High Performance Computing & Big Data Services

hpc.uni.lu

hpc@uni.lu

@ULHPC

LUXEMBOURG
LET'S MAKE IT HAPPEN

UL HPC School 2019

Overview & Challenges of the UL HPC facility at the EuroHPC Horizon

Prof. P. Bouvry, Dr. S. Varrette

V. Plugaru, S. Peter, H. Cartiaux, C. Parisot,

Dr. F. Pinel and Dr. E. Kieffer

University of Luxembourg (UL), Luxembourg

https://hpc.uni.lu

June. 20th, 2019, Belval

Uni.lu HPC School 2019

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

- **9th edition** of this training...
  - ↪ started in 2014
  - ↪ This one is the **long** version
    - ✓ 2-days event
    - ✓ Parallel sessions, feat. **basic & advanced** tutorials
- **Requirement**:
  - ↪ your favorite laptop with your favorite OS
    - ✓ Linux / Mac OS preferred, but Windows accepted
  - ↪ basic knowledge in Linux command line
  - ↪ ability to take notes (Markdown etc.)

Online UL HPC Tutorials

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

Uni.lu HPC School 2019

- Several **brand new tutorials**
  - ↪ **PS5c**: Scalable Science III (Advanced - GPU programming)
  - ↪ **PS10b**: Python II (Advanced) Parallel ML & Evol. Comp.
  - ↪ **PS12**: Bioinformatics workflows with snakemake and conda
  - ↪ **PS13b**: Deep learning II (Advanced - distributed) Horovod
  - ↪ **PS14**: Optimizers (CPLex, Gurobi)

Uni.lu HPC School 2019

- Several **brand new tutorials**
  - ↪ **PS5c**: Scalable Science III (Advanced - GPU programming)
  - ↪ **PS10b**: Python II (Advanced) Parallel ML & Evol. Comp.
  - ↪ **PS12**: Bioinformatics workflows with snakemake and conda
  - ↪ **PS13b**: Deep learning II (Advanced - distributed) Horovod
  - ↪ **PS14**: Optimizers (CPLex, Gurobi)

- Special Guest for this edition: **Filip Kucerak**
  - ↪ Winner of the EU Contest for Young Scientists EUCYS18

Uni.lu HPC School 2019

- Several **brand new tutorials**
  - ↪ **PS5c**: Scalable Science III (Advanced - GPU programming)
  - ↪ **PS10b**: Python II (Advanced) Parallel ML & Evol. Comp.
  - ↪ **PS12**: Bioinformatics workflows with snakemake and conda
  - ↪ **PS13b**: Deep learning II (Advanced - distributed) Horovod
  - ↪ **PS14**: Optimizers (CPLex, Gurobi)

- Special Guest for this edition: **Filip Kucerak**
  - ↪ Winner of the EU Contest for Young Scientists EUCYS18

- **Next edition** planned for **Nov., 2019**
  - ↪ in Belval
  - ↪ Short 1-day event, adressing basic tutorials

Agenda Day 1: June 20th, 2019

Day 1	Main Track (MSA **4.540**)	Speaker
9h00-9h45	PS1: Preliminaries (SSH)	H. Cartiaux, C. Parisot
9h45-10h15	**Coffee break**	
10h15-11h30	**Keynote**: Overview and Challenges of the UL HPC Facility	S. Varrette
11h30-12h30	PS2: Getting Started 2.0	H. Cartiaux
12h30-13h30	**LUNCH**	
13h30-15h00	PS4a: Monitoring & Profiling I (basics)	H. Cartiaux
15h00-16h30	PS5a: Scalable Science I (Basics - OpenMP/MPI jobs)	S. Varrette
16h30-17h00	**Coffee break**	
17h00-18h30	PS6: Big Data Applications (batch, stream, hybrid)	S. Varrette

Day 1	Parallel Track (MSA **4.140**)	Speaker
9h00-11h30		
11h30-12h30	PS3: (Advanced) Job scheduling with SLURM	C.Parisot
12h30-13h30	**LUNCH**	
13h30-15h00	PS4b: Monitoring & Profiling II (Advanced)	V.Plugaru,X.Besseron
15h00-16h30	PS5b: Scalable Science II (Computational Physics, Chemistry...)	V. Plugaru
16h30-17h00	**Coffee break (in 4.540)**	
17h00-18h30	PS7: Scientific computing using MATLAB	V. Plugaru

PS = *Practical Session using your laptop*

Agenda Day 2: June 21th, 2019

Day 2	Main Track (MSA **4.540**)	Speaker
9h00-10h00	PS8: HPC Containers with Singularity	V. Plugaru
10h00-10h30	**Coffee break**	
10h30-11h15	**Keynote/PS9**: Data management (backup, decommissioning...)	S. Peter
11h15-12h00	**User Session**	
12h00-13h00	**LUNCH**	
13h00-14h15	PS10a: Python I (Basic) Prototyping with Python	C. Parisot, S.Peter
14h15-15h45	PS10b: Python II (Advanced) SCOOP, DEAP (ipyparallel, scikit-learn)	E. Kieffer
15h45-16h15	**Coffee break**	
16h15-17h30	PS13a: Deep Learning I (Basics) Keras/Tensorflow CPU/GPU	F. Pinel
17h30-18h30	PS13b: Deep learning II (Advanced - distributed) Horovod	V. Plugaru, F.Pinel

Day 2	Parallel Track (MSA **4.510**)	Speaker
9h00-10h00	PS5c: Scalable Science III (Advanced - GPU programming)	F. Pinel
10h00-10h30	**Coffee break (in 4.540)**	
10h30-11h15	**Keynote (in 4.540)**	
11h15-12h00	**User Session (in 4.540)**	
12h00-13h00	**LUNCH**	
13h00-14h15	PS11: R - statistical computing	A. Ginolhac, R. Krause
14h15-15h45	PS12: Bioinformatics workflows with snakemake and conda	S. Peter
15h45-16h15	**Coffee break (in 4.540)**	
16h15-17h30	PS14: Optimizers (CPLex, Gurobi)	E. Kieffer
17h30-18h30		

Summary

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** ...

Why HPC and BD ?

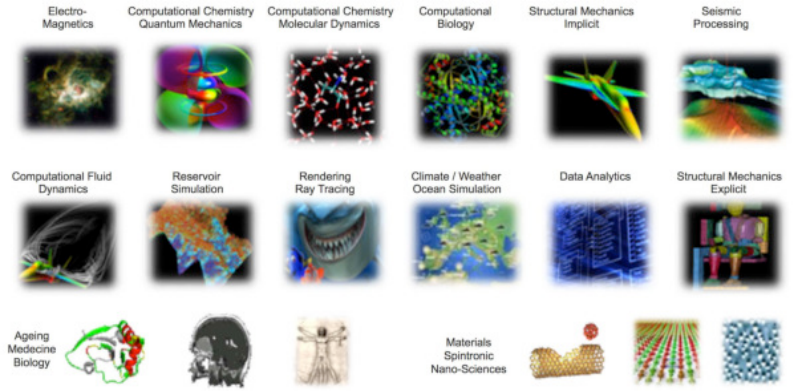

HPC: **H**igh **P**erformance **C**omputing
BD: **B**ig **D**ata

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.

Why HPC and BD ?

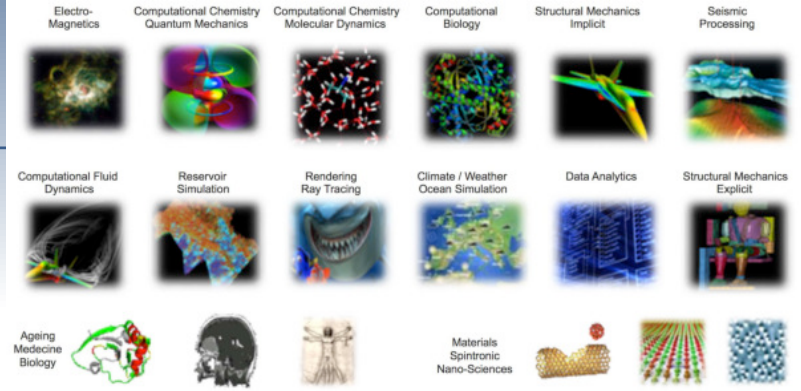

HPC: **H**igh **P**erformance **C**omputing
BD: **B**ig **D**ata

- Essential tools for **Science, Society and Industry**
  - ↪ **Data driven economy context**
  - ↪ All scientific disciplines are becoming computational today
    - ✓ require very high computing power, handle **huge** volumes of data
- **Industry, SMEs** increasingly relying on HPC
  - ↪ to invent innovative solutions
  - ↪ . . . while reducing cost & decreasing time to market

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.

Why HPC and BD ?

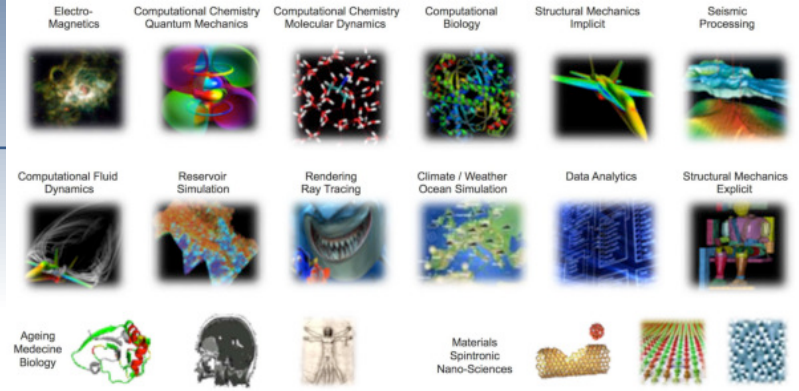

HPC: **H**igh **P**erformance **C**omputing
BD: **B**ig **D**ata

- Essential tools for **Science, Society and Industry**
  - ↪ **Data driven economy context**
  - ↪ All scientific disciplines are becoming computational today
    - ✓ require very high computing power, handle **huge** volumes of data
- **Industry, SMEs** increasingly relying on HPC
  - ↪ to invent innovative solutions
  - ↪ . . . while reducing cost & decreasing time to market
- HPC = **global race** (strategic priority) - EU takes up the challenge:
  - ↪ PRACE / EuroHPC / IPCEI on HPC and Big Data (BD) Applications

