-gLite_CDH@system-info def root@oplasim30.cern.ch:root
tycoon_stop [mp]tycoon-gLite_CDH@system-info def root@oplasim30.cern.ch:root@tycoonstop
tycoon_stop [mp]tycoon-gLite_CDH@system-info def root@oplasim30.cern.ch:root@tycoonstop
tycoon_ssh root@oplasim30.cern.ch "root@tycoonstop"
ssh root@oplasim24-2.cern.ch "top@tycoonstop"
tycoon host create_account oplasim21.cern.ch 1 "0 dak:1.200.200" file_system=File (/var/lib/tycoon/raid/ra33/Model/Default.ec2)
tycoon_stop [mp]tycoon-gLite_CDH@system-info def root@oplasim21.cern.ch:root
tycoon_stop [mp]tycoon-gLite_CDH@system-info def root@oplasim21.cern.ch:root@tycoonstop
tycoon_ssh root@oplasim21.cern.ch "root@tycoonstop"
ssh root@oplasim24-5.cern.ch "top@tycoonstop"
tycoon host create_account oplasim31.cern.ch 1 "0 dak:1.200.200" file_system=File (/var/lib/tycoon/raid/ra33/Model/Default.ec2)
tycoon_stop [mp]tycoon-gLite_CDH@system-info def root@oplasim31.cern.ch:root
tycoon_stop [mp]tycoon-gLite_CDH@system-info def root@oplasim31.cern.ch:root@tycoonstop
  
```

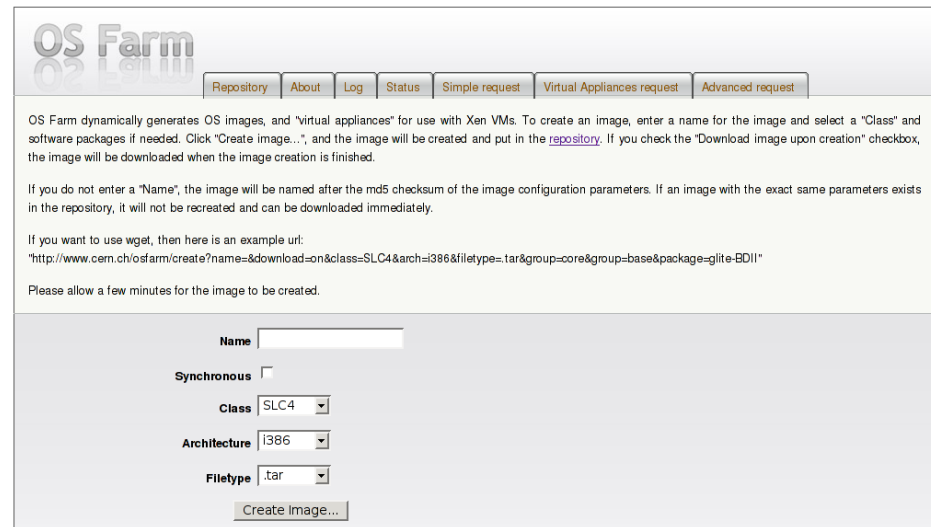


- Scheduling (X. Grehant)
 - Tycoon (J.M. Dana)
 - Virtualization (H. Bjerke)
-

- Benchmarks run on paravirtualized and hardware-assisted virtualization platforms
 - point to strengths and weaknesses in hypervisors



- VM images generated using a layered cache
 - Core layer is instantaneous, using copy-on-write
 - Supports Debian and Red Hat based distributions
- Contextualization - customizes images according to deployment context
- Web service interface w/
example Java client
- XML image descriptions



The screenshot shows the OS Farm web interface. At the top, there is a navigation menu with links: Repository, About, Log, Status, Simple request, Virtual Appliances request, and Advanced request. Below the menu, there is a text block explaining the service: "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." Below this text, there is a note: "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." Then, there is an example URL: "http://www.cern.ch/osfarm/create?name=&download-on&class=SLC4&arch=i386&filetype=tar&group=core&group=base&package=glite-BDII". Below the URL, there is a button labeled "Create Image...".

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&group=base&package=glite-BDII"

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

Name

Synchronous

Class

Architecture

Filetype

- Most VM images are relatively similar
 - Transfer only the delta between images
- Efficiency close to hypothetical max (infinite CPU power)
- Integration with OS Farm

