Uni.lu HPC School 2020

Overview and Challenges of the ULHPC facility at the EuroHPC Horizon

Dr. S. Varrette & UL HPC Team

University of Luxembourg (UL)

https://hpc.uni.lu

10th ULHPC School

Dec. 15-16, 2020, Luxembourg

Uni.lu HPC School 2020

https://hpc.uni.lu/hpc-school/

- **10th edition** of this training initiated in 2014
  - ↪ This one is the **long remote** version (COVID regulation)
    - ✓ 2-days, semi-parallel sessions, feat. **basic & advanced** tutorials
  - ↪ **New:** *dedicated* Slack workspace for live support
    - ✓ ulhpc-school-2020.slack.com
- **Requirement** (outside coffee or tea)
  - ↪ your favorite laptop with your favorite OS
    - ✓ Linux / Mac OS preferred, but Windows accepted
  - ↪ basic knowledge in Linux command line and ability to take notes (Markdown etc.)

Online UL HPC Tutorials

https://ulhpc-tutorials.rtfd.io/

ULHPC Technical Documentation

https://hpc.uni.lu https://hpc-docs.uni.lu

Uni.lu HPC School 2020 Agenda

Day 1	Tuesday, December 15, 2020		
	Main Track		
		Description	***Speaker***
08:30 - 09:00	**Remote setup and registration**		
09:00 - 10:30	Keynote	Welcome Overview and Challenges of the UL HPC Facility at the EuroHPC Horizon	S. Varrette
10:30 - 11:00	**BREAK**		
11h00 - 12:30	PS1	Preliminaries (SSH - OpenOnDemand)	T. Valette, A. Olloh
12:30 - 14:00	**LUNCH**		
14:00 - 16:00	PS2	Getting Started 2.0 SLURM, performance engineering and basic launchers	H. Cartiaux
16:00 - 16:30	**BREAK**		
16:30 - 18:00	PS3	HPC Management of Sequential and Embarrassingly parallel jobs	S. Varrette

Uni.lu HPC School 2020 Agenda

Day 2	Wednesday, December 16, 2020						
	Main Track			**Parallel Session**			
		Description	***Speaker***			***Description***	***Speaker***
09:00 - 11:00	PS4	Big Data Analytics	S. Varrette	09:00 - 11:00	PS5	Scalable Science with OpenMP/MPI	E. Krishnasamy
11:00 - 11h30	**BREAK**						
11:30 - 12:30	Keynote	Data management (backup, security, ...)	S. Peter				
12:30 - 14:00	**LUNCH**			12:30 - 14:00	**LUNCH**		
14:00 - 15:00	PS6	HPC Containers with Singularity	E. Kieffer	14:00 - 15:00	PS7a	Introduction to GPU programming with CUDA (Part I)	F. Pinel
SHORT BREAK							
15:15 - 16:15	PS8	Advanced distributed computing with Python	E. Kieffer	15:15 - 16:15	PS7b	Introduction to GPU programming with CUDA (Part II)	L. Koutsantonis
SHORT BREAK							
16:30 - 18:00	PS10	R - statistical computing	A. Ginolhac	16:30 - 18:00	PS9	Multi-GPU Training of Neural Networks	C. Hundt (Nvidia)
SHORT BREAK							
18:15 - 18:30	**Closing Remarks / Take Away messages**		S. Varrette				

Uni.lu HPC School 2020 Contributors

... in alphabetical order

Hyacinthe Cartiaux
Infra. & HPC Arch. Engineer

Dr. Aurelien Ginohac
Research Scientist

Dr. Christian Hundt
NVidia AI/DL Solutions Architect

Dr. Emmanuel Kieffer
Research Scientist

Dr. Ezhilmathi Krishnasamy
Postdoctoral Researcher

Dr. Loizos Koutsantonis
Postdoctoral Researcher

Abatcha Olloh
Infra. & HPC Arch. Engineer

Sarah Peter
Infra. & Arch. Engineer

Dr. Frederic Pinel
Research Scientist

Teddy Valette
Infra. & HPC Arch. Engineer

Dr. Sebastien Varrette
Research Scientist

... and additional help (Survey, session tests)

Arlyne Vandeventer
Project Manager

Uni.lu HPC School 2020 Contributors

... in alphabetical order

Hyacinthe Cartiaux
Infra. & HPC Arch. Engineer

Dr. Aurelien Ginohac
Research Scientist

Dr. Christian Hundt
NVidia AI/DL Solutions Architect

Dr. Emmanuel Kieffer
Research Scientist

Dr. Ezhilmathi Krishnasamy
Postdoctoral Researcher

Dr. Loizos Koutsantonis
Postdoctoral Researcher

Abatcha Olloh
Infra. & HPC Arch. Engineer

Sarah Peter
Infra. & Arch. Engineer

Dr. Frederic Pinel
Research Scientist

Teddy Valette
Infra. & HPC Arch. Engineer

Dr. Sebastien Varrette
Research Scientist

... and additional help (Survey, session tests)

Arlyne Vandeventer
Project Manager

Summary

Summary

Prerequisites: Metrics

- **HPC**: **H**igh **P**erformance **C**omputing — **BD**: **B**ig **D**ata

Main HPC/BD Performance Metrics

- **Computing Capacity**: often measured in **flops** (or **flop/s**)
  - ↪ **Fl**oating **p**oint **o**perations per **s**econds (often in DP)
  - ↪ **GFlops** $= 10^{9}$ **TFlops** $= 10^{12}$ **PFlops** $= 10^{15}$ **EFlops** $= 10^{18}$

Prerequisites: Metrics

- **HPC**: **H**igh **P**erformance **C**omputing — **BD**: **B**ig **D**ata

Main HPC/BD Performance Metrics

- **Computing Capacity**: often measured in **flops** (or **flop/s**)
  - ↪ **Fl**oating **p**oint **o**perations per **s**econds (often in DP)
  - ↪ **GFlops** $= 10^9$ **TFlops** $= 10^{12}$ **PFlops** $= 10^{15}$ **EFlops** $= 10^{18}$
- **Storage Capacity**: measured in multiples of **bytes** = 8 **bits**
  - ↪ **GB** $= 10^9$ bytes **TB** $= 10^{12}$ **PB** $= 10^{15}$ **EB** $= 10^{18}$
  - ↪ **GiB** $= 1024^3$ bytes **TiB** $= 1024^4$ **PiB** $= 1024^5$ **EiB** $= 1024^6$
- Transfer rate on a medium measured in **Mb/s** or **MB/s**
- Other metrics: Sequential vs Random **R/W speed**, **IOPS** ...

HPC, Big Data & Cloud

- **HPC**: **H**igh **P**erformance **C**omputing
- **BD**: **B**ig **D**ata
- **Cloud Computing**:
  - ↪ Network access to a shared pool of configurable computing resources* which is:
    - ✓ Ubiquitous, Convenient, On-demand

HPC, Big Data & Cloud

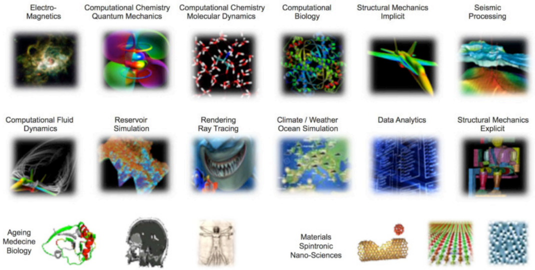

- **HPC**: **H**igh **P**erformance **C**omputing
- **BD**: **B**ig **D**ata
- **Cloud Computing**:
  - ↪ Network access to a shared pool of configurable computing resources* which is:
    - ✓ Ubiquitous, Convenient, On-demand
- All scientific disciplines are becoming **computational** today
  - ↪ Modern scientific discovery requires very high computing power, handles huge data volumes
  - ↪ cf. J. Rifkin report: "3rd Industrial Revolution Strategy for the Gd. Duchy of Luxembourg"
  - ↪ **Research Projects**, **Industry** and **SMEs** are increasingly relying on computing resources
    - ✓ . . . to invent innovative solutions while reducing cost and decreasing time to market
- All are nowadays essential tools for **Research, Science, Society and Industry**

HPC, Big Data & Cloud

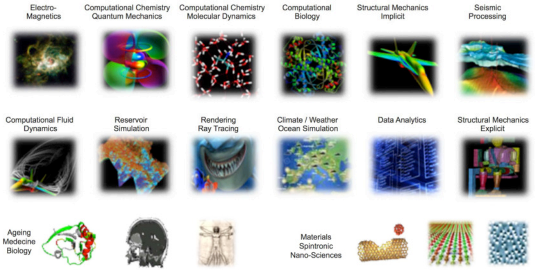

- **HPC**: **H**igh **P**erformance **C**omputing
- **BD**: **B**ig **D**ata
- **Cloud Computing**:
  - ↪ Network access to a shared pool of configurable computing resources* which is:
    - ✓ Ubiquitous, Convenient, On-demand
- All scientific disciplines are becoming **computational** today
  - ↪ Modern scientific discovery requires very high computing power, handles huge data volumes
  - ↪ cf. J. Rifkin report: "3rd Industrial Revolution Strategy for the Gd. Duchy of Luxembourg"
  - ↪ **Research Projects**, **Industry** and **SMEs** are increasingly relying on computing resources
    - ✓ ... to invent innovative solutions while reducing cost and decreasing time to market
- All are nowadays essential tools for **Research, Science, Society and Industry**

Andy Grant, Head of Big Data and HPC, Atos UK&I

To out-compete you must out-compute

Increasing competition, heightened customer expectations and shortening product development cycles are forcing the pace of acceleration across all industries.

Different Needs for Different Domains

Material Science & Engineering

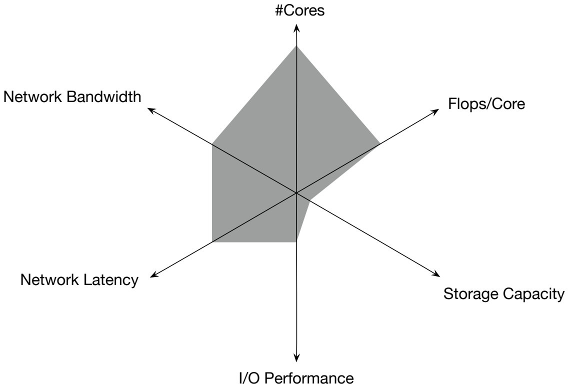

Different Needs for Different Domains

Biomedical Industry / Life Sciences

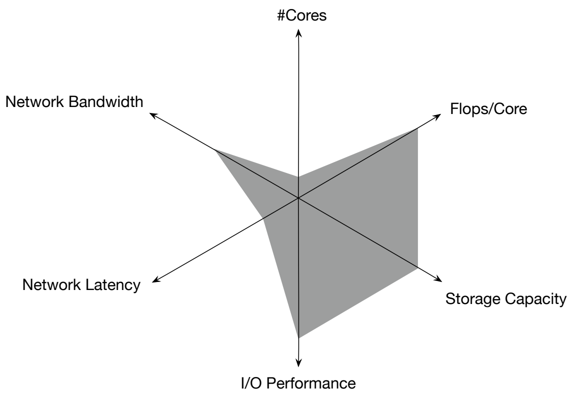

Different Needs for Different Domains

Deep Learning / Cognitive Computing

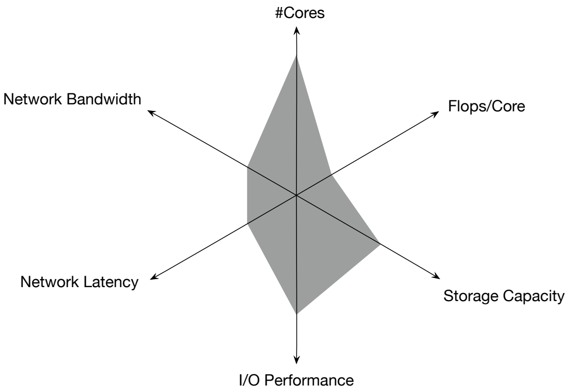

Different Needs for Different Domains

IoT, FinTech

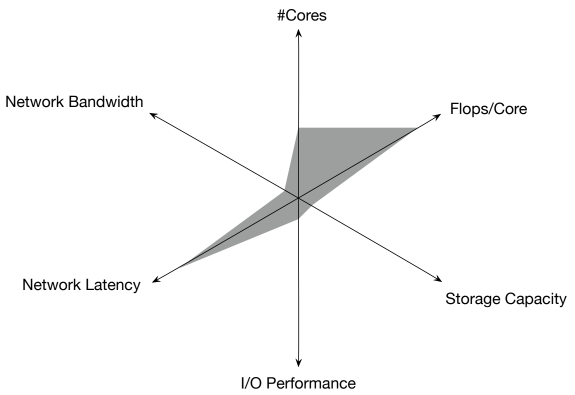

Different Needs for Different Domains

ALL Research Computing Domains

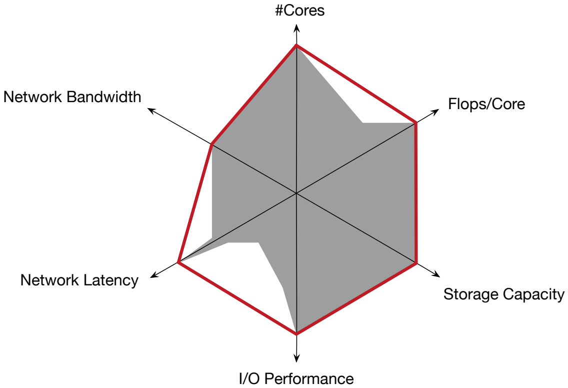

Different Needs for Different Domains

ALL Research Computing Domains

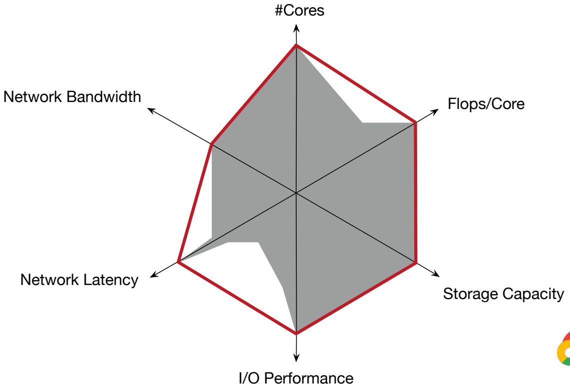

Different Needs for Different Domains

ALL Research Computing Domains

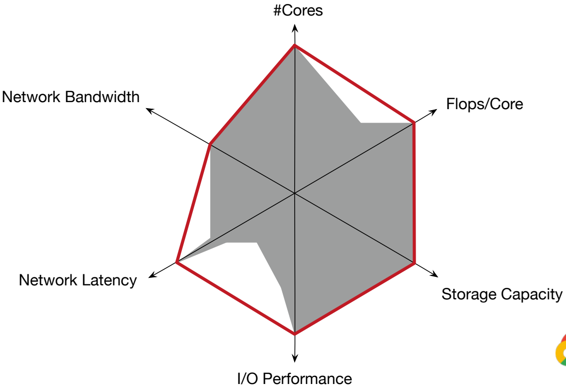

EuroHPC Joint Undertaking

PRACE

amazon web services

Microsoft Azure

Google Cloud

High Performance Computing & Big Data Services

hpc.uni.lu

hpc@uni.lu

@ULHPC

LUXEMBOURG LET'S MAKE IT HAPPEN

Computing for Researchers: Laptop

- **Regular PC / Local Laptop / Workstation**
  ↪ **Native OS** (Windows, Linux, Mac etc.)