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 HPC Needs per Domains

Material Science & Engineering

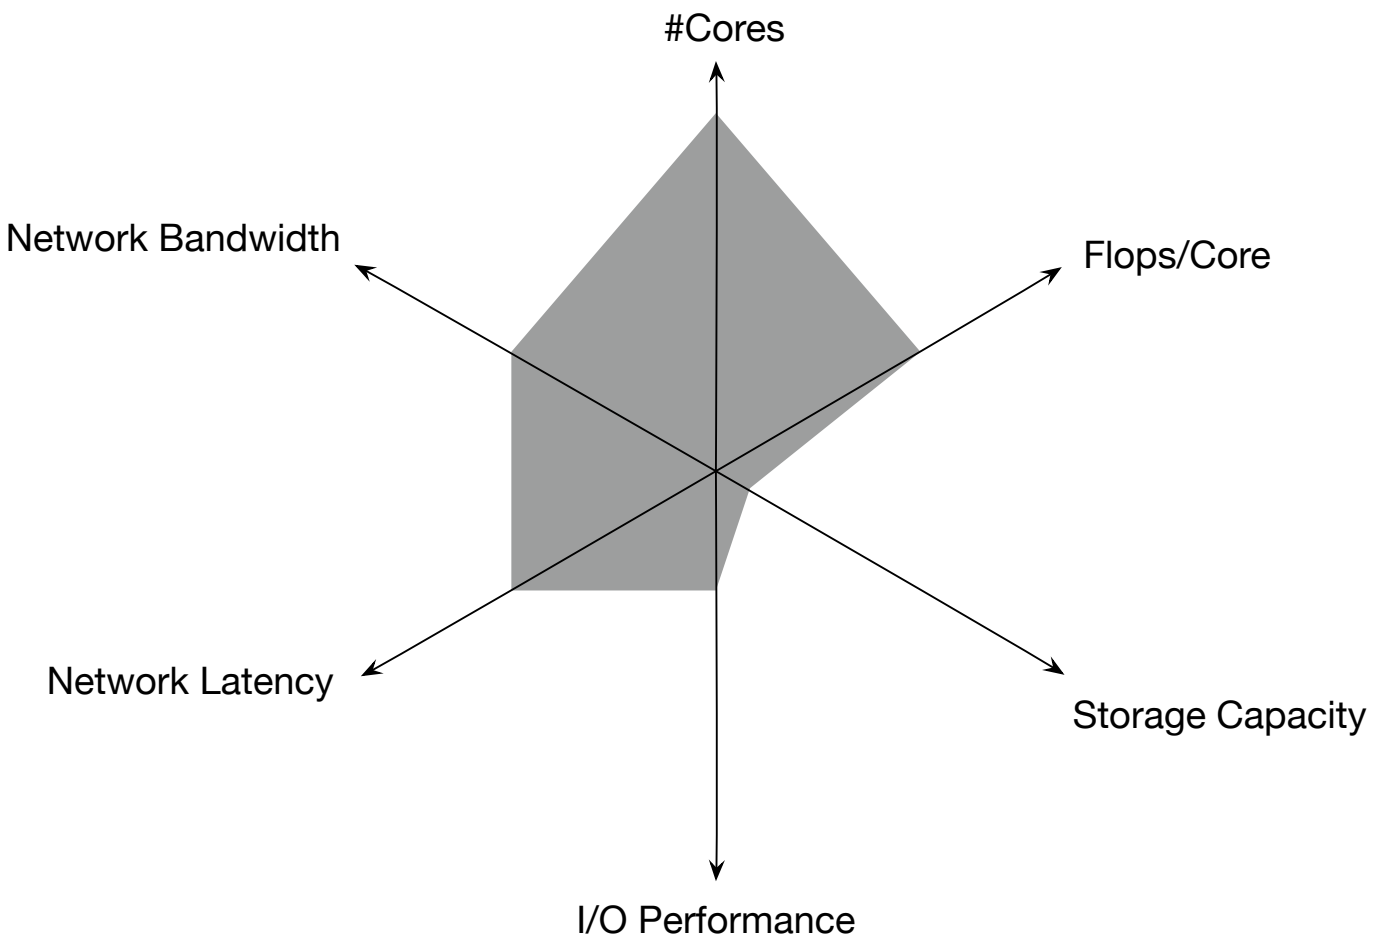

Different HPC Needs per Domains

Biomedical Industry / Life Sciences

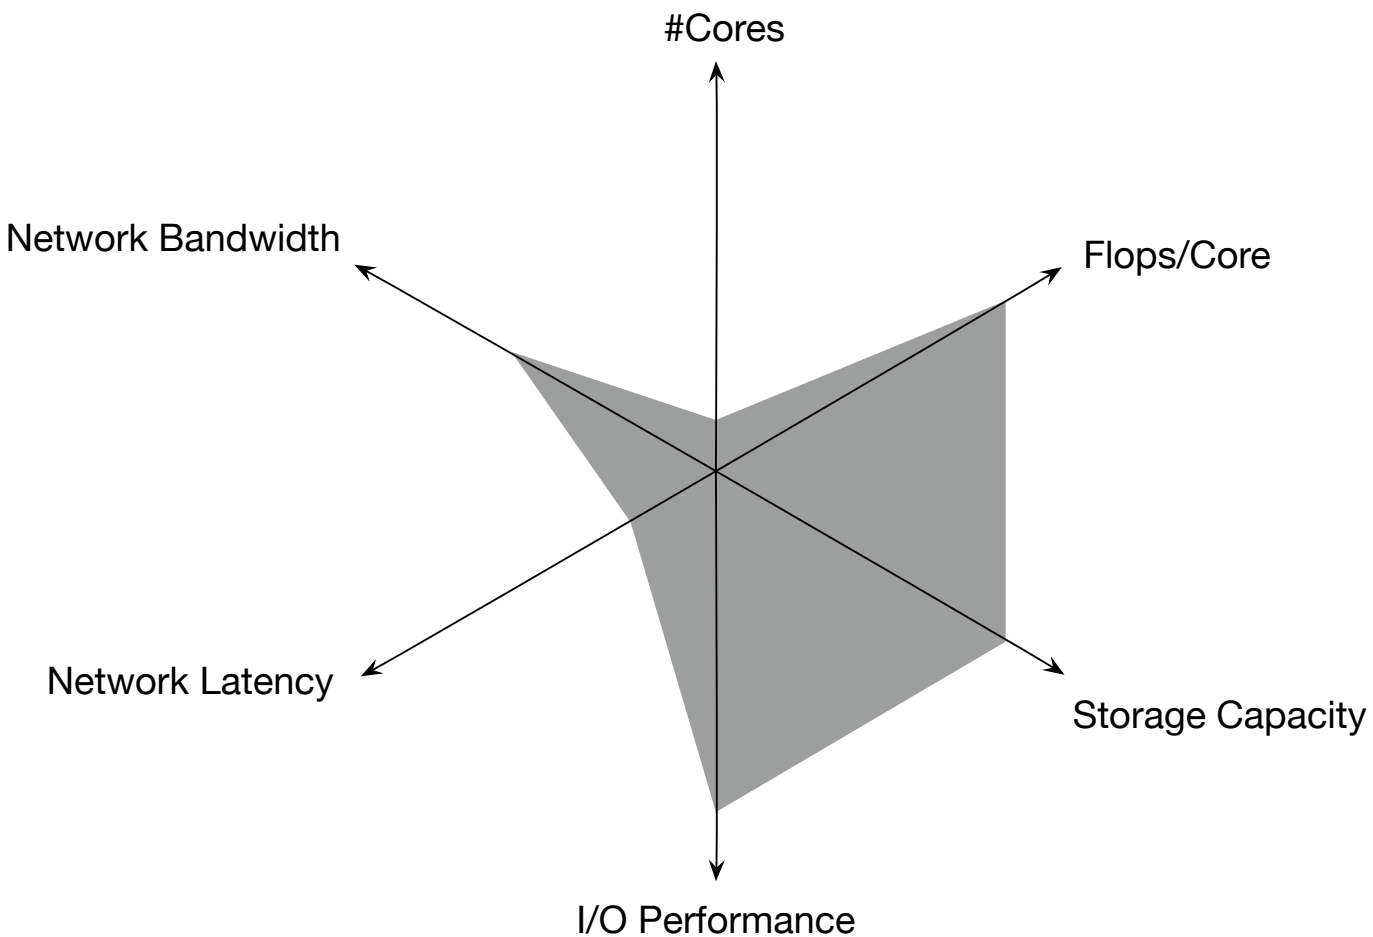

Different HPC Needs per Domains

Deep Learning / Cognitive Computing

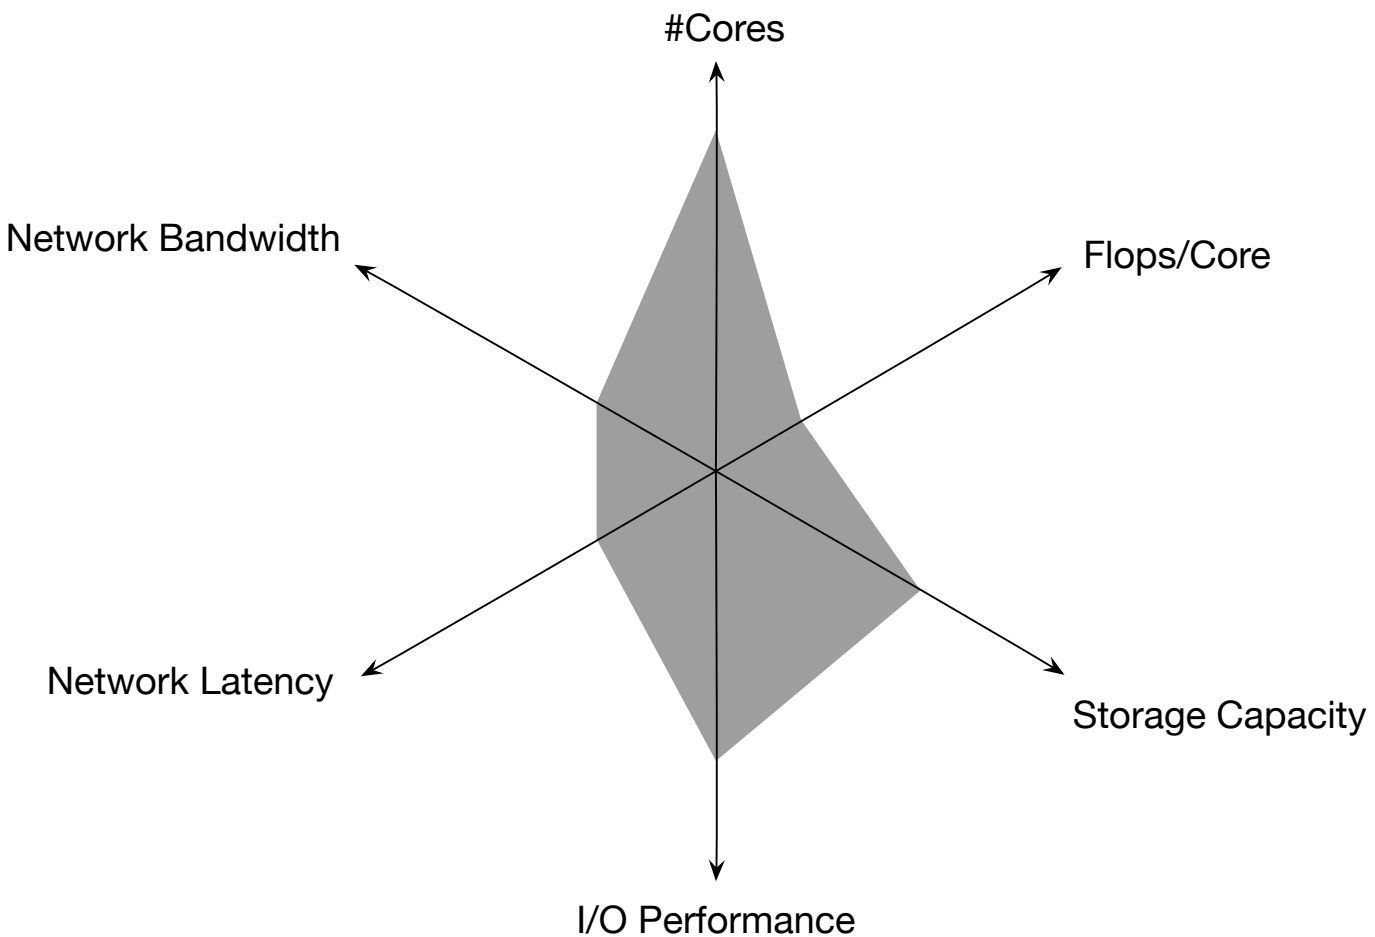

Different HPC Needs per Domains

IoT, FinTech

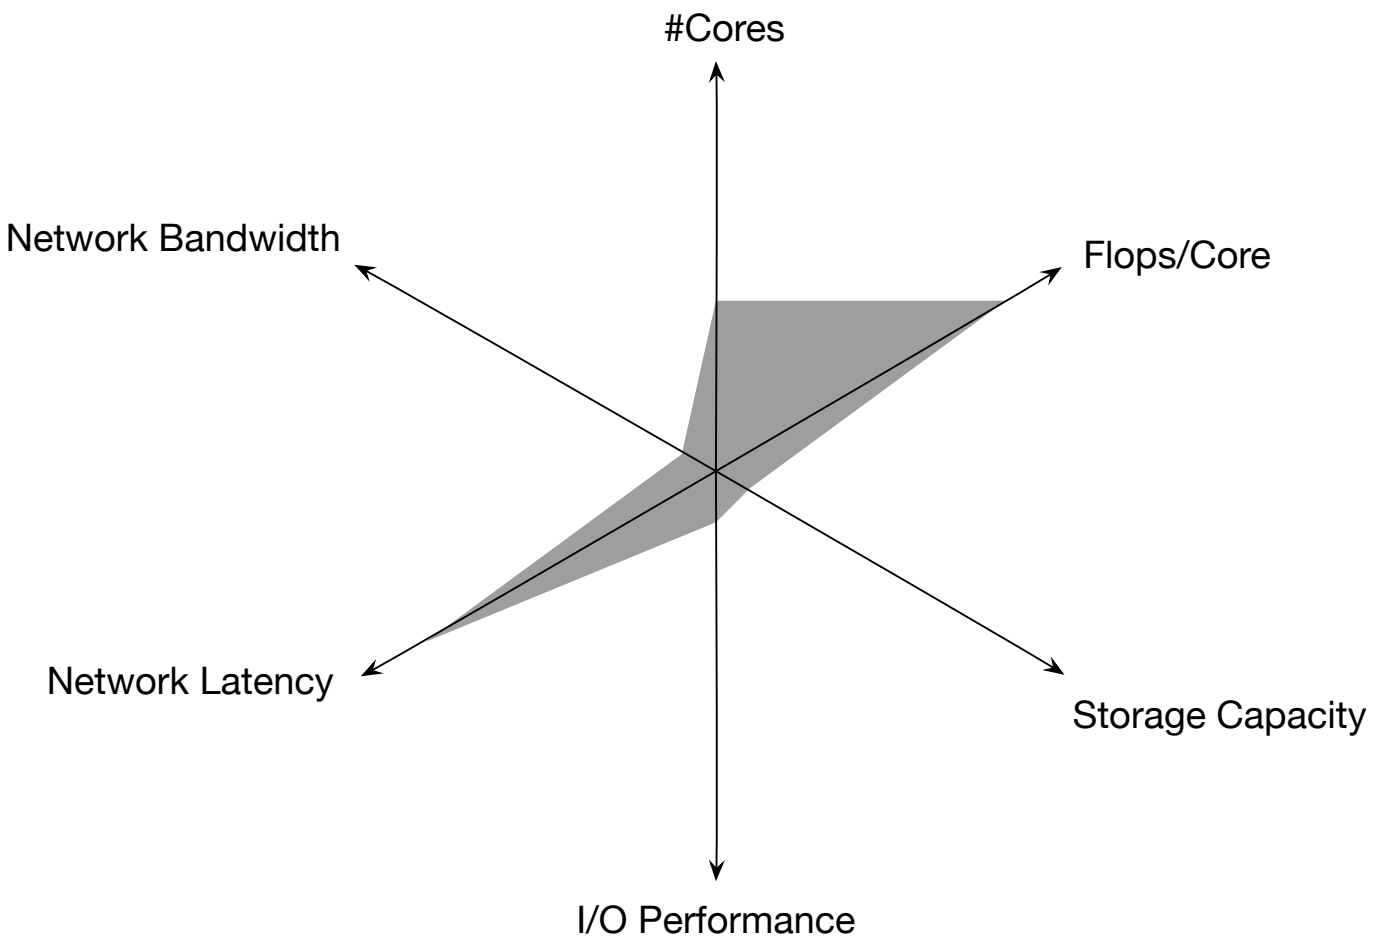

Different HPC Needs per Domains

ALL Research Computing Domains

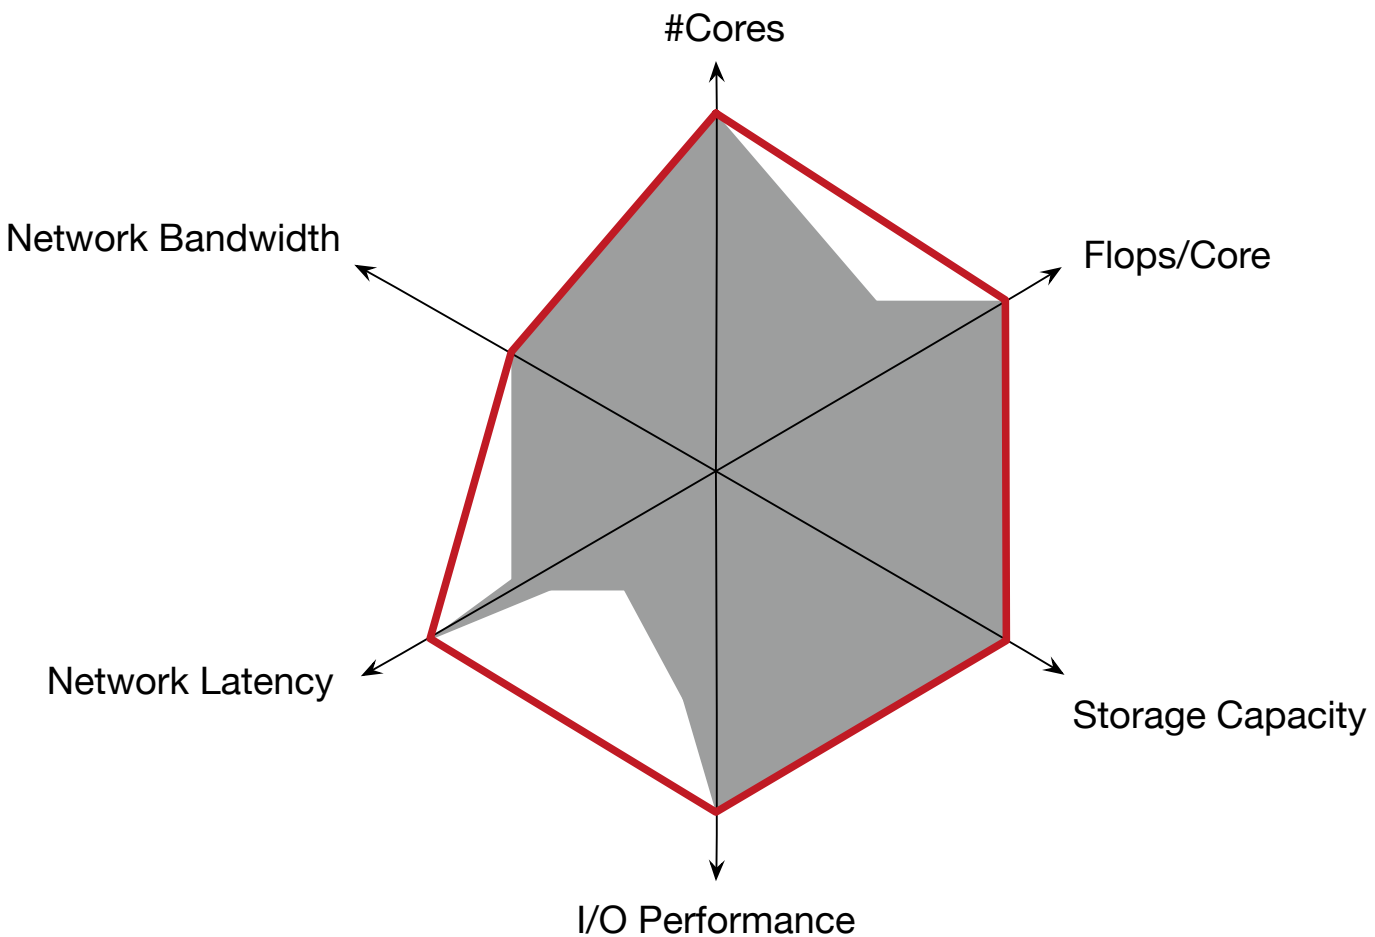

Computing for Researchers: Laptop

- **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 / environmen
    - ✓ Ex: Xen, VMWare ESXi, KVM, VirtualBox
    - ✓ Non-negligible Performance loss: $\geq 20\%$

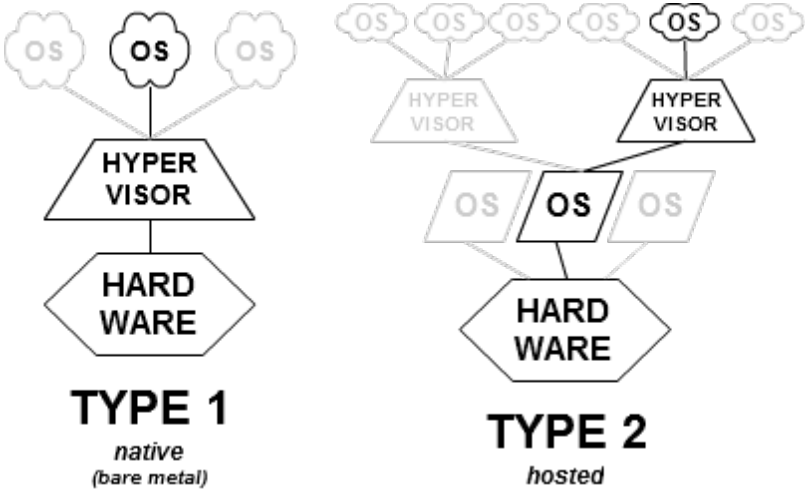

Computing for Researchers: Laptop

- **Regular PC / Local Laptop / Workstation**
  - ↪ **Native OS** (Windows, Linux, Mac etc.)
- Virtualized OS (**VM**) through an **hypervisor**
  - ↪ *Hypervisor*: core virtualization engine / environmen
    - ✓ 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

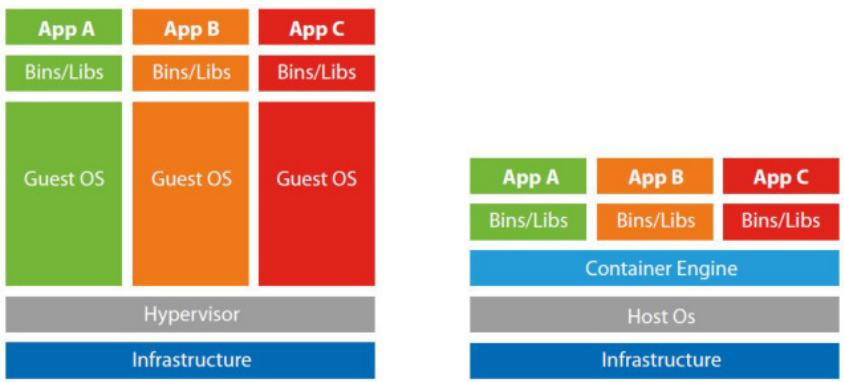

Computing for Researchers: Cloud

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

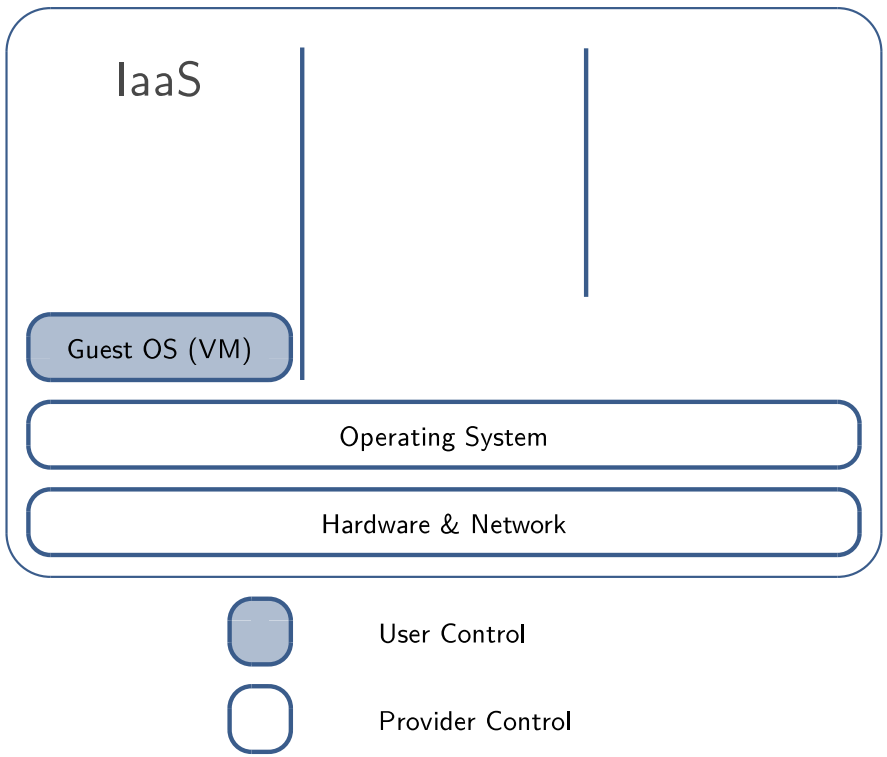

Computing for Researchers: Cloud

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

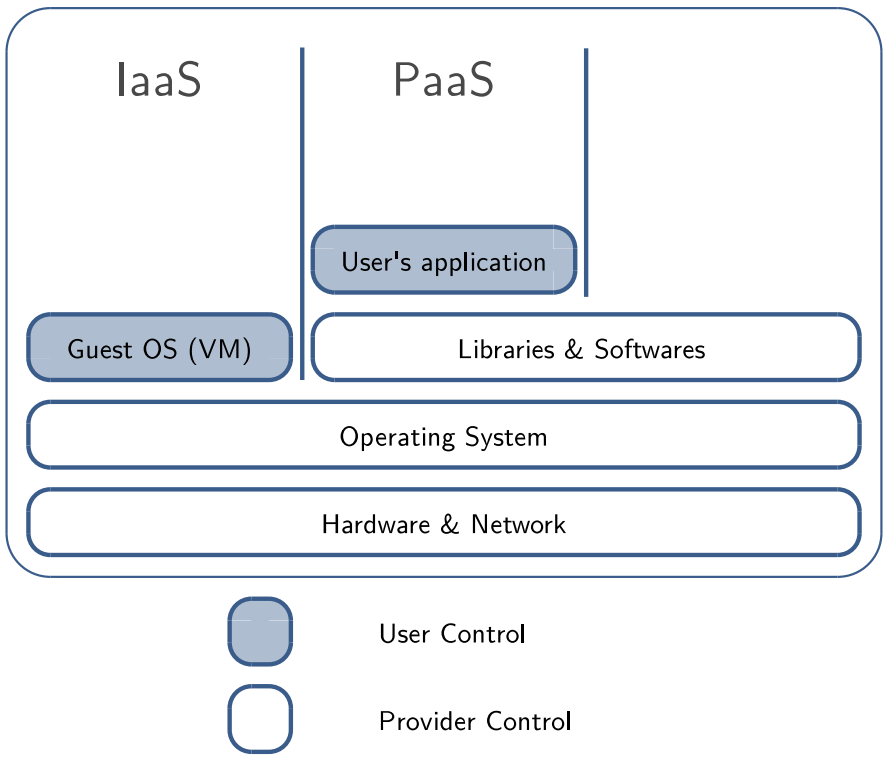

Computing for Researchers: Cloud

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

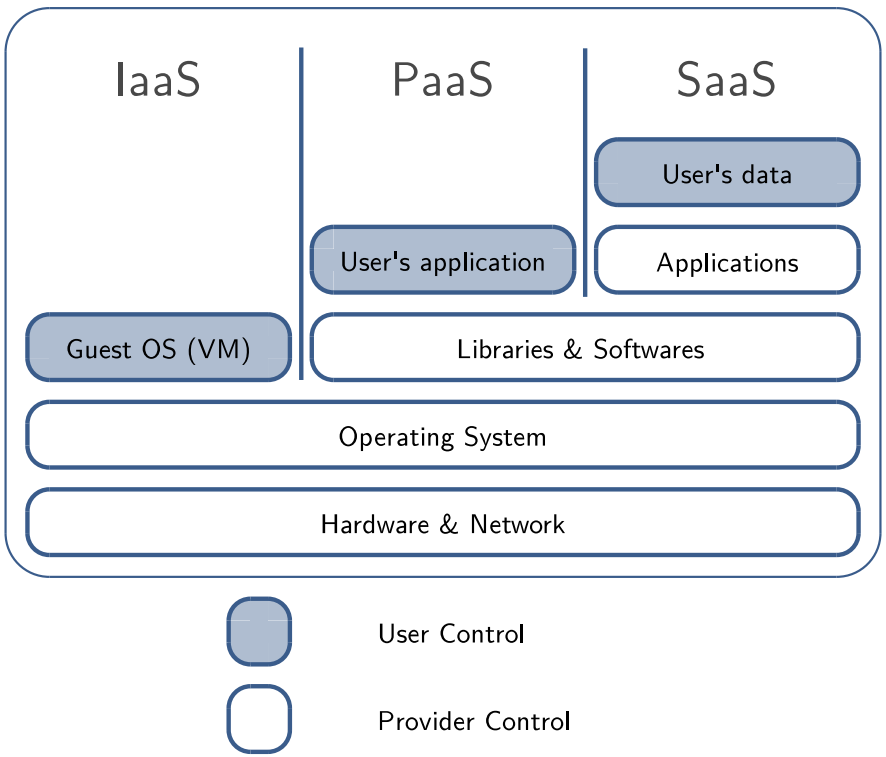

Computing for Researchers: Cloud

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

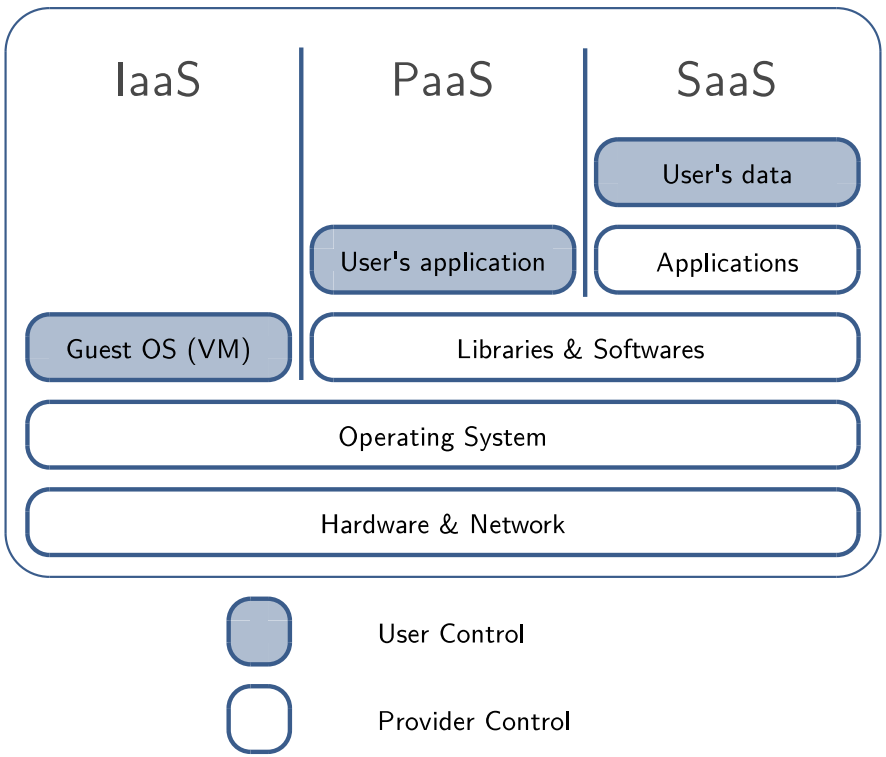

Computing for Researchers: HPC

- 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**

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**

YET...

