UL HPC School 2017

Closing Remarks: Take Away Messages

Prof. Pascal Bouvry, Dr. Sebastien Varrette

and the UL HPC Team

June. 13th, 2017, MSA 4.510

University of Luxembourg (UL), Luxembourg

http://hpc.uni.lu

Agenda Day 1: June 12th, 2017

Time	**Main Track (MSA 4.510)**
9h00 – 10h00	**PS1**: Getting Started on the UL HPC platform
10h00 – 10h30	**Coffee break**
10h30 – 12h30	**Overview and Challenges of the UL HPC Facility at the Belval and EuroHPC Horizon**
12h30 – 13h30	**LUNCH**
13h30 – 15h30	**PS2**: HPC workflow with sequential jobs (test cases on GROMACS, Java and Python)
15h30 – 16h00	**Coffee break**
16h00 – 17h00	**PS4a**: UL HPC Monitoring in practice: why, what, how, where to look
17h00 – 18h30	**PS5**: HPC workflow with Parallel/Distributed jobs

Time	**Advanced Parallel Track (MSA 4.520)**
9h00 – 10h00	
10h00 – 10h30	**Coffee break**
10h30 – 12h30	**Overview and Challenges of the UL HPC Facility at the Belval and EuroHPC Horizon**
12h30 – 13h30	**LUNCH**
13h30 – 15h30	**PS3**: Advanced Scheduling (Slurm, OAR) and Software Customization
15h30 – 16h00	**Coffee break**
16h00 – 17h00	**PS4b**: Debugging, profiling and performance analysis
17h00 – 18h30	**PS6**: Bioinformatics workflows and applications

PS = *Practical Session using your laptop*

Agenda Day 2: June 12th, 2017

Time	Main Track (MSA 4.510)
9h00 – 10h30	PS7: Big Data Applications
10h30 – 11h00	**Coffee break**
11h00 – 12h30	Users' session: UL HPC experiences
12h30 – 13h30	**LUNCH**
13h30 – 15h00	PS9: [Advanced] Prototyping with Python
15h30 – 16h00	**Coffee break**
16h00 – 17h30	PS10: R - statistical computing
17h30 – 18h30	Closing Keynote: Take Away Messages

Time	Advanced Parallel Track (MSA 4.520)
9h00 – 10h30	PS8: MATLAB (interactive, passive, sequential, checkpointing and parallel)
10h30 – 11h00	**Coffee break**
11h00 – 12h30	
12h30 – 13h30	**LUNCH**
13h30 – 15h00	PS11: Multi-Physics workdflows (CFD / MD / Chemistry Applications)
15h30 – 16h00	**Coffee break**
16h00 – 17h30	PS12: Virtualization
17h30 – 18h30	

PS = *Practical Session using your laptop*

Thanks of Attending this 5th Edition !

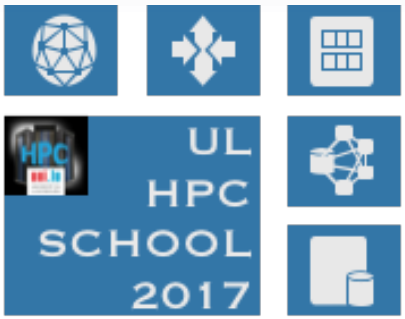

- https://hpc.uni.lu/hpc-school/
  - ↪ more than **70 registered attendees**
  - ↪ we **hope you found it useful**
  - ↪ **Feedback / comments welcome!**
    - ✓ hpc-sysadmins@uni.lu
    - ✓ Improvement suggestion / pull requests etc.

Thanks of Attending this 5th Edition !

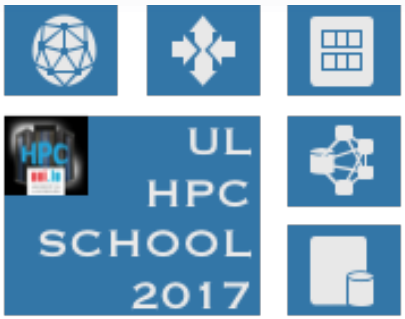

- https://hpc.uni.lu/hpc-school/
  - ↪ more than **70 registered attendees**
  - ↪ we **hope you found it useful**
  - ↪ **Feedback / comments welcome!**
    - ✓ hpc-sysadmins@uni.lu
    - ✓ Improvement suggestion / pull requests etc.