Computing for Researchers: Laptop

- **Regular PC / Local Laptop / Workstation**
  - ↪ **Native OS** (Windows, Linux, Mac etc.)
- Virtualized OS (**VM**) through an **hypervisor**
  - ↪ *Hypervisor*: core virtualization engine / environment
    - ✓ Ex: Xen, VMWare ESXi, KVM, VirtualBox – **Non-negligible Performance loss**: $\geq 20\%$

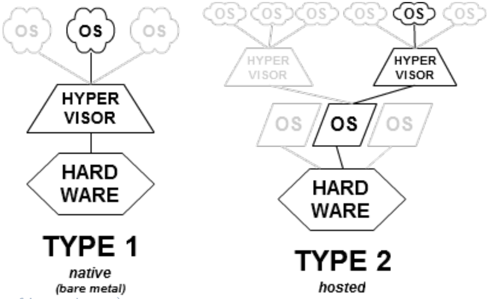

Computing for Researchers: Laptop

- **Regular PC / Local Laptop / Workstation**
  - ↪ **Native OS** (Windows, Linux, Mac etc.)
- Virtualized OS (**VM**) through an **hypervisor**
  - ↪ *Hypervisor*: core virtualization engine / environment
    - ✓ Ex: Xen, VMWare ESXi, KVM, VirtualBox – **Non-negligible Performance loss**: $\geq 20\%$
- **Container**-based Virtualization
  - ↪ similar to VMs, **yet** containers **share** the system kernel of the host with others
    - ✓ Ex: Docker, Singularity, Shifter
  - ↪ Orchestration: Kubernetes

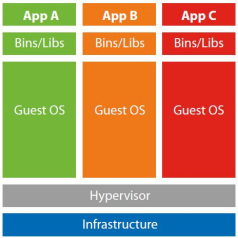

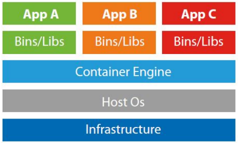

Computing for Researchers: Laptop

- **Regular PC / Local Laptop / Workstation**
  - ↪ **Native OS** (Windows, Linux, Mac etc.)
- Virtualized OS (**VM**) through an **hypervisor**
  - ↪ *Hypervisor*: core virtualization engine / environment
    - ✓ Ex: Xen, VMWare ESXi, KVM, VirtualBox – **Non-negligible Performance loss**: $\geq 20\%$
- **Container**-based Virtualization
  - ↪ similar to VMs, **yet** containers **share** the system kernel of the host with others
    - ✓ Ex: Docker, Singularity, Shifter
  - ↪ Orchestration: Kubernetes

HPC-compliant Containers

- No docker for security reasons
- **Singularity**, Sarus

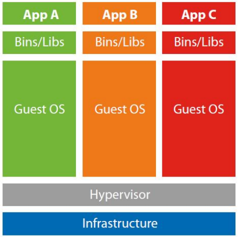

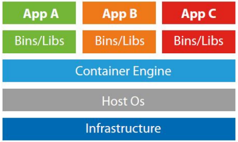

Computing for Researchers: Cloud

- **Cloud Computing**
  - ↪ access to shared *(generally virtualized)* resources
  - ↪ pay-per-use approach
  - ↪ **Infrastructure** as a Service (**IaaS**)

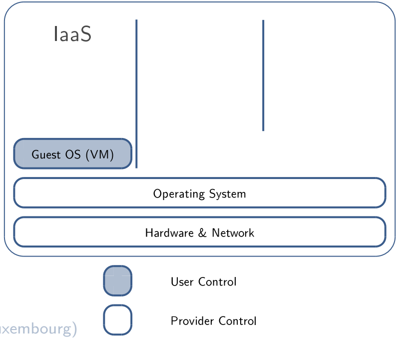

Computing for Researchers: Cloud

- **Cloud Computing**
  - ↪ access to shared *(generally virtualized)* resources
  - ↪ pay-per-use approach
  - ↪ **P**latform as a Service (**P**aaS)

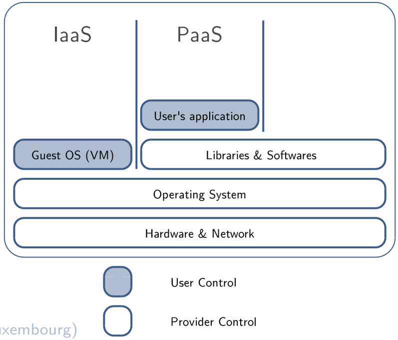

Computing for Researchers: Cloud

- **Cloud Computing**
  - ↪ access to shared *(generally virtualized)* resources
  - ↪ pay-per-use approach
  - ↪ **Software** as a Service (**S**aaS)

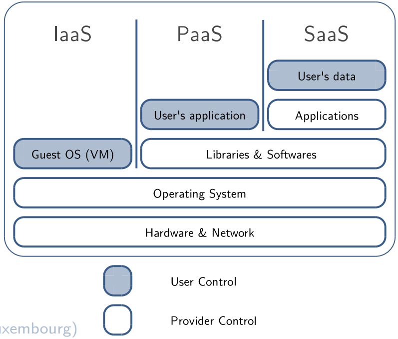

Computing for Researchers: Cloud

- **Cloud Computing**
  - ↪ access to shared *(generally virtualized)* resources
  - ↪ pay-per-use approach
  - ↪ **XXX** as a Service (<**X**>aaS)

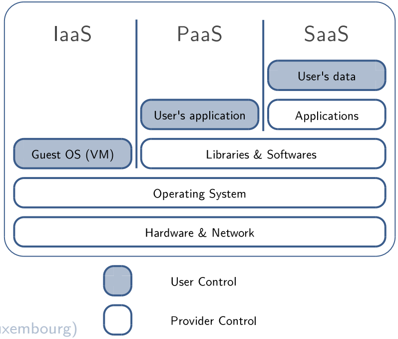

Computing for Researchers: HPC

- High Performance Computing (HPC) platforms
  ↪ For **Speedup**, **Scalability** and **Faster Time to Solution**

Computing for Researchers: HPC

- High Performance Computing (HPC) platforms
  - ↪ For **Speedup**, **Scalability** and **Faster Time to Solution**
  - ↪ Available peta-scale facilities in Luxembourg
    - ✓ Tier 0/1: MeluXina (part of EuroHPC network)

Computing for Researchers: HPC

- High Performance Computing (HPC) platforms
  - ↪ For **Speedup**, **Scalability** and **Faster Time to Solution**
  - ↪ Available peta-scale facilities in Luxembourg
    - ✓ Tier 0/1: MeluXina (part of EuroHPC network)
    - ✓ Tier 2: Uni.lu HPC

Computing for Researchers: HPC

- High Performance Computing (HPC) platforms
  - ↪ For **Speedup**, **Scalability** and **Faster Time to Solution**
  - ↪ Available peta-scale facilities in Luxembourg
    - ✓ Tier 0/1: MeluXina (part of EuroHPC network)
    - ✓ Tier 2: Uni.lu HPC

YET...

$$\textbf{PC} \neq \textbf{Cloud} \neq \textbf{HPC}$$

Computing for Researchers: HPC

- High Performance Computing (HPC) platforms
  - ↪ For **Speedup**, **Scalability** and **Faster Time to Solution**
  - ↪ Available peta-scale facilities in Luxembourg
    - ✓ Tier 0/1: MeluXina (part of EuroHPC network)
    - ✓ Tier 2: Uni.lu HPC

YET...

PC ≠ Cloud ≠ HPC

- HPC ≃ Formula 1
  - ↪ relies on ultra efficient hardware / interconnect (IB HDR. . .)
  - ↪ . . . when Cloud has to stay standard (10 GbE etc. . .)
- **Does not mean the 3 approaches cannot work together**

HPC Computing Hardware

Base

- **CPU** (Central Processing Unit) *Highest software flexibility*
  - ↪ High performance across all computational domains
  - ↪ Ex: Intel Core i9-10900K *(Q2'20)* $R_{peak} \simeq$ 1,18 TFlops (DP)
    - ✓ 10 cores @3.7GHz (14nm, 125W, $\simeq$ 7 billion transistors) + integ. graphics

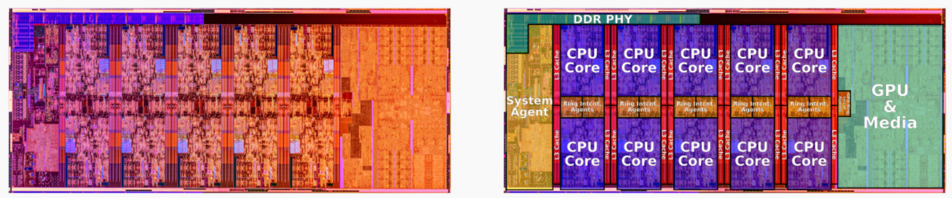

Intel Comet Lake die (2020)

HPC Computing Hardware

Base

- **CPU** (Central Processing Unit) *Highest software flexibility*
  - ↪ High performance across all computational domains
  - ↪ Ex: Intel Core i9-10900K *(Q2'20)* $R_{peak} \simeq$ 1,18 TFlops (DP)
    - ✓ 10 cores @3.7GHz (14nm, 125W, $\simeq$ 7 billion transistors) + integ. graphics

Accelerators

- **GPU** (Graphics Processing Unit): *Ideal for ML/DL workloads*
  - ↪ Ex: Nvidia Tesla A100 *(Q1'20)* $R_{peak} \simeq$ 9.7 TFlops (DP)
    - ✓ 6912 cores @ 1.41GHz (7nm, 400W, 54,2 billion transistors)

HPC Computing Hardware

Base

- **CPU** (Central Processing Unit) *Highest software flexibility*
  - ↪ High performance across all computational domains
  - ↪ Ex: Intel Core i9-10900K *(Q2'20)* $R_{peak} \simeq$ 1,18 TFlops (DP)
    - ✓ 10 cores @3.7GHz (14nm, 125W, $\simeq$ 7 billion transistors) + integ. graphics

Accelerators

- **GPU** (Graphics Processing Unit): *Ideal for ML/DL workloads*
  - ↪ Ex: Nvidia Tesla A100 *(Q1'20)* $R_{peak} \simeq$ 9.7 TFlops (DP)
    - ✓ 6912 cores @ 1.41GHz (7nm, 400W, 54,2 billion transistors)
- ~~Intel MIC (Many Integrated Core) Accelerator~~
- **ASIC** (Application-Specific Integrated Circuits), **FPGA** (Field Programmable Gate Array)
  - ↪ least software flexibility, highest performance for specialized problems
    - ✓ Ex: AI, Mining, Sequencing...

$\Longrightarrow$ toward hybrid platforms w. DL enabled accelerators

HPC Components: Local Memory

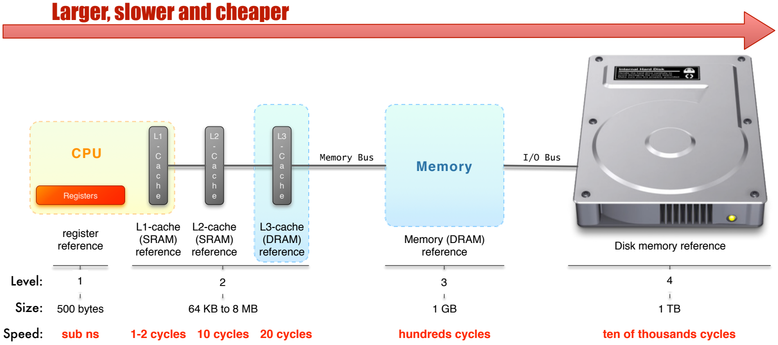

- SSD (SATA3) R/W: 550 MB/s; 100000 IOPS **450 €/TB**
- HDD (SATA3 @ 7,2 krpm) R/W: 227 MB/s; 85 IOPS **54 €/TB**

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	40μs to 300μs
10 Gigabit Ethernet	10 Gb/s	1.25 GB/s	4μs to 5μs
Infiniband QDR	40 Gb/s	5 GB/s	1.29μs to 2.6μs
Infiniband EDR	100 Gb/s	12.5 GB/s	0.61μs to 1.3μs
Infiniband HDR	200 Gb/s	25 GB/s	0.5μs to 1.1μs
100 Gigabit Ethernet	100 Gb/s	1.25 GB/s	30μs
Intel Omnipath	100 Gb/s	12.5 GB/s	0.9μs

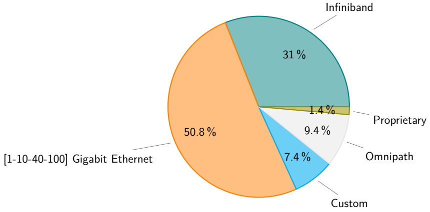

[Source : www.top500.org, Nov 2020]

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	40μs to 300μs
10 Gigabit Ethernet	10 Gb/s	1.25 GB/s	4μs to 5μs
Infiniband QDR	40 Gb/s	5 GB/s	1.29μs to 2.6μs
Infiniband EDR	100 Gb/s	12.5 GB/s	0.61μs to 1.3μs
Infiniband HDR	200 Gb/s	25 GB/s	0.5μs to 1.1μs
100 Gigabit Ethernet	100 Gb/s	1.25 GB/s	30μs
Intel Omnipath	100 Gb/s	12.5 GB/s	0.9μs

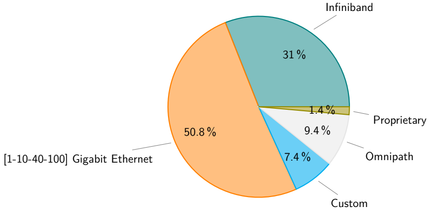

[Source : www.top500.org, Nov 2020]

Network Topologies

- **Direct** vs. **Indirect** interconnect
  - ↪ *direct*: each network node attaches to at least one compute node
  - ↪ *indirect*: compute nodes attached at the edge of the network only
    - ✓ many routers only connect to other routers.

Network Topologies

- **Direct** vs. **Indirect** interconnect
  - ↪ *direct*: each network node attaches to at least one compute node
  - ↪ *indirect*: compute nodes attached at the edge of the network only
    - ✓ many routers only connect to other routers.

Main HPC Topologies

- **CLOS Network / Fat-Trees** [Indirect]
  - ↪ can be fully non-blocking (1:1) or blocking (x:1)
  - ↪ typically enables **best performance**
    - ✓ Non blocking bandwidth, lowest network latency

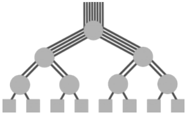

Network Topologies

- **Direct** vs. **Indirect** interconnect
  - ↪ *direct*: each network node attaches to at least one compute node
  - ↪ *indirect*: compute nodes attached at the edge of the network only
    - ✓ many routers only connect to other routers.

Main HPC Topologies

- **CLOS Network / Fat-Trees** [Indirect]
  - ↪ can be fully non-blocking (1:1) or blocking (x:1)
  - ↪ typically enables **best performance**
    - ✓ Non blocking bandwidth, lowest network latency

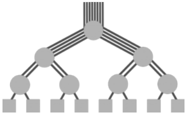

- **Mesh or 3D-torus** [Direct]
  - ↪ Blocking network, cost-effective for systems at scale
  - ↪ Great performance solutions for applications with locality
  - ↪ Simple expansion for future growth

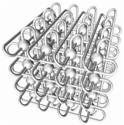

HPC Components: Operating System

- Exclusively Linux-based (**really** 100%)
  - ↪ Note: Used to be Unix before
  - ↪ **better to become familiar with Linux environments**
    - ✓ interaction can be done from **ANY** OS
- Reasons:
  - ↪ stability
  - ↪ development flexibility

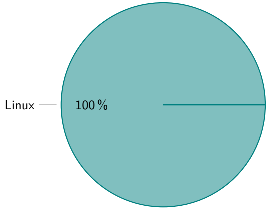

[Source : www.top500.org, Nov 2020]

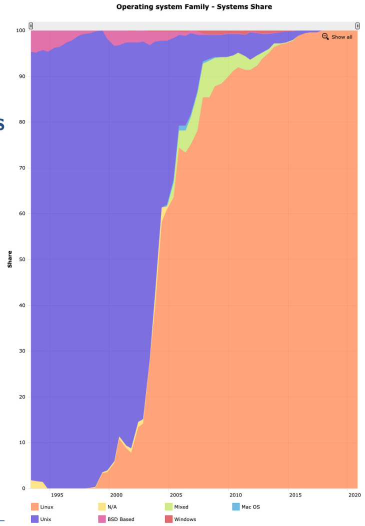

HPC Components: Software Stack

- **Remote connection to the platform** SSH
- **Identity Management / SSO**: LDAP, Kerberos, IPA...
- **Resource management**: job/batch scheduler
  ↪ SLURM, OAR, PBS, MOAB/Torque...
- (Automatic) **Node Deployment**:
  ↪ FAI, Kickstart, Puppet, Chef, Ansible, Kadeploy...
- (Automatic) **User Software Management**:
  ↪ Easybuild, Environment Modules, LMod
- **Platform Monitoring**:
  ↪ Nagios, Icinga, Ganglia, Foreman, Cacti, Alerta...

[Big]Data Management: Disk Encl.

- $\simeq$ **150 K€**/encl. - 60-84 disks (4U)
  ↪ incl. redundant (i.e. 2) RAID controllers (master/slave)

[Big]Data Management: FS Summary

- **File System (FS)**: Logical manner to *store, organize & access* data
  - ↪ (local) **Disk FS** : `FAT32`, `NTFS`, `HFS+`, `ext4`, `{x,z,btr}fs`...
  - ↪ **Networked FS**: `NFS`, `CIFS/SMB`, `AFP`
  - ↪ **Parallel/Distributed FS**: `SpectrumScale/GPFS`, `Lustre`
    - ✓ typical FS for HPC / HTC (High Throughput Computing)

[Big]Data Management: FS Summary

- **File System (FS)**: Logical manner to *store, organize & access* data
  - ↪ (local) **Disk FS** : `FAT32`, `NTFS`, `HFS+`, `ext4`, `{x,z,btr}fs`...
  - ↪ **Networked FS**: `NFS`, `CIFS/SMB`, `AFP`
  - ↪ **Parallel/Distributed FS**: `SpectrumScale/GPFS`, `Lustre`
    - ✓ typical FS for HPC / HTC (High Throughput Computing)

Main Characteristic of Parallel/Distributed File Systems

Capacity and Performance increase with #servers

[Big]Data Management: FS Summary

- **File System (FS)**: Logical manner to *store, organize & access* data
  - ↪ (local) **Disk FS** : `FAT32, NTFS, HFS+, ext4, {x,z,btr}fs...`
  - ↪ **Networked FS**: `NFS, CIFS/SMB, AFP`
  - ↪ **Parallel/Distributed FS**: `SpectrumScale/GPFS, Lustre`
    - ✓ typical FS for HPC / HTC (High Throughput Computing)

Main Characteristic of Parallel/Distributed File Systems

Capacity and Performance increase with #servers

Name	Type	Read* [GB/s]	Write* [GB/s]
`ext4`	Disk FS	0.426	0.212
`nfs`	Networked FS	0.381	0.090
`gpfs` (iris/aion)	Parallel/Distributed FS	**11.25**	**9,46**
`lustre` (iris/aion)	Parallel/Distributed FS	**12.88**	**10,07**

* maximum **random** read/write, per IOZone or IOR measures, using concurrent nodes for networked FS.

HPC Components: Data Center

Definition (Data Center)

- Facility to house computer systems & components
  - ↪ Basic storage component: **rack** (height: 42 RU)

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

HPC Components: Data Center

Definition (Data Center)

- Facility to house computer systems & components
  - ↪ Basic storage component: **rack** (height: 42 RU)

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

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

- Power/Heat dissipation per rack:
  - ↪ HPC **computing** racks: **30-120 kW**
  - ↪ **Storage** racks: **15 kW**
  - ↪ **Interconnect** racks: **5 kW**
- Various **Cooling** Technology
  - ↪ Airflow with In-Row cooling

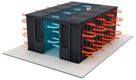

  - ↪ Direct-Liquid Cooling, Immersion...