PC $\neq$ Cloud $\neq$ HPC

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

Summary

HPC Computing Hardware

Base

- **CPU** (Central Processing Unit) *Highest software flexibility*
  - ↪ High performance across all computational domains
  - ↪ Ex: Intel Core i9-9900K *(Q4'18)* $R_{peak} \simeq 922$ GFlops (DP)
    - ✓ 8 cores @3.6GHz (14nm, 95W, $\simeq$ 3.5 billion transistors) + integ. graphics

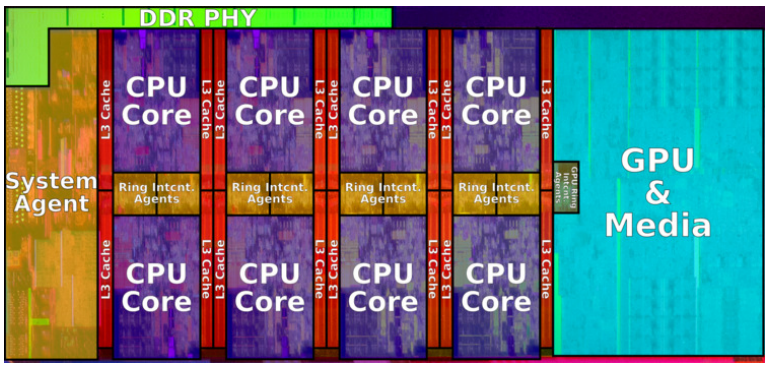

Intel Coffee Lake die

HPC Computing Hardware

Base

- **CPU** (Central Processing Unit) — *Highest software flexibility*
  - ↪ High performance across all computational domains
  - ↪ Ex: Intel Core i9-9900K *(Q4'18)* — $R_{peak} \simeq 922$ GFlops (DP)
    - ✓ 8 cores @3.6GHz (14nm, 95W, $\simeq$ 3.5 billion transistors) + integ. graphics

Accelerators

- **GPU** (Graphics Processing Unit): — *Ideal for ML/DL workloads*
  - ↪ Ex: Nvidia Tesla V100 SXM2 *(Q2'17)* — $R_{peak} \simeq 7.8$ TFlops (DP)
    - ✓ 5120 cores @ 1.3GHz (12nm, 250W, 21 billion transistors)

HPC Computing Hardware

Base

- **CPU** (Central Processing Unit) — *Highest software flexibility*
  - ↪ High performance across all computational domains
  - ↪ Ex: Intel Core i9-9900K *(Q4'18)* — $R_{peak} \simeq 922$ GFlops (DP)
    - ✓ 8 cores @3.6GHz (14nm, 95W, $\simeq$ 3.5 billion transistors) + integ. graphics

Accelerators

- **GPU** (Graphics Processing Unit): — *Ideal for ML/DL workloads*
  - ↪ Ex: Nvidia Tesla V100 SXM2 *(Q2'17)* — $R_{peak} \simeq 7.8$ TFlops (DP)
    - ✓ 5120 cores @ 1.3GHz (12nm, 250W, 21 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

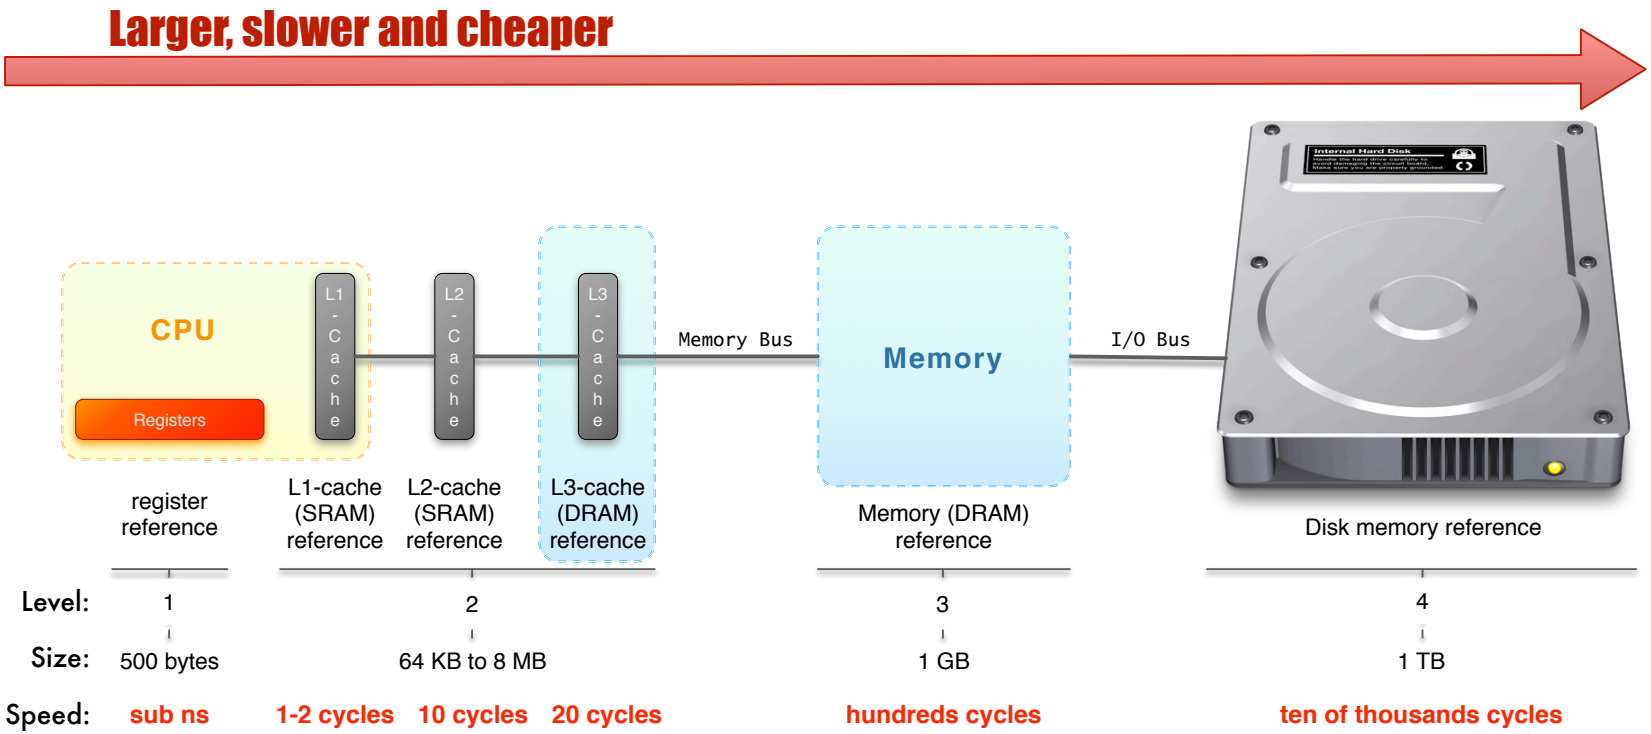

- SSD (SATA3) R/W: 550 MB/s; 100000 IOPS **450 €/TB**
- HDD (SATA3 @ 7,2 krpm) R/W: 227 MB/s; 85 IOPS **54 €/TB**

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

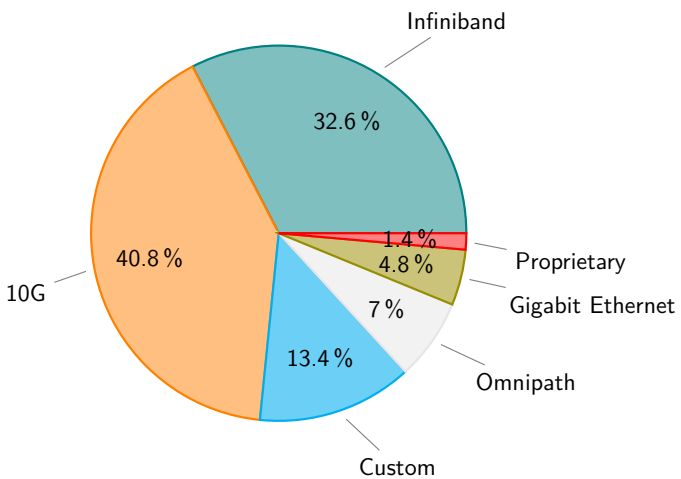

[Source : www.top500.org, Nov. 2017]

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

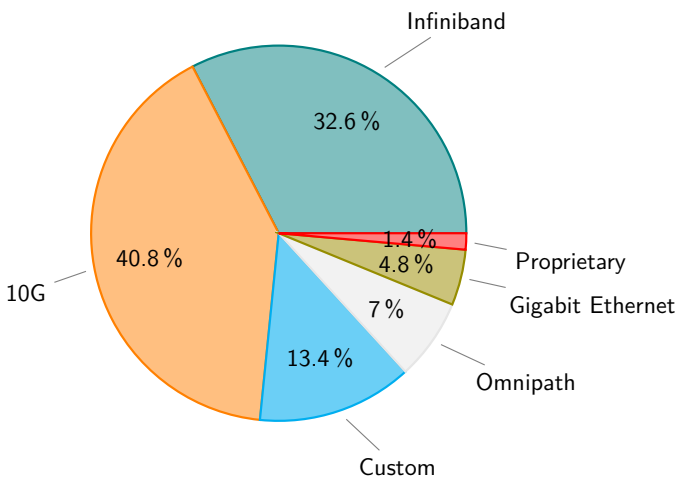

[Source : www.top500.org, Nov. 2017]

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

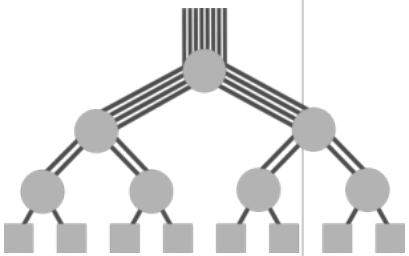

Network Topologies

- **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

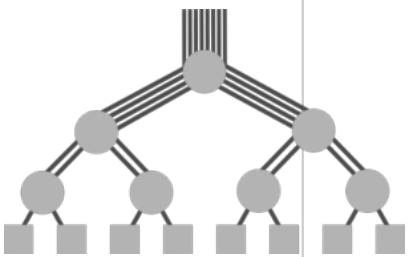

- **Mesh or 3D-torus** [Direct]
  - ↪ Blocking network, cost-effective for systems at scale
  - ↪ Great performance solutions for applications with locality
  - ↪ Simple expansion for future growth

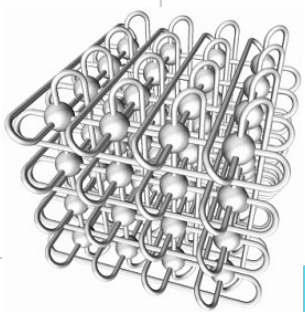

HPC Components: Operating System

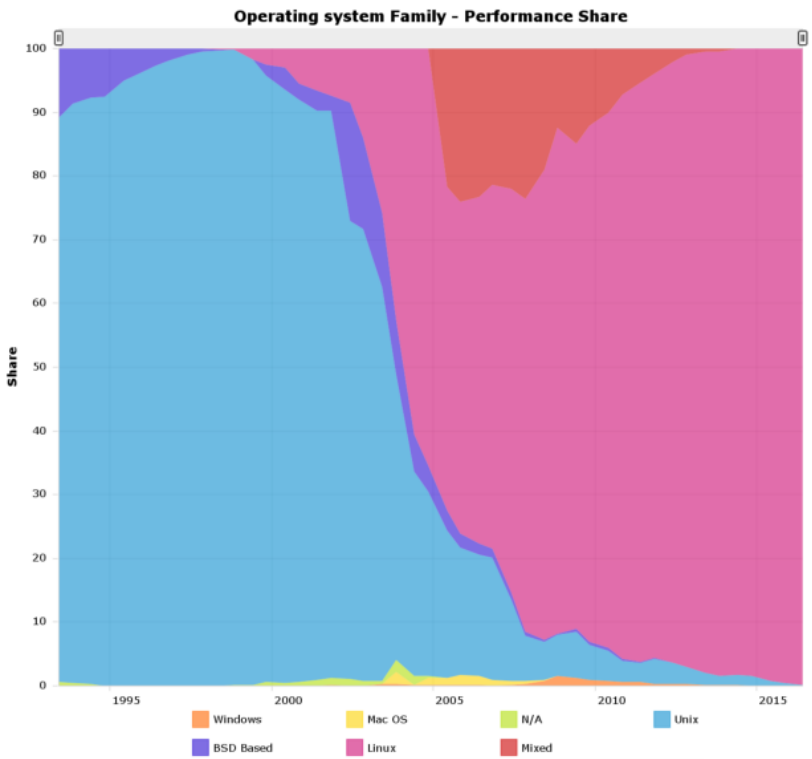

- Exclusively Linux-based (**really** 100%)
- Reasons:
  - ↪ stability
  - ↪ development flexibility

[Source : www.top500.org, Nov 2017]

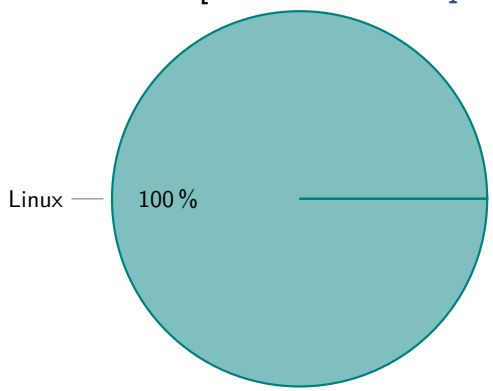

HPC Components: Software Stack

- **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.

- $\simeq$ **120 K€** - enclosure - 48-60 disks (4U)
  - $\hookrightarrow$ 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)	Parallel/Distributed FS	**11.25**	**9,46**
`lustre` (iris)	Parallel/Distributed FS	**12.88**	**10,07**
`gpfs` (gaia)	Parallel/Distributed FS	7.74	6.524
`lustre` (gaia)	Parallel/Distributed FS	4.5	2.956

* 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 and associated components
  - ↪ Basic storage component: **rack** (height: 42 RU)

HPC Components: Data Center

Definition (Data Center)

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

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

- Power/Heat dissipation per rack:
  - ↪ HPC **computing** racks: **30-120 kW**
  - ↪ **Storage** racks: **15 kW**
  - ↪ **Interconnect** racks: **5 kW**
- Various **Cooling** Technology
  - ↪ Airflow
  - ↪ Direct-Liquid Cooling, Immersion...

Power Usage Effectiveness

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

Summary

Summary

Univ. of Luxembourg & HPC

- *With regards to HPC*, Univ. of Luxembourg offers:
  - ↪ **People**
    - ✓ Domain experts
    - ✓ Computational and data scientists
    - ✓ Specialists in parallel algorithmics
  - ↪ **Services**
    - ✓ HPC clusters and management team
    - ✓ IT team (SIU)
    - ✓ Infrastructure team in collab. w. Fonds Belval
  - ↪ **Infrastructure**
    - ✓ Data center and a set of high-end clusters
  - ↪ **Education & Training**

High Performance Computing @ UL

- **Started in 2007**
  - ↪ under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ expert Uni.lu HPC team
    - ✓ *S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot* + **(new 2019) F. Pinel, E. Kieffer**
    - ✓ . . . and multiple domain experts per RU
  - ↪ Largest HPC facility in Luxembourg. . .
    - ✓ until EuroHPC Meluxina in production

https://hpc.uni.lu

High Performance Computing @ UL

- **Started in 2007**
  - ↪ under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ expert Uni.lu HPC team
    - ✓ *S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot* + **(new 2019) F. Pinel, E. Kieffer**
    - ✓ . . . and multiple domain experts per RU
  - ↪ Largest HPC facility in Luxembourg. . .
    - ✓ until EuroHPC Meluxina in production

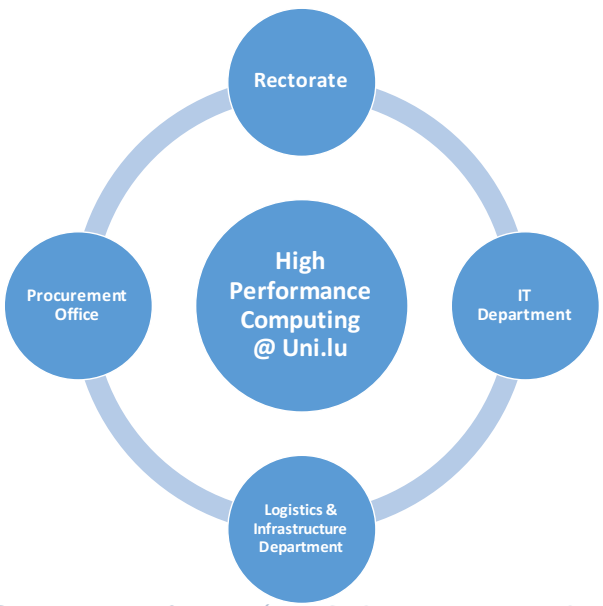

High Performance Computing @ UL

- **Started in 2007**
  - ↪ under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ expert Uni.lu HPC team
    - ✓ *S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot* + **(new 2019) F. Pinel, E. Kieffer**
    - ✓ . . . and multiple domain experts per RU
  - ↪ Largest HPC facility in Luxembourg. . .
    - ✓ until EuroHPC Meluxina in production

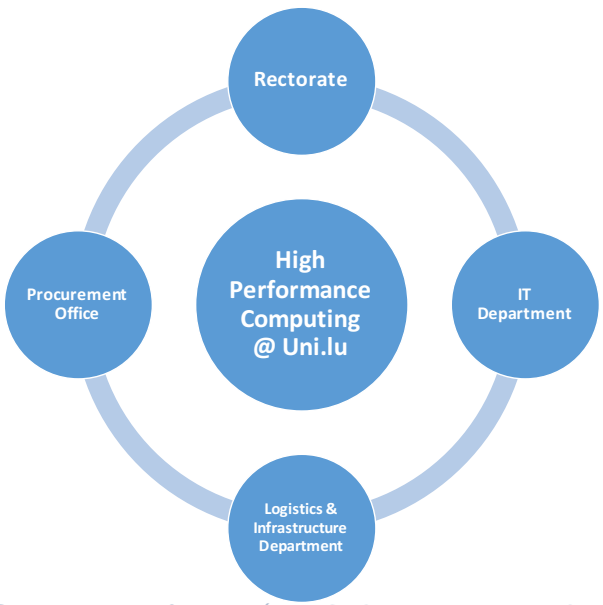

HPC/Computing Capacity*

1263.322 TFlops

(incl. 825.02 GPU TFlops)

*: including 6 new GPU nodes

HTC/Storage Capacity

9852.4 TB storage

UL HPC Beneficiaries

23 computational domains accelerated on UL HPC

- for the UL Faculties, Research Units and Interdisciplinary Centres
  - ↪ incl. LCSB, SnT... and now C2DH thematics
  - ↪ UL **strategic research priorities**
    - ✓ computational sciences, finance (fintech)
    - ✓ systems biomedicine, security, reliability and trust

- UL HPC features special systems targeting specific workloads:
  - ↪ **Machine Learning & AI**: GPU accelerators
    - ✓ ~~10 Tesla K40 + 16 Tesla K80 + 24 Tesla M20x~~: **~~76 GPU Tflops~~**
    - ✓ **Q4 2018**: 18*4 V100 (part of RFP 180027): **561 GPU Tflops**
    - ✓ **Q2 2019**: 6*4 V100 (part of RFP 180027): **187 GPU Tflops**
  - ↪ **BigData analytics & data driven science**:
    - ✓ 4 Large memory/SMP systems with 3 TB RAM

Accelerating UL Research

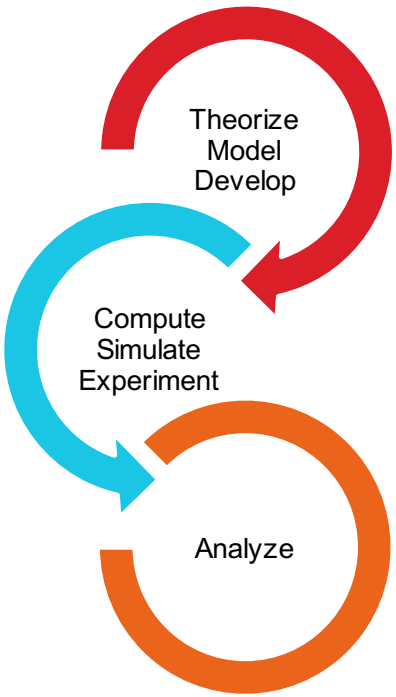

- **over 210 software packages** available for researchers
  ↪ via Environment modules/LMod from Easybuild

Domain	Software
Compiler Toolchains	(2018a) FOSS, Intel, PGI
MPI suites	OpenMPI, Intel MPI, MVAPICH2
Machine Learning	PyTorch, TensorFlow, Keras, 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	Allinea/ARM Forge & Perf Reports, Python, Go, Rust...
Container systems	Singularity
Visualisation	ParaView, OpenCV, XCS portal
...	

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

Uni.lu HPC Team

Prof. Pascal Bouvry
Senior advisor for the president as regards the HPC strategy
Leader of PCO Group, Head Uni.lu HPC

Dr. Sébastien Varrette
Research Scientist, Deputy head Uni.lu HPC

Valentin Plugaru, MSc.
R&D Specialist, Senior HPC Architect

Sarah Peter, MSc.
R&D Specialist, HPC/LCSB Support Liaison

Hyacinthe Cartiaux
HPC System administrator

Clément Parisot, MSc.
HPC System administrator

New 2019 Team members:

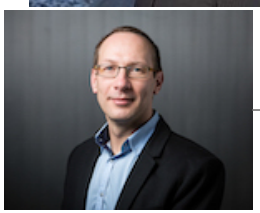

Dr. Frederic Pinel
Research Scientist,
Coordinator NVidia Joint AI Lab

Dr. Emmanuel Kieffer
Postdoctoral Researsher

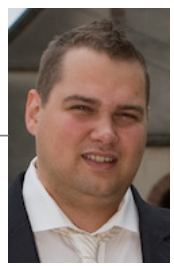

. . . and computational scientists / domain experts from across ALL the University

