





# Summary

- 1 Overview of the Main HPC Components**
- 2 The UL HPC platform**
  - Overview
  - Platform Deployment with FAI and Puppet
  - Monitoring
  - Statistics & Milestones
- 3 UL HPC in Practice: Toward an [Efficient] Win-Win Usage**
  - General considerations
  - SSH Access
  - The OAR Batch Scheduler
  - Available Software and Environment Modules
- 4 Last Challenges**
  - Memory bottleneck
  - Effective Storage and Memory Management
  - Reporting problems and Promoting the Platform



# Summary

- 1 Overview of the Main HPC Components**
- 2 The UL HPC platform**
  - Overview
  - Platform Deployment with FAI and Puppet
  - Monitoring
  - Statistics & Milestones
- 3 UL HPC in Practice: Toward an [Efficient] Win-Win Usage**
  - General considerations
  - SSH Access
  - The OAR Batch Scheduler
  - Available Software and Environment Modules
- 4 Last Challenges**
  - Memory bottleneck
  - Effective Storage and Memory Management
  - Reporting problems and Promoting the Platform

# HPC Components: [GP]CPU

## CPU

- Always multi-core
- Ex: Intel Core i7-970 (July 2010)  $R_{peak} \simeq 100$  GFlops (DP)  
↪ 6 cores @ 3.2GHz (32nm, 130W, 1170 millions transistors)

## GPU / GPGPU


- Always multi-core, optimized for vector processing
- Ex: Nvidia Tesla C2050 (July 2010)  $R_{peak} \simeq 515$  GFlops (DP)  
↪ 448 cores @ 1.15GHz

**$\simeq 10$  Gflops for 50 €**



# HPC Components: Local Memory

**Larger, slower and cheaper**



- SSD R/W: 560 MB/s; 85000 IOps

**1500 €/TB**


- HDD (SATA @ 7,2 krpm) R/W: 100 MB/s; 190 IOps

**150 €/TB**

## HPC Components: Interconnect

- **latency**: time to send a minimal (0 byte) message from A to B
- **bandwidth**: max amount of data communicated per unit of time


Technology	Effective Bandwidth		Latency
Gigabit Ethernet	1 Gb/s	125 MB/s	40 $\mu$ s to 300 $\mu$ s
Myrinet (Myri-10G)	9.6 Gb/s	1.2 GB/s	2.3 $\mu$ s
10 Gigabit Ethernet	10 Gb/s	1.25 GB/s	4 $\mu$ s to 5 $\mu$ s
Infiniband QDR	40 Gb/s	5 GB/s	1.29 $\mu$ s to 2.6 $\mu$ s
SGI NUMalink	60 Gb/s	7.5 GB/s	1 $\mu$ s



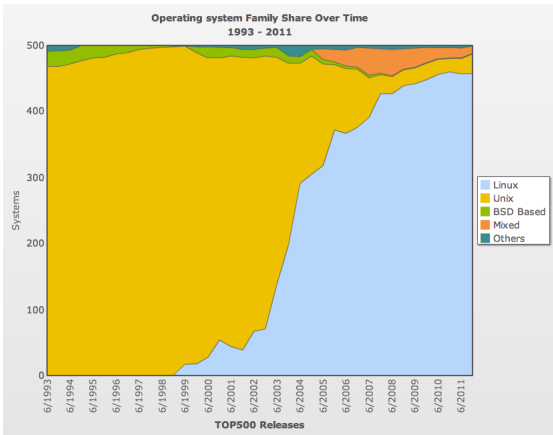
## HPC Components: Interconnect

- **latency**: time to send a minimal (0 byte) message from A to B
- **bandwidth**: max amount of data communicated per unit of time

Technology	Effective Bandwidth		Latency
Gigabit Ethernet	1 Gb/s	125 MB/s	40 $\mu$ s to 300 $\mu$ s
Myrinet (Myri-10G)	9.6 Gb/s	1.2 GB/s	2.3 $\mu$ s
10 Gigabit Ethernet	10 Gb/s	1.25 GB/s	4 $\mu$ s to 5 $\mu$ s
Infiniband QDR	40 Gb/s	5 GB/s	1.29 $\mu$ s to 2.6 $\mu$ s
SGI NUMalink	60 Gb/s	7.5 GB/s	1 $\mu$ s



# HPC Components: Operating System



- Mainly Linux-based OS (91.4%) (Top500, Nov 2011)
- ... or Unix based (6%)
- Reasons:
  - ↳ stability
  - ↳ prone to devels



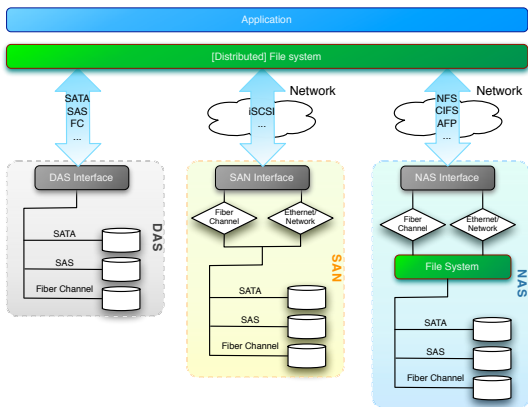


## HPC Components: Software Stack

- **Remote connection to the platform:** SSH
- **User SSO:** NIS or OpenLDAP-based
- **Resource management:** job/batch scheduler
  - ↪ OAR, PBS, Torque, MOAB Cluster Suite
- **(Automatic) Node Deployment:**
  - ↪ FAI (Fully Automatic Installation), Kickstart, Puppet, Chef, Kadeploy etc.
- **Platform Monitoring:** Nagios, Ganglia, Cacti etc.
- (eventually) **Accounting:**
  - ↪ oarnodeaccounting, Gold allocation manager etc.


# HPC Components: Data Management

## Storage architectural classes & I/O layers




# HPC Components: Data Management

## RAID standard levels



# HPC Components: Data Management


## RAID combined levels






# HPC Components: Data Management

## RAID combined levels



# HPC Components: Data Management

## RAID combined levels



- Software vs. **Hardware** RAID management
- RAID Controller card performances differs!
  - ↪ Basic (low cost): 300 MB/s; Advanced (expensive): 1,5 GB/s



# HPC Components: Data Management

## File Systems

- Logical manner to store, organize, manipulate and access data.
- **Disk file systems:** FAT32, NTFS, HFS, ext3, ext4, xfs...
- **Network file systems:** NFS, SMB
- **Distributed parallel file systems:** HPC target
  - ↪ data are striped over multiple servers for high performance.
  - ↪ generally add robust failover and recovery mechanisms
  - ↪ Ex: Lustre, GPFS, FhGFS, GlusterFS...

- HPC storage make use of high density **disk enclosures**
  - ↪ includes [redundant] RAID controllers



# HPC Components: Data Center

## Definition (Data Center)

Facility to house computer systems and associated components

↔ Basic storage component: **rack** (height: 42 RU)

# HPC Components: Data Center

## Definition (Data Center)

Facility to house computer systems and associated components

↪ Basic storage component: **rack** (height: 42 RU)

**Challenges:** Power (UPS, battery), Cooling, Fire protection, Security

- Power/Heat dissipation per rack:
  - ↪ 'HPC' (computing) racks: 30-40 kW
  - ↪ 'Storage' racks: 15 kW
  - ↪ 'Interconnect' racks: 5 kW

## Power Usage Effectiveness

$$PUE = \frac{\text{Total facility power}}{\text{IT equipment power}}$$

# HPC Components: Data Center





## HPC Components: Summary

HPC platforms involves:

- A data center / server room carefully designed
- Computing elements: CPU/GPGPU
- Interconnect elements
- Storage elements: HDD/SDD, disk enclosure,  
↳ disks are virtually aggregated by RAID/LUNs/FS
- A flexible software stack
- **Above all:** expert system administrators...