Summary

University of Luxembourg & HPC

- *With regards to HPC*, University of Luxembourg offers:
  - ↪ **People**
    - ✓ **Domain experts**, Computational and Data scientists
    - ✓ Specialists in parallel algorithmics

High Performance
Computing &
Big Data Services

hpc.uni.lu
hpc@uni.lu
@ULHPC

University of Luxembourg & HPC

- *With regards to HPC*, University of Luxembourg offers:
  - ↪ **People**
    - ✓ **Domain experts**, Computational and Data scientists
    - ✓ Specialists in parallel algorithmics
  - ↪ **Services**
    - ✓ **HPC Compute & Data services** (HPC for research)
    - ✓ IT services (SIU)

University of Luxembourg & HPC

- *With regards to HPC*, University of Luxembourg offers:
  - ↪ **People**
    - ✓ **Domain experts**, Computational and Data scientists
    - ✓ Specialists in parallel algorithmics
  - ↪ **Services**
    - ✓ **HPC Compute & Data services** (HPC for research)
    - ✓ IT services (SIU)
  - ↪ **Infrastructure**
    - ✓ **State-of-the-art HPC systems**, **2.7 PFlops** compute capacity
    - ✓ Highly capable Data Center (*Centre De Calcul* **CDC**)
    - ✓ Cutting-edge energy-efficient Direct Liquid Cooling capability

University of Luxembourg & HPC

- *With regards to HPC*, University of Luxembourg offers:
  - ↪ **People**
    - ✓ **Domain experts**, Computational and Data scientists
    - ✓ Specialists in parallel algorithmics
  - ↪ **Services**
    - ✓ **HPC Compute & Data services** (HPC for research)
    - ✓ IT services (SIU)
  - ↪ **Infrastructure**
    - ✓ **State-of-the-art HPC systems**, **2.7 PFlops** compute capacity
    - ✓ Highly capable Data Center (*Centre De Calcul* **CDC**)
    - ✓ Cutting-edge energy-efficient Direct Liquid Cooling capability
  - ↪ **Education & Training**
    - ✓ MICS Parallel and Grid Computing lecture, Bi-annual HPC School
    - ✓ **Technology Transfer HPC workshops & seminars**

      ... in collaboration with UL / National HPC Competence Center)

High Performance Computing @ UL

- **Started in 2007** under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ 2nd Largest HPC facility in Luxembourg. . .
    - ✓ after EuroHPC MeluXina (≥ 15 PFlops) system

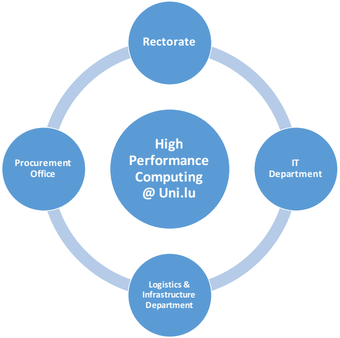

https://hpc.uni.lu/

HPC/Computing Capacity

2794.23 TFlops

(incl. 748.8 GPU TFlops)

Shared Storage Capacity

10713.4 TB storage

High Performance Computing @ UL

UL HPC Cluster (2020)

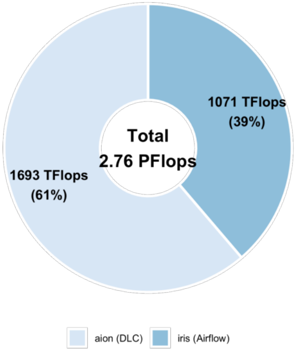

UL HPC Storage FileSystems (2020)

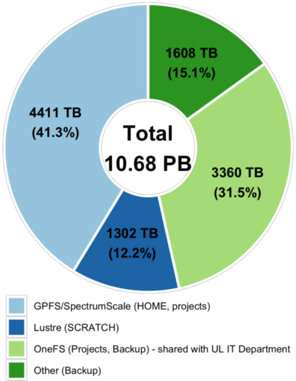

HPC in Luxembourg and Around in EU

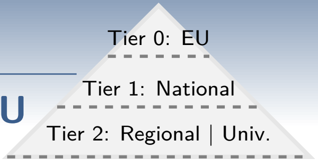

Country	System(s)	Type	Institute	(CPU) #Nodes	#Cores	#[GPU]Accelerators	R_{peak}	Shared Storage
	MeluXina (2021)	Euro-HPC **Petascale** *Tier 0/1 (EU,Nat)*	LuxProvide	824	≃ 88 000	764 NVidia A100	**17.57 PF**	≃ 20 PB
Luxembourg	aion,iris	*Tier 2 (Univ)*	**Uni.lu HPC**	**552**	46896	**96** NVidia V100	**2.79 PF**	**10.71 PB**
		Tier 2 (local)	LIST	40	1280	8 Nvidia V100	0.126 PF	0.58 PB
	TGCC (Joliot-Curie)	*Tier 0 (EU)*	GENCI/CEA	4808	430 448	828 Xeon Phi, 128 NVidia V100	22.26 PF	35PB
France	JeanZay	*Tier 1 (Nat.)*	GENCI/Idris	1 528	61 120	1292 NVidia V100	14.97 PF	31.2 PB
	ROMEO	*Tier 2 (Reg.)*	Univ. Reims	115	3 220	280 NVidia P100	1.75 PF	0.634
	Vlaams	*Tier 1 (Nat.)*	VSC	988	27 664	n/a	1.63 PF	1.3PB
Belgium	zenobe	*Tier 1 (Nat.)*	Cenaero	584	14 016	4 NVidia K40	0.41 PF	0.356PB
	Hortense	*Tier 2 (Reg.)*	Gent Univ.	n/a	≃ 40 000	88 NVidia V100	3.3PF	3PB
	JUWELS	*Tier 0 (EU)*	JSC	2571	122 768	224 Nvidia V100	12.3 PF	130.3PB
	JURECA	*Tier 0 (EU)*	JSC	3524	156 736	1640 Xeon Phi	7.24 PF	(as above)
Germany	Hawk	*Tier 0 (EU)*	HLRS, Univ. Stuttgart	5632	720 896	n/a	26 PF	≃25PB
	SuperMUC-NG	*Tier 0 (EU)*	LRZ, Munich	6480	311 040	n/a	26.9 PF	70.16PB
	CLAIX-2018	*Tier 2 (Univ)*	Univ. Aachen	1307	61 200	108 Nvidia V100	4.11 PF	3PB
Bulgaria	PetaSC (2021)	Euro-HPC **Petascale** *Tier 0/1 (EU,Nat)*	SofiaTech	n/a	n/a	n/a	**≃ 4 PF**	n/a
Czech Republic	Barbora	*Tier 1 (Nat.)*	IT4Innovation	201	7232	32 NVidia V100	0.85 PF	≃ 1PB
	EURO_IT4I (2021)	Euro-HPC **Petascale** *Tier 0/1 (EU,Nat)*	IT4Innovation	826	≃ 100K	560 NVidia A100	**15.2 PF**	1PB
Finland	LUMI (2021)	Euro-HPC **Pre-exascale** *Tier 0 (EU)*	CSC	n/a	≃ 200K (LUMI-C)	n/a	**552 PF**	127PB
	Marconi-A3	*Tier 0 (EU)*	Cineca	3216	154 368	n/a	10.37 PF	10PB
Italy	Galileo	*Tier 1 (Nat.)*	Cineca	1022	36792	n/a	1.35 PF	1.92PB
	Leonardo (2021)	Euro-HPC **Pre-exascale** *Tier 0 (EU)*	Cineca	4992	n/a	13824 Nvidia A100	**322.6 PF**	100PB
Portugal	Deucalion (2021)	Euro-HPC **Petascale** *Tier 0/1 (EU,Nat)*	MACC	n/a	n/a	n/a	**≃ 10 PF**	n/a
Slovenia	VEGA (2021)	Euro-HPC **Petascale** *Tier 0/1 (EU,Nat)*	Maribor SC	960	122,8K	240 NVidia A100	**10.1 PF**	24 PB
Spain	MareNostrum 4	*Tier 0 (EU)*	BSC	3456	165 888	n/a	11.15 PF	14PB
	MareNostrum 5 (2021)	Euro-HPC **Pre-exascale** *Tier 0 (EU)*	BSC	n/a	n/a	n/a	**≃ 200 PF**	n/a
Switzerland	Piz-Daint	*Tier 0 (EU)*	CSCS, ETH Zürich	7517	387 872	5704 NVidia P100	29.34 PF	8.8PB

Uni.lu HPC Users

- **1518** registered HPC Users
  - ↪ 23 computational domains accelerated on UL HPC

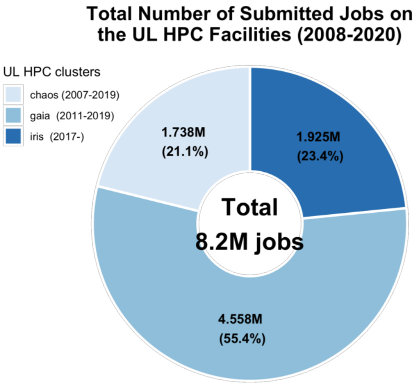

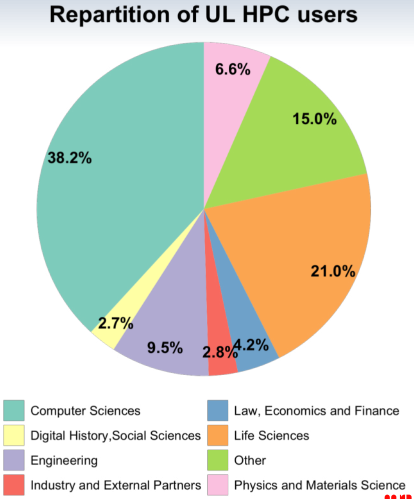

Uni.lu HPC Cumulative Usage

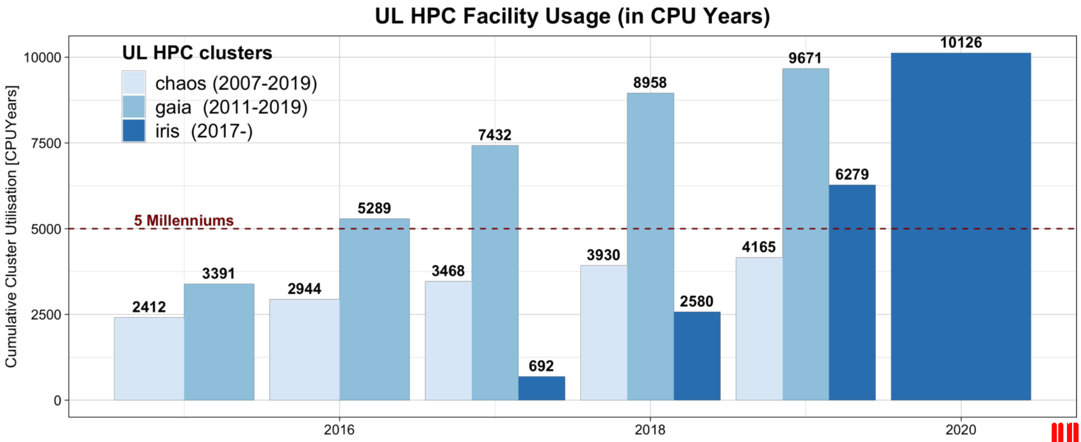

Accelerating UL Research - User Software Sets

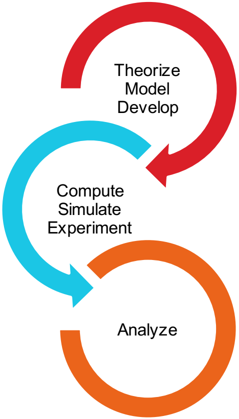

- **Over 230 software packages** available for researchers
  - ↪ software environment generated using Easybuild / LMod
  - ↪ containerized applications delivered with Singularity system

Domain	**2019 Software environment**
Compiler Toolchains	FOSS (GCC), Intel, PGI
MPI suites	OpenMPI, Intel MPI
Machine Learning	PyTorch, TensorFlow, Keras, Horovod, Apache Spark...
Math & Optimization	Matlab, Mathematica, R, CPLEX, Gurobi...
Physics & Chemistry	GROMACS, QuantumESPRESSO, ABINIT, NAMD, VASP...
Bioinformatics	SAMtools, BLAST+, ABySS, mpiBLAST, TopHat, Bowtie2...
Computer aided engineering	ANSYS, ABAQUS, OpenFOAM...
General purpose	ARM Forge & Perf Reports, Python, Go, Rust, Julia...
Container systems	Singularity
Visualisation	ParaView, OpenCV, VMD, VisIT
Supporting libraries	numerical (arpack-ng, cuDNN), data (HDF5, netCDF)...
...	

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

UL HPC Governance & Pillars

- Rectorate
  - Head HPC for Research Prof. P. Bouvry
    - Deputy head HPC for research Dr. S. Varrette
    - Research Computing
      - HPC User Engagement & L1/L2Support
        - Scientific Computing High Level Support
        - Compute & Data Service Support
        - Industry & Business Support
        - Web Portals & Documentation
        - HPC Tickets & Accounts
        - Uni.lu Faculty and ICs Liaison
        - Public Research Centres Liaison
      - Scientific Software & Libraries
        - Toolchains, debuggers, programming languages
        - Bioinformatics, biology and biomedical
        - Computational science
        - AI, DL, BigData analytics
        - High-level mathematical software
        - Performance evaluation & Benchmarks
        - Visualization
        - GPU accelerated software
      - HPC Operations and Supercomputing Services
        - Compute Services
        - Network , Monitoring and Security Services
        - Storage, Data & Backup Services
        - Resource Allocation & Scheduling
        - Project & Identity Management, Accounting
        - Data Center & Infrastructure Operations
        - DevOps - CI/CD
        - Disaster Recovery
    - HPC R&D / Training
      - HPC Applications & Middleware optimization
      - Energy Efficiency & Hybrid computing optimization
      - Emerging Technologies Edge and Fog Computing
      - AI / Machine Learning
      - Technology Watch
      - Training
      - Dissemination
    - Strategic Developments, Partnership
      - HPC Service Agreements & Consulting
      - Partnership & Business Services
      - Reporting and Auditing
      - Security & Data Protection
      - National HPC Coordination
      - HPC Competence Center
      - NVidia AI Technology Center
      - EU HPC Projects (EuroHPC, PRACE[-6IP], ETP4HPC, Grid5000, SLICE…)
      - International HPC Cooperation (E-READI…)
    - Administration & Information
      - Communications, Media & Event Support
      - Financial Project Management & Control
      - HPC Procurement & Inventory
      - Human Ressources
      - Licences & Maintenance Contract management
      - Project Coordination

UL HPC Core Team

Prof. Stephane Pallage
Rector

Uni.lu HPC Team

Prof. Pascal Bouvry
Head Uni.lu HPC*
* *Temporary mandate: Half-time CEO LuxProvide*

Dr. Sebastien Varrette
Research Scientist,
Deputy Head, Uni.lu HPC

Research Computing & HPC Operations
* *Acting Head during mandate: Dr. S. Varrette*

Hyacinthe Cartiaux
Infrastructure and
HPC Architecture Engineer

Abatcha Olloh
Infrastructure and
HPC Architecture Engineer

Teddy Valette
Infrastructure and
HPC Architecture Engineer

Sarah Peter
Infrastructure & Architecture Engineer
LCSB BioCore sysadmins manager

HPC Research & Trainings
* *Acting Deputy Head during mandate: Dr. F. Pinel*

Dr. Frederic Pinel
Research Scientist,
Coordinator NVidia Joint AI Lab

Dr. Emmanuel Kieffer
Research Scientist

Dr. Ezhilmathi Krishnasamy
Postdoctoral Researcher,
Coordinator H2020 PRACE-6IP

Dr. Loizos Koutsantonis
Postdoctoral Researcher, EuroCC

N/A
Postdoctoral Researcher, EuroCC

Strategic Developments Partnership

Arlyne Vandeventer
Project Manager
EuroHPC Comp. Center (EuroCC)

Administration & Information

UL HPC Core Team

Prof. Stephane Pallage
Rector

Uni.lu HPC Team

Prof. Pascal Bouvry
Head Uni.lu HPC*
** Temporary mandate: Half-time CEO LuxProvide*

Dr. Sebastien Varrette
Research Scientist,
Deputy Head, Uni.lu HPC

Research Computing & HPC Operations
** Acting Head during mandate: Dr. S. Varrette*

Hyacinthe Cartiaux
Infrastructure and
HPC Architecture Engineer

Abatcha Olloh
Infrastructure and
HPC Architecture Engineer

Teddy Valette
Infrastructure and
HPC Architecture Engineer

Sarah Peter
Infrastructure & Architecture Engineer
LCSB BioCore sysadmins manager

HPC Research & Trainings
** Acting Deputy Head during mandate: Dr. F. Pinel*

Dr. Frederic Pinel
Research Scientist,
Coordinator NVidia Joint AI Lab

Dr. Emmanuel Kieffer
Research Scientist

Dr. Ezhilmathi Krishnasamy
Postdoctoral Researcher,
Coordinator H2020 PRACE-6IP

Dr. Loizos Koutsantonis
Postdoctoral Researcher, EuroCC

N/A
Postdoctoral Researcher, EuroCC

Strategic Developments Partnership

Arlyne Vandeventer
Project Manager
EuroHPC Comp. Center (EuroCC)

Administration & Information

. . . and domain experts across ALL the University

Uni.lu Data Center

Belval Campus

Centre De Calcul (CDC)