Sites / Data centers

Kirchberg	Belval
CS.43, AS. 28	Biotech I, CDC/MSA

2 sites, $\geq$ 4 server rooms

Sites / Data centers

Kirchberg

CS.43, AS. 28

Belval

Biotech I, CDC/MSA

2 sites, $\geq$ 4 server rooms

UL HPC: General cluster organization

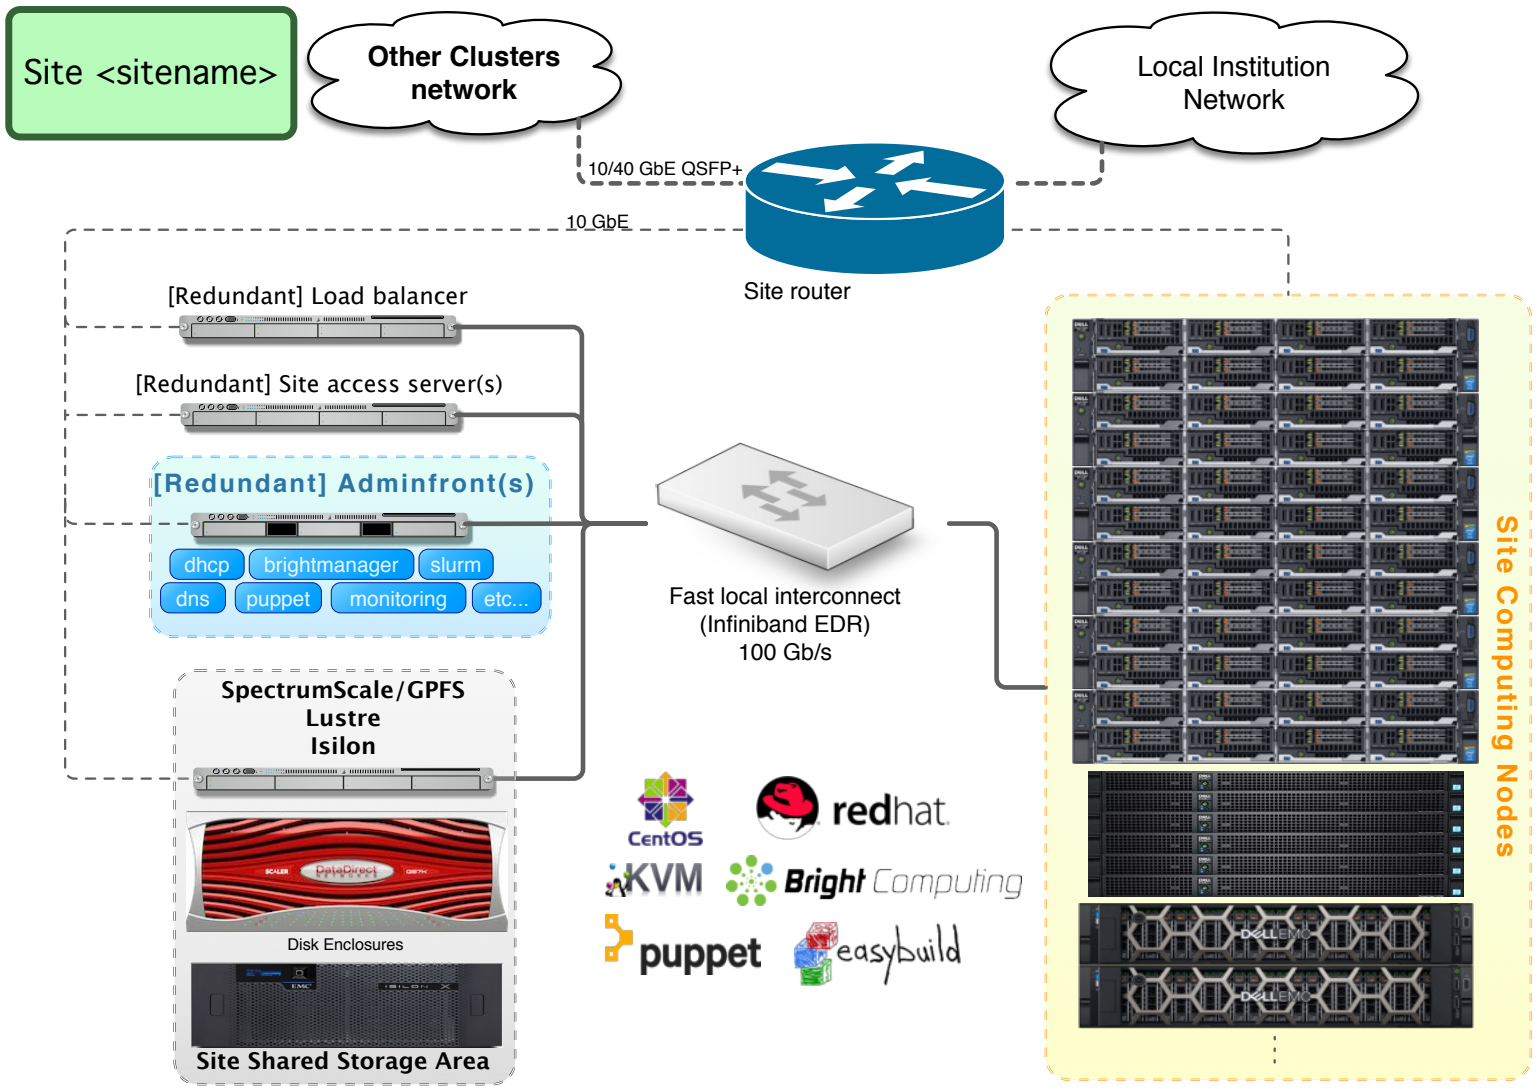

The flagship `iris` cluster

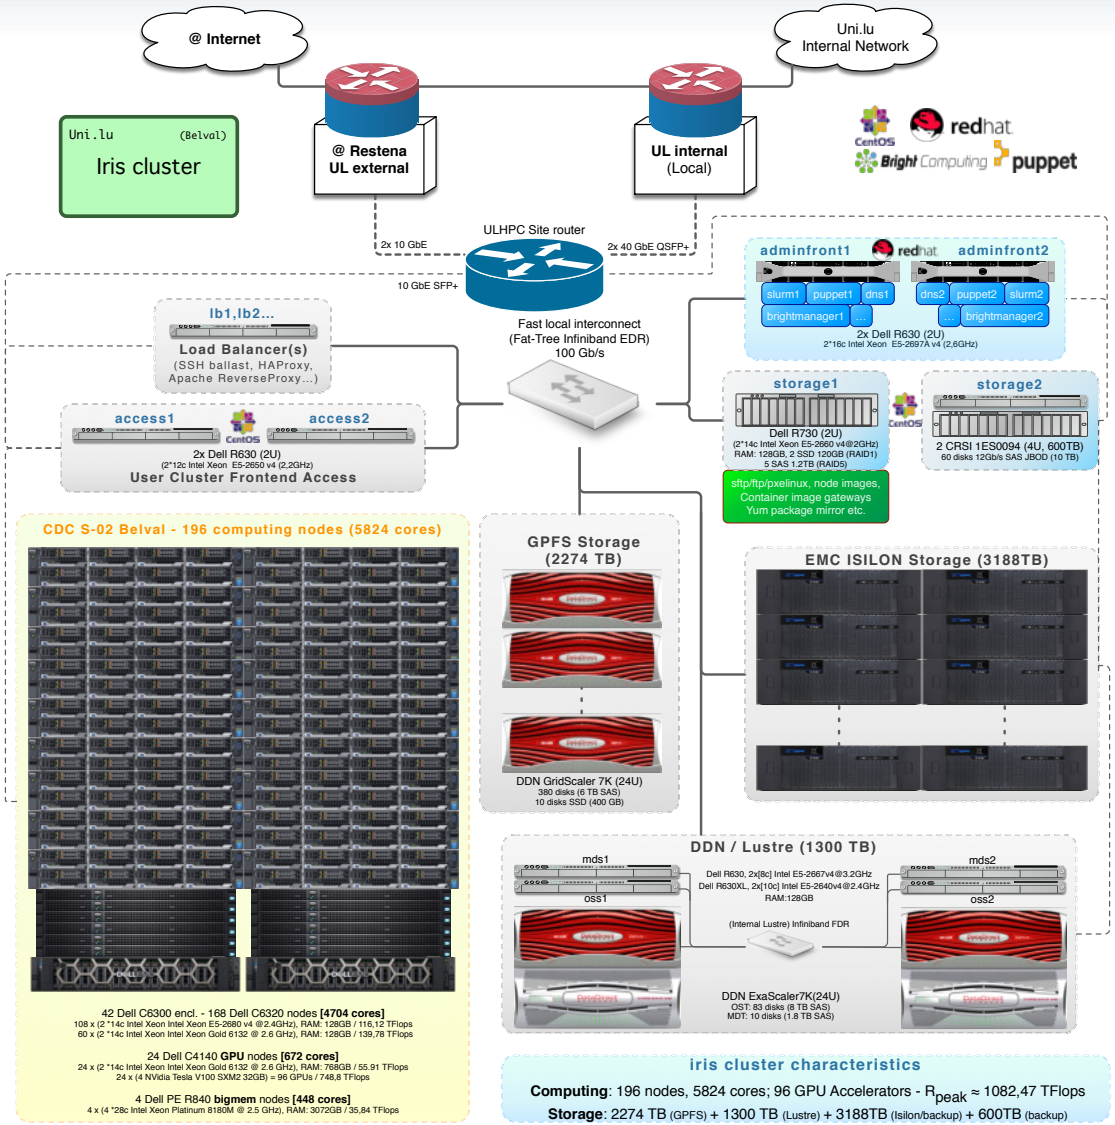

UL HPC Computing capacity

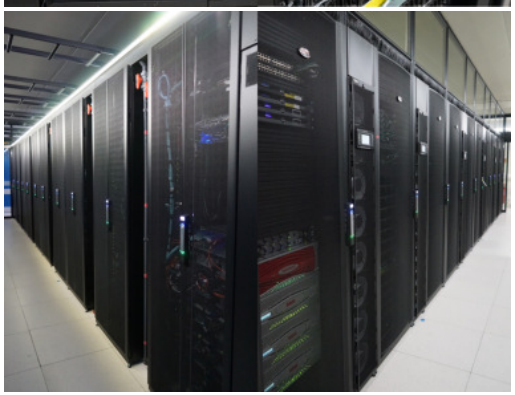

~~5 clusters / 2 sites~~

1 cluster / 1 site

1263.322 TFlops

(incl. 825.02 GPU TFlops)

690 nodes

11252 CPU cores

(+ 612224 GPU cores)

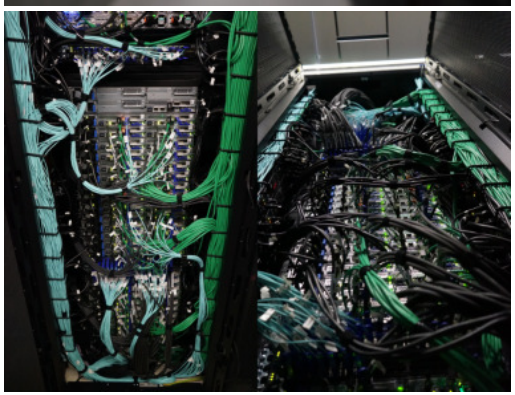

- IB interconnect
- Fat tree topo. in general

UL HPC Computing Clusters

Cluster	Location	**#N**	**#C**	R_{peak} [TFlops]	GPU R_{peak} [TFlops]
`iris`	CDC S-01	196	5824	347.65	748.8
`gaia`*	BT1	273	3440	69.296	76.22
`chaos`*	Kirchberg	81	1120	14.495	0
`g5k`	Kirchberg	38	368	4.48	0
`nyx`* (experimental)	BT1	102	500	2.381	0
	TOTAL:	**690**	**11252**	**438.302**	**+ 825.02 TFlops**

*: *Deprecated end-2019!!*

Uni.lu HPC Total Computing Capacity:
1263.322 TFlops

UL HPC - Detailed Computing Nodes

	#N	#C	R_{peak}
Uni.lu HPC TOTAL:	**690**	**11252**	**1263.322 TFlops**
			(incl. 825.02 GPU TFlops)

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
			Per node: 4x NVIDIA Tesla V100 SXM2 16 or 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**
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**
Testing cluster:						
nyx, viridis, pyro...	2012	Dell	Intel Xeon E5-2420@1.9GHz 1 × 6C,32GB	2	12	0.091 TFlops
	2013	Viridis	ARM A9 Cortex@1.1GHz 1 × 4C,4GB	96	384	0.422 TFlops
	2015	Dell	Intel Xeon E5-2630Lv2@2.4GHz 2 × 6C,32GB	2	24	0.460 TFlops
	2015	Dell	Intel Xeon E5-2660v2@2.2GHz 2 × 10C,32GB	4	80	1.408 TFlops
			`nyx/viridis` TOTAL:	**102**	**500**	**2.381 TFlops**

Decommissioned clusters (Planned END-2019)

Cluster	Start	End	Vendors	#N	#C	R_{peak}
gaia	[2011-2016]	2019	Bull/Dell/HPE/SGi/Delta	**273**	**3440**	**69.296 TFlops**
			Intel architecture (Gulftown, Sandy/Ivy Bridge, Haswell)			
			incl. 21 GPU nodes (2 to 4 NVidia Tesla per node)	50 GPUs	120704	**+76.22 GPU TFlops**
chaos	[2010-2012]	2019	Dell/HPE	**81**	**1120**	**14.495 TFlops**

UL HPC Storage capacity

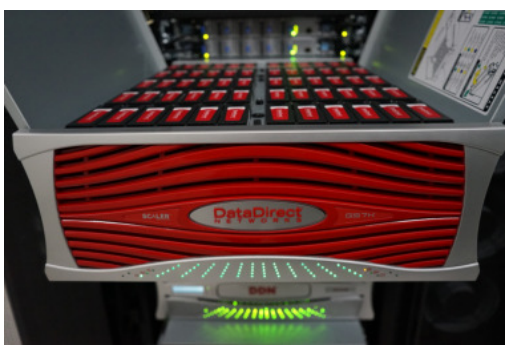

9852.4 TB (incl. 1020TB for Backup)

2425 disks

- 4 distributed/parallel FS
  - ↪ GPFS : 3244 TB
  - ↪ Lustre: 1940 TB
  - ↪ OneFS: 3188 TB...

UL HPC Shared Storage Capacities

Cluster	GPFS	Lustre	Other	Backup
iris	2284	1280	6/3188[2]	600
gaia[1]	960	660	0/3188[2]	240
chaos[1]	0	0	180	180
g5k	0	0	32.4	0
nyx[1] (experimental)	0	0	242	0
TOTAL:	**3244 TB**	**1940 TB**	**3648.4 TB**	**1020 TB**

[1]: *Deprecated end-2019!!*

[2]: *Common Isilon/OneFS shared storage mounted on gaia and iris*

Uni.lu HPC Total Storage Capacity:
9852.4 TB

UL HPC Software Stack

- **Operating System**: **Linux** CentOS 7 (`iris`), ~~Debian 8 (others)~~
- **Remote connection to the platform**: SSH
- **User SSO**: IPA, OpenLDAP
- **Resource management**: job/batch scheduler: **Slurm**(`iris`), OAR
- **(Automatic) Computing Node Deployment**:
  - ↪ ~~FAI (Fully Automatic Installation)(`gaia`, `chaos` clusters)~~
  - ↪ Bright Cluster Manager (`iris`)
  - ↪ Puppet
  - ↪ Kadeploy
- **Platform Monitoring**:
  - ↪ OAR Monika/Drawgantt, Ganglia, Allinea Perf Report, SLURM
  - ↪ Icinga, NetXMS, PuppetBoard etc.
- **Commercial Softwares**:
  - ↪ ANSYS, ABAQUS, MATLAB, Intel Cluster Studio XE, Allinea DDT, Stata etc.

The case of Grid'5000

http://www.grid5000.fr

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

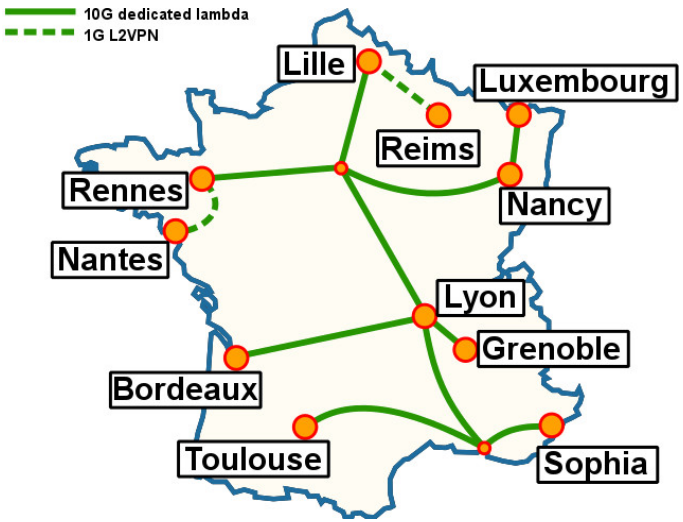

- 8 sites, 7 in France
  - ↪ **Abroad**: Luxembourg
  - ↪ Total: **12326** cores over **31** clusters
- 1-10GbE / Infiniband
  - ↪ **10Gb/s dedicated** between all sites
- Unique software stack
  - ↪ **kadeploy, kavlan, kwapi**

The case of Grid'5000

http://www.grid5000.fr

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

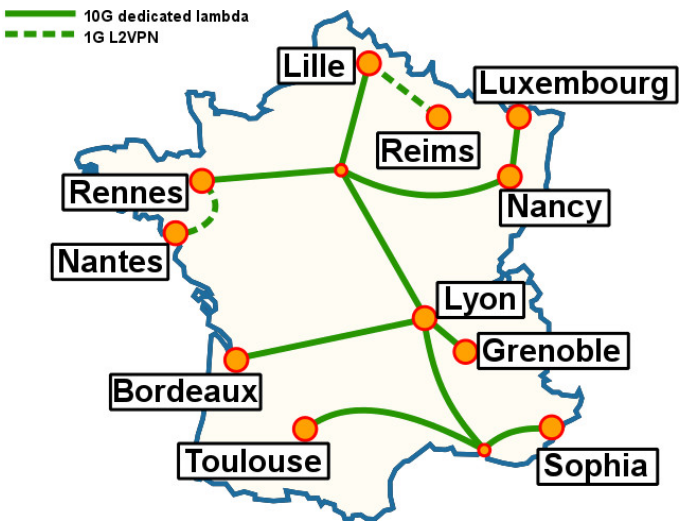

- 8 sites, 7 in France
  - ↪ **Abroad**: Luxembourg
  - ↪ Total: **12326** cores over **31** clusters
- 1-10GbE / Infiniband
  - ↪ **10Gb/s dedicated** between all sites
- Unique software stack
  - ↪ **kadeploy, kavlan, kwapi**

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

Uni.lu HPC Users

- **496** registered HPC Users

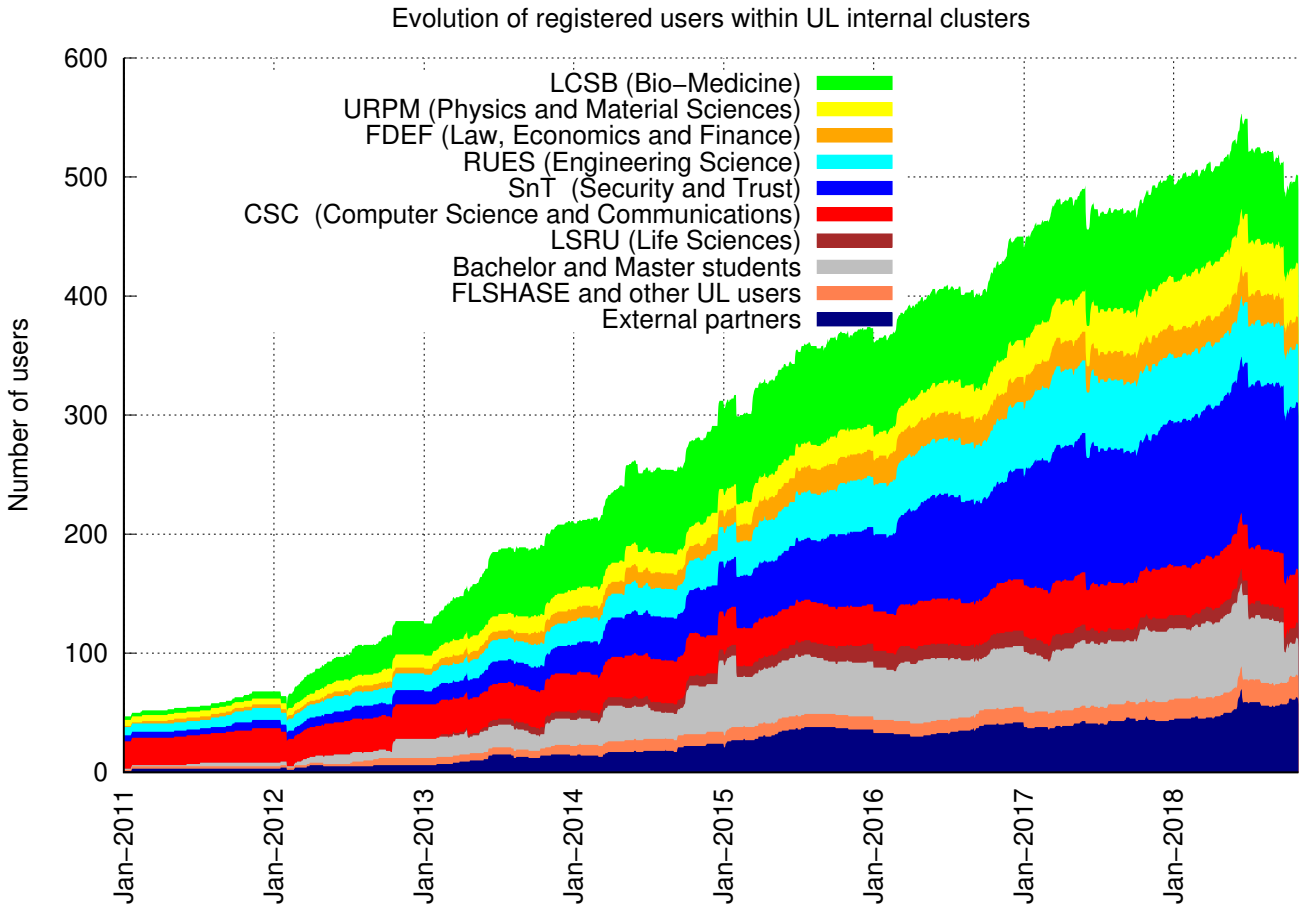

Past Year Job Statistics (June 2018 → June 2019)

Cluster	Allocated [CPU Hours]	Typical Command
`iris`	10,997,387	`sreport -t hours cluster Utilization start=[...]`
`gaia`	12,298,255	`oarstat --accounting "2018-06-01, 2019-06-01" [...]`
`chaos`	3,925,224	
TOTAL	42,153,043 CPU Hours	

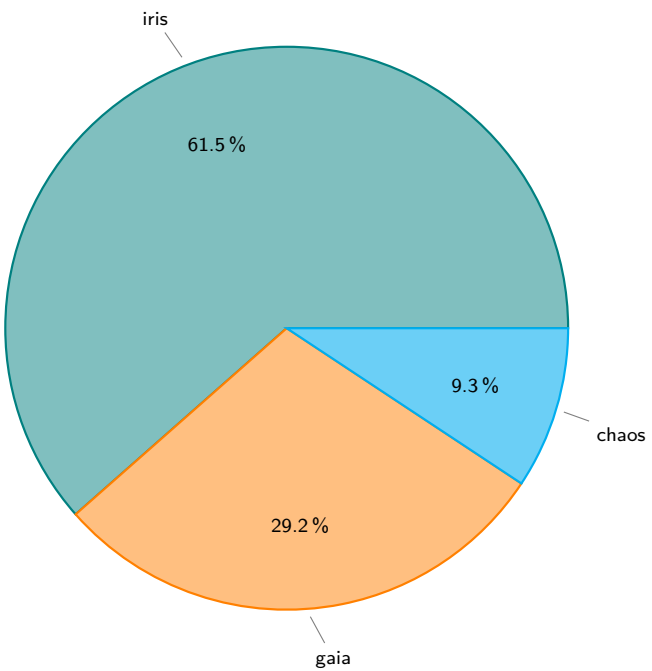

Past Year Job Statistics (June 2018 → June 2019)

Cluster	Allocated [CPU Hours]	Typical Command
`iris`	10,997,387	`sreport -t hours cluster Utilization start=[...]`
`gaia`	12,298,255	`oarstat --accounting "2018-06-01, 2019-06-01" [...]`
`chaos`	3,925,224	
TOTAL	42,153,043 CPU Hours	**... aka 4612 CPU Years**

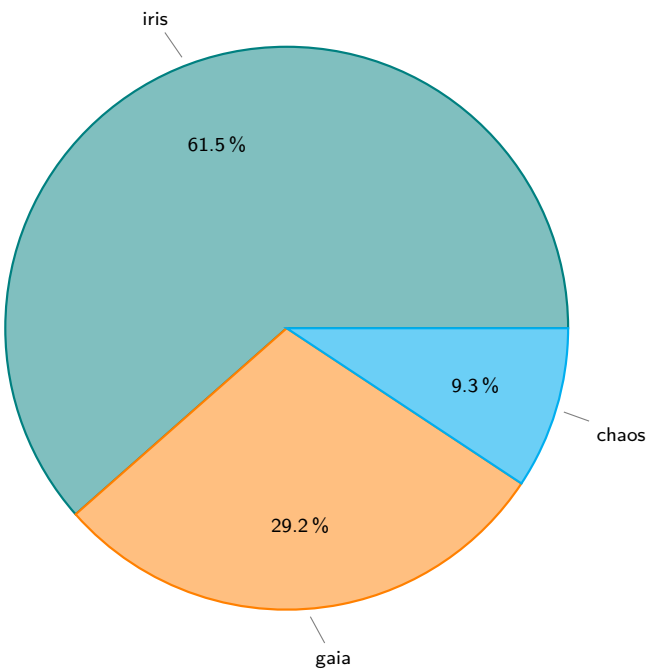

Past Year Job Statistics

(June 2018 → June 2019)