# Summary

- 1 Overview of the Main HPC Components
- 2 **The UL HPC platform**
  - Overview
  - Platform Deployment with FAI and Puppet
  - Monitoring
  - Statistics & Milestones
- 3 UL HPC in Practice: Toward an [Efficient] Win-Win Usage
  - General considerations
  - SSH Access
  - The OAR Batch Scheduler
  - Available Software and Environment Modules
- 4 **Last Challenges**
  - Memory bottleneck
  - Effective Storage and Memory Management
  - Reporting problems and Promoting the Platform





The screenshot shows the homepage of the HPC @ Uni.lu website. At the top, there is a navigation bar with links for 'Systems', 'For Users', 'Live Status', 'Blog/News', and 'About'. A search bar is also present. The main content area features a 'Welcome to the HPC @ Uni.lu platform!' section with a paragraph about the website's purpose and a quote from the Council on Competitiveness. Below this is a large image of a server room with a caption identifying it as the server room at the LCSB building in Belval. To the right, there are sections for 'Recent Posts' and 'GitHub Repos'. At the bottom, there are sections for 'Featured Systems', 'Platform Status', and 'User Docs'. A 'Tweets' section is also visible on the right side.

**Welcome to the HPC @ Uni.lu platform !**  
This is the official website of HPC @ Uni.lu platform, which assembles information about the computing clusters operated by the University of Luxembourg and the organization running them.  
The country that out-computes will be the one that out-competes.  
— The Council on Competitiveness

**Recent Posts**

- SFS & nocell4
- FOSDEM 2014
- Almas Press Release
- New 10 interconnect and new notes on Gaia
- Quick configuration guide for the Infiniband switch Mellanox Vt606

**GitHub Repos**

- tutorials
- qualif
- dorfies
- launcher-scripts
- reports
- vindis-bootstrap
- ganglia\_infiniband\_module

**Featured Systems**  
We currently operate a total of 371 computing nodes (3908 cores, 43.751 TFlops) and a shared storage capacity of 934.4 TB (+ 360 TB for backup).

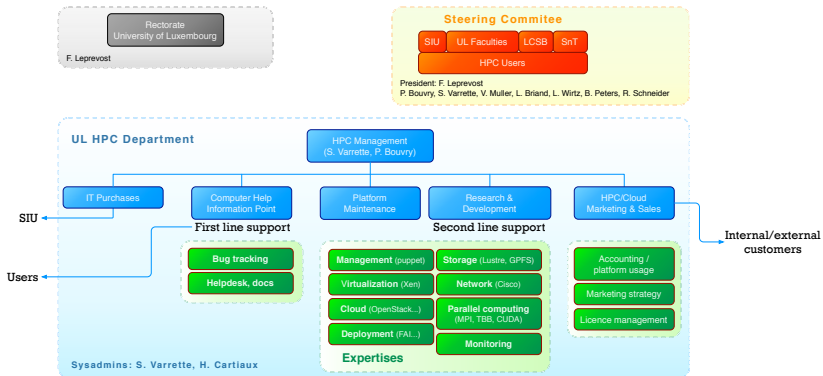
**Platform Status**  
Several tools report in live the current status of our systems.  
[Check them out!](#)

**User Docs**  
We took the time to make the HPC documentation as complete as possible. Please make sure you read it carefully.

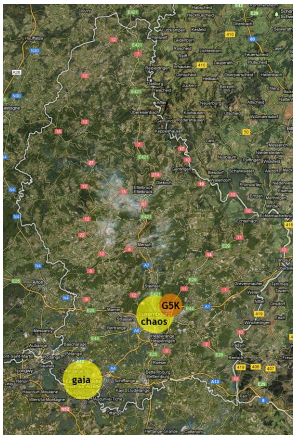
**Tweets**

HPC UGent @HPCUGent 1 Feb  
Xavier Bessaon presents his talk on Automated Testing of Installed Software, with an MPI focus #HPC #FOSDEM pic.twitter.com/odQ84fAeM  
13 Retweeted by ULHPC

# Organizational Chart




## UL HPC platforms at a glance (2014)




- 2 geographical sites, 3 server rooms
- 4 clusters: chaos+gaia, granduc, nyx.
  - ↳ 368 nodes, 3880 cores, 43.204 TFlops
  - ↳ incl. 18 dual [GP]GPU nodes
  - ↳ 1996.4 TB (raw) shared storage
    - incl. 1.3 PB on 7 backup enclosures
- 2 sysadmins `hpc-sysadmins@uni.lu`
- 5,749,432€ (Cumul. HW Investment) since 2007
  - ↳ Hardware acquisition only
  - ↳ 3,487,029€ (excluding server rooms)
- Open-Source software stack
  - ↳ SSH, LDAP, OAR, Puppet, Modules...

## HPC server rooms

- **2009** CS.43 (Kirchberg campus) 14 racks, 100 m<sup>2</sup>, ≈ 800,000€



- **2011** LCSB 6<sup>th</sup> floor (Belval) 14 racks, 112 m<sup>2</sup>, ≈ 1,100,000€





# UL HPC Computing Nodes

	Date	Vendor	Proc. Description	#N	#C	R <sub>peak</sub>
chaos	2010	HP	Intel Xeon L5640@2.26GHz 2 × 6C,24GB	32	384	3.472 TFlops
	2011	Dell	Intel Xeon L5640@2.26GHz 2 × 6C,24GB	16	192	1.736 TFlops
	2012	Dell	Intel Xeon X7560@2,26GHz 4 × 6C, 1TB	1	32	0.289 TFlops
	2012	Dell	Intel Xeon E5-2660@2.2GHz 2 × 8C,32GB	16	256	4.506 TFlops
	2012	HP	Intel Xeon E5-2660@2.2GHz 2 × 8C,32GB	16	256	4.506 TFlops
<b>chaos TOTAL:</b>				<b>81</b>	<b>1120</b>	<b>17.026 TFlops</b>

gaia	2011	Bull	Intel Xeon L5640@2.26GHz 2 × 6C,48GB	72	864	7.811 TFlops
	2012	Dell	Intel Xeon E5-4640@2.4GHz 4 × 8C, 1TB	1	32	0.307 TFlops
	2012	Bull	Intel Xeon E7-4850@2GHz 16 × 10C,1TB	1	160	1.280 TFlops
	2013	Dell	Intel Xeon E5-2660@2.2GHz 2 × 8C,64GB	5	80	1.408 TFlops
	2013	Bull	Intel Xeon X5670@2.93GHz 2 × 6C,48GB	40	480	5.626 TFlops
	2013	Bull	Intel Xeon X5675@3.07GHz 2 × 6C,48GB	32	384	4.746 TFlops
<b>gaia TOTAL:</b>				<b>151</b>	<b>2000</b>	<b>21.178 TFlops</b>

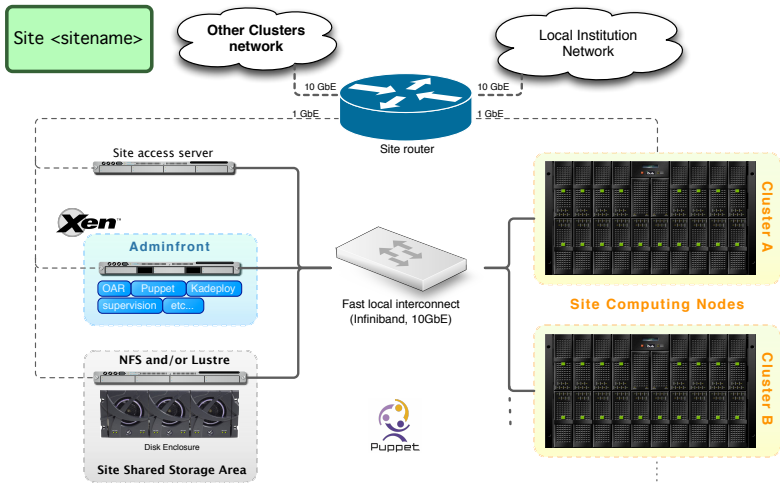
g5k	2008	Dell	Intel Xeon L5335@2GHz 2 × 4C,16GB	22	176	1.408 TFlops
	2012	Dell	Intel Xeon E5-2630L@2GHz 2 × 6C,24GB	16	192	1.536 TFlops
<b>granduc/petitprince TOTAL:</b>				<b>38</b>	<b>368</b>	<b>4.48 TFlops</b>