- Power generation station for HPC floor:
  - ↪ up to **3 MW of electrical power**
  - ↪ **2.4 MW of cold water** at a 12-18°C regime
    - ✓ used for traditional Airflow with In-Row cooling.
  - ↪ Separate hot water circuit (between 30 and 40°C)
    - ✓ used for Direct Liquid Cooling (DLC): `aion`

Location	Cooling	Usage
CDC S-02-001	Airflow	**Future extension**
CDC S-02-002	Airflow	**Future extension**
CDC S-02-003	DLC	**Future extension** - High Density/Energy efficient HPC
CDC S-02-004	**DLC**	**High Density/Energy efficient HPC**: `aion`
CDC S-02-005	**Airflow**	**Storage / Traditional HPC**: `iris` and common equipment

UL HPC Supercomputers: General Architecture

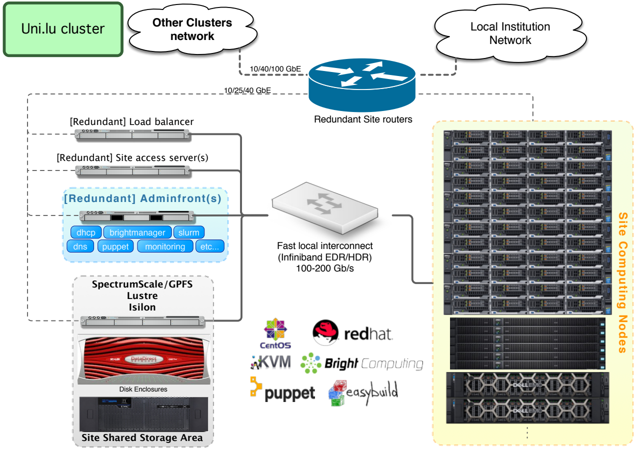

UL HPC Supercomputers: `iris` cluster

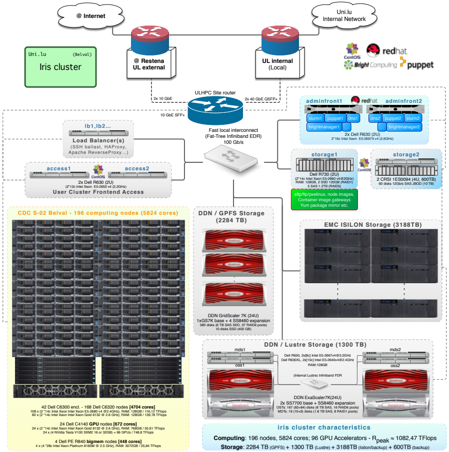

- **Dell/Intel** supercomputer, Air-flow cooling
  - ↪ 196 compute nodes
    - ✓ 5824 compute cores
    - ✓ Total 52224 GB RAM
  - ↪ R_{peak}: **1,072 PetaFLOP/s**
- Fast InfiniBand (IB) EDR network
  - ↪ **Fat-Tree** Topology — blocking factor 1:1.5

Rack ID	Purpose	Description
D02	Network	Interconnect equipment
D04	Management	Management servers, Interconnect
D05	Compute	`iris-[001-056]`, interconnect
D07	Compute	`iris-[057-112]`, interconnect
D09	Compute	`iris-[113-168]`, interconnect
D11	Compute	`iris-[169-177,191-193](gpu)`, `iris-[187-188](bigmem)`
D12	Compute	`iris-[178-186,194-196](gpu)`, `iris-[189-190](bigmem)`

UL HPC Supercomputers: aion cluster

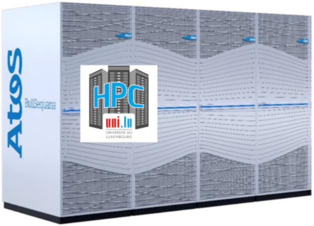

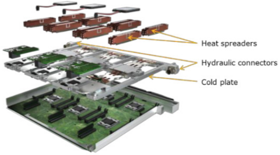

- **Atos/AMD** supercomputer, DLC cooling
  - ↪ 4 BullSequana XH2000 adjacent racks
  - ↪ 318 compute nodes
    - ✓ 40704 compute cores
    - ✓ Total 81408 GB RAM
  - ↪ R_{peak}: **1,693 PetaFLOP/s**
- Fast InfiniBand (IB) HDR network
  - ↪ **Fat-Tree** Topology

blocking factor 1:2

	Rack 1	Rack 2	Rack 3	Rack 4	**TOTAL**
Weight [kg]	1872,4	1830,2	1830,2	1824,2	**7357 kg**
#X2410 Rome Blade	28	26	26	26	**106**
#Compute Nodes	84	78	78	78	**318**
#Compute Cores	10752	9984	9984	9984	**40704**
R_{peak} **[TFlops]**	447,28 TF	415,33 TF	415,33 TF	415,33 TF	**1693.29 TF**

UL HPC Computing capacity

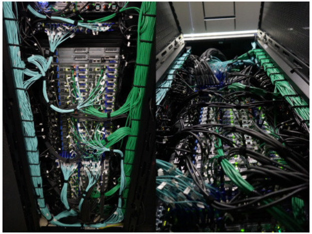

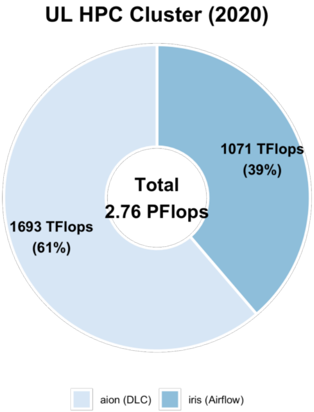

UL HPC - Detailed Computing Nodes

	#N	#C	R_{peak}
Uni.lu HPC TOTAL:	**552**	**46896**	**2794.23 TFlops**
			(incl. 748.8 GPU TFlops)

Cluster	Date	Vendor	Proc. Description		#N	#C	R_{peak}
aion	2020*	Atos	AMD EPYC 7H12 @2.6 GHz	2 × 64c, 256GB	318	40704	1693,29 TFlops
				aion **TOTAL:**	**318**	**40704**	**1693.3 TFlops**

** installation delayed due to global COVID crisis*

Cluster	Date	Vendor	Proc. Description		#N	#C	R_{peak}
iris	2017	Dell	Intel Xeon E5-2680 v4@2.4GHz	2 × 14C,128GB	108	3024	116,12 TFlops
	2018	Dell	Intel Xeon Gold 6132 @ 2.6 GHz	2 × 14C,128GB	60	1680	139,78 TFlops
	2018	Dell	Intel Xeon Gold 6132 @ 2.6 GHz	2 × 14C,768GB	24	672	55,91 TFlops
	2019		*Per node*: 4x NVIDIA Tesla V100 SXM2 16/32GB		96 GPUs	491520	748,8 GPU TFlops
	2018	Dell	Intel Xeon Platinum 8180M @ 2.5 GHz	4 × 28C,3072GB	4	448	35,84 TFlops
				iris **TOTAL:**	**196**	**5824**	**347.65 TFlops**
					96 GPUs	491520	**+748.8 GPU Tflops**

Cluster	Date	Vendor	Proc. Description		#N	#C	R_{peak}
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	3.072 TFlops
				granduc/petitprince **TOTAL:**	**38**	**368**	**4.48 TFlops**

Fast Local Interconnect Network

- HPC interconnect technologies nowadays divided into three categories
  1. Ethernet: dominant interconnect standard yet underlying protocol has inherent limitations
     - ✓ preventing low-latency deployments expected in real HPC environment
  2. InfiniBand: predominant interconnect technology in the HPC market
  3. Vendor specific interconnects: Cray/HPC Slingshot, Intel Omni-Path, Bull BXI...

- **On ULHPC**: **InfiniBand (IB)** solution
  - ↪ `iris`: Infiniband (IB) **EDR** Fabric in a **Fat-Tree** Topology
  - ↪ `aion`: Infiniband (IB) **HDR100** Fabric in a **Fat-Tree** Topology

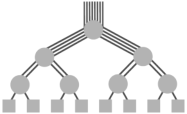

- Up/Down InfiniBand Routing Algorithm
  - ↪ super-set of Fat-Tree with a tracker mode (allow each node to have dedicated route)
  - ↪ well adapted to IO traffic patterns

Fast Local Infiniband Interconnect Network

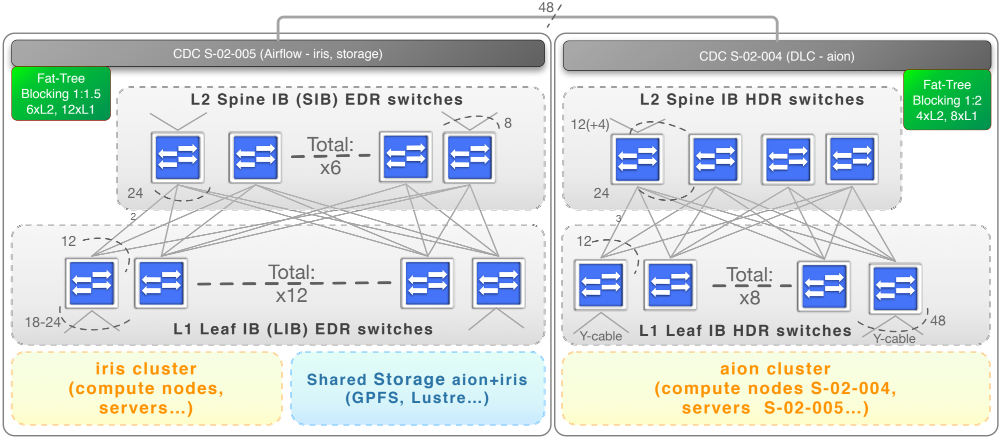

Ethernet Network

- High-bandwidth and low-latency network: local Fast IB interconnect network
  - ↪ support efficient HPC and Big data workloads
- Flexibility of Ethernet-based networks still required

UL HPC Ethernet network

- **2-layers** topology
  - ↪ Upper level: **Gateway Layer**
    - ✓ routing, switching features, network isolation and filtering (ACL) rules
    - ✓ meant to interconnect only switches.
    - ✓ allows to interface the University network (LAN/WAN)
  - ↪ bottom level: **Switching Layer**
    - ✓ composed by [stacked] core switches as well as the TOR (Top-the-rack) switches,
    - ✓ meant to interface HPC servers and compute nodes

UL HPC Storage Systems

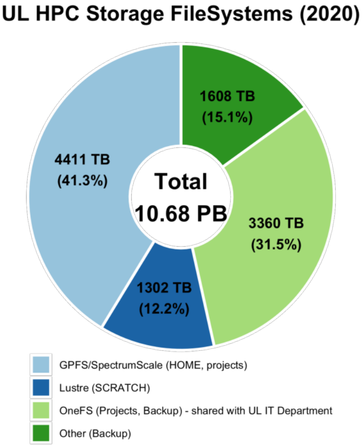

UL HPC Software Stack

Operating System: Linux CentOS/Redhat

- **User Single Sign-on**: Redhat IdM/IPA
- **Remote connection & data transfer**: SSH/SFTP
  ↪ **User Portal**: Open OnDemand
- **Scheduler/Resource management**: Slurm
- **(Automatic) Server / Compute Node Deployment**:
  ↪ BlueBanquise, Bright Cluster Manager, Ansible, Puppet and Kadeploy
- **Virtualization and Container Framework**: KVM, Singularity
- **Platform Monitoring** (User level): Ganglia, SlurmWeb, OpenOndemand...
- **ISV software**:
  ↪ ABAQUS, ANSYS, MATLAB, Mathematica, Gurobi Optimizer, Intel Cluster Studio XE, ARM Forge & Perf. Report, Stata, ...

The case of Grid'5000

http://www.grid5000.fr

- Large scale nation wide infrastructure
  ↪ for large scale parallel and distributed computing research.

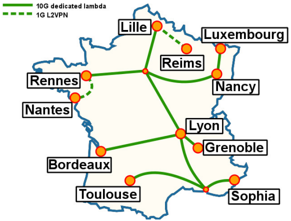

- 8 sites, 7 in France (1 site **Abroad**: Luxembourg)
  ↪ Total: **12326** cores over **31** clusters
- 1-10GbE / Infiniband
  ↪ **10Gb/s dedicated** between all sites
- Unique software stack
  ↪ **kadeploy, kavlan, kwapi**

The case of Grid'5000

http://www.grid5000.fr

- Large scale nation wide infrastructure
  - ↪ for large scale parallel and distributed computing research.

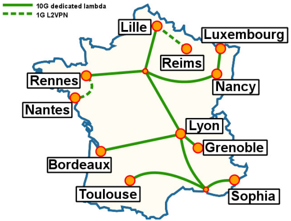

- 8 sites, 7 in France (1 site **Abroad**: Luxembourg)
  - ↪ Total: **12326** cores over **31** clusters
- 1-10GbE / Infiniband
  - ↪ **10Gb/s dedicated** between all sites
- Unique software stack
  - ↪ **kadeploy, kavlan, kwapi**

- **Out of scope for this talk**
  - ↪ General information: https://hpc.uni.lu/g5k
  - ↪ Grid'5000 website and documentation: https://www.grid5000.fr

Summary

Notable Events/Activities (2019)

... since last HPC School 2019!

- **June-July 2019**: MeluXina proposal accepted (EuroHPC call for Petascale Supercomputers)
  ↪ elaborated by Min. of Economy consortium, incl. LuxConnect, ULHPC, LIST, JSC, Partec
  - ✓ July 24, 2019: creation of LuxProvide S.A., hosting entity managing MeluXina supercomputer
  - ✓ Initiate RFI preparation phase of the tender for MeluXina acquisition (concluded Jan 2020)

Notable Events/Activities (2019)

... since last HPC School 2019!

- **July-Sept 2019**: Preparation of Uni.lu tender **RFP 190027: New DLC Cluster Aion**
  - ↪ 116 pages (new template), 430 criteria evaluated, **Target Budget ≃ 3 M€(HT)**
    - ✓ Lot 1: DLC Computing cluster `aion`
    - ✓ Lot 2: Adaptation and extension of the existing High-Performance Storage systems
    - ✓ Lot 3: Adaptation of the network (Ethernet and IB)
  - ↪ **Sept 11, 2019**: official public release
    - ✓ TED72/2019-608787 / PMP Portal n°1901442

University of Luxembourg RFP 190027 – HPC Cluster Aion

TENDER FILE

RFP 190027 - TENDER FOR [THE ACQUISITION OF A COMPLEMENTARY HIGH PERFORMANCE COMPUTING (HPC) CLUSTER 'AION']

To be executed in the interests of **the University of Luxembourg**

The opening of the Submission of Tenders will be on **21 October 2019 at 11:00 AM** via the Public Procurement Portal. The Tender File contains to 116 pages numbered from 1 to 116, including annexes.

This Tender File was prepared by:
The HPC Team of the University of Luxembourg (UL HPC), led by Prof. Pascal Bouvry and Dr. Sebastien Varrette

The tendering is done:	☐ by core works	☒ by general contractor

The start of services is planned for **16 March 2020.**
The Contract has a duration of **4 years** with support maintenance and warranty potentially extendable for two (2) periods of twelve (12) months at the request of the University.

	To be completed by the Tenderer	Reserved for the Contracting Authority
The amount of the tender:	.... €	.... €
VAT :	.... €	.... €
Total:	.... €	.... €

Name of the Tenderer: ....
Address: ....
Telephone: ....
Fax: ....
Email: ....
Government authorisation number: ....

Commitment form:
By signing the commitment form, the service provider confirms that it is familiar with all of the documents in the Tender File governing the services provided in compliance with the conditions of the Tender File, with the aforementioned staff, in accordance with professional standards, within the deadline and at the Tender price.

Unique signature and stamp of the Tenderer	Reserved for the Contracting Authority

N.B. In case of a joint venture and/or in case of a general contractor, the forms listed in 7.1, 7.2, 7.3, must also be completed.

1 / 116

RFP 190027 - TENDER EVALUATION - OVERVIEW

100%

ID		Criteria	Weighting criteria [%]	Weighting sub-criteria [%]
1		**Presentation and organisation of the Tender**	**5%**	
	1.1	Clarity and conciseness of the Tender		20%
	1.2	Summary of the Tender in a maximum of 4 pages by lot (Equipment, Maintenance, Service)		10%
	1.3	Completeness of the proposal		20%
	1.4	Advantage (s) accorded to the University		30%
	1.5	Proposed methods to avoid a possible conflict of interest		20%
				100%
2		**Management of resources and transition**	**5%**	
	2.1	CV of the named account manager and the replacement for the named account manager		10%
	2.2	CV of the service delivery manager and his replacement		5%
	2.3	CVs of the technicians/engineers		15%
	2.4	CVs of the replacement technicians/engineers		15%
	2.5	Personal references		10%
	2.6	Continuous training of the technicians/engineers		20%
	2.7	Transfer of knowledge		25%
				100%
3		**Tenderer's similar experience**	**10%**	
	3.1	Similar projects		40%
	3.2	Expertise in all of the fields concerned by this Contract		60%
				100%
4		**Tenderer's ability**	**5%**	
	4,1	Professional and support services organization		80%
	4,2	Number of qualified staff within the implementation team		20%
				100%
5		**Technical value and compliance of the IT equipment**	**50%**	
	5.1	General performance / General technical quality of the offer		5%
	5.2	Functional performance		5%
	5.3	General Hosting performance		8%
	5.4	Technical performance Lot 1 (DLC Computing cluster Aion)		28%
	5.5	Technical performance Lot 2 (Adaptation and extension of the existing High-Performance Storage systems)		10%
	5.6	Technical performance Lot 3 (Adaptation of the network (Ethernet and IB))		10%
	5.7	Management performance		15%
	5.8	Service & Support performance		13%
	5.9	Solution Validation performance		6%
				100%
6		**Financial value and competitive pricing offered**	**25%**	
	6.1	Total amount of the Tender		35%
	6.2	Relative Price (excl. VAT) within Lot 1 per (theoretical) perf. [€/TFlop]		25%
	6.3	Relative Price (excl. VAT) within Lot 1 per computing core		5%
	6.4	Relative Price (excl. VAT) within Lot 2 per (effective) TB [€/TB]		15%
	6.5	Relative Price (excl. VAT) within Lot 3 per GB/s of Infiniband network bisection bandwidth		5%
	6.6	Relative Price (excl. VAT) within Lot 3 per TB/s of Ethernet equipment aggregated bandwidth		5%
	6.7	[Binding] Price for an additional year of service (warranty and support) (year 5)		5%
	6.8	[Binding] Price for a second additional year of service (warranty and support) (year 6)		5%
				100%

Notable Events/Activities (2019)

... since last HPC School 2019!