Cluster	Allocated [CPU Hours]	Typical Command
`iris`	10,997,387	`sreport -t hours cluster Utilization start=[...]`
`gaia`	12,298,255	`oarstat --accounting "2018-06-01, 2019-06-01" [...]`
`chaos`	3,925,224	
TOTAL	42,153,043 CPU Hours	**... aka 4612 CPU Years**

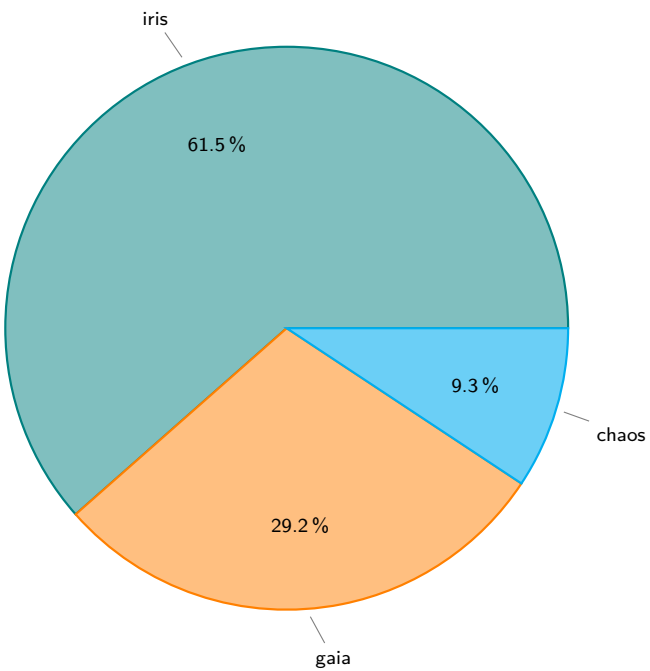

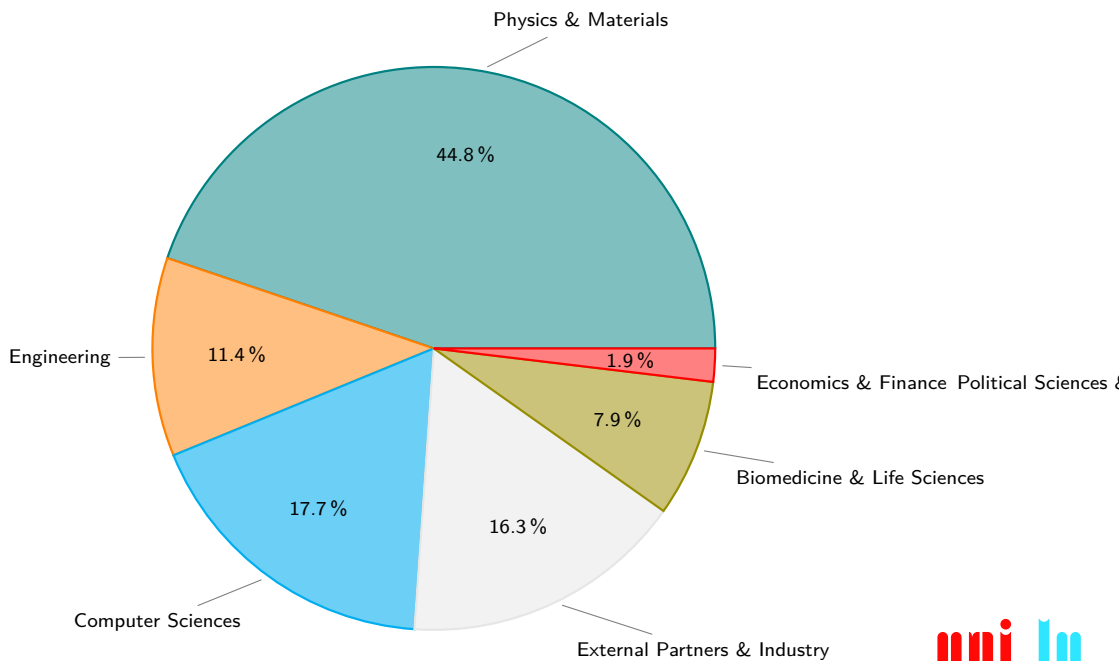

Cluster Utilization (`iris`)

- Availability since Jan 2019: **96.59%**

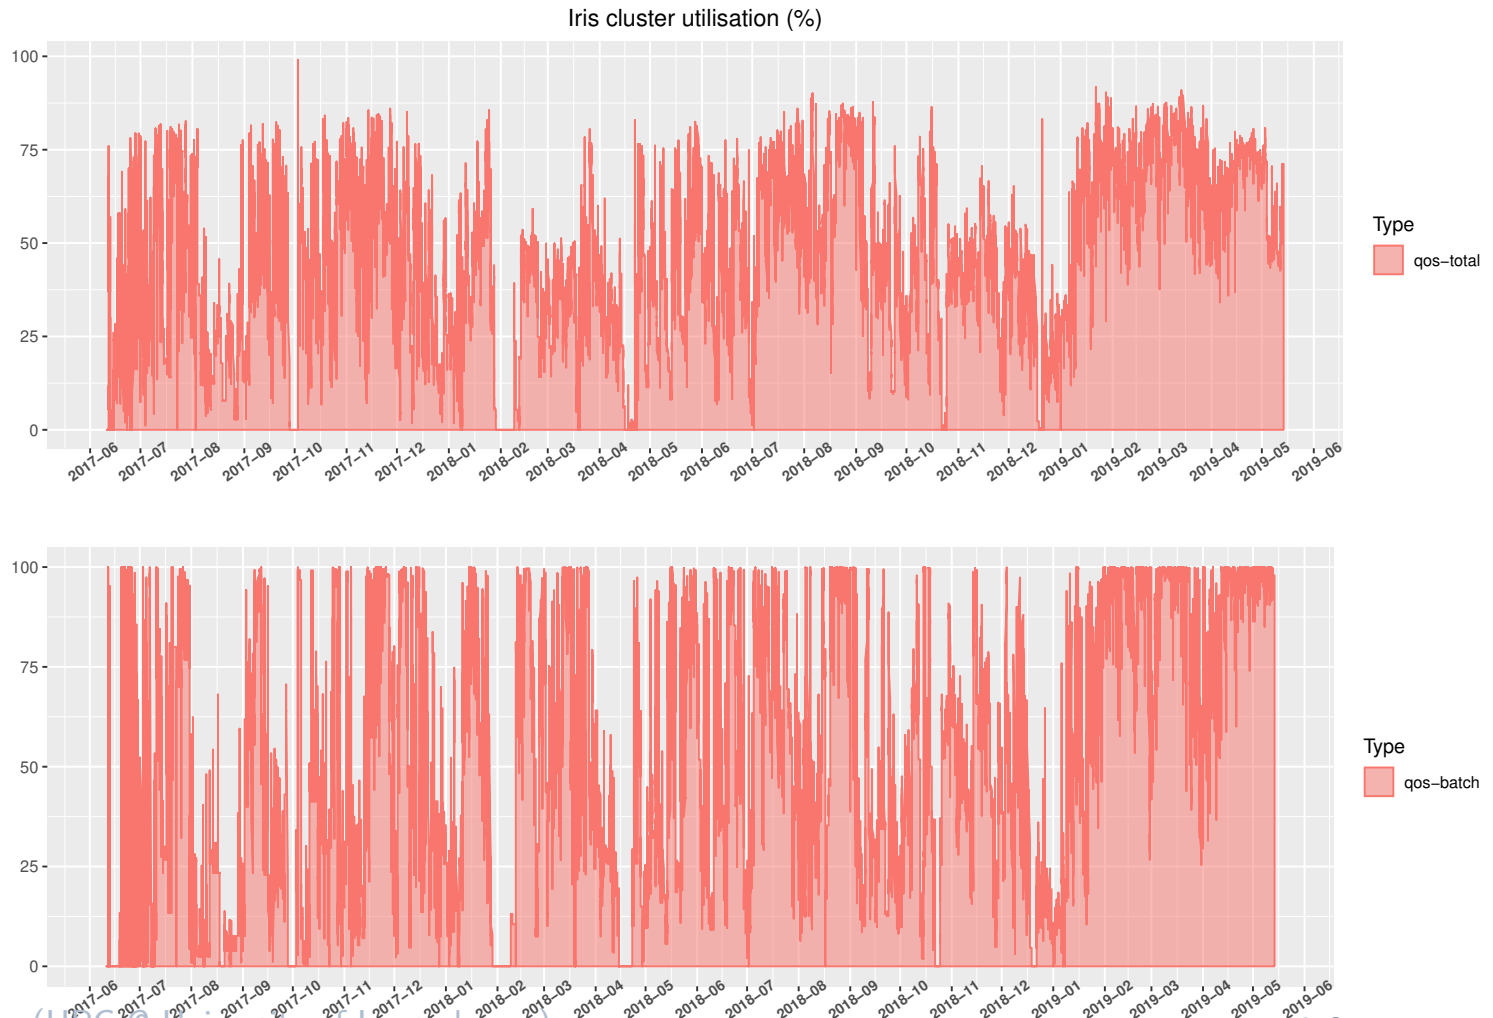

Summary

Computing nodes Management

Node deployment by FAI/Bright Manager

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

Fully Automatic Installation

Bright Computing

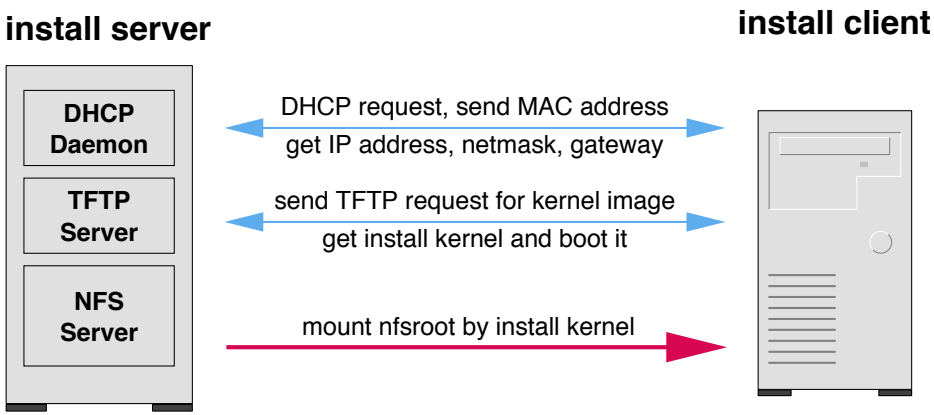

Computing nodes Management

Node deployment by FAI/Bright Manager

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

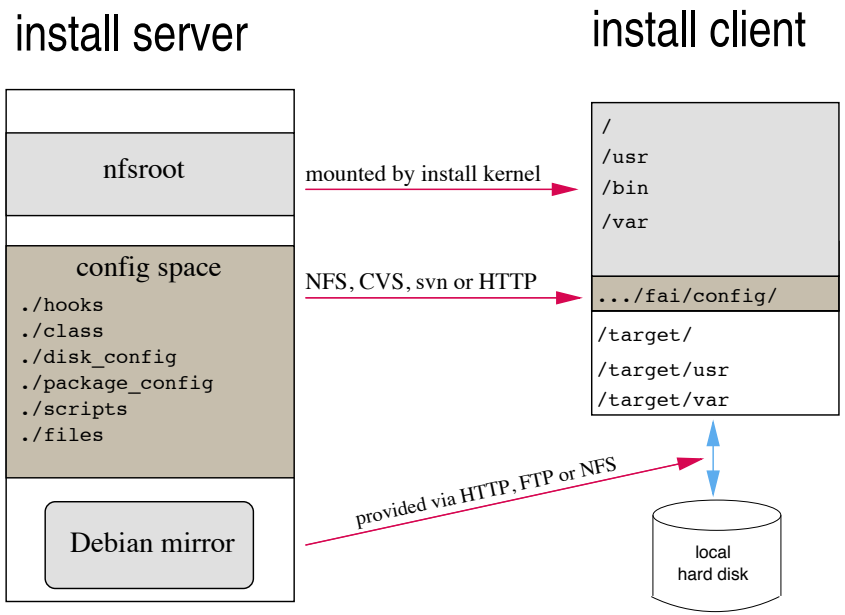

Computing nodes Management

Node deployment by FAI/Bright Manager

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

Computing nodes Management

Node deployment by FAI/Bright Manager

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

Average (full) reinstallation time: $\simeq$ 600s

IT Serv[er|ice] Management: Puppet

Server/Service configuration by Puppet

- **IT Automation** for configuration management
  - ↪ idempotent, agent/master OR stand-alone
  - ↪ cross-platform through Puppet Resource Abstraction Layer (RAL)
  - ↪ git-based workflow with r10k (*role & profiles* workflow)
  - ↪ PKI-based security (X.509)
- DevOps tool of choice for configuration management
  - ↪ Reusable modules
  - ↪ per-environment hierarchy lookup with hiera

http://puppetlabs.com

https://forge.puppet.com/

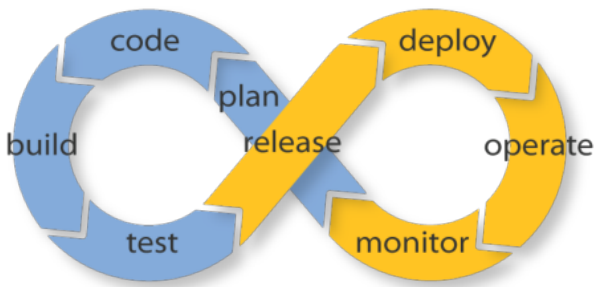

Endless Possibilities: DevOps can create an infinite loop of release and feedback for all your code and deployment targets.

IT Serv[er|ice] Management: Puppet

Server/Service configuration by Puppet

- **IT Automation** for configuration management
  - ↪ idempotent, agent/master OR stand-alone
  - ↪ cross-platform through Puppet Resource Abstraction Layer (RAL)
  - ↪ git-based workflow with r10k (*role & profiles* workflow)
  - ↪ PKI-based security (X.509)
- DevOps tool of choice for configuration management
  - ↪ Reusable modules
  - ↪ per-environment hierarchy lookup with hiera

`http://puppetlabs.com`

`https://forge.puppet.com/`

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

General Puppet Infrastructure

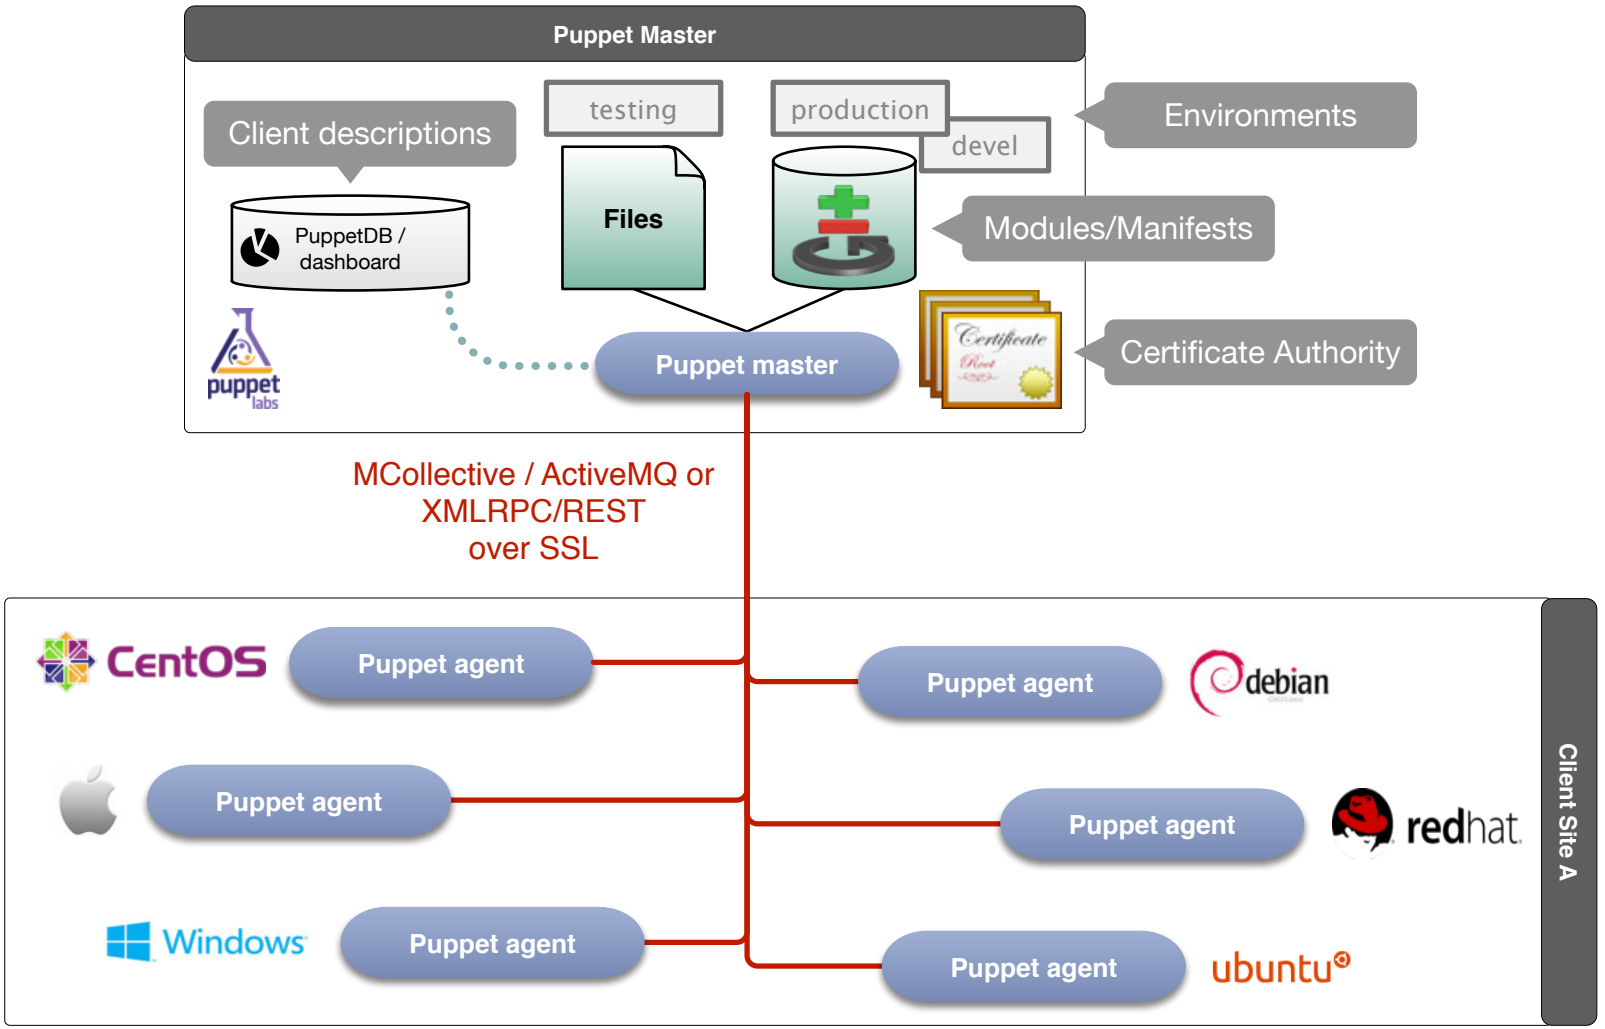

Software/Modules Management

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

- Based on Environment Modules / LMod
  - ↪ convenient way to dynamically change the users environment `$PATH`
  - ↪ permits to easily load software through `module` command
- Currently on UL HPC:
  - ↪ > **210 software packages**, in *multiple* versions, within **18 categ.**
  - ↪ reworked software set for `iris` cluster and now deployed everywhere
    - ✓ RESIF v2.0, allowing [real] semantic versioning of released builds
  - ↪ hierarchical organization **Ex**: `toolchain/{foss,intel}`

```
$> module avail                                    # List available modules
```

```
$> module load <category>/<software>[/<version>]
```

Software/Modules Management

- Key `module` variable: `$MODULEPATH` / where to look for modules
  ↪ altered with `module use <path>`. **Ex**:

```
export EASYBUILD_PREFIX=$HOME/.local/easybuild
export LOCAL_MODULES=$EASYBUILD_PREFIX/modules/all
module use $LOCAL_MODULES
```

Software/Modules Management

- Key `module` variable: `$MODULEPATH` / where to look for modules
  ↪ altered with `module use <path>`. **Ex**:

```
export EASYBUILD_PREFIX=$HOME/.local/easybuild
export LOCAL_MODULES=$EASYBUILD_PREFIX/modules/all
module use $LOCAL_MODULES
```

Main modules commands:

Command	Description
`module avail`	Lists all the modules which are available to be loaded
`module spider <pattern>`	Search for among available modules **(Lmod only)**
`module load <mod1> [mod2...]`	Load a module
`module unload <module>`	Unload a module
`module list`	List loaded modules
`module purge`	Unload all modules (purge)
`module display <module>`	Display what a module does
`module use <path>`	Prepend the directory to the MODULEPATH environment variable
`module unuse <path>`	Remove the directory from the MODULEPATH environment variable

Software/Modules Management

http://hpcugent.github.io/easybuild/

- Easybuild: open-source framework to (automatically) build scientific SW
- **Why?**: *"Could you please install this software on the cluster?"*
  - ↪ Scientific software is often **difficult** to build
    - ✓ non-standard build tools / incomplete build procedures
    - ✓ hardcoded parameters and/or poor/outdated documentation
  - ↪ EasyBuild helps to facilitate this task
    - ✓ **consistent** software **build and installation** framework
    - ✓ includes testing step that helps validate builds
    - ✓ **automatically generates LMod modulefiles**

```
$> module use $LOCAL_MODULES
$> module load tools/EasyBuild
# Search for recipes for a given software
$> eb -S Spark
$> eb Spark-2.4.0-Hadoop-2.7-Java-1.8.eb -Dr # Dry-run install
$> eb Spark-2.4.0-Hadoop-2.7-Java-1.8.eb -r
```

Software/Modules Management

http://hpcugent.github.io/easybuild/

- Easybuild: open-source framework to (automatically) build scientific SW
- **Why?**: *"Could you please install this software on the cluster?"*
  - ↪ Scientific software is often **difficult** to build
    - ✓ non-standard build tools / incomplete build procedures
    - ✓ hardcoded parameters and/or poor/outdated de
  - ↪ EasyBuild helps to facilitate this t
    - ✓ **consistent** soft
    - ✓
    - ✓

Also covered in a Uni.lu HPC tutorial!

http://ulhpc-tutorials.rtfd.io/en/latest/tools/easybuild/

```
$> module use
$> module loa
# Search for
$> eb -S Spark
$> eb Spark-2.4.0-Hadoop-2.7-Java-1.8.eb -Dr # Dry-run install
$> eb Spark-2.4.0-Hadoop-2.7-Java-1.8.eb -r
```