Thanks for helping us testing the new `iris` cluster!

Thanks of Attending this 5th Edition !

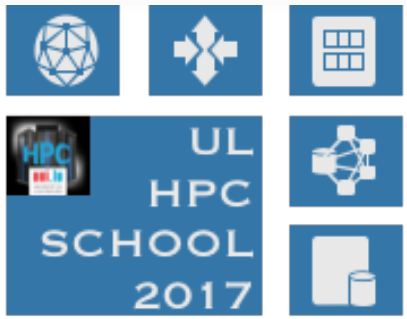

- https://hpc.uni.lu/hpc-school/
  - ↪ more than **70 registered attendees**
  - ↪ we **hope you found it useful**
  - ↪ **Feedback / comments welcome!**
    - ✓ hpc-sysadmins@uni.lu
    - ✓ Improvement suggestion / pull requests etc.

Thanks for helping us testing the new `iris` cluster!

- Something *confusing*? *badly presented*? topics *not covered*?
  - ↪ **Tell us**! . . . we will do our best to have it covered
  - ↪ suggest you own well-made tutorials. . .
    - ✓ . . . and (if you're not shy) present it!
- **Looking forward meeting you in Nov, 2017**

A Few Take Away Messages

PC $\neq$ HPC $\neq$ Cloud

- does not mean they cannot work in synergy

Context	Local PC	HPC
Sequential	$T_1(\text{local}) = 100$	$T_1(\text{hpc}) = \mathbf{120s}$
Parallel/Distributed	$T_2(\text{local}) = 70\text{s}$	$T_2(\text{hpc}) = 80\text{s}$
		$T_8(\text{hpc}) = \mathbf{60s}$

- **Parallel/Distributed** runs **DO NOT COME FOR FREE**
  - ↪ runs **will be sequential** even if you reserve ≥ 2 cores/nodes
  - ↪ you have to **explicitly** adapt your jobs to benefit from the multi-cores/nodes

[New] Tools for you

- Find a nice **editor** `vim`, `emacs`, `spacemacs...`
- **Git everywhere**
  - ↪ Git[Lab,Hub...] supports SSH deploy key
- SSH is your new friend (as the UL HPC platform)
  - ↪ as `rsync` etc.

For more information...

PhD Seminar: IT/Dev[op]s Army Knives Tools for the Researcher

- (advanced) Gnu Screen is your new friend
  - ↪ see ULHPC tutorial
- Easybuild, Spack for HPC software configuration
  - ↪ Use HomeBrew / Linuxbrew on your laptop

Know What You Are Doing...

- Always **check what you are doing on the platform!**
  - ↪ common pitfalls : **Out of memory**
  - ↪ tools for you: `htop`, Ganglia, `valgrind`, `ddt`, `map`...
  - ↪ good practice: **benchmark** your code
    - ✓ helps to anticipate on **appropriate walltime**
- Many **[advanced] debugging/profiling tools** are available for you
  - ↪ use it! we pay for it ;-)
- **write launchers files**! `{sbatch | oarsub -S} <launcher.sh>`
  - ↪ use `#OAR` or `#SBATCH` comments to define default reservation

OAR Pattern	Description
`enclosure=N`	number of enclosure
`nodes=N`	number of nodes
`core=N`	number of cores
`walltime=hh:mm:ss`	job's max duration

sbatch/srun option	Description
`-N $n`	Number of nodes
`--ntasks-per-node=$tn`	Tasks per node
`-c $ct`	Cores per task (multithreading)
`-n $nt`	Total # of tasks
`--mem=$mGB`	Memory per node
`--time=d-hh:mm:ss`	Request job walltime