Testing cluster:

nyx	2012	Dell	Intel Xeon E5-2420@1.9GHz 1 × 6C,32GB	2	12	0.091 TFlops
viridis	2013	Viridis	ARM A9 Cortex@1.1GHz 1 × 4C,4GB	96	384	0.422 TFlops
<b>nyx/viridis TOTAL:</b>				<b>98</b>	<b>392</b>	<b>0.52 TFlops</b>

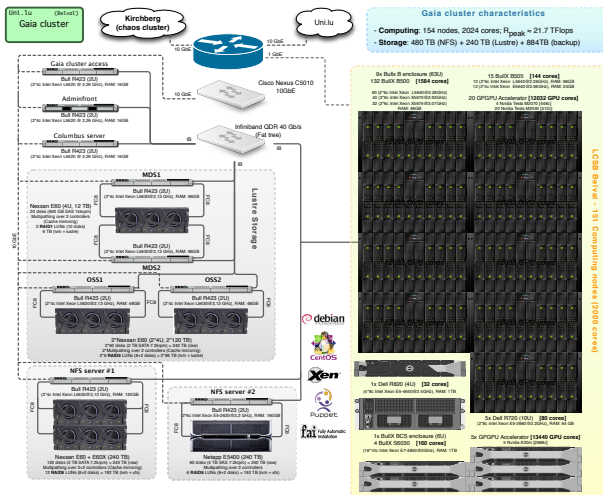
**TOTAL: 368 nodes, 3880 cores, 43.204 TFlops**

# UL HPC: General cluster organization






## Ex: The gaia cluster





## Ex: Some racks of the gaia cluster





# UL HPC Software Stack Characteristics

- **Operating System:** Linux Debian (CentOS on storage servers)
- **Remote connection to the platform:** SSH
- **User SSO:** OpenLDAP-based
- **Resource management:** job/batch scheduler: OAR
- **(Automatic) Computing Node Deployment:**
  - ↪ FAI (Fully Automatic Installation) (chaos,gaia,nyx only)
  - ↪ Puppet
  - ↪ Kadeploy (granduc,petitprince/Grid5000 only)
- **Platform Monitoring:** OAR Monika, OAR Drawgantt, Ganglia, Nagios, Puppet Dashboard etc.
- **Commercial Softwares:**
  - ↪ Intel Cluster Studio XE, TotalView, Allinea DDT, Stata etc.



# HPC in the Grande region and Around


Country	Name/Institute	#Cores	TFlops	TB	FTEs
			$R_{peak}$	Storage	Manpower
Luxembourg	UL CRP GL	3880	43.204	1996.4	2
		800	6.21	144	1.5
France	TGCC Curie, CEA	77184	1667.2	5000	n/a
	LORIA, Nancy	3724	29.79	82	5.05
	ROMEO, UCR, Reims	564	4.128	15	2
Germany	Juqueen, Juelich	393216	5033.2	448	n/a
	MPI, RZG	2556	14.1	n/a	5
	URZ, (bwGrid), Heidelberg	1140	10.125	32	9
Belgium	UGent, VCS	4320	54.541	82	n/a
	CECI, UMONS/UCL	2576	25.108	156	> 4
UK	Darwin, Cambridge Univ	9728	202.3	20	n/a
	Legion, UCLondon	5632	45.056	192	6
Spain	MareNostrum, BCS	33664	700.2	1900	14

# The case of Grid'5000



- Large scale nation wide infrastructure

↪ for large scale parallel and distributed computing research.



- 10 sites in France
- **Abroad:** Luxembourg, Porto Allegre
- Total: **7782** cores over **26** clusters
- 1-10GbE / Myrinet / Infiniband
  - ↪ **10Gb/s dedicated** between all sites
- Unique software stack
  - ↪ **kadeploy, kavlan, storage5k**



## Grid'5000 Usage

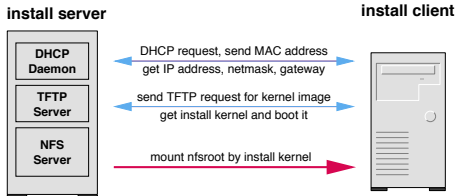
- Out of scope for this talk
- **Grid'5000 school 2014**
  - ↪ Organized in Lyon, from June 16th to June 19th 2014
  - ↪ More info: [G5K School website](#)
- General information: <https://hpc.uni.lu/g5k>
- Grid'5000 website and documentation: <https://www.grid5000.fr>

# Computing nodes Management

## Node deployment by FAI

<http://fai-project.org/>

- Boot via network card (PXE)
  - ↪ ensure a running diskless Linux OS

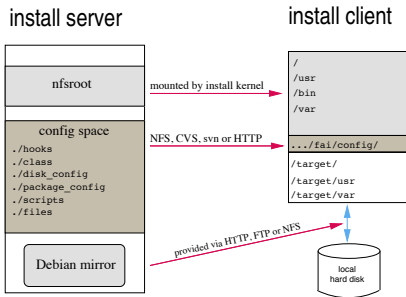


# Computing nodes Management

## Node deployment by FAI

<http://fai-project.org/>

- Boot via network card (PXE)
  - ↳ ensure a running diskless Linux OS
- Get configuration data (NFS)





# Computing nodes Management

## Node deployment by FAI

<http://fai-project.org/>

- Boot via network card (PXE)
  - ↳ ensure a running diskless Linux OS
- Get configuration data (NFS)
- Run the installation
  - ↳ partition local hard disks and create filesystems
  - ↳ install software using apt-get command
  - ↳ configure OS and additional software
  - ↳ save log files to install server, then reboot new system







# Computing nodes Management

## Node deployment by FAI

<http://fai-project.org/>

- Boot via network card (PXE)
  - ↪ ensure a running diskless Linux OS
- Get configuration data (NFS)
- Run the installation
  - ↪ partition local hard disks and create filesystems
  - ↪ install software using apt-get command
  - ↪ configure OS and additional software
  - ↪ save log files to install server, then reboot new system



**Average reinstallation time:  $\simeq$  500s**



# Platform Management: Puppet

## Server/Service configuration by Puppet

<http://puppetlabs.com>



- *Configuration management made easy*
- Automates service deployment / sysadmin tasks
- 3 components:
  - ↔ declarative language (manifests)
  - ↔ client/server model
  - ↔ Git-based workflow + Rakefile + Capistrano
- Hierarchical PKI



# Platform Management: Puppet

## Server/Service configuration by Puppet

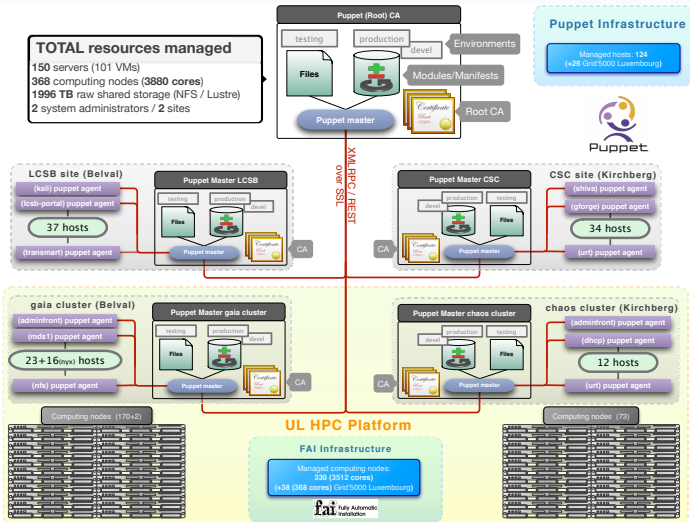
<http://puppetlabs.com>




- *Configuration management made easy*
- Automates service deployment / sysadmin tasks
- 3 components:
  - ↔ declarative language (manifests)
  - ↔ client/server model
  - ↔ Git-based workflow + Rakefile + Capistrano
- Hierarchical PKI