- **July-Sept 2019**: Preparation of Uni.lu tender **RFP 190027: New DLC Cluster Aion**
  - ↪ 116 pages (new template), 430 criteria evaluated, **Target Budget ≃ 3 M€(HT)**
    - ✓ Lot 1: DLC Computing cluster `aion`
    - ✓ Lot 2: Adaptation and extension of the existing High-Performance Storage systems
    - ✓ Lot 3: Adaptation of the network (Ethernet and IB)
  - ↪ **Sept 11, 2019**: official public release
    - ✓ TED72/2019-608787 / PMP Portal n°1901442

- **July-Aug 2019**: PRACE Summer of HPC 2019
  - ↪ S. Mahon (Ireland), M. Stringher (Italy) welcome in ULHPC Team
    - ✓ perf. evaluation deep learning frameworks and energy efficiency

Notable Events/Activities (2019)

... since last HPC School 2019!

- **July-Sept 2019**: Preparation of Uni.lu tender **RFP 190027: New DLC Cluster Aion**
  - ↪ 116 pages (new template), 430 criteria evaluated, **Target Budget ≃ 3 M€(HT)**
    - ✓ Lot 1: DLC Computing cluster `aion`
    - ✓ Lot 2: Adaptation and extension of the existing High-Performance Storage systems
    - ✓ Lot 3: Adaptation of the network (Ethernet and IB)
  - ↪ **Sept 11, 2019**: official public release
    - ✓ TED72/2019-608787 / PMP Portal n°1901442
- **July-Aug 2019**: PRACE Summer of HPC 2019
  - ↪ S. Mahon (Ireland), M. Stringher (Italy) welcome in ULHPC Team
    - ✓ perf. evaluation deep learning frameworks and energy efficiency
- **Sept 2019**: Changes in ULHPC Team
  - ↪ Clément Parisot left to tackle a new challenge in Nancy
  - ↪ Dr. Ezhilmathi Krishnasamy joined (Postdoc researcher)
    - ✓ Coordinator PRACE-6IP project

Notable Events/Activities (2019)

... since last HPC School 2019!

- **Oct-Nov 2019**:
  - ↪ *(transparent)* Migration of ULHPC Identity Management, HA RedHat IdM/IPA-based

Notable Events/Activities (2019)

... since last HPC School 2019!

- **Oct-Nov 2019**:
  - ↪ *(transparent)* Migration of ULHPC Identity Management, HA RedHat IdM/IPA-based
  - ↪ **Oct 29, 2019**: Official opening of RFP 190027 (Aion supercomputer) proposals by vendors
    - ✓ initiate offers analysis by ULHPC team with procurement/legal dept. (concluded Dec. 2019)

Notable Events/Activities (2019)

... since last HPC School 2019!

- **Oct-Nov 2019**:
  - ↪ *(transparent)* Migration of ULHPC Identity Management, HA RedHat IdM/IPA-based
  - ↪ **Oct 29, 2019**: Official opening of RFP 190027 (Aion supercomputer) proposals by vendors
    - ✓ initiate offers analysis by ULHPC team with procurement/legal dept. (concluded Dec. 2019)
  - ↪ **35th PRACE Council meeting** held in Luxembourg (Neimenster Abbey) (Nov 27-28, 2019)
    - ✓ organized by P. Bouvry, S. Varrette and M. Martin
    - ✓ IAC, SSC, General assembly meetings, EuroHPC updates by T. Skordas, G Kalbe & H. Zeisel

Notable Events/Activities (2019)

... since last HPC School 2019!

- **Oct-Nov 2019**:
  - ↪ *(transparent)* Migration of ULHPC Identity Management, HA RedHat IdM/IPA-based
  - ↪ **Oct 29, 2019**: Official opening of RFP 190027 (Aion supercomputer) proposals by vendors
    - ✓ initiate offers analysis by ULHPC team with procurement/legal dept. (concluded Dec. 2019)
  - ↪ **35th PRACE Council meeting** held in Luxembourg (Neimenster Abbey) (Nov 27-28, 2019)
    - ✓ organized by P. Bouvry, S. Varrette and M. Martin
    - ✓ IAC, SSC, General assembly meetings, EuroHPC updates by T. Skordas, G Kalbe & H. Zeisel
  - ↪ **2nd EU-ASEAN HPC Coordination Group Meeting** (Nov 29-30, 2019)
    - ✓ EC (DG RTD) EU-ASEAN HPC Mapping study *(S. Varrette, I.F. Sulaiman)* endorsed


Notable Events/Activities (2019)

... since last HPC School 2019!

- **Oct-Nov 2019**:
  - ↪ *(transparent)* Migration of ULHPC Identity Management, HA RedHat IdM/IPA-based
  - ↪ **Oct 29, 2019**: Official opening of RFP 190027 (Aion supercomputer) proposals by vendors
    - ✓ initiate offers analysis by ULHPC team with procurement/legal dept. (concluded Dec. 2019)
  - ↪ **35th PRACE Council meeting** held in Luxembourg (Neimenster Abbey) (Nov 27-28, 2019)
    - ✓ organized by P. Bouvry, S. Varrette and M. Martin
    - ✓ IAC, SSC, General assembly meetings, EuroHPC updates by T. Skordas, G Kalbe & H. Zeisel
  - ↪ **2nd EU-ASEAN HPC Coordination Group Meeting** (Nov 29-30, 2019)
    - ✓ EC (DG RTD) EU-ASEAN HPC Mapping study *(S. Varrette, I.F. Sulaiman)* endorsed
- **Dec 2019**
  - ↪ ULHPC Software Environment **2019a** released
    - ✓ became production starting Jan 15, 2020

Notable Events/Activities (2019)

... since last HPC School 2019!

- **Oct-Nov 2019**:
  - ↪ *(transparent)* Migration of ULHPC Identity Management, HA RedHat IdM/IPA-based
  - ↪ **Oct 29, 2019**: Official opening of RFP 190027 (Aion supercomputer) proposals by vendors
    - ✓ initiate offers analysis by ULHPC team with procurement/legal dept. (concluded Dec. 2019)
  - ↪ **35th PRACE Council meeting** held in Luxembourg (Neimenster Abbey) (Nov 27-28, 2019)
    - ✓ organized by P. Bouvry, S. Varrette and M. Martin
    - ✓ IAC, SSC, General assembly meetings, EuroHPC updates by T. Skordas, G Kalbe & H. Zeisel
  - ↪ **2nd EU-ASEAN HPC Coordination Group Meeting** (Nov 29-30, 2019)
    - ✓ EC (DG RTD) EU-ASEAN HPC Mapping study *(S. Varrette, I.F. Sulaiman)* endorsed
- **Dec 2019**
  - ↪ ULHPC Software Environment **2019a** released
    - ✓ became production starting Jan 15, 2020
  - ↪ Software Carpentry Life Science Workshop
    - ✓ practical application of Python, conda and panda, by S. Peter

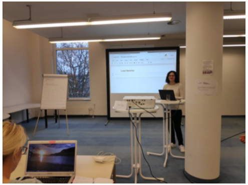

Notable Events/Activities (2019)

... since last HPC School 2019!

- **Dec 2019**
  - ↪ gaia and chaos clusters officially **decommissioned** (process initiated since Feb 2019)
    - ✓ After 8 (resp. 12) years of good & faithful service
    - ✓ 6.2 million jobs were processed, cumulating **13,8 MILLENIUM of CPU Time usage**

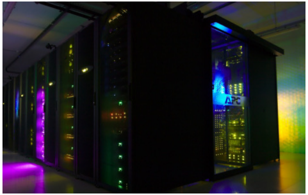

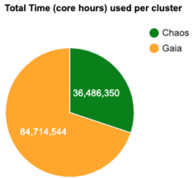

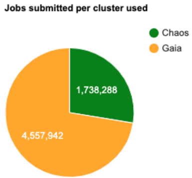

Notable Events/Activities (2019)

... since last HPC School 2019!

- **Dec 2019**
  - ↪ `gaia` and `chaos` clusters officially **decommissioned** (process initiated since Feb 2019)
    - ✓ After 8 (resp. 12) years of good & faithful service
    - ✓ 6.2 million jobs were processed, cumulating **13,8 MILLENIUM of CPU Time usage**
  - ↪ **RFP 190027 attributed to Atos**
    - ✓ **318 compute nodes**, 2xAMD Epyc 7H12 @ 2.6GHz (2x64c), 256GB RAM
    - ✓ Fast Interconnect: **IB HDR** Fabric in a Fat tree topology (2:1 blocking)
    - ✓ **GPFS extension** by 1720TB/1560TiB to reach a total of **4.41 PB**
    - ✓ Adaptation of the network to bind `iris` and `aion` island (Ethernet and IB)
    - ✓ Delivery and Installation planned April 2020

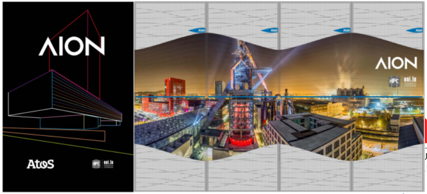

Notable Events/Activities (2019)

... since last HPC School 2019!

- **Dec 2019**
  - ↪ gaia and chaos clusters officially **decommissioned** (process initiated since Feb 2019)
    - ✓ After 8 (resp. 12) years of good & faithful service
    - ✓ 6.2 million jobs were processed, cumulating **13,8 MILLENIUM of CPU Time usage**
  - ↪ **RFP 190027 attributed to Atos**
    - ✓ **318 compute nodes**, 2xAMD Epyc 7H12 @ 2.6GHz (2x64c), 256GB RAM
    - ✓ Fast Interconnect: **IB HDR** Fabric in a Fat tree topology (2:1 blocking)
    - ✓ **GPFS extension** by 1720TB/1560TiB to reach a total of **4.41 PB**
    - ✓ Adaptation of the network to bind iris and aion island (Ethernet and IB)
    - ✓ Delivery and Installation ~~planned April 2020~~ **started Dec 2020 (COVID crisis impact)**

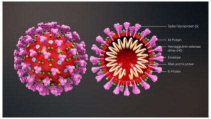

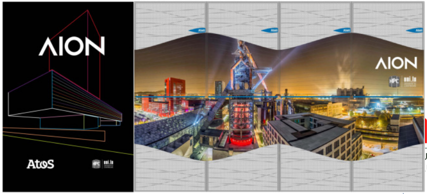

Notable Events/Activities (2020)

... since last HPC School 2019!

- **Jan-Feb 2020**
  - ↪ Kickoff meeting with Atos team for Aion deployment
  - ↪ Updated Acceptable Use Policy (AUP) 2.0

University of Luxembourg High Performance Computing Acceptable Use Policy

Introduction

The University of Luxembourg operates since 2007 a large academic HPC facility which remains the reference implementation within the country, offering a cutting-edge research infrastructure to Luxembourg public research while serving as edge access to the upcoming Euro-HPC Luxembourg supercomputer. Special focus was laid on the development of large computing power combined with huge data storage capacity to accelerate the research performed in intensive computing and large-scale data analytics (Big Data).

The University extends access to its HPC resources (including facilities, services and HPC experts) to its students, staff, research partners (including scientific staff of national public organizations and external partners for the duration of joint research projects) and to industrial partners.

All users of UL HPC resources and PIs must abide by the following policies.

Computing systems use

Use of UL HPC computing resources should be used only for work directly related to the projects for which the resources were requested and granted, and primarily to advance University's missions of education and research. Use of UL HPC computing resources for personal activities is prohibited.

The UL HPC Team maintains up-to-date documentation on its computational resources and their proper use, and provides regular training and constant support to users. Users assume the responsibility for following the documentation, training sessions and best practice guides in order to understand the proper and considerate use of the UL HPC computing resources.

Notable Events/Activities (2020)

... since last HPC School 2019!

- **Jan-Feb 2020**
  - ↪ Kickoff meeting with Atos team for Aion deployment
  - ↪ Updated Acceptable Use Policy (AUP) 2.0
  - ↪ New **HPC Cost Model**, [FNR] Funded projects & Externals
    - ✓ Approved by Rectorate on July 7, 2020

High Performance Computing Resource Allocations for Research Projects and External Partners

OVERVIEW AND GUIDELINES

Executive Summary

With the advent of the technological revolution and the digital transformation that made all scientific disciplines becoming computational nowadays, High-Performance Computing (HPC) is increasingly identified as a strategic asset and enabler to accelerate the research performed in all areas requiring intensive computing and large-scale Big Data analytic capabilities.
The University of Luxembourg (UL) operates since 2007 a large academic HPC facility that remains the reference implementation within the country, offering a cutting-edge research infrastructure to Luxembourg public research while serving as edge access to the upcoming Euro-HPC Luxembourg supercomputer. Special focus was laid on the development of large computing power combined with huge data storage capacity to accelerate the research performed in intensive computing and large-scale data analytic (Big Data).
For more details: see hpc.uni.lu (Main contacts: Prof. Pascal Bouvry (Head), Dr. Sébastien Varrette (Deputy head), HPC for research).

The University extends access to its HPC resources (i.e., facility and expert HPC consultants) to the scientific staff of national public organizations and external partners for the duration of joint research projects under the conditions defined in this document.

Version 1.0 - Approved by UL rectorate in July 7, 2020

Notable Events/Activities (2020)

... since last HPC School 2019!

- **Jan-Feb 2020**
  - ↪ Kickoff meeting with Atos team for Aion deployment
  - ↪ Updated Acceptable Use Policy (AUP) 2.0
  - ↪ New **HPC Cost Model**, [FNR] Funded projects & Externals
    - ✓ Approved by Rectorate on July 7, 2020
- **Mar-June 2020**
  - ↪ Changes in ULHPC Team
    - ✓ V. Plugaru left us to join LuxProvide as CTO — Mar 2020
    - ✓ Prof. P. Bouvry started temporary mandate as CEO LuxProvide — June 2020

Notable Events/Activities (2020)

... since last HPC School 2019!

- **Jan-Feb 2020**
  - ↪ Kickoff meeting with Atos team for Aion deployment
  - ↪ Updated Acceptable Use Policy (AUP) 2.0
  - ↪ New **HPC Cost Model**, [FNR] Funded projects & Externals
    - ✓ Approved by Rectorate on July 7, 2020
- **Mar-June 2020**
  - ↪ Changes in ULHPC Team
    - ✓ V. Plugaru left us to join LuxProvide as CTO
    - ✓ Prof. P. Bouvry started temporary mandate as CEO LuxProvide
  - ↪ **COVID-19 pandemic and global lockdown**
    - ✓ University's supercomputer supports fight against COVID-19

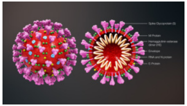

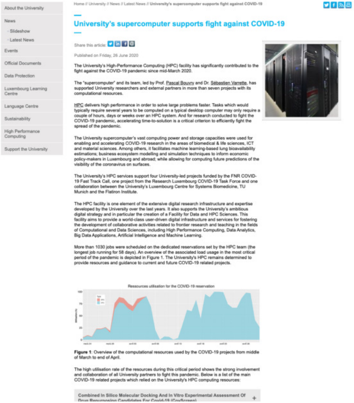

About the University
News
Slideshow
Latest News
Events
Official Documents
Data Protection
Luxembourg Learning Centre
Language Centre
Sustainability
High Performance Computing
Support the University

Home // University // News // Latest News // University's supercomputer supports fight against COVID-19

University's supercomputer supports fight against COVID-19

Share this article:

Published on Friday, 26 June 2020

The University's High-Performance Computing (HPC) facility has significantly contributed to the fight against the COVID-19 pandemic since mid-March 2020.

The "supercomputer" and its team, led by Prof. Pascal Bouvry and Dr. Sébastien Varrette, has supported University researchers and external partners in more than seven projects with its computational resources.

HPC delivers high performance in order to solve large problems faster. Tasks which would typically require several years to be computed on a typical desktop computer may only require a couple of hours, days or weeks over an HPC system. And for research conducted to fight the COVID-19 pandemic, accelerating time-to-solution is a critical criterion to efficiently fight the spread of the pandemic.

The University supercomputer's vast computing power and storage capacities were used for enabling and accelerating COVID-19 research in the areas of biomedical & life sciences, ICT and material sciences. Among others, it facilitates machine learning-based lung bioavailability estimations; business ecosystem modelling and simulation techniques to inform economic policy-makers in Luxembourg and abroad; while allowing for computing future predictions of the visibility of the coronavirus on surfaces.

The University's HPC services support four University-led projects funded by the FNR COVID-19 Fast Track Call, one project from the Research Luxembourg COVID-19 Task Force and one collaboration between the University's Luxembourg Centre for Systems Biomedicine, TU Munich and the Flatiron Institute.