Software/Modules Management

- **RESIF**: Revolutionary EasyBuild-based Software Installation Framework
  - ↪ Automatic Management of **software sets**
  - ↪ Fully automates software builds and supports all available toolchains
  - ↪ Clean (hierarchical) modules layout to facilitate its usage
  - ↪ (**incoming**) per-release module, per-ISA builds,

```
MODULEPATH=/opt/apps/resif/data/{devel,production,stable,testing}/default/modules/all/
```

RESIF: Revolutionary EasyBuild-based Software Installation Framework

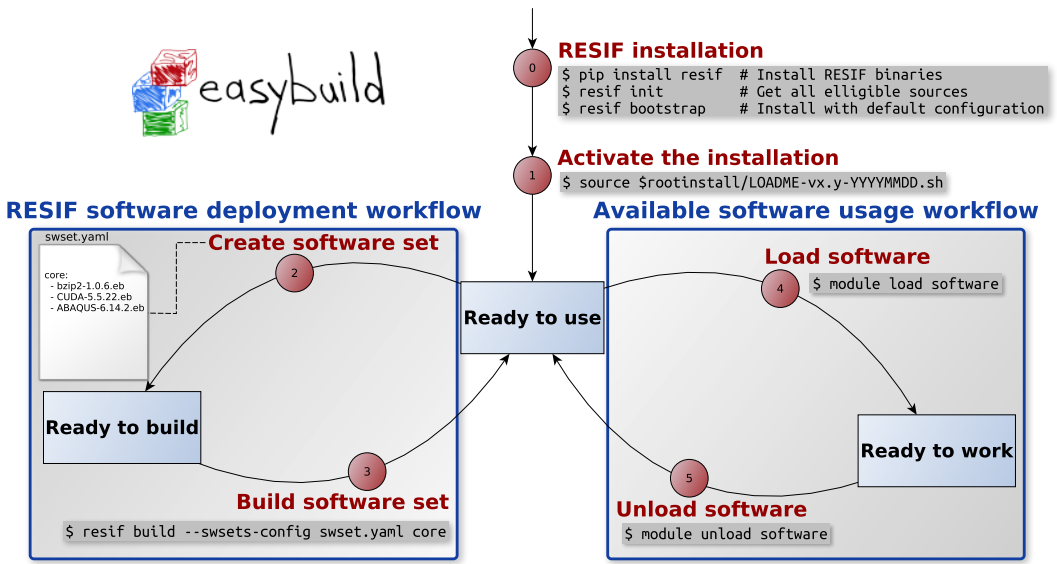

Platform Monitoring

- **General Live Status** `http://hpc.uni.lu/status/overview.html`

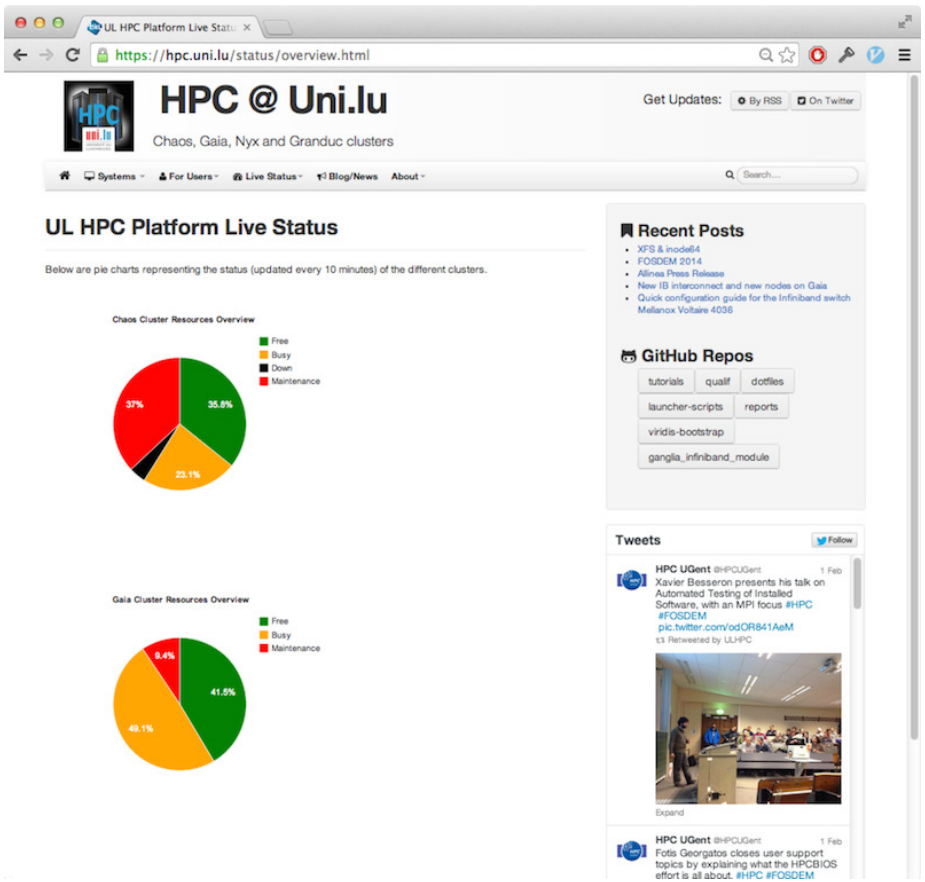

Platform Monitoring

- **Ganglia**

`http://hpc.uni.lu/{gaia,chaos,g5k,iris}/ganglia`

Platform Monitoring

- **SLURM-Web**

http://hpc.uni.lu/iris/slurm/

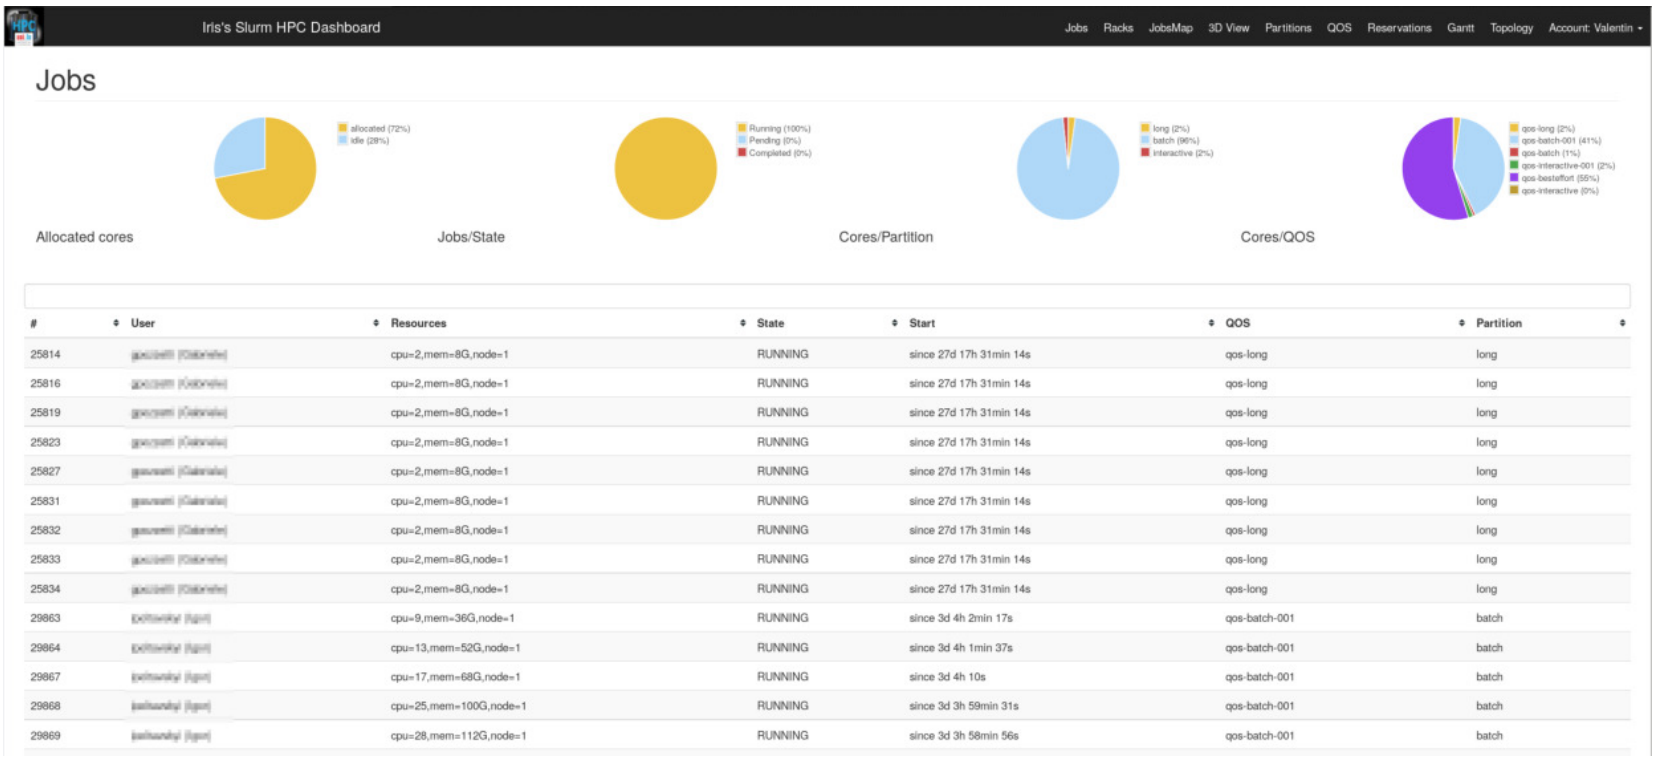

#	User	Resources	State	Start	QOS	Partition
25814	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25816	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25819	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25823	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25827	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25831	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25832	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25833	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25834	[illegible]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
29863	[illegible]	cpu=9,mem=36G,node=1	RUNNING	since 3d 4h 2min 17s	qos-batch-001	batch
29864	[illegible]	cpu=13,mem=52G,node=1	RUNNING	since 3d 4h 1min 37s	qos-batch-001	batch
29867	[illegible]	cpu=17,mem=68G,node=1	RUNNING	since 3d 4h 10s	qos-batch-001	batch
29868	[illegible]	cpu=25,mem=100G,node=1	RUNNING	since 3d 3h 59min 31s	qos-batch-001	batch
29869	[illegible]	cpu=28,mem=112G,node=1	RUNNING	since 3d 3h 58min 56s	qos-batch-001	batch

Platform Monitoring

- **Internal Monitoring**

Icinga / Puppet / NetXMS (networking)

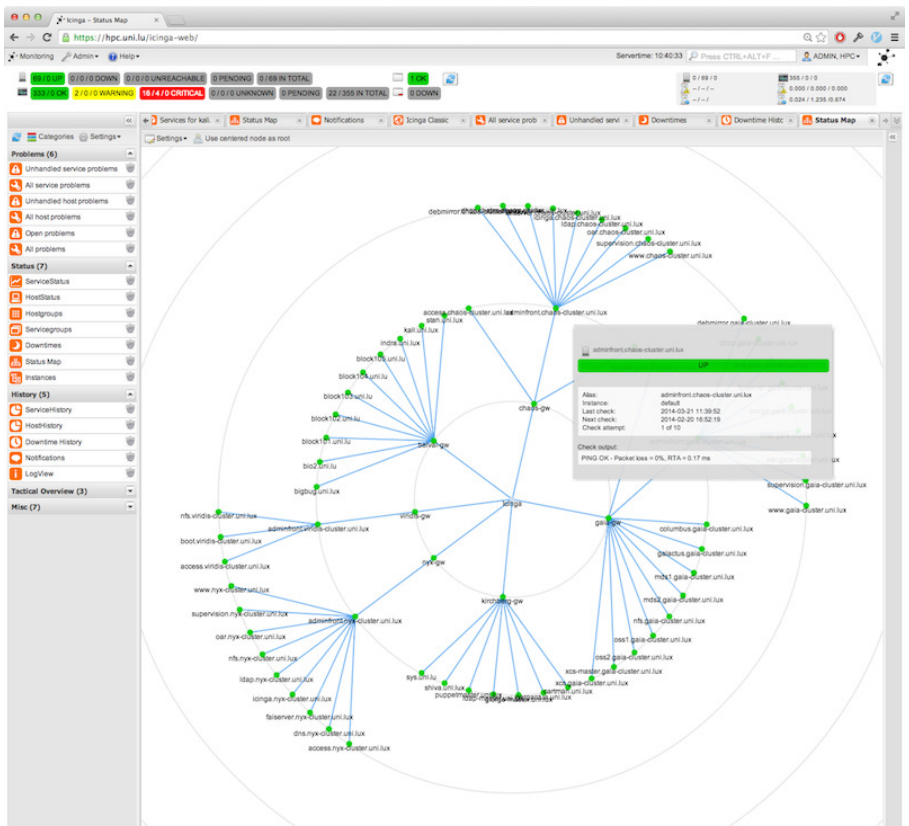

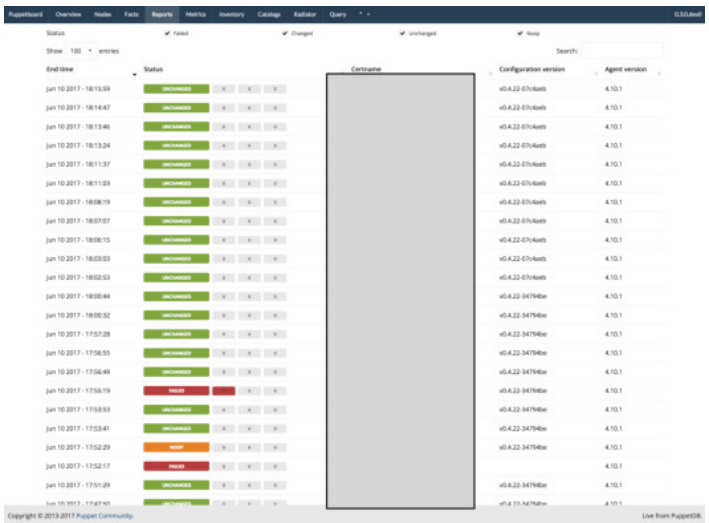

Platform Monitoring

- **Internal Monitoring**

[Disk] Enclosure status

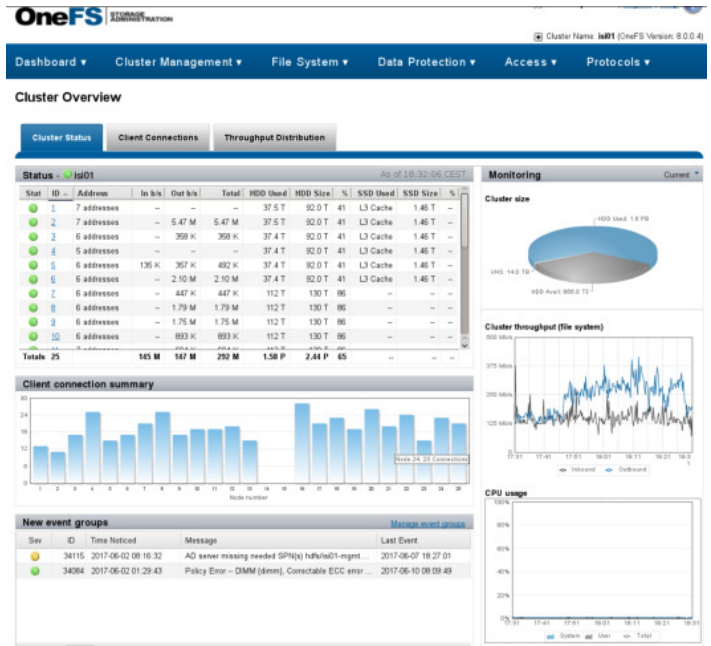

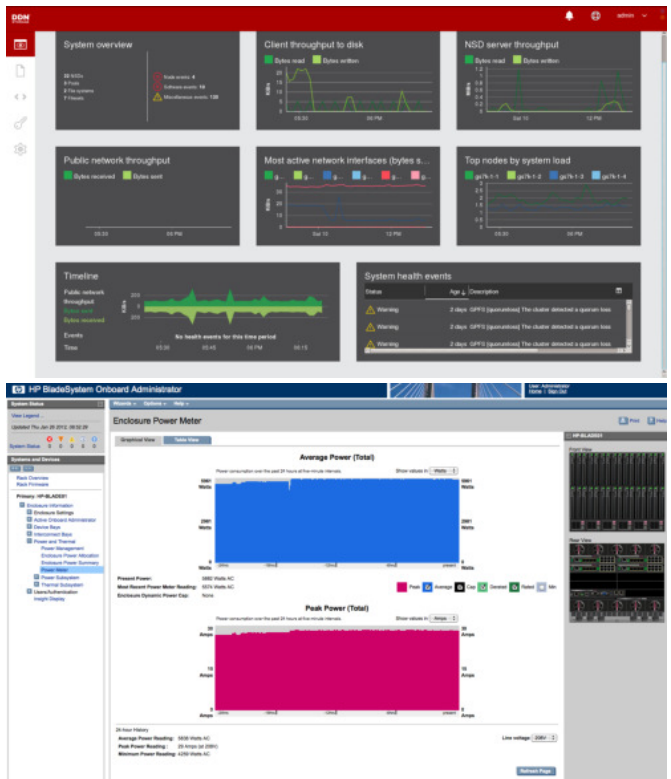

Summary

Past Achievements / Technical

- **2018/2019: new GPU/Bigmem nodes in `iris`**
  - ↪ RFP 180027: **22 GPU/Bigmem nodes (initially)**,
    - ✓ **18 GPU nodes** x 4 Nvidia V100 SXM2 ~~32GB~~ 16GB
    - ✓ 4 x 4 Intel Xeon Platinum 8180M@2.5 GHz [4x28c, 3TB RAM]
  - ↪ Deployment in Dec. 2018
    - ✓ Wrong GPU cards delivered (16GB instead of 32GB)
    - ✓ HPL CUDA on 18 GPU nodes: 282.4 TFlops (52.65% efficiency)
  - ↪ Long compensation transactions until Apr. 2019
    - ✓ ... leading to delivery of 6 new GPU nodes (with 32GB cards)




Past Achievements / Technical

- **2018/2019: new GPU/Bigmem nodes in `iris`**
  - ↪ **24 GPU nodes** x 4 Nvidia V100 SXM2 {16,32}GB +748.8 GPU TFlops
    - ✓ Measured *HPL* performances: 378 TFlops. . .
      . . . yet meant for AI-oriented workflow
    - ✓ part of `gpu` partition.
    - ✓ You **HAVE TO** explicitly reserve the GPU cards: `--gres:gpu:<N>`

Past Achievements / Technical

- **2018/2019: new GPU/Bigmem nodes in `iris`**
  - ↪ **24 GPU nodes** x 4 Nvidia V100 SXM2 {16,32}GB +748.8 GPU TFlops
    - ✓ Measured *HPL* performances: 378 TFlops. . .
      . . . yet meant for AI-oriented workflow
    - ✓ part of `gpu` partition.
    - ✓ You **HAVE TO** explicitly reserve the GPU cards: `--gres:gpu:<N>`
  - ↪ **4 Bigmem nodes** 4x28c i.e. 112 cores per node
    - ✓ part of the `bigmem` partition

Compute nodes	Features	SBATCH option	sbatch/srun command line
`iris-[169-186]`	skylake,volta	`#SBATCH -p gpu`	`sbatch -p gpu --gres=gpu:[1-4]`
`iris-[191-196]`	skylake,volta,volta32		`[...]`
`iris-[187-190]`	skylake	`#SBATCH -p bigmem`	`sbatch -p bigmem`

Past Achievements / Technical

- **2018/2019: new GPU/Bigmem nodes in `iris`**
  - ↪ **24 GPU nodes** x 4 Nvidia V100 SXM2 {16,32}GB +748.8 GPU TFlops
    - ✓ Measured *HPL* performances: 378 TFlops. . .
      . . . yet meant for AI-oriented workflow
    - ✓ part of `gpu` partition.
    - ✓ You **HAVE TO** explicitly reserve the GPU cards: `--gres:gpu:<N>`
  - ↪ **4 Bigmem nodes** 4x28c i.e. 112 cores per node
    - ✓ part of the `bigmem` partition

Compute nodes	Features	SBATCH option	sbatch/srun command line
`iris-[169-186]`	skylake,volta	`#SBATCH -p gpu`	`sbatch -p gpu --gres=gpu:[1-4]`
`iris-[191-196]`	skylake,volta,volta32		`[...]`
`iris-[187-190]`	skylake	`#SBATCH -p bigmem`	`sbatch -p bigmem`

Past Achievements / Technical

- **2018/2019: new GPU/Bigmem nodes in `iris`**
  - ↪ **24 GPU nodes** x 4 Nvidia V100 SXM2 {16,32}GB +748.8 GPU TFlops
    - ✓ Measured *HPL* performances: 378 TFlops
    - ... yet meant for AI-oriented
    - ✓ part of gpu parti
    - ✓ X

Learn More w. Uni.lu HPC Tutorials!

- PS5c: Scalable Science III (GPU programming)
- PS6: Big Data Applications
- PS13{a,b}: Deep Learning I & II

ir		sbatch -p gpu --gres=gpu:[1-4]
ir		[...]
iri	#SBATCH -p bigmem	sbatch -p bigmem

Past Achievements

```
$> srun -p gpu --gres=gpu:2 --ntasks-per-node 2  --ntasks-per-socket 1 -c 14 --pty bash
$> nvidia-smi
Wed Jun 19 12:04:20 2019
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 418.40.04    Driver Version: 418.40.04    CUDA Version: 10.1     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla V100-SXM2...  On   | 00000000:1A:00.0 Off |                    0 |
| N/A   39C    P0    43W / 300W |      0MiB / 32480MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   1  Tesla V100-SXM2...  On   | 00000000:1C:00.0 Off |                    0 |
| N/A   36C    P0    42W / 300W |      0MiB / 32480MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
```

Past Achievements

LE GOUVERNEMENT
DU GRAND-DUCHÉ DE LUXEMBOURG

NVIDIA.

- **Jan 2019**: Collaboration agreement Luxembourg/NVidia
  - ↪ creation of a **joint AI Lab in Luxembourg**
  - ↪ Luxembourg consortium: SMC, Uni.lu & LIST
    - ✓ Uni.lu: ULHPC Team, LCSB, SnT
    - ✓ **Coordinator**: Dr. F. Pinel
  - ↪ official gouvernement communication

Partnership with NVIDIA Boosts Research in Artificial Intelligence

Share this article:

Published on Thursday, 31 January 2019

Luxembourg's research community, the Luxembourg Government, and NVIDIA announced today the creation of a joint AI laboratory in Luxembourg. This national AI collaboration, the first with NVIDIA in Europe, is a milestone towards working together to solving society's most important challenges using artificial intelligence and high-performance computing.