**Average server installation/configuration time:  $\simeq$  3-6 min**

# Platform Management: Puppet



# Puppet Performances



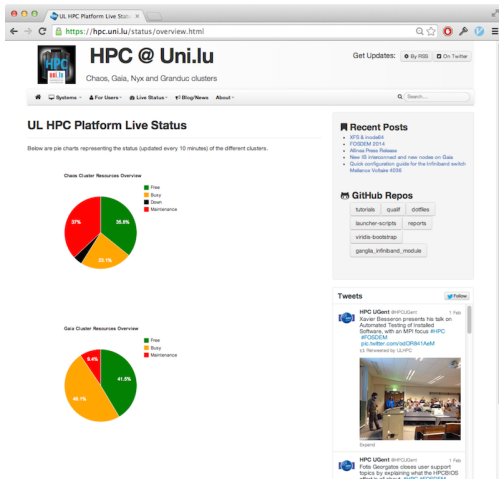
- between 161s and 364s to completely **bootstrap** a virgin node  
 ↳ between 20s and 31s to later check/correct the config

Now proposed as an IT service to external consumers

# Platform Monitoring

## General Live Status

<http://hpc.uni.lu/status/overview.html>

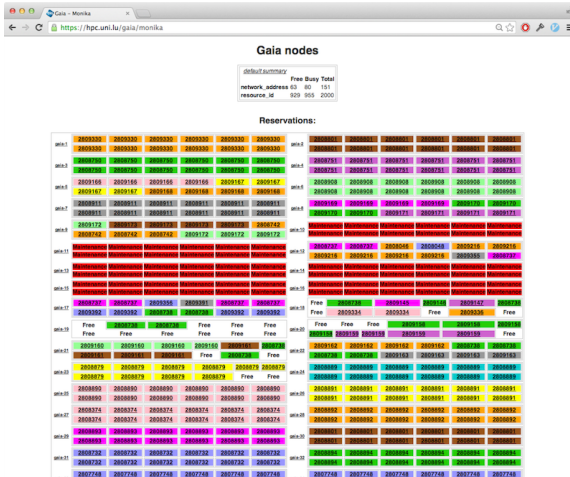




# Platform Monitoring

## Monika

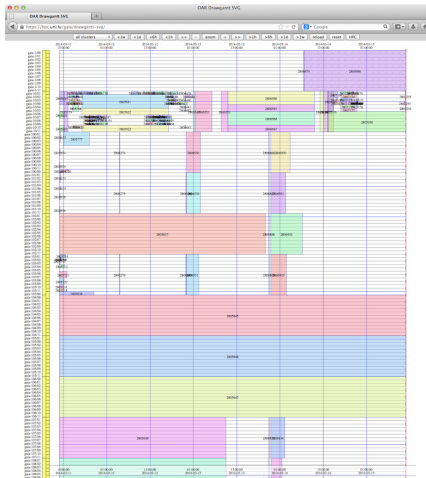
<http://hpc.uni.lu/{chaos,gaia,g5k}/monika>



# Platform Monitoring

## Drawgantt

<http://hpc.uni.lu/{chaos,gaia,g5k}/drawgantt>





# Platform Monitoring

## Ganglia

<http://hpc.uni.lu/{chaos,gaia,g5k}/ganglia>





# Platform Monitoring

## CDash

<http://cdash.uni.lu/>

The screenshot shows the CDash web interface for 'UL-HPC-Testing'. The page includes a navigation bar with 'Dashboard', 'Calendar', 'Previous', 'Current', and 'Project' tabs. Below the navigation bar, there is a table of build results. The table has columns for 'Site', 'Build Name', 'Update', 'Configure', 'Build', 'Test', and 'Build Time'. The 'Test' column is further divided into 'Net Rtn', 'Fail', and 'Pass' sub-columns. The table lists various build configurations for different clusters (Chaos and Gaia) and MPI modules, with status indicators (green for success, red for failure) in the 'Test' sub-columns.

Site	Build Name	Update		Configure		Build		Test			Build Time
		Files	Error	Warn	Error	Warn	Net Rtn	Fail	Pass		
Chaos cluster	MPI Module MPICH2_1.1-GCC-4.8.1	0	0	0	0	0	0	9	4		9 hours ago
Gaia cluster	MPI Module MPICH2_1.1-GCC-4.8.1	0	0	0	0	0	0	9	4		9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.3-iccfort-2011.13.367	0	0	0	0	0	0	9	4		9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.3-iccfort-2011.13.367	0	0	0	0	0	0	9	4		9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.4-ClangGCC-1.1.3	0	0	0	0	0	0	9	4		9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.4-ClangGCC-1.1.3	0	0	0	0	0	0	9	4		9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.4-GCC-4.8.4	0	0	0	0	0	0	9	4		9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.4-GCC-4.8.4	0	0	0	0	0	0	9	4		9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.4-GCC-4.7.2	0	0	0	0	0	0	9	4		9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.4-GCC-4.7.2	0	0	0	0	0	0	9	4		9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.5-GCC-4.7.2	0	0	0	0	0	0	9	4		9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.5-GCC-4.7.2	0	0	0	0	0	0	9	4		9 hours ago
Chaos cluster	MPI Module OpenMPI_1.7.3-gccuda-2.6.10	0	0	0	0	0	0	9	4		9 hours ago
Gaia cluster	MPI Module OpenMPI_1.7.3-gccuda-2.6.10	0	0	0	0	0	0	9	4		9 hours ago
Chaos cluster	MPI Module impi_3.2.2.006	0	0	0	0	0	5	5	3		9 hours ago
Gaia cluster	MPI Module impi_3.2.2.006	0	0	0	0	0	5	5	3		9 hours ago
Chaos cluster	MPI Module impi_4.0.0.028	0	0	0	0	0	5	5	3		9 hours ago
Gaia cluster	MPI Module impi_4.0.0.028	0	0	0	0	0	5	5	3		9 hours ago
Chaos cluster	MPI Module impi_4.0.0.028	0	0	0	0	0	5	5	3		9 hours ago




## Key numbers

- **368 nodes, 3880 cores, 43.204 TFlops**
- **1996.4 TB (raw) shared storage**
  - ↪ incl. 1.3 PB on 7 backup enclosures
- **5,749,432€ (Cumul. HW Investment)** since 2007
- 235 registered users
- Total **Asked** resources: **40537 years** since 2008
- Total **Used** resources: **2482 years** since 2008


# CPU-year usage since 2008

- **CPU-hour:** *work* done by a CPU in one hour of wall clock time

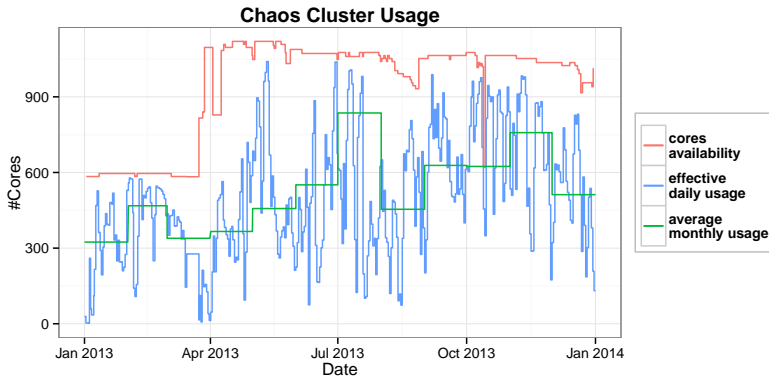


## CPU-year usage since 2008


- **CPU-hour:** *work* done by a CPU in one hour of wall clock time



# Chaos Cluster Usage in 2013



# Gaia Cluster Usage in 2013





## 2013 Milestones

- Creation of the HPC Steering Committee
- **2013 Budget: 715 963,04 €**
  - ↪ HPC: 54,5%      LCSB: 28,73%      CSC:5,62%      Other:11,15%
- Platform [continuous] evolution
  - ↪ Ex: +181 nodes (+1432 cores), +12,21 TFlops
  - ↪ Consolidation of shared storage stability (NFS, Lustre)
  - ↪ New website, Tutorials on Github etc.
  - ↪ **Job submission web-portal** (testing)
- Events / System Downtime