The HPC facility is one element of the extensive digital research infrastructure and expertise developed by the University over the last years. It also supports the University's ambitious digital strategy and in particular the creation of a Facility for Data and HPC Sciences. This facility aims to provide a world-class user-driven digital infrastructure and services for fostering the development of collaborative activities related to frontier research and teaching in the fields of Computational and Data Sciences, including High Performance Computing, Data Analytics, Big Data Applications, Artificial Intelligence and Machine Learning.

More than 1030 jobs were scheduled on the dedicated reservations set by the HPC team (the longest job running for 58 days). An overview of the associated load usage in the most critical period of the pandemic is depicted in Figure 1. The University's HPC remains determined to provide resources and guidance to current and future COVID-19 related projects.

Ressources utilisation for the COVID-19 reservation

Figure 1: Overview of the computational resources used by the COVID-19 projects from middle of March to end of April.

The high utilisation rate of the resources during this critical period shows the strong involvement and collaboration of all University partners to fight this pandemic. Below is a list of the main COVID-19 related projects which relied on the University's HPC computing resources:

Combined In Silico Molecular Docking And In Vitro Experimental Assessment Of Drug Repurposing Candidates For Covid-19 (CovScreen) +

Notable Events/Activities (2020)

... since last HPC School 2019!

- **April-June 2020**: Reworked model for HPC User Software management – **RESIF 3.0**
  - ↪ RESIF 2 with a complex workflow, not distributed, and too many "custom" easyconfigs
    - ✓ broken compliance & divergence from streamline `easybuilders/easybuild-easyconfigs`
    - ✓ never any contribution back to streamline Easybuild
  - ↪ **New model more robust**, **architecture optimized**, 91% reduction of custom EB
    - ✓ Module path: `/opt/apps/resif/<cluster>/<version>/<arch>/modules/all`
  - ↪ **2019b release** rebuild in Nov, **in testing mode until Jan. 31th, 2020**
    - ✓ 2020a in progress (required for epyc), pending Aion deployment

Name	**Type**	2019[a] (prod/old)	**2019b** (devel)	**2020a** (next)
`GCCCore`	compiler	8.2.0	8.3.0	9.3.0
`foss`	toolchain	2019a	2019b	2020a
`intel`	toolchain	2019a	2019b	2020a
Python		3.7.2 (and 2.7.15)	3.7.4 (and 2.7.16)	3.8.2

Notable Events/Activities (2020)

... since last HPC School 2019!

- **April-June 2020**: Reworked model for HPC User Software management – **RESIF 3.0**
  - ↪ RESIF 2 with a complex workflow, not distributed, and too many "custom" easyconfigs
    - ✓ broken compliance & divergence from streamline `easybuilders/easybuild-easyconfigs`
    - ✓ never any contribution back to streamline Easybuild
  - ↪ **New model more robust**, **architecture optimized**, 91% reduction of custom EB
    - ✓ Module path: `/opt/apps/resif/<cluster>/<version>/<arch>/modules/all`
  - ↪ **2019b release** rebuild in Nov, **in testing mode until Jan. 31th, 2020**
    - ✓ 2020a in progress (required for epyc), pending Aion deployment

```
# (new) 2019b software set - iris cluster
unset MODULEPATH
module use /opt/apps/resif/iris/2019b/broadwell/modules/all
# OR (when appropriate) skylake/GPU-specialized builds **ONLY** on skylake/GPU nodes
module use /opt/apps/resif/iris/2019b/skylake/modules/all
module use /opt/apps/resif/iris/2019b/gpu/modules/all
```

Notable Events/Activities (2020)

... since last HPC School 2019!

- **April-June 2020**: Reworked model for HPC User Software management – **RESIF 3.0**
  - ↪ RESIF 2 with a complex workflow, not distributed, and too many "custom" easyconfigs
    - ✓ broken compliance & divergence from streamline `easybuilders/easybuild-easyconfigs`
    - ✓ never any contribution back to streamline Easybuild
  - ↪ **New model more robust**, **architecture optimized**, 91% reduction of custom EB
    - ✓ Module path: `/opt/apps/resif/<cluster>/<version>/<arch>/modules/all`
  - ↪ **2019b release** rebuild in Nov, **in testing mode until Jan. 31th, 2020**
    - ✓ 2020a in progress (required for epyc), pending Aion deployment

```
# INCOMING in 2021: (new) 2020a software set - aion cluster
module use /opt/apps/resif/aion/2020a/epyc/modules/all
# iris cluster
module use /opt/apps/resif/iris/2020a/{broadwell,skylake,gpu}/modules/all
```

Notable Events/Activities (2020)

... since last HPC School 2019!

- **June-Sept 2020**:
  - ↪ **Slurm configuration update 2.0** **Effective since Oct. 24, 2020**
    - ✓ Global accounting and billing implementing new HPC Cost Model and `aion` integration
    - ✓ Updated model for Fairshare, Account Hierarchy and Limits for Iris/Aion

Notable Events/Activities (2020)

... since last HPC School 2019!

- **June-Sept 2020**:
  - ↪ **Slurm configuration update 2.0** — **Effective since Oct. 24, 2020**
    - ✓ Global accounting and billing implementing new HPC Cost Model and `aion` integration
    - ✓ Updated model for Fairshare, Account Hierarchy and Limits for Iris/Aion
  - ↪ **New equipment for G5K** and transfer to CDC

Notable Events/Activities (2020)

... since last HPC School 2019!

- **June-Sept 2020**:
  - ↪ **Slurm configuration update 2.0** **Effective since Oct. 24, 2020**
    - ✓ Global accounting and billing implementing new HPC Cost Model and `aion` integration
    - ✓ Updated model for Fairshare, Account Hierarchy and Limits for Iris/Aion
  - ↪ **New equipment for G5K** and transfer to CDC
  - ↪ *Partial* delivery `Aion` (**Ethernet** interconnect equipment)
    - ✓ **CDC HPC Ethernet Network 2.0** setup

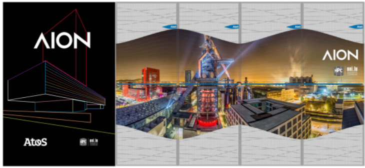

Notable Events/Activities (2020)

... since last HPC School 2019!

- **June-Sept 2020**:
  - ↪ **Slurm configuration update 2.0** — **Effective since Oct. 24, 2020**
    - ✓ Global accounting and billing implementing new HPC Cost Model and `aion` integration
    - ✓ Updated model for Fairshare, Account Hierarchy and Limits for Iris/Aion
  - ↪ **New equipment for G5K** and transfer to CDC
  - ↪ *Partial* delivery `Aion` (**Ethernet** interconnect equipment)
    - ✓ **CDC HPC Ethernet Network 2.0** setup
  - ↪ Support tracker **migration to Uni.lu Service NOW**
    - ✓ `hpc-tracker.uni.lu` portal decommissionned
    - ✓ **1178** issues/support tickets managed since 2013 (8 years)
    - ✓ **half L3/L4 support level**,
      ... **YET** too many un-tracked effort by direct mails answers
    - ✓ ULHPC users requested to use Service Now (HPC)
    - ✓ **Effective since Oct 5, 2020**

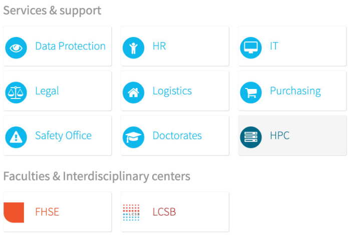

Notable Events/Activities (2020)

... since last HPC School 2019!

- **Sept-Dec 2020**:
  - ↪ Changes in ULHPC Team:
    - ✓ T. Valette and A. Olloh joined the HPC Team as Infr. & HPC Engineers — Sept 2020
    - ✓ A. Moinier-Vandeventer *(Project Manager)*, Dr. L. Koutsantonis *(PostDoc)* EuroCC — Oct 2020

Notable Events/Activities (2020)

... since last HPC School 2019!

- **Sept-Dec 2020**:
  - ↪ Changes in ULHPC Team:
    - ✓ T. Valette and A. Olloh joined the HPC Team as Infr. & HPC Engineers — Sept 2020
    - ✓ A. Moinier-Vandeventer *(Project Manager)*, Dr. L. Koutsantonis *(PostDoc)* EuroCC — Oct 2020
  - ↪ **ULHPC Websites 2.0 and Technical Documentation** — production release postponed EOY

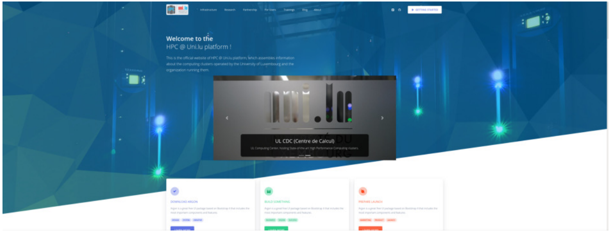

Notable Events/Activities (2020)

... since last HPC School 2019!

- **Sept-Dec 2020**:
  - ↪ Changes in ULHPC Team:
    - ✓ T. Valette and A. Olloh joined the HPC Team as Infr. & HPC Engineers — Sept 2020
    - ✓ A. Moinier-Vandeventer *(Project Manager)*, Dr. L. Koutsantonis *(PostDoc)* EuroCC — Oct 2020
  - ↪ **ULHPC Websites 2.0 and Technical Documentation** — production release postponed EOY

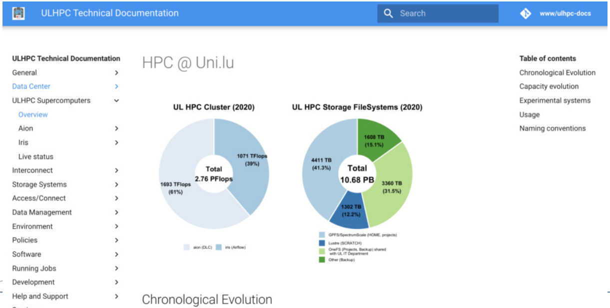

Notable Events/Activities (2020)

... since last HPC School 2019!

- **Sept-Dec 2020**:
  - ↪ Changes in ULHPC Team:
    - ✓ T. Valette and A. Olloh joined the HPC Team as Infr. & HPC Engineers — Sept 2020
    - ✓ A. Moinier-Vandeventer *(Project Manager)*, Dr. L. Koutsantonis *(PostDoc)* EuroCC — Oct 2020
  - ↪ **ULHPC Websites 2.0 and Technical Documentation** — production release postponed EOY
  - ↪ *Partial* delivery Aion (**Servers, Storage equipment**) — Nov 2020
    - ✓ Main **final** delivery & installation expected in the coming weeks — Jan 2021

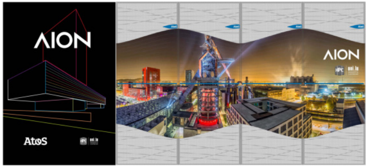




Summary

General Guidelines

Acceptable Use Policy (AUP) 2.0

Uni.lu-HPC-Facilities_Acceptable-Use-Policy_v2.0.pdf

- **UL HPC is a shared** *(and expansive)* **facility: you must practice good citizenship**
  - ↪ **Users are accountable for their actions**
    - ✓ Users are allowed **one account per person** - **user credentials sharing is strictly prohibited**
    - ✓ Use of UL HPC computing resources for personal activities is prohibited
    - ✓ limit activities that may impact the system for other users.
  - ↪ **Do not abuse the shared filesystems**
    - ✓ Avoid too many simultaneous file transfers
    - ✓ regularly clean your directories from useless files
  - ↪ **Do not run programs or I/O bound processes on the login nodes**
  - ↪ Plan large scale experiments during night-time or week-ends
- Resource allocation is done on a **fair-share** principle, with **no guarantee** of being satisfied

General Guidelines

Acceptable Use Policy (AUP) 2.0

- **Data Use / GDPR**
  - ↪ **You** are responsible to ensure the appropriate level of protection, backup & integrity checks
    - ✓ Data Authors/generators/owners are responsible for its correct categorization as sensitive/non-sensitive
    - ✓ Owners of sensitive information are responsible for its secure handling, transmission, processing, storage, and disposal on the UL HPC systems
    - ✓ Data Protection inquiries can be directed to the Uni.lu Data Protection Officer
  - ↪ We make **no guarantee** against loss of data
- We provide [project] **usage report** to user/PI **on-demand** and *(by default)* on a **yearly basis**

- For **ALL** publications having results produced using the UL HPC Facility
  - ↪ **Acknowledge** the UL HPC facility and **cite** reference ULHPC article
    - ✓ using official banner
  - ↪ Tag your publication upon registration on ORBiLu.

ULHPC Websites 2.0 and Documentation

Main Website

`hpc.uni.lu`

ULHPC Tutorials

`ulhpc-tutorials.rtfd.io`

ULHPC Technical Docs

`hpc-docs.uni.lu`

ULHPC HelpDesk

`hpc.uni.lu/support`

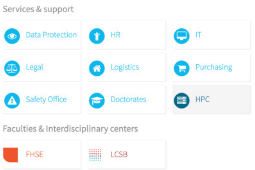

- **Fallback Support**:
  - ↪ `hpc-team@uni.lu`
  - ↪ ULHPC Community: `hpc-users@uni.lu`
    - ✓ moderated

Compute Nodes / Storage Environment

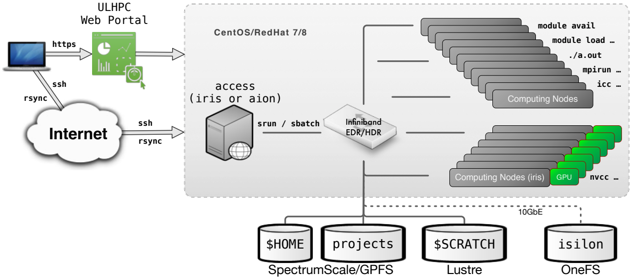

- **Storage usage**: `df-ulhpc [-i]`
  - ↪ `$HOME`: regular backup policy
  - ↪ `$SCRATCH` **NO** backup & purged
    - ✓ 60 days retention policy
  - ↪ Project quotas attached to group
    - ✓ **not** *(default)* `clusterusers` group
    - ✓ Commands writing in project dir: `sg <group> -c "<command>"`
- **LMod/Environment modules**
  - ↪ **Not** on access, **only** on compute nodes

Directory	FileSystem	Max size	Max #files	Backup
`$HOME` (iris)	GPFS	500 GB	1.000.000	YES
`$SCRATCH`	Lustre	10 TB	1.000.000	**NO**
Project	GPFS	*per request*		PARTIALLY (`/backup` subdir)
Project	OneFS	*per request*		PARTIALLY

Reporting Problems

- **First checks**
  1. My issue is probably documented https://hpc-docs.uni.lu
  2. An event is on-going: **check ULHPC Live status page** https://hpc.uni.lu/live-status/motd/
     - ✓ Planned maintenance are announced *at least* 2 weeks in advance
     - ✓ The proper SSH banner is displayed during **planned** downtime
  3. check the state of your nodes
     - ✓ `{ scontrol show job <jobid> | sjoin <jobid>}; htop` *on active jobs*
     - ✓ `{ slist <jobid> | sacct [-X] -j <jobid> -l }` *post-mortem*

ULHPC Helpdesk / Support Ticket Portal

Reporting Problems

- **First checks**
  1. My issue is probably documented — https://hpc-docs.uni.lu
  2. An event is on-going: **check ULHPC Live status page** — https://hpc.uni.lu/live-status/motd/
     - ✓ Planned maintenance are announced *at least* 2 weeks in advance
     - ✓ The proper SSH banner is displayed during **planned** downtime
  3. check the state of your nodes
     - ✓ `{ scontrol show job <jobid> | sjoin <jobid>}; htop` — *on active jobs*
     - ✓ `{ slist <jobid> | sacct [-X] -j <jobid> -l }` — *post-mortem*

- **ONLY NOW**, consider the following depending on the severity:
  - ↪ Open an new issue on https://hpc.uni.lu/support — **(preferred)**
    - ✓ Uni.lu Service Now Helpdesk Portal: relies on **Uni.lu (≠ ULHPC) credentials**
  - ↪ Mail (only now) us — `hpc-team@uni.lu`
  - ↪ **Ask the help of other users** — `hpc-users@uni.lu`

Reporting Problems

- **First checks**
  1. My issue is probably documented `https://hpc-docs.uni.lu`
  2. An event is on-going: **check ULHPC Live status page** `https://hpc.uni.lu/live-status/motd/`
     - ✓ Planned maintenance are announced *at least* 2 weeks in advance
     - ✓ The proper SSH banner is displayed during **planned** downtime
  3. check the state of your nodes
     - ✓ `{ scontrol show job <jobid> | sjoin <jobid>}; htop` *on active jobs*
     - ✓ `{ slist <jobid> | sacct [-X] -j <jobid> -l }` *post-mortem*

- **ONLY NOW**, consider the following depending on the severity:
  - ↪ Open an new issue on `https://hpc.uni.lu/support` **(preferred)**
    - ✓ Uni.lu Service Now Helpdesk Portal: relies on **Uni.lu (≠ ULHPC) credentials**
  - ↪ Mail (only now) us `hpc-team@uni.lu`
  - ↪ **Ask the help of other users** `hpc-users@uni.lu`

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

Summary

Interactive Jobs

```
# BEFORE
srun -p interactive --qos qos-interactive -C {broadwell|skylake} [...] --pty bash`
# AFTER -- match feature name with target partition
srun -p interactive --qos debug -C {batch,gpu,bigmem} [...] --pty bash
```

- **Before**: guaranteed access to interative jobs on regular nodes even if `batch` partition full
  ↪ **YET** no way to use `qos-interactive` for GPU/bigmem
    - ✓ default node category QOS/partition used, inherits from default limits
    - ✓ `srun -p gpu --qos qos-gpu -G 4 [...] --pty bash` can stay 5 days in a screen
- **After**: no guarantee if partition is full **YET** backfilling and priority ensure first served