NVIDIA, the University of Luxembourg's Interdisciplinary Centre for Security, Reliability and Trust (SnT), its Luxembourg Centre for Systems Biomedicine (LCSB), its High-Performance Computing Team, and the Luxembourg Institute for Science and Technology (LIST) have initially committed to dedicate a team of six people to work on joint projects. Each institution will gain better access to cutting-edge NVIDIA hardware and software, for use in tackling their most difficult research problems. Enabled by the government's "Digital Luxembourg" initiative, the joint AI laboratory will encourage research in diverse fields.

"This partnership will give our researchers access to cutting-edge equipment," said University of Luxembourg Rector Stéphane Pallage. "But most importantly, it will enable them to tap into NVIDIA's experience in addressing real-world problems using artificial intelligence. This will allow them to both identify new application areas and to push ahead with existing work, from our use of drones for automated airplane and bridge inspections to the analysis of genomes and mobile health sensor data. We are also particularly excited about the interdisciplinary nature of the partnership. Tackling real-world problems involves bringing together experts from across disciplines."

Initially bootstrapped by the research community, the founders plan to open up the AI laboratory to Luxembourg's industry partners and startup community. There is a growing demand for AI-related technologies in Luxembourg, including in finance, healthcare, space resources research and more.

Learn more here

Past Achievements

- **Jan/Feb 2019**: First **HPC Service contracts** with Industry
  - ↪ **coordinator** Uni.lu HPC Team: V. Plugaru
  - ↪ Signed agreements so far:
    - ✓ Arcelor Mittal, Ceratizit
  - ↪ Yearly contract, tied to `iris` only,
    - ✓ cannot exceed 10% of the platform capacity
  - ↪ Weighted charging
    - ✓ **Pre-defined Max computing capacity**
    - ✓ Slurm-based accounting
    - ✓ Shared project directory
  - ↪ Expert-level service and support
    - ✓ initial 4 hour kickstart free of charge
  - ↪ Billing at three intervals
    - ✓ signature - 6 m - 1Y

```
sreport -t hours cluster AccountUtilizationByUser Accounts=<Company> [...]
```

Past Achievements

University grants ArcelorMittal and CERATIZIT access to Supercomputer

Share this article:

Published on Tuesday, 21 May 2019

The University of Luxembourg has entered service agreements with Luxembourg-based companies ArcelorMittal and CERATIZIT.

For the first time, this contract will give access to commercial companies to the University's supercomputing capacities. Initially, the projects will run for a duration of one year.

"The service contracts have three main pillars: The companies are granted computing time in our High-Performance Computing (HPC) facilities, they can store the project-related data on our systems, and they have access to our HPC expert group." explains Prof. Pascal Bouvry, leading the University High Performance Computing service.

At the University, the HPC capabilities are gaining importance as a tool for research and teaching activities enabled by a large investment in hardware equipment. With storage capacities of over 10 PetaBytes and a computing capacity of 1 PetaFlops, the University boasts one of the most potent supercomputing infrastructures in the country. While mostly reserved for research, it allocates up to ten percent of its HPC capacities for projects of external partners.

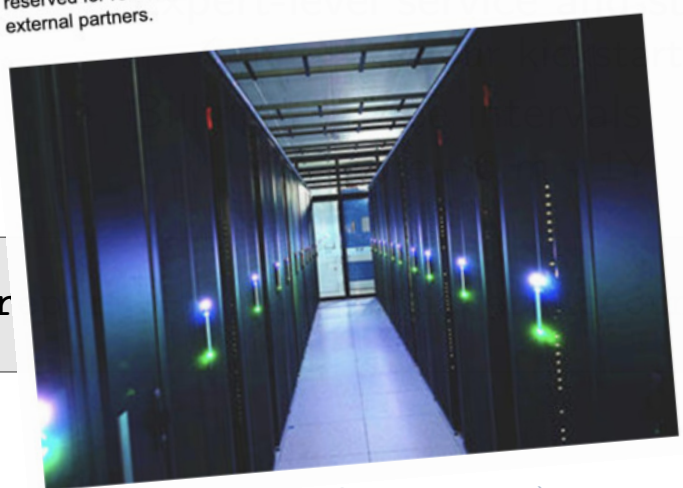

...ith Industry

```
sr                                                    ounts=<Company> [...]
```

Past Achievements

- **Feb/May 2019**: National HPC Competence Center / **Meluxina**
  - ↪ to answer the EuroHPC JU call for Petascale Supercomputers
  - ↪ More details later in this talk

Past Achievements

- **Feb/May 2019**: National HPC Competence Center / **Meluxina**
  - ↪ to answer the EuroHPC JU call for Petascale Supercomputers
  - ↪ More details later in this talk

- **`chaos` and `gaia` will be DECOMMISSIONNED end-2019**
  - ↪ prepare the transition to `iris`!!!

Past Achievements

- **Feb/May 2019**: National HPC Competence Center / **Meluxina**
  - ↪ to answer the EuroHPC JU call for Petascale Supercomputers
  - ↪ More details later in this talk

- **`chaos` and `gaia` will be DECOMMISSIONNED end-2019**
  - ↪ prepare the transition to `iris`!!!

- **Feb-now 2019**: Initiate `gaia`/`chaos` decommissioning
  - ↪ Preparation of the policy, advertisement for Data transfer
  - ↪ High bandwidth link upgrade checks `gaia` (BT1) / `iris` (CDC)
  - ↪ Background project transfer in progress (inode update, Isilon etc.)

```
# Within a screen session
$> rsync --bwlimit=10m --rsh='ssh -p 8022' --exclude="/.local" \
     --exclude="/.cache" -avzP . access-iris.uni.lu:~/gaia_home/
```

Past Achievements

- **Feb/May 2019**: National HPC Competence Center / **Meluxina**
  - ↪ to answer the EuroHPC JU call for Petascale Supercomputers
  - ↪ More details later in this talk

- **`chaos` and `gaia` will be DECOMMISSIONNED end-2019**
  - ↪ prepare the transition to `iris`!!!

- **Feb-now 2019**: Initiate `gaia`/`cl`...

Learn More with Tomorrow's keynote

- Keynote/PS9: Data management (backup, decommissioning...)

```
... --bwlimit=10m --rsh='ssh -p 8022' --exclude="/.local" \
--exclude="/.cache" -avzP . access-iris.uni.lu:~/gaia_home/
```

Past Achievements

- **Mar 2019**: Gitlab 2.0 infrastructure
  - ↪ Now on Gitlab Enterprise Edition
  - ↪ featuring Docker registry etc.
  - ↪ C. Parisot & H. Cartiaux

Past Achievements

- **Mar 2019**: Gitlab 2.0 infrastructure
  - ↪ Now on Gitlab Enterprise Edition
  - ↪ featuring Docker registry etc.
  - ↪ C. Parisot & H. Cartiaux

- **Mar,May 2019**: New HPC Team members
  - ↪ **Dr. F. Pinel** (Research Scientist)
    - ✓ Coordinator NVidia Joint AI Lab
    - ✓ advanced HPC/research support on GPU/AI workflows.
  - ↪ **Dr. E. Kieffer** (Postdoc researcher)
    - ✓ Expert in Bi-level optimization and scalable science with Python
    - ✓ advanced HPC and research support to the HPC community.
  - ↪ **Note**: one more Postdoc arriving Aug 2019
    - ✓ PRACE-6IP Project

Past Achievements / Technical

- Continuous `iris` setup improvement
  - ↪ consolidated **SLURM** configuration
  - ↪ consolidation of the monitoring etc.
  - ↪ **Continuous OS / software modules / security Upgrade**
- **Storage**:
  - ↪ **GPFS** `iris`: CES service recurrent degradation, pending resolution
    - ✓ short maintenance planned mid-july (2 days) to mitigate this issue

Again...

- `chaos` **and** `gaia` **will be DECOMMISSIONNED end-2019**
  - ↪ prepare the transition to `iris`!!!

Past Events

- **HPC Trainings**
  - ↪ 8th Uni.lu HPC School 2018 (bis) Nov 2018
    - ✓ Luxembourg Learning Centre, 1 day event
  - ↪ Data Sciences pilot training for EC employees (Digital Pole) Feb 2019
  - ↪ SCCamp 2019, Madrid, Jun 2019
    - ✓ Trainers: Dr. X. Besseron, C. Parisot

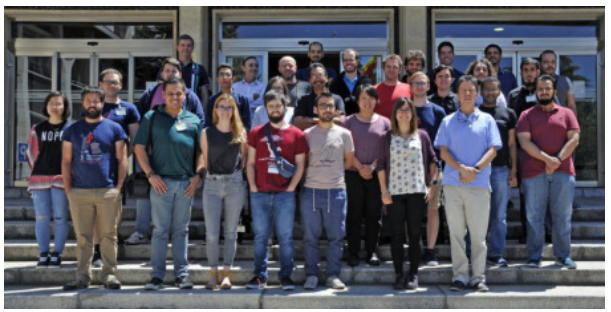

Past Events

- **HPC Trainings**
  - ↪ 8th Uni.lu HPC School 2018 (bis) — Nov 2018
    - ✓ Luxembourg Learning Centre, 1 day event
  - ↪ Data Sciences pilot training for EC employees (Digital Pole) — Feb 2019
  - ↪ SCCamp 2019, Madrid, — Jun 2019
    - ✓ Trainers: Dr. X. Besseron, C. Parisot

- **HPC conferences and exhibitions**:
  - ↪ SC'18, Dallas, US — Nov 2018
  - ↪ FOSDEM, Brussels, — Feb 2019
  - ↪ SCAsia'19, Singapore — Mar 2019
  - ↪ EuroHPC Summit Week'19, Poznan — May 2019
  - ↪ ISC'19, Frankfurt, DE — Jun 2018

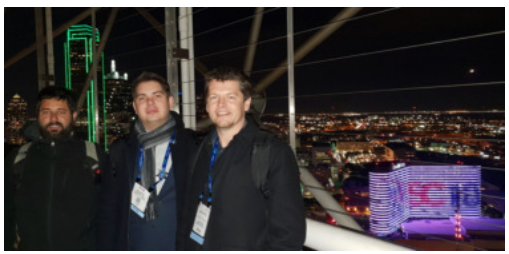

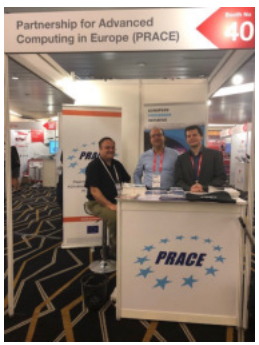

Past Events

- **HPC EU/International Coordination**
  - ↪ EU PRACE council meetings (every 3 months)
    - ✓ Delegate: Prof. Pascal Bouvry
    - ✓ Advisor: Dr. Sebastien Varrette
  - ↪ Support for development of ThaiSC
    - ✓ National Supercomputer Centre in Thailand / NECTEC
  - ↪ EC Enhanced Regional EU-ASEAN Dialogue Instrument (E-READI)
    - ✓ 1st EU-ASEAN HPC Coordination Group Meeting Mar 2019

Uni.lu CDC (Centre de Calcul)

- **Toward Energy-Efficient HPC enabling DLC**
  - ↪ 2x500 m^2 deployed since 2015, one floor for HPC developments

Location	Cooling	Usage	Max Capacity [kW]
CDC S-02-001	Airflow	*Future extension*	280 kW (120 m2)
CDC S-02-002	Airflow	Storage / Traditional HPC /Cloud/FPGA	280 kW (88 m2)
CDC S-02-003	**DLC**	**High Density/Energy efficient HPC**	**1050 kW (90 m2)**
CDC S-02-004	**DLC**	**High Density/Energy efficient HPC**	**1050 kW (92 m2)**
CDC S-02-005	Airflow	Storage / Traditional HPC (`iris` cluster)	300 kW (128 m2)

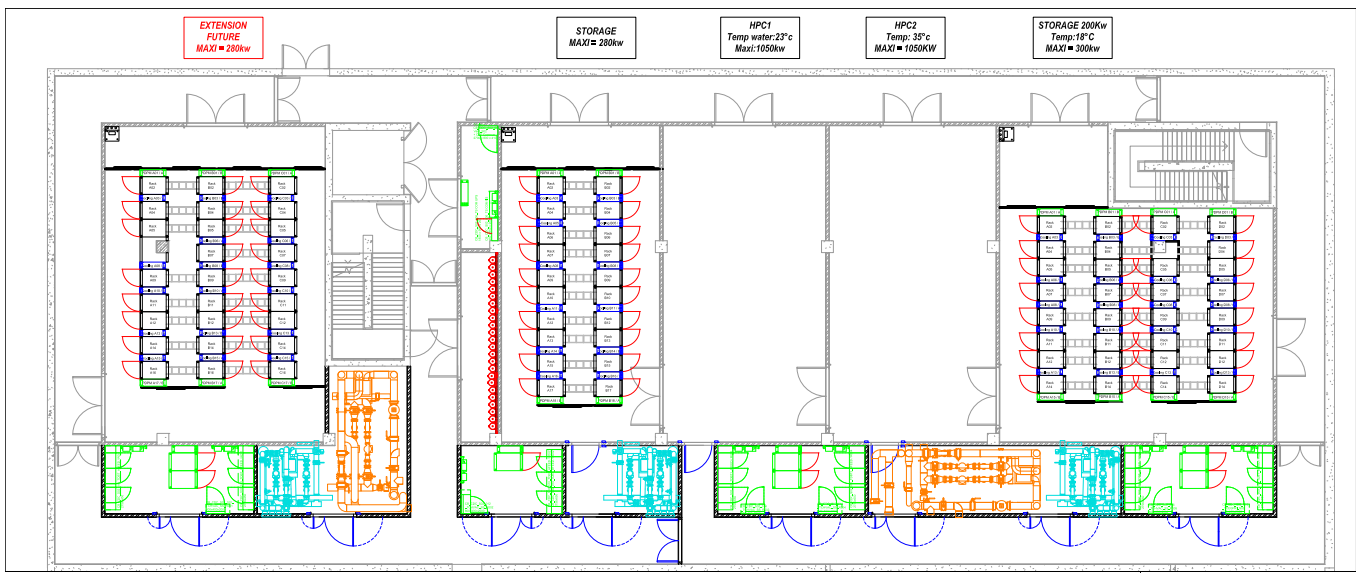

Next actions (Uni.lu HPC)

- **Short** maintenance (2d) Mid July
  - ↪ DDN intervention to tackle the recurrent GPFS issue
- **Big** maintenance End August (1w)
  - ↪ Just too many items to list here...

RFP 190027

- Public tender, **new `aion` cluster**,
  - ↪ *DLC*-based, to be deployed in CDC S-02-004
  - ↪ Lot 1: DLC racks/cells,
    - ✓ should compensate partially the decommissioned nodes
    - ✓ obj: at least doubling #nodes in `iris`
  - ↪ Lot 2: `iris` storage extension (central)
    - ✓ keep same performances despite combined `iris`+`aion` access
  - ↪ Lot 3: interconnect adaptation
    - ✓ `iris` FT is saturated

Summary

Summary

General Guidelines

- The UL HPC is a ***shared*** resource
  - ↪ hundreds of users may be logged on at one time
  - ↪ hundreds of jobs may be running on all compute nodes,
- All users must practice ***good citizenship***
  - ↪ 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 on the login nodes**
  - ↪ Plan large scale experiments during night-time or week-ends
    - ✓ **no more than 120 cores** during working day and working hours

General Guidelines

- The UL HPC is a ***shared*** resource
  - ↪ hundreds of users may be logged on at one time
  - ↪ hundreds of jobs may be running on all compute nodes,
- All users must practice ***good citizenship***
  - ↪ 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 on the login nodes**
  - ↪ Plan large scale experiments during night-time or week-ends
    - ✓ **no more than 120 cores** during working day and working hours

- For **ALL** publications having results produced using the UL HPC
  - ↪ Acknowledge / cite the UL HPC facility (using official banner)
  - ↪ Tag your publication upon registration on ORBiLu.

Summary

Compute Nodes Environment

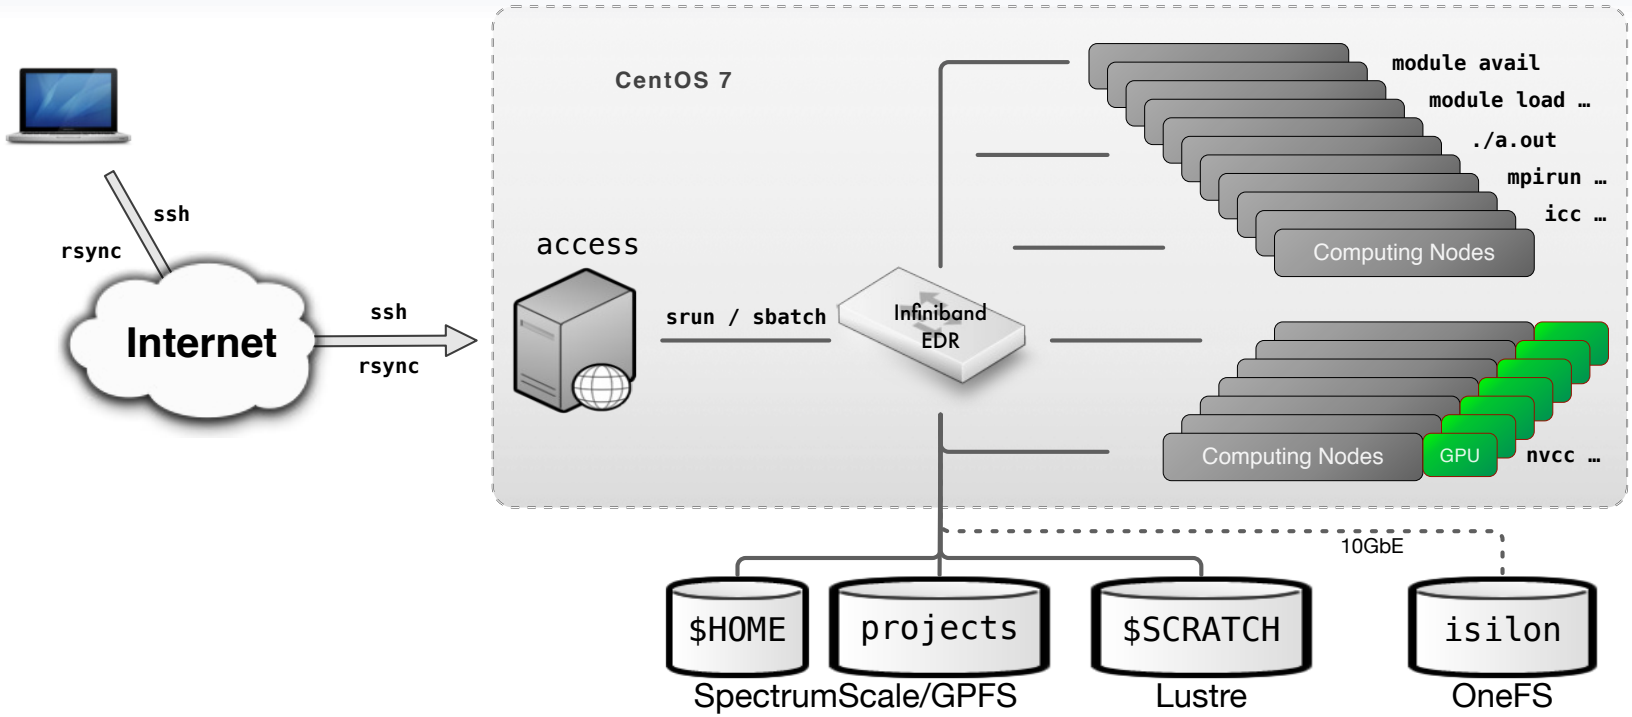

- *OS*: ~~Debian 8~~ / CentOS 7
- **Storage usage**: `df-ulhpc`
- **Env. modules**: `modules`
  - ↪ **Not** available on frontends
  - ↪ ***Only*** on compute nodes

Directory	Max size	Max #files	Backup
`$HOME` (`chaos,gaia`)	100 GB	1.000.000	YES
`$HOME` (`iris`)	500 GB	1.000.000	YES
`$WORK` (excl. `iris`)	3 TB		NO
`$SCRATCH`	per request		NO

Compute Nodes Environment

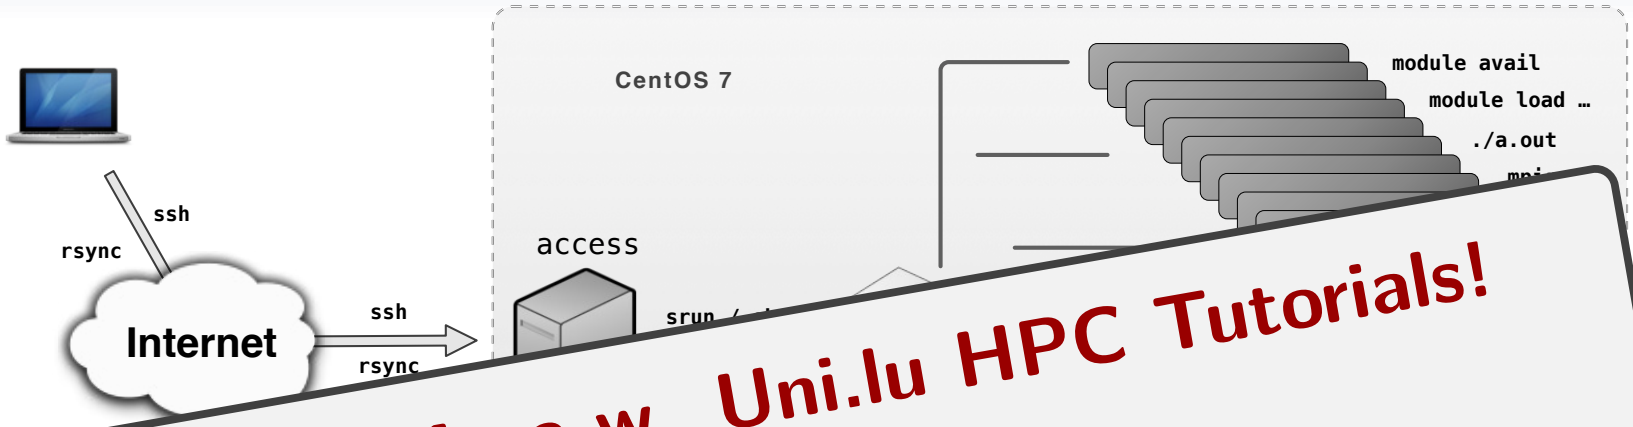

Learn More w. Uni.lu HPC Tutorials!

- PS2: Getting Started 2.0
- PS3: (Advanced) Job scheduling with SLURM
- Keynote/PS9: Data management (backup, decommissioning...)