Cluster	Yearly Downtime	Yearly Uptime
chaos	70h	<b>99.201%</b>
gaia	401h	<b>95.422%</b>
g5k	320h	<b>96.347%</b>





## 2014 Milestones

- 30% budget drop
  - ↪ new budget repartition in progress
  - ↪ **unavoidable**: new cost-model (computing & storage)
- 2014 budget fully dedicated to storage component
  - ↪ BigNAS RFP (1 PB storage) – shared SIU / HPC / LCSB
  - ↪ Lustre consolidation
- PRACE observer status

### Incoming milestones (~~2015~~ 2016)

- Belval *Centre De Calcul* (CDC)
  - ↪ 5 new server rooms (3 storage, 2 HPC)
  - ↪ Pending discussions with Fond Belval to re-justify everything



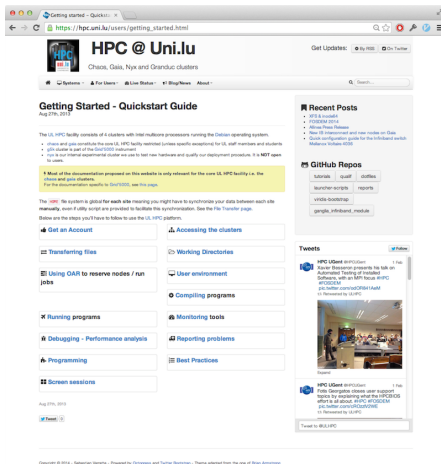
# Summary

- 1 Overview of the Main HPC Components
- 2 The UL HPC platform
  - Overview
  - Platform Deployment with FAI and Puppet
  - Monitoring
  - Statistics & Milestones
- 3 **UL HPC in Practice: Toward an [Efficient] Win-Win Usage**
  - General considerations
  - SSH Access
  - The OAR Batch Scheduler
  - Available Software and Environment Modules
- 4 Last Challenges
  - Memory bottleneck
  - Effective Storage and Memory Management
  - Reporting problems and Promoting the Platform

## General Consideration

- The platform is **\*restricted\*** to UL members and is **\*shared\***
- Everyone should be civic-minded.
  - ↳ Just **avoid** the following behavior: (or you'll be banned)  
*"My work is the most important: I use all the resources for 1 month"*
  - ↳ regularly clean your homedir from useless files
- Plan large scale experiments during night-time or week-ends
  - ↳ try not to use more than 40 computing cores during working day
  - ↳ ... or use the 'besteffort' queue

... aka the rtfm paradigm



The screenshot shows a web browser displaying the 'Getting Started - Quickstart Guide' page for HPC @ Uni.lu. The page is dated August 27th, 2013, and provides a comprehensive list of links for users to explore, including sections for 'Getting an Account', 'Accessing the clusters', 'Transferring files', 'Working Directories', 'Using OAR to reserve nodes / run jobs', 'User environment', 'Compiling programs', 'Running programs', 'Monitoring tools', 'Debugging - Performance analysis', 'Reporting problems', 'Programming', and 'Best Practices'. The page also features a 'Recent Posts' section with links to 'HPC @ Uni.lu' and 'HPC @ Uni.lu' and a 'Tweets' section with a tweet from @ULHPC. The page is powered by Octopress and Twitter Bootstrap.



## User Account

**Get an account:** [https://hpc.uni.lu/get\\_an\\_account](https://hpc.uni.lu/get_an_account)

With your account, you'll get:

- Access to the UL HPC wiki <http://hpc.uni.lu/>
- Access to the UL HPC bug tracker <http://hpc-tracker.uni.lu/>
- Subscribed to the mailing lists `hpc-{users,platform}@uni.lu`
  - ↪ raise questions and concerns. Help us to make it a community!
  - ↪ notification of platform maintenance on [hpc-platform@uni.lu](mailto:hpc-platform@uni.lu)
- A nice way to reach workstation in the internal UL network (ProxyCommand)



## Typical Workflow on UL HPC resources


- 1 Connect to the frontend of a site/cluster `ssh`
- 2 (eventually) synchronize you code `scp/rsync/svn/git`
- 3 (eventually) Reserve a few interactive resources `oarsub -I`
  - ↪ (eventually) Configure the resources `kadeploy`
  - ↪ (eventually) Prepare your experiments `gcc/icc/mpicc/javac/...`
  - ↪ Test your experiment on small size problem `mpirun/java/bash....`
  - ↪ Free the resources
- 4 Reserve some resources `oarsub`
- 5 Run your experiment via a launcher script `bash/python/perl/ruby...`
- 6 Grab the results `scp/rsync`
- 7 Free the resources

## UL HPC access

- **\*Restricted\*** to **SSH** connection **with public key authentication**

↪ on a non-standard port (8022)

*limits kiddie script scans/dictionary's attacks*





## UL HPC SSH access

~/.ssh/config

```
Host chaos-cluster
  Hostname      access-chaos.uni.lu
Host gaia-cluster
  Hostname      access-gaia.uni.lu
Host *-cluster
  User          login
  Port          8022
  ForwardAgent no

Host myworkstation
  User          localadmin
  Hostname      myworkstation.uni.lux
Host *.ext_ul
  ProxyCommand ssh -q gaia-cluster "nc -q 0 %h %p"
```

```
$> ssh {chaos,gaia}-cluster
$> ssh myworkstation
```

When @ Home:

```
$> ssh myworkstation.ext_ul
```

### More on this in PS1

- Getting Started

### ● Transferring data...

```
$> rsync -avzu /devel/myproject chaos-cluster:
(gaia)$> gaia_sync_home *
(chaos)$> chaos_sync_home devel/
```





## UL HPC SSH access (Windows)

- Download all the Putty tools
  - ↪ Extract them in an easy-to-find place, such as C:\Putty
- Run Pageant and load your SSH private key
- Run Putty (connection type: SSH)
  - ↪ Host Name: `access-{chaos,gaia}.uni.lu`
  - ↪ Port: 8022
  - ↪ Saved session: {Chaos,Gaia}
  - ↪ in Category:Connection:Data :
    - Auto-login username: your login
- Transferring data...
  - ↪ Download `cwRsync`

## UL HPC SSH access (Windows)

- Access to an internal workstation (ProxyCommand via Gaia)

- ↪ Host Name: Workstation IP or hostname

- ↪ Port: 22

- ↪ Saved session: MyWorkstation

- ↪ in Category:Connection:Proxy :

- Proxy type: Local
- Proxy hostname: access-gaia.uni.lu
- Port: 8022
- Username: your login
- Telnet command:

```
C:\Putty\plink -P %proxyport %user@%proxyhost nc -q 0 %host %port
```

- ↪ in Category:Connection:Data :

- Auto-login username: your login on your workstation

# UL HPC resource manager: OAR

## The OAR Batch Scheduler

<http://oar.imag.fr>

- Versatile resource and task manager

- ↳ schedule **jobs** for users on the cluster **resource**
- ↳ OAR resource = a node or part of it (CPU/core)
- ↳ OAR job = execution time (**walltime**) on a set of resources