Node Type	Slurm command	Helper script
regular	`srun -p interactive --qos debug -C batch [-C {broadwell,skylake}] [...] --pty bash`	`si [...]`
gpu	`srun -p interactive --qos debug -C gpu [-C volta[32]] -G 1 [...] --pty bash`	`si-gpu [...]`
bigmem	`srun -p interactive --qos debug -C bigmem [...] --pty bash`	`si-bigmem [...]`

Regular Jobs

- **NO MORE `qos-*` QOS**
  - ↪ **ALL slurm launchers to review to remove/adapt QOS attributes**
  - ↪ all default to `normal` QOS, **except** CRP/externals who default to `low`
  - ↪ thus no need to precise, except to access higher priority QOS if allowed
    - ✓ Ex: `#SBATCH --qos high`
- **NEW: Add `-A <project|lecture>` account when appropriate!**
  - ↪ Non-default L3 meta-account used:
    - ✓ project name `<project>`
    - ✓ lecture/course name: `<lecture>`

```
   #SBATCH -p batch              #SBATCH -p gpu              #SBATCH -p bigmem
-- #SBATCH --qos qos-batch    -- #SBATCH --qos qos-gpu    -- #SBATCH --qos qos-bigmem
++ #SBATCH -A <project>       ++ #SBATCH -A <project>     ++ #SBATCH -A <project>
```

Regular Jobs

- Relatively similar as before, **YET** now restricted to **Max 2 days** / Max 64 nodes
  - ↪ walltime reduction would have affected 1.22% of the jobs completed since July, 1st 2020
  - ↪ default QOS induced by the `job_submit.lua` plugin as before
  - ↪ enforce precision of project/training account (`-A <account>`)

Node Type	Slurm command
regular	`sbatch [-A <project>] -p batch [--qos {high,urgent}] [-C {broadwell,skylake}] [...]`
gpu	`sbatch [-A <project>] -p gpu [--qos {high,urgent}] [-C volta[32]] -G 1 [...]`
bigmem	`sbatch [-A <project>] -p bigmem [--qos {high,urgent}] [...]`

- Slurm Federation configuration between `iris` and `aion`
  - ↪ ensures global policy (coherent job ID, global scheduling, etc.) within ULHPC systems
  - ↪ easily submit jobs from one cluster to another `-M, --cluster aion|iris`

```
# Ex (from iris): try first on iris, then on aion
sbatch -p batch -M iris,aion [...]
```

Long Jobs

```
# BEFORE - only on regular nodes
sbatch -p long --qos qos-long [...]
# AFTER -- select target partition to bypass default walltime restrictions
sbatch -p {batch | gpu | bigmem} --qos long [...]
```

- **Before**: extended Max walltime (`MaxWall`) set to **30 days, restricted to regular nodes**
  - ↪ Max 6 nodes, Max 2 nodes per Job, Max 10 Jobs per User
  - ↪ No way to run long jobs on GPU or Large-Memory nodes
- **After**: extended Max walltime (`MaxWall`) set to **14 days** EuroHPC/PRACE Recommendations
  - ↪ Max 6 nodes, Max 2 nodes per Job, Max **1** Job per User

Node Type	Slurm command
regular	`sbatch [-A <project>] -p batch --qos long [-C {broadwell,skylake}] [...]`
gpu	`sbatch [-A <project>] -p gpu --qos long [-C volta[32]] -G 1 [...]`
bigmem	`sbatch [-A <project>] -p bigmem --qos long [...]`

Other Misc Changes

- (complex) Depth-Oblivious Fairshare ⟹ **Fair tree** Algorithm
- Special **preemptible QOS** kept for **best-effort Jobs** **YET** renamed: ~~qos-~~besteffort
  - ↪ `sbatch -p {batch | gpu | bigmem} --qos besteffort [...]`
- **NO MORE dedicated QOS** `qos-batch-00*` but global **restricted** `high`**(priority) QOS**
  - ↪ Incentives for User groups/Projects contributing to the HPC budget line
    - ✓ **Updated every year** based on past funding amount and depreciation (default: 12 months)
    - ✓ Affect raw share for the L2/L3 account

$$FundingScore(Year) = \left\lfloor \alpha_{\text{level}} \times \frac{Investment(Year-1)}{100 \times \#months} \right\rfloor$$

- **Restricted** `urgent` **QOS** for ultra-high priority jobs (Ex: covid-19)
- End-User raw-share increased based on past year efficiency
  - ↪ Efficiency Score for L4 user, **Average Wall-time Accuracy (WRA)**

Account Hierarchy 2.0

- Every user job runs under a group account
  - ↪ granting access to specific QOS levels.
  - ↪ default raw share for accounts: 1

- **L1**: Organization Level: UL, CRPs, Externals, Projects, Trainings
  - ↪ guarantee 80% of the shares for core UL activities
- **L2**: Organizational Unit (Faculty, ICs, External partner, Funding program. . .)
  - ↪ Raw share depends on outdegree and past year funding
- **L3**: Principal Investigator (PIs), Projects, Course
  - ↪ Raw share depends on past year funding
  - ↪ Eventually restricted **only** to projects and courses
- **L4**: End User (ULHPC login)
  - ↪ Raw share based on efficiency score

Account Hierarchy 2.0

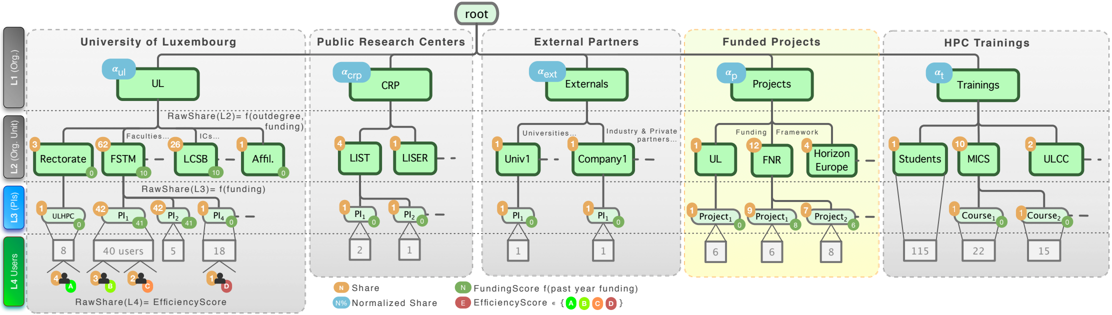

Account Hierarchy 2.0

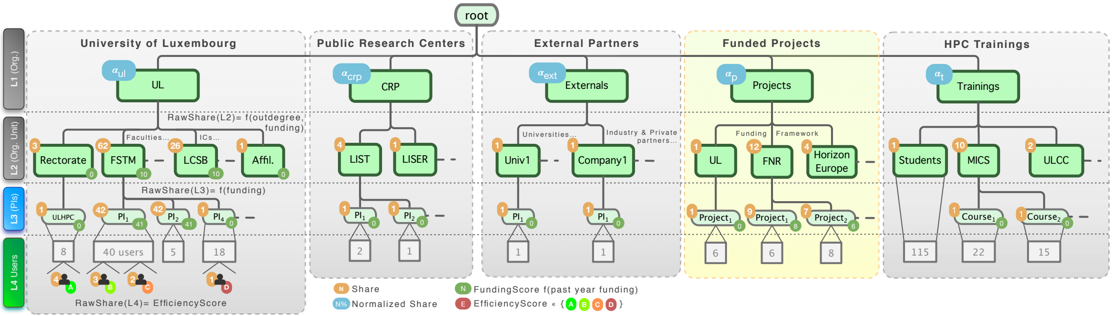


```
# L1,L2 or L3 account /!\ ADAPT <name> accordingly
sacctmgr show association tree where accounts=<name> format=account,share
# End user (L4)
sacctmgr show association where users=$USER  format=account,User,share,Partition,QOS
```

Efficiency Score (L4)

- **Updated every year based on past jobs efficiency.**
  - ↪ Similar notion of "nutri-score': A(very good - 3), B (good: 2), C (bad, 1), D(very bad - 0)
- Proposed Metric for **user** U: **Average Wall-time Accuracy (WRA)** (higher the better)
  - ↪ Defined for a given time period (past year)

```
sacct -u <U> -X -S <start> -E <end> [...] # --format User,JobID,state,time,elapsed
```

  - ↪ Reduction for N `COMPLETED` jobs:

$$S_{\text{efficiency}}(U, Year) = \frac{1}{N} \sum_{JobID \in (U, Year)} \frac{T_{\text{elapsed}}(JobID)}{T_{\text{asked}}(JobID)}$$

- Default thresholds

Score	Avg. WRA
A	$S_{\text{efficiency}} \geq 75\%$
B	$50\% \leq S_{\text{efficiency}} < 75\%$
C	$25\% \leq S_{\text{efficiency}} < 50\%$
D	$S_{\text{efficiency}} < 25\%$

- **WIP**: integrate other efficiency metrics (CPU, mem, GPU efficiency)

Job Priority, Fairsharing and Fair Tree

- **Fairsharing**: way of ensuring that users get their appropriate portion of a system
  - ↪ **Share**: portion of the system users have been granted.
  - ↪ **Usage**: amount of the system users have actually **used**.
  - ↪ **Fairshare score**: value the system calculates based off of user's usage.
    - ✓ difference between the portion of the computing resource that has been promised and the amount of resources that has been consumed
  - ↪ **Priority score**: priority assigned based off of the user's fairshare score.
- ULHPC Slurm configuration with Multifactor Priority Plugin and **Fair tree** algorithm
  - ↪ rooted plane tree (rooted ordered tree) being created then sorted by Level Fairshare
  - ↪ All users from a higher priority account receive a higher fair share factor than all users from a lower priority account

```
$> sshare -l                                                    # See Level FS
```

ULHPC Job Prioritization Factors

- **Age**: length of time a job has been waiting (PD state) in the queue
- **Fairshare**: difference between the portion of the computing resource that has been promised and the amount of resources that has been consumed
- **Partition**: factor associated with each node partition
  - ↪ Ex: privilege `interactive` over `batch`
- **QOS** A factor associated with each Quality Of Service (`low` ⟶ `urgent`)

```
Job_priority =
    PriorityWeightAge       * age_factor +
    PriorityWeightFairshare * fair-share_factor+
    PriorityWeightPartition * partition_factor +
    PriorityWeightQOS       * QOS_factor +
    - nice_factor
```

ULHPC Job Prioritization Factors

- **Age**: length of time a job has been waiting (PD state) in the queue
- **Fairshare**: difference between the portion of the computing resource that has been promised and the amount of resources that has been consumed
- **Partition**: factor associated with each node partition
  - ↪ Ex: privilege `interactive` over `batch`
- **QOS** A factor associated with each Quality Of Service (`low` ⟶ `urgent`)

```
Job_priority =
    PriorityWeightAge         * age_factor +
    PriorityWeightFairshare * fair-share_factor+
    PriorityWeightPartition * partition_factor +
    PriorityWeightQOS         * QOS_factor +
    - nice_factor
```

```
# Show current weights
sprio -w
# List pending jobs, sorted by jobid
sprio [-n]
# List pending jobs, sorted by  priority
sprio [-n] -S+Y
sprio [-n] | sort -k 3 -n
sprio [-n] -l | sort -k 4 -n
```

Fairshare Factor and Job billing

- Utilization of the University computational resources is charged in Service Unit (SU)
  - ↪ 1 SU $\simeq$ 1 hour on 1 physical processor core on regular computing node
  - ↪ Usage charged **0,03€ per SU (VAT excluded)** (external partners, funded projects etc.)
- A Job is characterized (and thus billed) according to the following elements:
  - ↪ T_{exec}: Execution time (in hours)
  - ↪ N_{Nodes}: number of computing nodes, and **per node**:
    - ✓ N_{cores}: number of CPU cores allocated per node
    - ✓ Mem: memory size allocated per node, in GB
    - ✓ N_{gpus}: number of GPU allocated per node
  - ↪ associated weighted factors $\alpha_{cpu}, \alpha_{mem}, \alpha_{GPU}$ defined as `TRESBillingWeight` in Slurm
    - ✓ account for consumed resources other than just CPUs
    - ✓ taken into account in fairshare factor
    - ✓ α_{cpu}: normalized relative perf. of CPU processor core (reference: skylake 73,6 GFlops/core)
    - ✓ α_{mem}: inverse of the average available memory size per core
    - ✓ α_{GPU}: weight per GPU accelerator

Fairshare Factor and Job billing

Number of SU associated to a job

$$N_{\text{Nodes}} \times [\alpha_{cpu} \times N_{\text{cores}} + \alpha_{mem} \times Mem + \alpha_{gpu} \times N_{\text{gpus}}] \times T_{\text{exec}}$$

Cluster	Node Type	Partition	#Cores/node	CPU	α_{cpu}	α_{mem}	α_{GPU}
Iris, Aion	Regular	`interactive`	28/128	n/a	0	0	0
Iris	Regular	`batch`	28	`broadwell`	1.0*	$\frac{1}{4} = 0,25$	0
Iris	Regular	`batch`	28	`skylake`	1.0	$\frac{1}{4} = 0,25$	0
Iris	GPU	`gpu`	28	`skylake`	1.0	$\frac{1}{27}$	50
Iris	Large-Mem	`bigmem`	112	`skylake`	1.0	$\frac{1}{27}$	0
Aion	Regular	`batch`	128	`epyc`	0,57	$\frac{1}{1.75}$	0

```
#  Billing rate for running job <jobID>
scontrol show job <jobID> | grep -i billing
# Billing rate for completed job <jobID>
sacct -X --format=AllocTRES%50,Elapsed -j <jobID>
```

Fairshare Factor and Job billing

Number of SU associated to a job

$$N_{\text{Nodes}} \times [\alpha_{cpu} \times N_{\text{cores}} + \alpha_{mem} \times Mem + \alpha_{gpu} \times N_{\text{gpus}}] \times T_{\text{exec}}$$

Cluster	Node Type	Partition	#Cores/node	CPU	α_{cpu}	α_{mem}	α_{GPU}
Iris, Aion	Regular	`interactive`	28/128	`n/a`	0	0	0
Iris	Regular	`batch`	28	`broadwell`	1.0*	$\frac{1}{4} = 0,25$	0
Iris	Regular	`batch`	28	`skylake`	1.0	$\frac{1}{4} = 0,25$	0
Iris	GPU	`gpu`	28	`skylake`	1.0	$\frac{1}{27}$	50
Iris	Large-Mem	`bigmem`	112	`skylake`	1.0	$\frac{1}{27}$	0
Aion	Regular	`batch`	128	`epyc`	0,57	$\frac{1}{1.75}$	0

- Continuous use of **2 regular skylake nodes** (56 cores, 224GB Memory) on `iris` cluster
  - ↪ 28 cores per node, 4 GigaByte RAM per core i.e., 112GB per node
  - ↪ **For 30 days**: 2 nodes$\times[\alpha_{cpu} \times 28 + \alpha_{mem} \times 4 \times 28 + \alpha_{gpu} \times 0]\times$ 30 days$\times$ 24 hours
    - ✓ Total: $2 \times [(1.0 + \frac{1}{4} \times 4) \times 28] \times 720 = 80640$ SU = **2419,2€ VAT excluded**

Fairshare Factor and Job billing

Number of SU associated to a job

$$N_{\text{Nodes}} \times [\alpha_{cpu} \times N_{\text{cores}} + \alpha_{mem} \times Mem + \alpha_{gpu} \times N_{\text{gpus}}] \times T_{\text{exec}}$$

Cluster	Node Type	Partition	#Cores/node	CPU	α_{cpu}	α_{mem}	α_{GPU}
Iris, Aion	Regular	`interactive`	28/128	n/a	0	0	0
Iris	Regular	`batch`	28	`broadwell`	1.0*	$\frac{1}{4} = 0, 25$	0
Iris	Regular	`batch`	28	`skylake`	1.0	$\frac{1}{4} = 0, 25$	0
Iris	GPU	`gpu`	28	`skylake`	1.0	$\frac{1}{27}$	50
Iris	Large-Mem	`bigmem`	112	`skylake`	1.0	$\frac{1}{27}$	0
Aion	Regular	`batch`	128	`epyc`	0,57	$\frac{1}{1.75}$	0

- Continuous use of **2 regular epyc nodes** (256 cores, 448GB Memory) on `aion` cluster
  - ↪ 128 cores per node, 1,75 GigaByte RAM per core i.e., 224 GB per node
  - ↪ **For 30 days**: 2 nodes$\times[\alpha_{cpu} \times 128 + \alpha_{mem} \times 1.75 \times 128 + \alpha_{gpu} \times 0] \times$ 30 days$\times$ 24 hours
    - ✓ Total: $2 \times [(0.57 + \frac{1}{1.75} \times 1.75) \times 128] \times 720 = 289382{,}4$ SU = **8681,47€ VAT excluded**

Fairshare Factor and Job billing

Number of SU associated to a job

$$N_{\text{Nodes}} \times [\alpha_{cpu} \times N_{\text{cores}} + \alpha_{mem} \times Mem + \alpha_{gpu} \times N_{\text{gpus}}] \times T_{\text{exec}}$$

Cluster	Node Type	Partition	#Cores/node	CPU	α_{cpu}	α_{mem}	α_{GPU}
Iris, Aion	Regular	`interactive`	28/128	n/a	0	0	0
Iris	Regular	`batch`	28	`broadwell`	1.0*	$\frac{1}{4} = 0,25$	0
Iris	Regular	`batch`	28	`skylake`	1.0	$\frac{1}{4} = 0,25$	0
Iris	GPU	`gpu`	28	`skylake`	1.0	$\frac{1}{27}$	50
Iris	Large-Mem	`bigmem`	112	`skylake`	1.0	$\frac{1}{27}$	0
Aion	Regular	`batch`	128	`epyc`	0,57	$\frac{1}{1.75}$	0