Slurm vs. OAR Main Commands

Action	SLURM command	OAR Command
Submit passive/batch job	`sbatch [...] $script`	`oarsub [...] $script`
Start interactive job	`srun [...] --pty bash`	`oarsub -I [...]`
Queue status	`squeue`	`oarstat`
User job status	`squeue -u $user`	`oarstat -u $user`
Specific job status (detailed)	`scontrol show job $jobid`	`oarstat -f -j $jobid`
Job accounting status (detailed)	`sacct --job $jobid -l`	
Delete (running/waiting) job	`scancel $jobid`	`oardel $jobid`
Hold job	`scontrol hold $jobid`	`oarhold $jobid`
Resume held job	`scontrol release $jobid`	`oarresume $jobid`
Node list and their properties	`scontrol show nodes`	`oarnodes`

Know Your Ennemy...

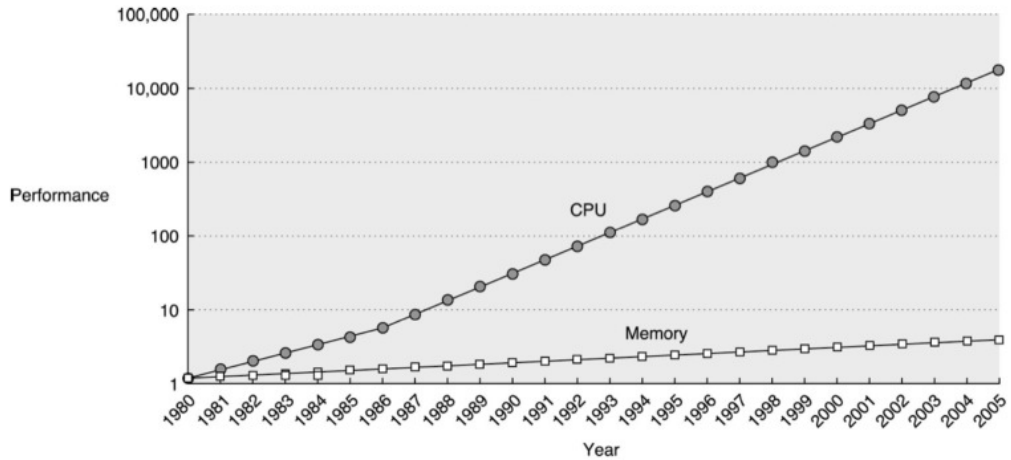

- A regular computing node have **at least 2GB/core RAM**
  - ↪ Do 12-24 runs fit in the memory?
  - ↪ If your job runs out of memory, it simply crashes
- Use fewer simultaneous runs if **really** needed!
  - ↪ **OR** request a big memory machine `{oarsub -t bigmem [...]`
    - ✓ 1TB to 4TB RAM
  - ↪ **OR (better)** explore parallization (MPI, OpenMP, pthreads)
  - ↪ **OR (better)** explore Big Data applications (Spark etc.)