# UL HPC resource manager: OAR

## The OAR Batch Scheduler

<http://oar.imag.fr>

- Versatile resource and task manager
  - ↳ schedule **jobs** for users on the cluster **resource**
  - ↳ OAR resource = a node or part of it (CPU/core)
  - ↳ OAR job = execution time (**walltime**) on a set of resources



OAR main features includes:

- **interactive vs. passive (aka. batch) jobs**
- **best effort jobs**: use more resource, accept their release any time
- **deploy jobs (Grid5000 only)**: deploy a customized OS environment
  - ↳ ... and have full (root) access to the resources
- **powerful resource filtering/matching**

## Main OAR commands

`oarsub` submit/reserve a job (by default: **1 core for 2 hours**)

`oardel` delete a submitted job

`oarnodes` shows the resources states

`oarstat` shows information about running or planned jobs

	Submission
<b>interactive</b>	<code>oarsub [options] -I</code>
<b>passive</b>	<code>oarsub [options] <b>scriptName</b></code>

- Each created job receive an identifier JobID
  - ↳ Default passive job log files: `OAR.JobID.std{out,err}`
- You can make a reservation with `-r "YYYY-MM-DD HH:MM:SS"`

## Main OAR commands

`oarsub` submit/reserve a job (by default: **1 core for 2 hours**)

`oardel` delete a submitted job

`oarnodes` shows the resources states

`oarstat` shows information about running or planned jobs

	Submission
<b>interactive</b>	<code>oarsub [options] -I</code>
<b>passive</b>	<code>oarsub [options] <b>scriptName</b></code>

- Each created job receive an identifier JobID
  - ↳ Default passive job log files: `OAR.JobID.std{out,err}`
- You can make a reservation with `-r "YYYY-MM-DD HH:MM:SS"`

Direct access to nodes by `ssh` is forbidden: use `oarsh` instead

## OAR job environment variables

Once a job is created, some environments variables are defined:

Variable	Description
<code>\$OAR_NODEFILE</code>	Filename which lists all reserved nodes for this job
<code>\$OAR_JOB_ID</code>	OAR job identifier
<code>\$OAR_RESOURCE_PROPERTIES_FILE</code>	Filename which lists all resources and their properties
<code>\$OAR_JOB_NAME</code>	Name of the job given by the "-n" option of oarsub
<code>\$OAR_PROJECT_NAME</code>	Job project name

Useful for MPI jobs for instance:

```
$> mpirun -machinefile $OAR_NODEFILE /path/to/myprog
```

... Or to collect how many cores are reserved per node:

```
$> cat $OAR_NODEFILE | uniq -c
```



## OAR job types

Job Type	Max Walltime (hour)	Max #active_jobs	Max #active_jobs_per_user
interactive	12:00:00	10000	5
default	120:00:00	30000	10
besteffort	9000:00:00	10000	1000

```
cf /etc/oar/admission_rules/*.conf
```

- **interactive**: useful to test / prepare an experiment
  - ↔ you get a shell on the first reserved resource
- **best-effort vs. default**: nearly unlimited constraints **YET**
  - ↔ a `besteffort` job can be killed as soon as a default job as no other place to go
  - ↔ enforce checkpointing (and/or idempotent) strategy



## Characterizing OAR resources

### Specifying wanted resources in a hierarchical manner

- Use the `-l` option of `oarsub`. Main constraints:

<code>enclosure=N</code>	number of enclosure
<code>nodes=N</code>	number of nodes
<code>core=N</code>	number of cores
<code>walltime=hh:mm:ss</code>	job's max duration

### Specifying OAR resource properties

- Use the `-p` option of `oarsub`: Syntax: `-p "property='value'"`

<code>gpu='{YES,NO}'</code>	has (or not) a GPU card
<code>host='fqdn'</code>	full hostname of the resource
<code>network_address='hostname'</code>	Short hostname of the resource
<b>(Chaos only)</b> <code>nodeclass='{k,b,h,d,r}'</code>	Class of node



## OAR (interactive) job examples

- 2 cores on 3 nodes (same enclosure) for 3h15:

Total: 6 cores

```
(frontend)$> oarsub -I -l /enclosure=1/nodes=3/core=2,walltime=3:15
```



## OAR (interactive) job examples

- 2 cores on 3 nodes (same enclosure) for 3h15: Total: 6 cores  

```
(frontend)$> oarsub -I -l /enclosure=1/nodes=3/core=2,walltime=3:15
```
- 4 cores on a GPU node for 8 hours Total: 4 cores  

```
(frontend)$> oarsub -I -l /core=4,walltime=8 -p "gpu='YES'"
```



## OAR (interactive) job examples

- 2 cores on 3 nodes (same enclosure) for 3h15: Total: 6 cores

```
(frontend)$> oarsub -I -l /enclosure=1/nodes=3/core=2,walltime=3:15
```

- 4 cores on a GPU node for 8 hours Total: 4 cores

```
(frontend)$> oarsub -I -l /core=4,walltime=8 -p "gpu='YES'"
```

- 2 nodes among the h-cluster1-\* nodes (Chaos only) Total: 24 cores

```
(frontend)$> oarsub -I -l nodes=2 -p "nodeclass='h'"
```



## OAR (interactive) job examples

- 2 cores on 3 nodes (same enclosure) for 3h15: Total: 6 cores

```
(frontend)$> oarsub -I -l /enclosure=1/nodes=3/core=2,walltime=3:15
```

- 4 cores on a GPU node for 8 hours Total: 4 cores

```
(frontend)$> oarsub -I -l /core=4,walltime=8 -p "gpu='YES'"
```

- 2 nodes among the h-cluster1-\* nodes (Chaos only) Total: 24 cores

```
(frontend)$> oarsub -I -l nodes=2 -p "nodeclass='h'"
```

- 4 cores on 2 GPU nodes + 20 cores on other nodes Total: 28 cores

```
$> oarsub -I -l "{gpu='YES'}/nodes=2/core=4+{gpu='NO'}/core=20"
```

## OAR (interactive) job examples

- 2 cores on 3 nodes (same enclosure) for 3h15: Total: 6 cores

```
(frontend)$> oarsub -I -l /enclosure=1/nodes=3/core=2,walltime=3:15
```
- 4 cores on a GPU node for 8 hours Total: 4 cores

```
(frontend)$> oarsub -I -l /core=4,walltime=8 -p "gpu='YES'"
```
- 2 nodes among the h-cluster1-\* nodes (Chaos only) Total: 24 cores

```
(frontend)$> oarsub -I -l nodes=2 -p "nodeclass='h'"
```
- 4 cores on 2 GPU nodes + 20 cores on other nodes Total: 28 cores

```
$> oarsub -I -l "{gpu='YES'}/nodes=2/core=4+{gpu='NO'}/core=20"
```
- A full big SMP node Total: 160 cores on gaia-74

```
$> oarsub -t bigsmp -I 1 node=1
```



## Some other useful features of OAR

### Connect to a running job

```
(frontend)$> oarsub -C JobID
```

### Cancel a job

```
(frontend)$> oardel JobID
```

### Status of a jobs

```
(frontend)$> oarstat -state -j JobID
```

### View the job

```
(frontend)$> oarstat
```

```
(frontend)$> oarstat -f -j JobID
```

### Get info on the nodes

```
(frontend)$> oarnodes
```

```
(frontend)$> oarnodes -l
```

```
(frontend)$> oarnodes -s
```