- Continuous use of **1 GPU nodes** (28 cores, 4 GPUs, 756GB Memory) on `iris` cluster
  - ↪ 28 cores per node, 4 GPUs per nodes, 27 GigaByte RAM per core, 756 GB per node
  - ↪ **For 30 days**: 1 node$\times[\alpha_{cpu} \times 28 + \alpha_{mem} \times 27 \times 28 + \alpha_{gpu} \times 4$ GPUS$]\times$ 30 days$\times$ 24 hours
    - ✓ Total: $1 \times [(1.0 + \frac{1}{27} \times 27) \times 28 + 50.0 \times 4] \times 720 = 184320$ SU $=$ **5529,6€ VAT excluded**

Fairshare Factor and Job billing

Number of SU associated to a job

$$N_{\text{Nodes}} \times [\alpha_{cpu} \times N_{\text{cores}} + \alpha_{mem} \times Mem + \alpha_{gpu} \times N_{\text{gpus}}] \times T_{\text{exec}}$$

Cluster	Node Type	Partition	#Cores/node	CPU	α_{cpu}	α_{mem}	α_{GPU}
Iris, Aion	Regular	interactive	28/128	n/a	0	0	0
Iris	Regular	batch	28	broadwell	1.0*	$\frac{1}{4} = 0,25$	0
Iris	Regular	batch	28	skylake	1.0	$\frac{1}{4} = 0,25$	0
Iris	GPU	gpu	28	skylake	1.0	$\frac{1}{27}$	50
Iris	Large-Mem	bigmem	112	skylake	1.0	$\frac{1}{27}$	0
Aion	Regular	batch	128	epyc	0,57	$\frac{1}{1.75}$	0

- Continuous use of **1 Large-Memory nodes** (112 cores, 3024GB Memory) on iris cluster
  - ↪ 112 cores per node, 27 GigaByte RAM per core i.e. 3024 GB per node
  - ↪ **For 30 days**: 1 node $\times[\alpha_{cpu} \times 112 + \alpha_{mem} \times 27 \times 112 + \alpha_{gpu} \times 0] \times$ 30 days $\times$ 24 hours
    - ✓ Total: $1 \times [(1.0 + \frac{1}{27} \times 27) \times 112] \times 720 = 161280$ SU = **4838,4€ VAT excluded**

Summary

European HPC strategy

- EU HPC strategy initiated in 2012
  ↪ implementation within H2020 program
- Based on three pillars:
  1. **HPC Infrastructure**: PRACE, GEANT,
  2. **HPC Technology**:
     ✓ ETP4HPC, European Processor Initiative (EPI)
  3. **Application expertise**:
     ✓ Centres of Excellence of Computing Applications (CoEs)

European HPC strategy

- EU HPC strategy initiated in 2012
  - ↪ implementation within H2020 program
- Based on three pillars:
  1. **HPC Infrastructure**: PRACE, GEANT, **EuroHPC**
  2. **HPC Technology**:
     - ✓ ETP4HPC, European Processor Initiative (EPI)
  3. **Application expertise**:
     - ✓ Centres of Excellence of Computing Applications (CoEs)
- **Significative** upgrade in **2018** of the EC Strategy on HPC

EuroHPC Joint Undertaking (JU)

Summary of the EU HPC funding efforts
[Source : ETP4HPC Handbook 2018]

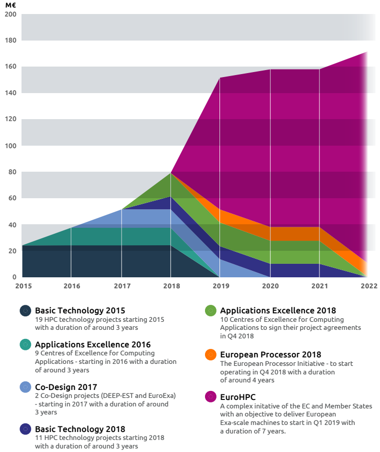

EU HPC Strategy Implementation

- **European Technology Platform (ETP) for HPC**
  - ↪ Industry-led forum feat. HPC stakeholders
  - ↪ Providing EU framework to define HPC research priorities/actions: SRA, HPC Handbook
    - ✓ UL part of ETP4HPC (2016-)

EU HPC Strategy Implementation

- **European Technology Platform (ETP) for HPC**
  - ↪ Industry-led forum feat. HPC stakeholders
  - ↪ Providing EU framework to define HPC research priorities/actions: SRA, HPC Handbook
    - ✓ UL part of ETP4HPC (2016-)
- **PRACE** - Partnership for Advanced Computing in Europe
  - ↪ Non-profit association, 25 member countries, now entering PRACE2/PRACE3
  - ↪ (Oct. 2017) **Luxembourg 25th country to join PRACE**
    - ✓ Official Delegate/Advisor (P. Bouvry/S. Varrette) from UL

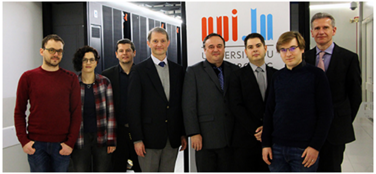

EU HPC Strategy Implementation

- **European High-Performance Computing Joint Undertaking**
  - ↪ EuroHPC JU effectively operational starting **Jan 1st, 2019**
    - ✓ administrative management from **Luxembourg**
  - ↪ **Budget 2018-2020 (Phase 1):** $\simeq$ **1.5 B€** (536 M€from EU)
  - ↪ Public and private members
    - ✓ EC, 32 MS, representatives from supercomputing/BD stakeholders
  - ↪ EU Objective with EuroHPC:
    - ✓ 5 **Petascale** systems (~~2020~~ 2021) (incl. MeluXina in Luxembourg)
    - ✓ 3 **Pre-exascale** systems (~~2020~~ 2021)
    - ✓ **2 exascale** systems (2022-2023)
    - ✓ Post-exascale system (2027)

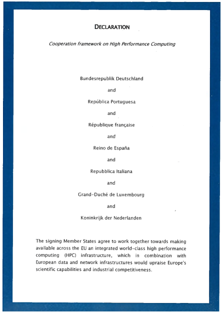

DECLARATION

Cooperation framework on High Performance Computing

Bundesrepublik Deutschland

and

República Portuguesa

and

République française

and

Reino de España

and

Repubblica Italiana

and

Grand-Duché de Luxembourg

and

Koninkrijk der Nederlanden

The signing Member States agree to work together towards making available across the EU an integrated world-class high performance computing (HPC) infrastructure, which in combination with European data and network infrastructures would upraise Europe's scientific capabilities and industrial competitiveness.

EU Tier-0 HPC systems	Total Capacity
PRACE	111.24 PFlops
EuroHPC {Peta,Pre-Exa}scale systems (as of Dec. 2020)	**917,9 PFlops**

EuroHPC Systems

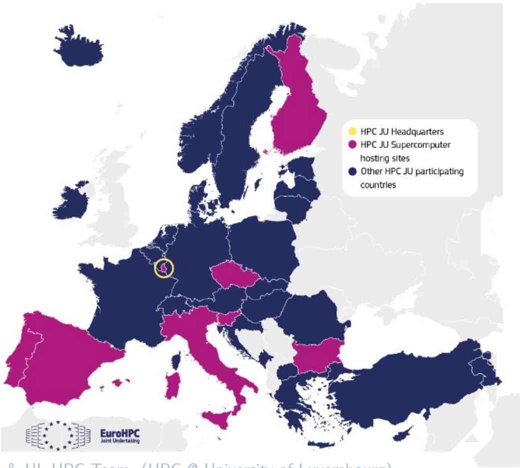

EuroHPC **Pre-exascale** Systems

Country	System	R_{peak}	Storage
Finland	LUMI	552 PF	127PB
Italy	Leonardo	322 PF	100PB
Spain	MareNostrum 5	$\geq$ 200 PF	

EuroHPC **Petascale** Systems

Country	System	R_{peak}	Storage
Luxembourg	MeluXina	17.57 PF	$\geq$ 20 PB
Czech Republic.	EUROITForl	15.2 PF	1PB
Portugal	Deucalion	$\geq$ 10 PF	
Slovenia	VEGA	10.1 PF	24 PB
Bulgaria	PetaSC	$\geq$ 4 PF	

PRACE/EuroHPC Ongoing Activities (Q4 2020)

- HPC in the lead for finding solutions for COVID-19 pandemics
  - ↪ PRACE specific fast track call for COVid-19 research
  - ↪ Exascale4Cov project
- **Newly appointed EuroHPC JU Executive Director**: Anders Dam Jensen
- EU-US cooperation (PRACE-XSEDE)
- Procurement for 8 supercomputers (3 pre-exa, 5 Peta) continued
  - ↪ EuroHPC 2020 Budget for MeluXina: 10,5 M€
- **EuroHPC Competence Center (EuroCC)**, CASTIEL

PRACE/EuroHPC Ongoing Activities (Q4 2020)

- HPC in the lead for finding solutions for COVID-19 pandemics
  - ↪ PRACE specific fast track call for COVid-19 research
  - ↪ Exascale4Cov project
- **Newly appointed EuroHPC JU Executive Director**: Anders Dam Jensen
- EU-US cooperation (PRACE-XSEDE)
- Procurement for 8 supercomputers (3 pre-exa, 5 Peta) continued
  - ↪ EuroHPC 2020 Budget for MeluXina: 10,5 M€
- **EuroHPC Competence Center (EuroCC)**, CASTIEL

- **Sept 2020**: EuroHPC Phase 2 (revised Regulation 2021-2027) announcement (EC)
  - ↪ State of the Union speech, by Ursula von der Leyen
  - ↪ **8 B€ investment in EU HPC / Digital sovereignty**
    - ✓ Next-generation exascale supercomputers
    - ✓ Quantum computers and hybrid computers
    - ✓ EU Cloud Gaia-X, a Federated Data Infrastructure for Europe. . .

EU HPC Strategy Implementation

- **European Processor Initiative** (EPI)

SiPearl

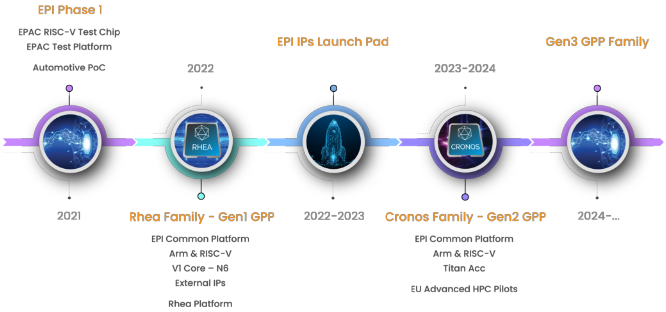

EU HPC Strategy Implementation

- **EuroHPC Competence Center (EuroCC)** and **CASTIEL** Projects
  - ↪ EU Project Management by HLRS (Stuttgart, Germany), 2Y project, Budget: **2M€**
  - ↪ **Objective**: **establish National Competence Centres (NCC) in HPC, HPDA and AI**
  - ↪ Luxembourg Consortium: LuxInnovation, LuxProvide, Univ. of Luxembourg
    - ✓ Uni.lu Coordinator: A. Vendeventer; Postdocs: L. Koutsantonis (one joining soon: Apr 2021)
  - ↪ Competence Centre Advisory Board (CAB):
    - ✓ Czech Republic (IT4I), Finland (CSC), Ireland (ICHEC), Norway (Sigma2), Sweden (ENCCS)

EU HPC Strategy Implementation

EURO

- **EuroHPC Competence Center (EuroCC)** and **CASTIEL** Projects
  - ↪ EU Project Management by HLRS (Stuttgart, Germany), 2Y project, Budget: **2M€**
  - ↪ **Objective**: **establish National Competence Centres (NCC) in HPC, HPDA and AI**
  - ↪ Luxembourg Consortium: LuxInnovation, LuxProvide, Univ. of Luxembourg
    - ✓ Uni.lu Coordinator: A. Vendeventer; Postdocs: L. Koutsantonis (one joining soon: Apr 2021)
  - ↪ Competence Centre Advisory Board (CAB):
    - ✓ Czech Republic (IT4I), Finland (CSC), Ireland (ICHEC), Norway (Sigma2), Sweden (ENCCS)

Uni.lu contributions to EuroCC and CASTIEL Projects

- Uni.lu **Task leader** on
  - ↪ Task 28.2: Training and Skills Development
  - ↪ Task 28.6: Facilitation of access to scientific and technical expertise and knowledge
- CASTIEL *"champions"*
  - ↪ “Training, Twinning, Mentoring” WG (S. Varrette/ F. Pinel)

Summary

Conclusion

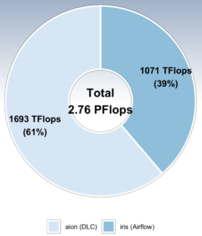

- Uni.lu HPC at the heart of the National Digital Strategy
  - ↪ Started in 2007, under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ **Expert UL HPC team** structured in 3 main pillars
    - ✓ **Research Computing & HPC Operations**: H. Cartiaux, T. Valette, A. Olloh. S. Peter *(Ext.)*
    - ✓ **Research & Trainings**: Dr. F Pinel, Dr. E. Kieffer, Dr. E. Krishnasamy, Dr. L. Koutsantonis
    - ✓ **Strategic Developments & Partnership**: A. Vandeventer
  - ↪ Several Computational scientists / domain experts across **ALL** the UL

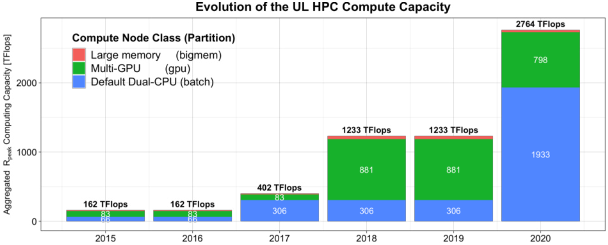

Conclusion

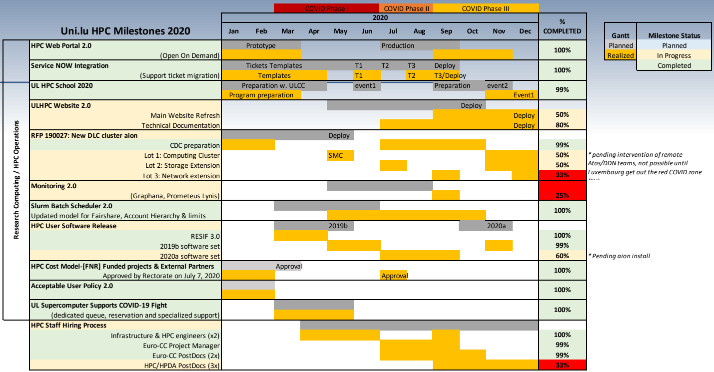

- **2020 very special year**
  - ↪ New joiners, departures
  - ↪ Role changes
  - ↪ Worldwide COVID crisis
    - ✓ 2020 HPC Budget impacted
      ... global UL strategy
    - ✓ delayed `aion` deployment
  - ↪ National/EuroHPC Developments in Luxembourg
- Yet nearly **ALL 2020 (ambitious) milestones completed** (modulo `aion` - GA: Q1 2021)
- **Several drastic changes implemented in 2020**
  - ↪ HPC Cost model implementation for funded research projects and external partners
  - ↪ AUP 2.0, Slurm Configuration 2.0, RESIF 3.0, OpenOnDemand Portal
  - ↪ Better HPC services to support research excellence
  - ↪ Enhance support (Service Now), documentation and training

Perspectives (2021)

- **Get out for COVID crisis!**
- **Aion Deployment and production release**

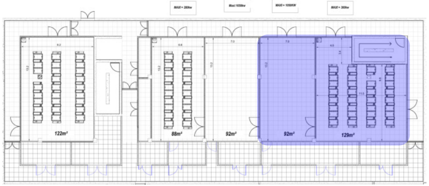

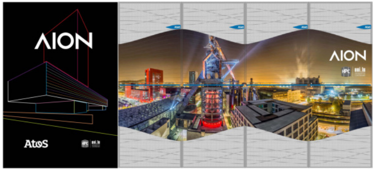

- Novel opportunities for **enhanced HPC partnership** (PRACE, EuroHPC, LuxProvide, EuroCC...)
  - ↪ consolidating National & European HPC ecosystem while preserving UL interest & expertise
  - ↪ consolidating UL expertise and leadership for reference HPC trainings at national level
  - ↪ sustain UL excellence in cutting-edge HPC/HPDA/AI research developments
- ULHPC to serve as edge access to National HPC Center (LuxProvide/MECO)
  - ↪ **YET** UL keep sovereignty and ownership of its internal HPC facility

Thank you for your attention...

Questions?

High Performance Computing @ Uni.lu

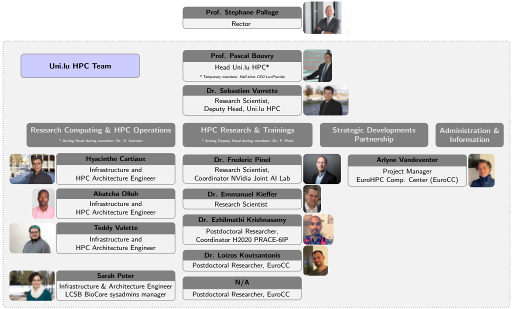

University of Luxembourg, Belval Campus
Maison du Nombre, 4th floor
2, avenue de l'Université
L-4365 Esch-sur-Alzette
mail: hpc@uni.lu

https://hpc.uni.lu