Storage Performances

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

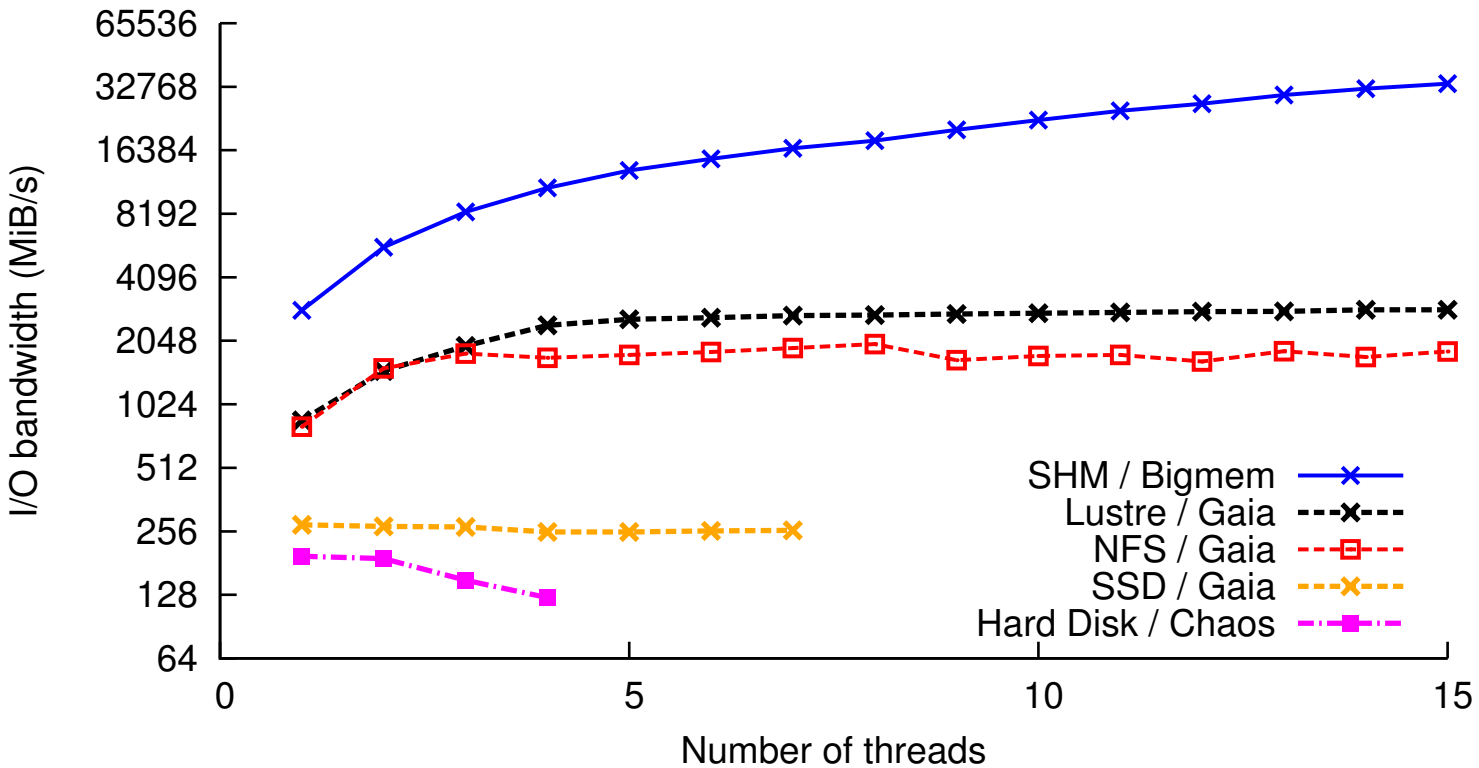

Storage Performances

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

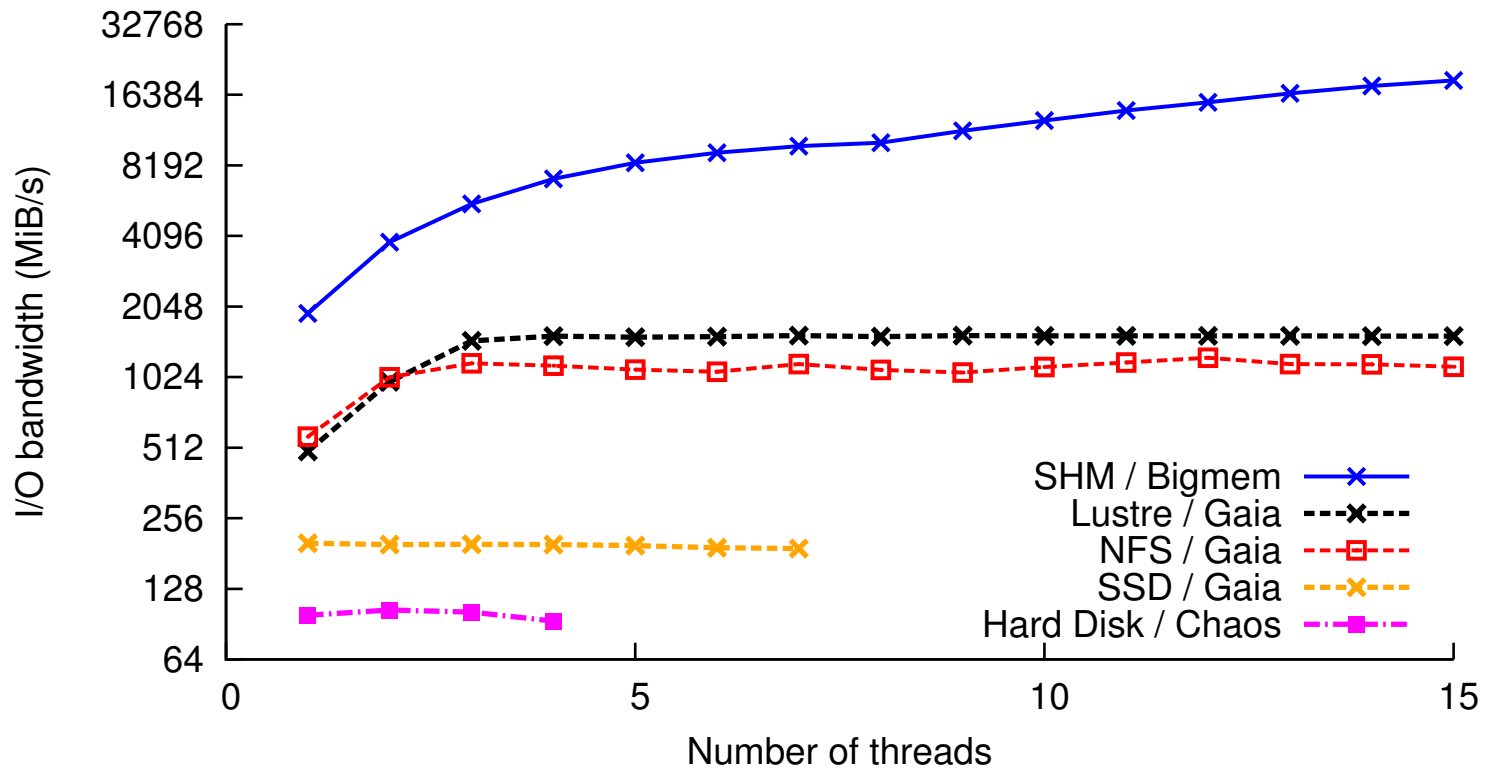

Understanding Your Storage Options

Where can I store and manipulate my data?

- **Shared** storage
  - ↪ `NFS` – **not scalable** $\simeq$ 1.5 GB/s (R) $\mathcal{O}$(100 TB)
  - ↪ `GPFS` – **scalable** $\simeq$ 10 GB/s (R) $\mathcal{O}$(1 PB)
  - ↪ `Lustre` – **scalable** $\simeq$ 5 GB/s (R) $\mathcal{O}$(0.5 PB)
- **Local** storage
  - ↪ local file system (`/tmp`) $\mathcal{O}$(200 GB)
    - ✓ over HDD $\simeq$ 100 MB/s
    - ✓ over SDD $\simeq$ 400 MB/s
  - ↪ RAM (`/dev/shm`) $\simeq$ 30 GB/s (R) $\mathcal{O}$(20 GB)

$\Rightarrow$ **In all cases:** small I/Os really **kill** storage performances

Happy Computing !!!

http://hpc.uni.lu

The UL HPC Team aka
Prof. Pascal Bouvry
Dr. Sebastien Varrette
Valentin Plugaru
Sarah Peter
Hyacinthe Cartiaux
Clément Parisot

Meet us at:
University of Luxembourg, Belval Campus
Maison du Nombre, 4th floor
2, avenue de l'Université
L-4365 Esch-sur-Alzette
mail: hpc@uni.lu

Spread the word...

http://hpc.uni.lu

Happy Computing !!!

http://hpc.uni.lu

The UL HPC Team aka
Prof. Pascal Bouvry
Dr. Sebastien Varrette
Valentin Plugaru
Sarah Peter
Hyacinthe Cartiaux
Clément Parisot

Meet us at:
University of Luxembourg, Belval Campus
Maison du Nombre, 4th floor
2, avenue de l'Université
L-4365 Esch-sur-Alzette
mail: hpc@uni.lu

Spread the word...

http://hpc.uni.lu