### Run a best-effort job

```
(frontend)$> oarsub -t besteffort ...
```





## Designing efficient OAR job launchers

### Resources/Example

<https://github.com/ULHPC/launcher-scripts>

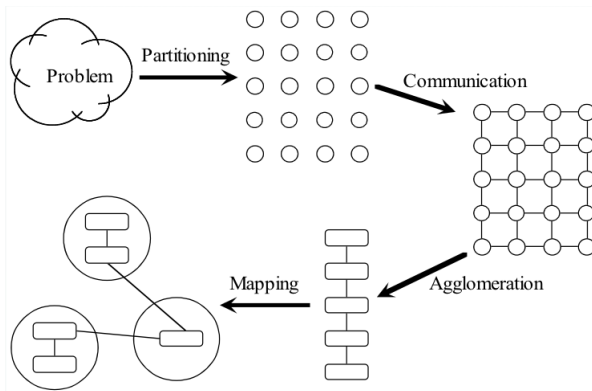


- UL HPC grant access to **parallel computing** resources
  - ↪ ideally: OpenMP/MPI/CUDA/OpenCL jobs
  - ↪ if serial jobs/tasks: run them efficiently
- Avoid to submit purely serial jobs to the OAR queue a
  - ↪ waste the computational power (11 out of 12 cores on gaia).
  - ↪ use whole nodes by running at least 12 serial runs at once
- **Key:** understand difference between **Task** and **OAR job**



# Designing efficient OAR job launchers

- Methodical Design of Parallel Programs



[Foster96] I. Foster, *Designing and Building Parallel Programs*. Addison Wesley, 1996.  
Available at: <http://www.mcs.anl.gov/dbpp>



## Serial tasks: BAD and NAIVE approach

```
#OAR -l nodes=1
#OAR -n BADSerial
#OAR -O BADSerial-%jobid%.log
#OAR -E BADSerial-%jobid%.log
if [ -f /etc/profile ]; then
    . /etc/profile
fi
# Now you can use: 'module load toto' or 'cd $WORK'
[...]
```



## Serial tasks: BAD and NAIVE approach

```
#OAR -l nodes=1
#OAR -n BADSerial
#OAR -O BADSerial-%jobid%.log
#OAR -E BADSerial-%jobid%.log
if [ -f /etc/profile ]; then
    . /etc/profile
fi
# Now you can use: 'module load toto' or 'cd $WORK'
[...]
```

```
# Example 1: run in sequence $TASK 1...$TASK $NB_TASKS
for i in `seq 1 $NB_TASKS`; do
    $TASK $i
done
# Example 2: For each line of $ARG_TASK_FILE, run in sequence
# $TASK <line1>... $TASK <lastline>
while read line; do
    $TASK $line
done < $ARG_TASK_FILE
```



## Serial tasks: A better approach

(fork & wait)


```
# Example 1: run in sequence $TASK 1...$TASK $NB_TASKS
for i in `seq 1 $NB_TASKS`; do
    $TASK $i &
done
wait
```

```
# Example 2: For each line of $ARG_TASK_FILE, run in sequence
# $TASK <line1>... $TASK <lastline>
while read line; do
    $TASK $line &
done < $ARG_TASK_FILE
fi
wait
```


## Serial tasks: A better approach

(fork & wait)


Different runs may not take the same time: **load imbalance**.



## Serial tasks with GNU Parallel



17 hours  
42% utilization



10 hours  
72% utilization



## Serial tasks with GNU Parallel

```
### Example 1: run in sequence $TASK 1...$TASK $NB_TASKS
# On a single node
seq $NB_TASKS | parallel -u -j 12 $TASK {}

# on many nodes
seq $NB_TASKS | parallel -tag -u -j 4 \\  
    -sshloginfile ${GP_SSHLOGINFILE}.task $TASK {}

### Example 2: For each line of $ARG_TASK_FILE, run in parallel
# $TASK <line1>... $TASK <lastline>
# On a single node
cat $ARG_TASK_FILE | parallel -u -j 12 -colsep ' ' $TASK {}

# on many nodes
cat $ARG_TASK_FILE | parallel -tag -u -j 4 \\  
    -sshloginfile ${GP_SSHLOGINFILE}.task -colsep ' ' $TASK {}
```



# Designing efficient OAR job launcher

- More information: **PS2**
  - ↪ HPC workflow with sequential jobs (C,python,java)



## Software Management

- Most default software comes from the OS
- Specific / optimized software available under Environment Modules

↪ <https://hpc.uni.lu/users/docs/modules.html>

↪ **To be used on cluster nodes.** Main commands:

```
(node)$> module help
(node)$> module avail |& less
(node)$> module avail |& grep -i '^gcc'
(node)$> module list
(node)$> module load modulename
```

- Compilation suites available: Intel CC, GCC

### Available software

<https://hpc.uni.lu/users/software/>

- **377 software packages**, in multiple versions



## MPI Suites available via module

### 1 OpenMPI

<http://www.open-mpi.org/>

```
(node)$> module load OpenMPI  
(node)$> make  
(node)$> mpirun -machinefile $OAR_NODEFILE /path/to/mpi_prog
```

### 2 MVAPICH2

<http://mvapich.cse.ohio-state.edu/overview/mvapich2>

```
(node)$> module purge  
(node)$> module load MVAPICH2  
(node)$> make clean && make  
(node)$> mpirun -machinefile $OAR_NODEFILE /path/to/mpi_prog
```

### 3 Intel Cluster Toolkit Compiler Edition (ictce for short):

```
(node)$> module purge  
(node)$> module load ictce  
(node)$> make clean && make  
(node)$> mpirun -machinefile $OAR_NODEFILE /path/to/mpi_prog
```



## Specific Software Guidelines

- **MPI:** see **PS3**
  - ↪ HPC workflow with MPI jobs. Ex: OSU MB/HPL
- **Parallel Debuggers:** see **PS4** and **PS8**
  - ↪ TotalView, Alinea
- **MATLAB:** see **PS5**
  - ↪ Using Matlab on the UL HPC Platform
- **R:** see **PS6**
  - ↪ Using R on the UL HPC Platform
- **ABYSS, Gromacs, Bowtie2 / TopHat, mpiBLAST:** see **PS7**
  - ↪ Bio-informatic softwares on the UL HPC Platform

## Software Management

### My favorite software is not installed on the cluster!

- Check if it does not exists via module
- **If not:** compile it in your home/work directory
  - ↳ using GNU stow <http://www.gnu.org/software/stow/>
  - ↳ Share it to others: consider EasyBuild / ModuleFile
- General workflow for programs based on **Autotools**
  - ↳ Get the software sources (version x.y.z)
  - ↳ Compile and install it in your home/work directory

```
(node)$> ./configure [options] -prefix=$BASEDIR/stow/mysoft.x.y.z
(node)$> make && make install
(node)$> cd $BASEDIR/stow && stow mysoft.x.y.z
```



# Summary

- 1 Overview of the Main HPC Components
- 2 The UL HPC platform
  - Overview
  - Platform Deployment with FAI and Puppet
  - Monitoring
  - Statistics & Milestones
- 3 UL HPC in Practice: Toward an [Efficient] Win-Win Usage
  - General considerations
  - SSH Access
  - The OAR Batch Scheduler
  - Available Software and Environment Modules
- 4 **Last Challenges**
  - Memory bottleneck
  - Effective Storage and Memory Management
  - Reporting problems and Promoting the Platform



## Memory bottleneck

- A regular computing node have at least 2GB/core RAM
  - ↪ Do 12-24 runs fit in the memory?
  - ↪ If your job runs out of memory, it simply crashes
- Use fewer simultaneous runs if **really** needed!
  - ↪ **OR** request a big memory machine (1TB RAM)  
`$> oarsub -t bigmem ...`
  - ↪ **Or** explore parallization (MPI, OpenMP, pthreads)

# Understanding Your Storage Options

## Where can I store and manipulate my data?

- **Shared storage**

- ↪ NFS – **not scalable**  $\simeq 1.5 \text{ GB/s (R)}$   $O(200 \text{ TB})$
- ↪ Lustre – **scalable**  $\simeq 3 \text{ GB/s (R)}$   $O(200 \text{ TB})$

- **Local storage**


- ↪ local file system (/tmp)  $O(200 \text{ GB})$ 
  - over HDD  $\simeq 100 \text{ MB/s}$
  - over SSD  $\simeq 400 \text{ MB/s}$
- ↪ RAM (/dev/shm)  $\simeq 30 \text{ GB/s (R)}$   $O(20 \text{ GB})$

⇒ **In all cases:** small I/Os really **kill** storage performances

# Storage performances

- Based on IOR or IOZone, reference I/O benchmarks

Read






# Storage performances

- Based on IOR or IOZone, reference I/O benchmarks

Write



# Speed Expectation on Data Transfer

<http://fasterdata.es.net/>

- How long to transfer **1 TB** of data across various speed networks?

Network	Time
10 Mbps	300 hrs (12.5 days)
100 Mbps	30 hrs
1 Gbps	3 hrs
10 Gbps	20 minutes

- **(Again)** small I/Os really **kill** performances
  - ↪ **Ex:** transferring 80 TB for the backup of ecosystem\_biology
  - ↪ same rack, 10Gb/s. 4 weeks → 63TB transfer...

# Speed Expectation on Data Transfer

<http://fasterdata.es.net/>

**Data set size**

<b>10PB</b>	<b>166.67 TB/sec</b>	<b>33.33 TB/sec</b>	<b>8.33 TB/sec</b>	<b>2.78 TB/sec</b>
<b>1PB</b>	<b>16.67 TB/sec</b>	<b>3.33 TB/sec</b>	<b>833.33 GB/sec</b>	<b>277.78 GB/sec</b>
<b>100TB</b>	<b>1.67 TB/sec</b>	<b>333.33 GB/sec</b>	<b>83.33 GB/sec</b>	<b>27.78 GB/sec</b>
<b>10TB</b>	<b>166.67 GB/sec</b>	<b>33.33 GB/sec</b>	<b>8.33 GB/sec</b>	<b>2.78 GB/sec</b>
<b>1TB</b>	<b>16.67 GB/sec</b>	<b>3.33 GB/sec</b>	<b>833.33 MB/sec</b>	<b>277.78 MB/sec</b>
<b>100GB</b>	<b>1.67 GB/sec</b>	<b>333.33 MB/sec</b>	<b>83.33 MB/sec</b>	<b>27.78 MB/sec</b>
<b>10GB</b>	<b>166.67 MB/sec</b>	<b>33.33 MB/sec</b>	<b>8.33 MB/sec</b>	<b>2.78 MB/sec</b>
<b>1GB</b>	<b>16.67 MB/sec</b>	<b>3.33 MB/sec</b>	<b>0.83 MB/sec</b>	<b>0.28 MB/sec</b>
<b>100MB</b>	<b>1.67 MB/sec</b>	<b>0.33 MB/sec</b>	<b>0.08 MB/sec</b>	<b>0.03 MB/sec</b>
	<b>1 Minute</b>	<b>5 Minutes</b>	<b>20 Minutes</b>	<b>1 Hour</b>
	<b>Time to transfer</b>			

**Legend:**

Requires less than 100Mbps throughput

Requires between 100Mbps and 10Gbps throughput

Requires between 10Gbps and 100Gbps throughput

Requires more than 100Gbps throughput

Note: Kilo, Mega, etc. are in SI units. E.g. 1KB is 1000 bytes, not 1024 bytes

# Speed Expectation on Data Transfer

<http://fasterdata.es.net/>

## Data set size

1XB	34.72 TB/sec	11.57 TB/sec	1.65 TB/sec	385.80 GB/sec
100PB	3.47 TB/sec	1.16 TB/sec	165.34 GB/sec	38.58 GB/sec
10PB	347.22 GB/sec	115.74 GB/sec	16.53 GB/sec	3.86 GB/sec
1PB	34.72 GB/sec	11.57 GB/sec	1.65 GB/sec	385.80 MB/sec
100TB	3.47 GB/sec	1.16 GB/sec	165.34 MB/sec	38.58 MB/sec
10TB	347.22 MB/sec	115.74 MB/sec	16.53 MB/sec	3.86 MB/sec
1TB	34.72 MB/sec	11.57 MB/sec	1.65 MB/sec	0.39 MB/sec
100GB	3.47 MB/sec	1.16 MB/sec	0.17 MB/sec	0.04 MB/sec
10GB	0.35 MB/sec	0.12 MB/sec	0.02 MB/sec	0.00 MB/sec
	8 Hours	24 Hours	7 Days	30 Days
	Time to transfer			

### Legend:

Requires less than 100Mbps throughput

Requires between 100Mbps and 10Gbps throughput

Requires between 10Gbps and 100Gbps throughput

Requires more than 100Gbps throughput

Note: Kilo, Mega, etc. are in SI units. E.g. 1KB is 1000 bytes, not 1024 bytes




# Fault Tolerance

- Cluster maintenance from time to time

# Fault Tolerance

- Cluster maintenance from time to time
- Reliability vs. Crash Faults in Distributed systems





# Fault Tolerance

- Cluster maintenance from time to time
- Reliability vs. Crash Faults in Distributed systems
- Fault Tolerance general strategy: checkpoint/rollback
  - ↪ assumes a way to save the state of your program
  - ↪ hints: OAR `-signal -checkpoint -idempotent...`, BLCR
  - ↪ combine best-effort jobs with checkpointing (<http://git.io/c-dn1A>)

# Reporting a problem

[https://hpc.uni.lu/users/docs/report\\_pbs.html](https://hpc.uni.lu/users/docs/report_pbs.html)

## ● First checks

- 1 My issue is probably documented User doc
- 2 An event is on-going cf mail from [hpc-platform@uni.lu](mailto:hpc-platform@uni.lu)
- 3 check the state of your nodes Ganglia (especially memory usage)

<http://tinyurl.com/ulhpc-ganglia-testcase>

## ● ONLY NOW consider the following depending on the severity

- ↪ Open an new issue on <http://hpc-tracker.uni.lu> (preferred)
- ↪ Mail (really all of) US [hpc-sysadmins@uni.lu](mailto:hpc-sysadmins@uni.lu)
- ↪ **Ask the help of other users** [hpc-users@uni.lu](mailto:hpc-users@uni.lu)

In all cases: **Carefully describe the problem and the context**





# Reporting your usage of the platform

<https://hpc.uni.lu/users/AUP.html>

- In your scientific publications:
  - ↪ **acknowledge** your usage of the UL HPC platform
  - ↪ cf **Acceptable Use Policy**

Acknowledgment: Experiments presented in this paper were carried out using the HPC facilities of University of Luxembourg~\cite{VBCG\_HPCS14}  
{\small -- see \url{http://hpc.uni.lu}}.

```
@InProceedings{VBCG_HPCS14,  
  author = {S. Varrette and P. Bouvry and H. Cartiaux and F. Georgatos},  
  title = {Management of an Academic HPC Cluster: The UL Experience},  
  booktitle = {Proc. of the 2014 Intl. Conf. on High Performance Computing \& Simulation (HPCS 2014)},  
  year = {2014},  
  OPTpages = {},  
  month = {July},  
  address = {Bologna, Italy },  
  publisher = {IEEE},  
}
```

# Thank you for your attention....

<http://hpc.uni.lu>

- 1 **Overview of the Main HPC Components**
- 2 **The UL HPC platform**
  - Overview
  - Platform Deployment with FAI and Puppet
  - Monitoring
  - Statistics & Milestones
- 3 **UL HPC in Practice: Toward an [Efficient] Win-Win Usage**
  - General considerations
  - SSH Access
  - The OAR Batch Scheduler
  - Available Software and Environment Modules
- 4 **Last Challenges**
  - Memory bottleneck
  - Effective Storage and Memory Management
  - Reporting problems and Promoting the Platform



## Contacts:

[hpc-sysadmins@uni.lu](mailto:hpc-sysadmins@uni.lu)