- C
- **St**
- **En** modules
  - ↪ **Not** available on frontends
  - ↪ ***Only*** on compute nodes

	Max size	Max #files	Backup
`$HOME` (chaos,gaia)	100 GB	1.000.000	YES
`$HOME` (iris)	500 GB	1.000.000	YES
`$WORK` (excl. iris)	3 TB		NO
`$SCRATCH`	per request		NO

Summary

Documentation

http://hpc.uni.lu/users/getting_started.html

... aka the **rtf$_{ine}$m** paradigm

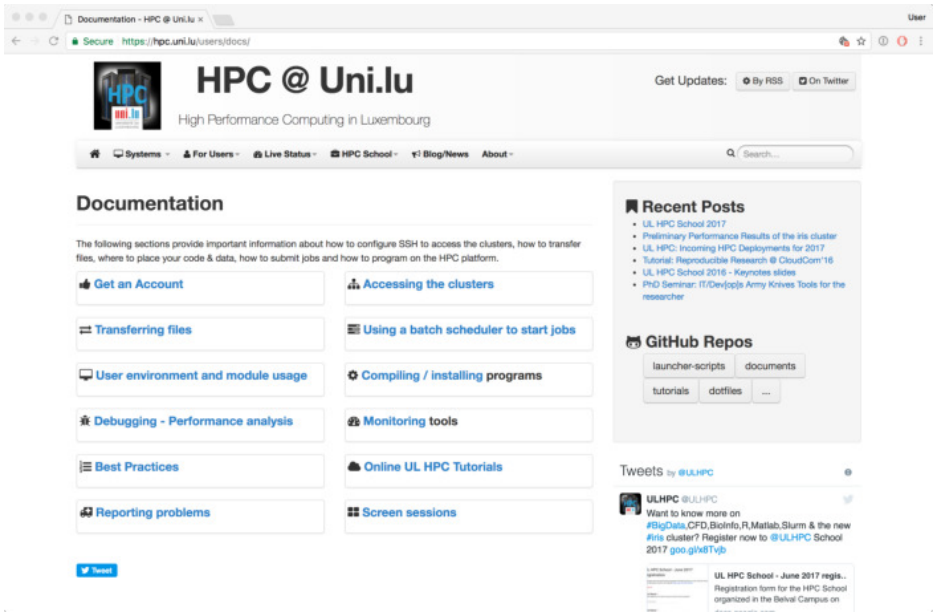

http://hpc.uni.lu

Reference Documentation

http://hpc.uni.lu/docs/

Online Tutorials

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

Documentation

http://hpc.uni.lu/users/getting_started.html

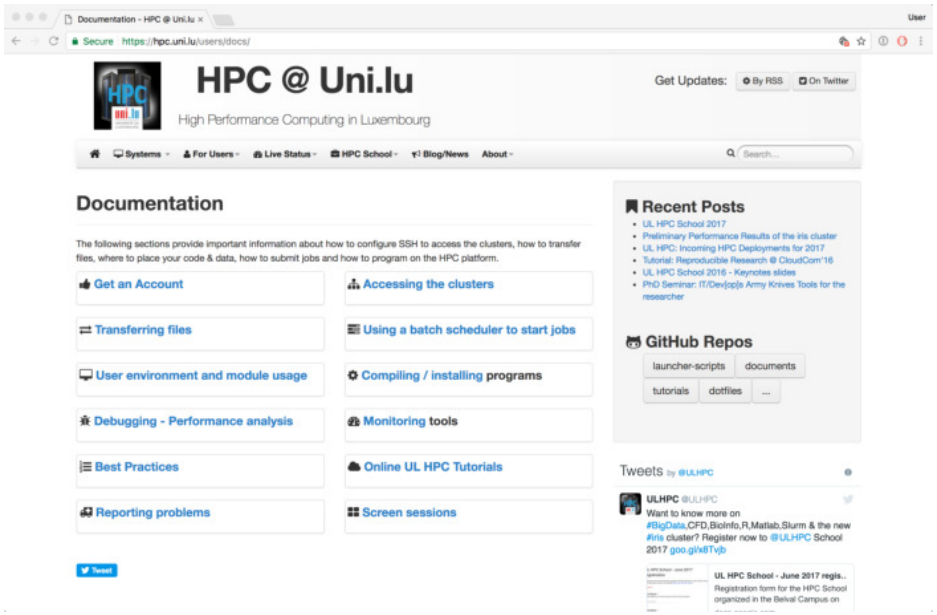

http://hpc.uni.lu

... aka the **rtf$_{ine}$m** paradigm

Reference Documentation

http://hpc.uni.lu/docs/

Online Tutorials

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

- UL HPC Ticketing System
  - ↪ https://hpc-tracker.uni.lu/
  - ↪ merge service.uni.lu ?
- Ask other users hpc-users@uni.lu
  - ↪ . . . or us hpc-sysadmins@uni.lu

Summary

Reporting Problems

- **First checks** https://hpc.uni.lu/users/docs/report_pbs.html
  1. My issue is probably documented — see User Doc
  2. An event is on-going — cf mail from `hpc-platform@uni.lu`
  3. check the state of your nodes
     - ✓ `{ oarsub -C <jobid> | ssh <node>}; htop` — *on active jobs*
     - ✓ `{ oarsub -f -j <jobid> }` — *post-mortem*
     - ✓ **iris**: `scontrol show job <jobid>` OR `sacct --job <jobid> -l`
     - ✓ Ganglia on your node(s) — https://hpc.uni.lu/status/ganglia.html

Reporting Problems

- **First checks** https://hpc.uni.lu/users/docs/report_pbs.html
  1. My issue is probably documented — see User Doc
  2. An event is on-going — cf mail from `hpc-platform@uni.lu`
  3. check the state of your nodes
     - ✓ `{ oarsub -C <jobid> | ssh <node>}; htop` — *on active jobs*
     - ✓ `{ oarsub -f -j <jobid> }` — *post-mortem*
     - ✓ **iris**: `scontrol show job <jobid>` OR `sacct --job <jobid> -l`
     - ✓ Ganglia on your node(s) https://hpc.uni.lu/status/ganglia.html

- **ONLY NOW**, consider the following depending on the severity:
  - ↪ Open an new issue on http://hpc-tracker.uni.lu (**preferred**)
  - ↪ Mail (only now) us `hpc-sysadmins@uni.lu`
  - ↪ **Ask the help of other users** `hpc-users@uni.lu`

Reporting Problems

- **First checks** https://hpc.uni.lu/users/docs/report_pbs.html
  1. My issue is probably documented — see User Doc
  2. An event is on-going — cf mail from `hpc-platform@uni.lu`
  3. check the state of your nodes
     - ✓ `{ oarsub -C <jobid> | ssh <node>}; htop` — *on active jobs*
     - ✓ `{ oarsub -f -j <jobid> }` — *post-mortem*
     - ✓ **iris**: `scontrol show job <jobid>` OR `sacct --job <jobid> -l`
     - ✓ Ganglia on your node(s) https://hpc.uni.lu/status/ganglia.html

- **ONLY NOW**, consider the following depending on the severity:
  - ↪ Open an new issue on http://hpc-tracker.uni.lu (**preferred**)
  - ↪ Mail (only now) us `hpc-sysadmins@uni.lu`
  - ↪ **Ask the help of other users** `hpc-users@uni.lu`

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

Reporting Obtained Results

- In your **scientific publications**: *as per Acceptable Use Policy (AUP)*
  - ↪ **acknowledge** your usage of the UL HPC platform
  - ↪ *(if possible)* **cite** the UL HPC paper `\cite{VBCG_HPCS14}`
- **More importantly**: add ULHPC Tag on your ORBi[lu] publication

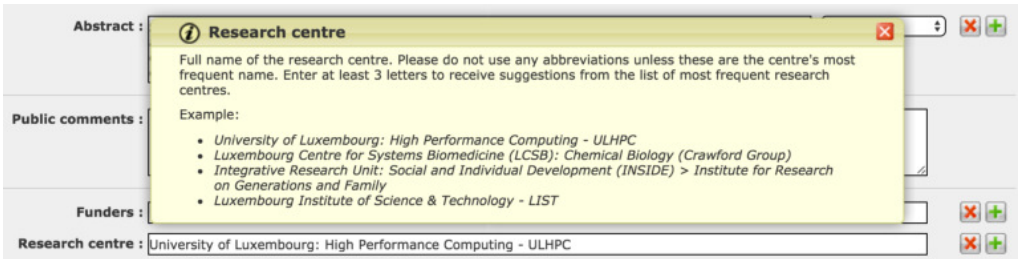


```
@InProceedings{VBCG_HPCS14,
   author =       {S. Varrette and P. Bouvry and H. Cartiaux and F. Georgatos},
   title =        {Management of an Academic HPC Cluster: The UL Experience},
   booktitle =    {Proc. of the 2014 Intl. Conf. on High Performance Computing \& Simulation (HPCS 2014)},
   year =         {2014},
   pages =        {959--967},
   month =        {July},
   address =      {Bologna, Italy},
   publisher =    {IEEE},
}
```

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
    2. **HPC Technology**: ETP4HPC, European Processor Initiative (EPI).
    3. **Application expertise**: Centres of Excellence of Computing Applications (CoEs)

- Junker Plan:
  - ↪ IPCEI on HPC and Big Data (BD) Applications (Nov. 2015)
    - ✓ Luxembourg (leader), France, Italy & Spain
    - ✓ Testbed around Personalized Medicine, Smart Space, Industry 4.0, Smart Manufacturing, New Materials, FinTech, Smart City...

European HPC strategy

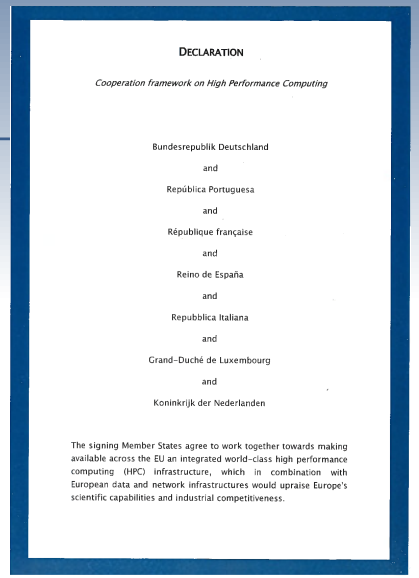
DECLARATION

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 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)

- Junker Plan:
  - ↪ IPCEI on HPC and Big Data (BD) Applications
    - ✓ Luxembourg (leader), France, Italy & Spain
    - ✓ Testbed around Personalized Medicine, Smart Space, Industry 4.0, Smart Manufacturing, New Materials, FinTech, Smart City...

- Significative upgrade in 2018 of the EC Strategy on HPC
  - ↪ EuroHPC Joint Undertaking (JU)

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
    - ✓ UL (P. Bouvry, S. Varrette, V.Plugaru) part of ETP4HPC (2016-)
    - ✓ See Strategic Research Agenda, 2017 European HPC Handbook...

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
    - ✓ UL (P. Bouvry, S. Varrette, V.Plugaru) part of ETP4HPC (2016-)
    - ✓ See Strategic Research Agenda, 2017 European HPC Handbook...

- **PRACE** - Partnership for Advanced Computing in Europe
  - ↪ Non-profit association, 25 member countries, now entering PRACE2
  - ↪ Providing access to **Five EU Tier-0** compute & data resources
  - ↪ (Oct. 2017) **Luxembourg 25th country to join PRACE**
    - ✓ Official Delegate/Advisor (P. Bouvry/S. Varrette) from UL

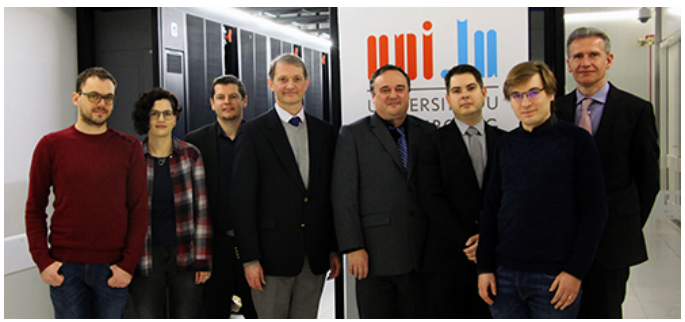

EU HPC Strategy Implementation

- **European High-Performance Computing Joint Undertaking**
  - ↪ EuroHPC JU effectively operational starting **Jan 1st, 2019**
    - ✓ administrative management from Luxembourg
  - ↪ Public and private members
    - ✓ EC, 28 MS, representatives from supercomputing/BD stakeholders
    - ✓ Governing Board (public members)
    - ✓ Industrial & Scientific Advisory Board (private members)
  - ↪ EU Objective with EuroHPC:
    - ✓ 5 **Petascale** systems (2020) (incl. Luxembourg)
    - ✓ 3 **Pre-exascale** systems (2020)
    - ✓ **2 exascale** systems (2022-2023)
    - ✓ Post-exascale system (2027)
- **European Processor Initiative (EPI)**
  - ↪ **120 M€** via Framework Partnership Agreement (FPA)

Meluxina EuroHPC Facility

- National HPC Competence Center / **Meluxina**
  - ↪ Luxembourg candidature to the EuroHPC JU call for Petascale Supercomputers
    - ✓ Piloted by Luxembourg Ministry of Economy
  - ↪ Proposal prepared by a **consortium lead by Luxconnect**
    - ✓ Uni.lu [HPC], LIST, JSC and Partec
  - ↪ Official press conference on **June 14, 2019** – Youtube video

MELUXINA
HIGH PERFORMANCE COMPUTING IN LUXEMBOURG

LE GOUVERNEMENT DU GRAND-DUCHÉ DE LUXEMBOURG
Ministère de l'Économie

LUXCONNECT

UNIVERSITÉ DU LUXEMBOURG

LIST

JÜLICH Forschungszentrum
JÜLICH SUPERCOMPUTING CENTRE

PARTEC CLUSTER COMPETENCE CENTER

Meluxina EuroHPC Facility

- Managed by **Luxconnect**, DC location: Bissen
  - ↪ 10 PFlops system

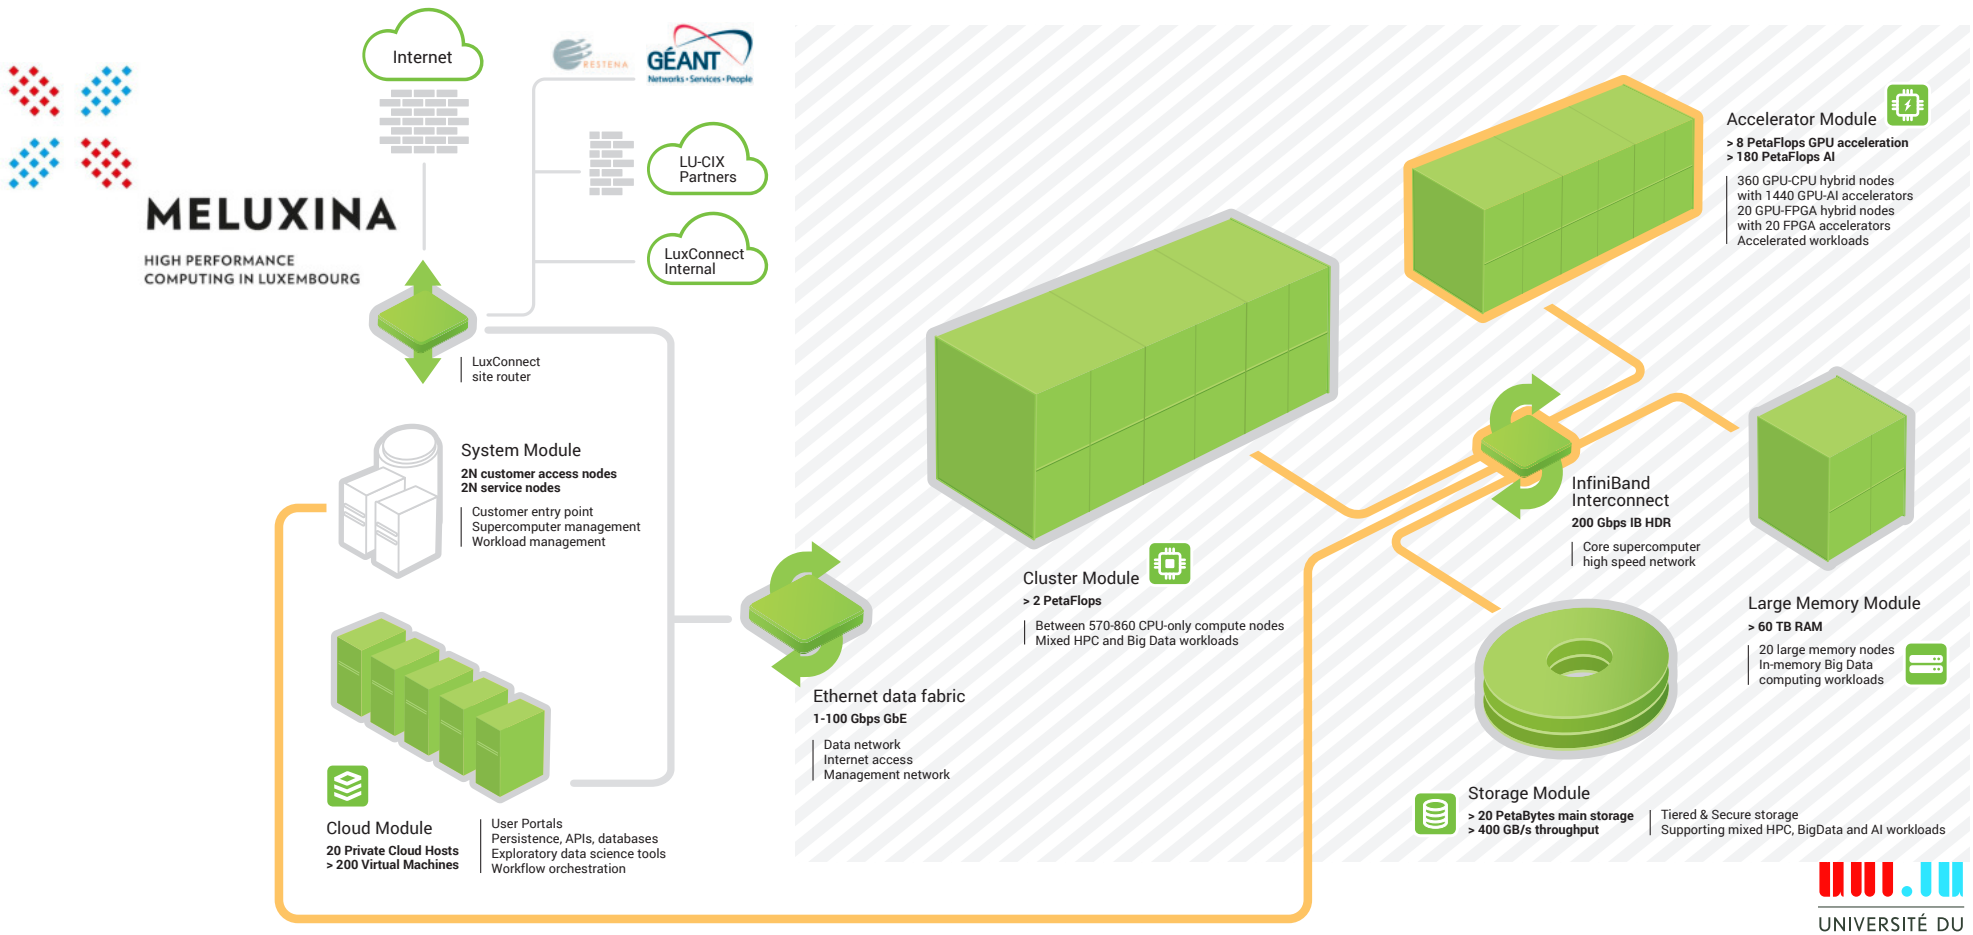

Summary

Conclusion

- **Luxembourg government priority on HPC**
  - ↪ sustained by University of Luxembourg HPC developments
    - ✓ started in 2007, under resp. of Prof P. Bouvry & Dr. S. Varrette
    - ✓ **expert UL HPC team**: *Dr. S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot* + ***2019 newcomers: Dr. F. Pinel, Dr. E. Kieffer, Dr. NN (Aug 2019)***
    - ✓ computational scientists / domain experts across ALL the UL
  - ↪ Uni.lu HPC (as of 2019): **1263.322 TFlops, 9852.4TB (shared)**

Conclusion

- **Luxembourg government priority on HPC**
  - ↪ sustained by University of Luxembourg HPC developments
    - ✓ started in 2007, under resp. of Prof P. Bouvry & Dr. S. Varrette
    - ✓ **expert UL HPC team**: *Dr. S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot* + ***2019 newcomers: Dr. F. Pinel, Dr. E. Kieffer, Dr. NN (Aug 2019)***
    - ✓ computational scientists / domain experts across ALL the UL
  - ↪ Uni.lu HPC (as of 2019): **1263.322 TFlops, 9852.4TB (shared)**

Several On-going Strategic HPC efforts in Europe...

- . . . **in which UL *(HPC)* is involved** . . .
  - ↪ ETP4HPC, EU COST Action NESUS etc.
  - ↪ PRACE - Official representative for Luxembourg from UL
    - ✓ Delegate: Prof. Pascal Bouvry; Advisor: Dr. Sebastien Varrette
  - ↪ NVidia Joint AI Lab since 2019
    - ✓ Corrdinator: Dr. Frederic Pinel
  - ↪ EuroHPC: **Meluxina by LuxConnect**
    - ✓ *National HPC-BD Competence Center*
    - ✓ Uni.lu HPC Team to collaborate **YET** conditions to formalize

Incoming Milestones in MSA CDC-02

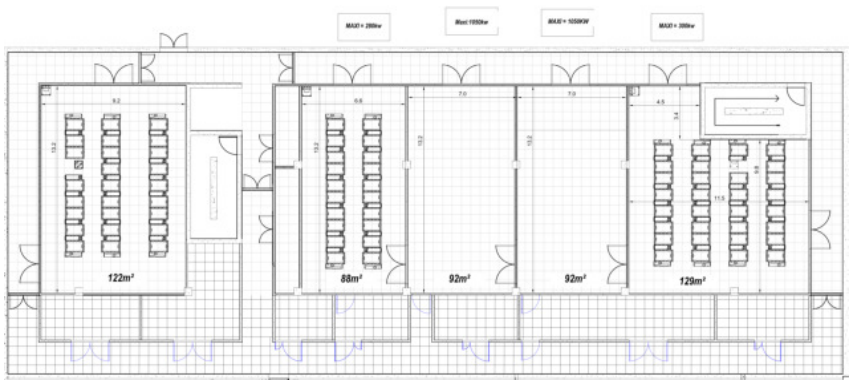

- $\simeq$ **1050kW** per **HPC** room
  - ↪ Direct Liquid Cooling (DLC)
- $\simeq$ **300kW** per **storage** room
  - ↪ rooms 1, 2 & 5
  - ↪ Air-flow storage / HPC

Incoming Milestones in MSA CDC-02

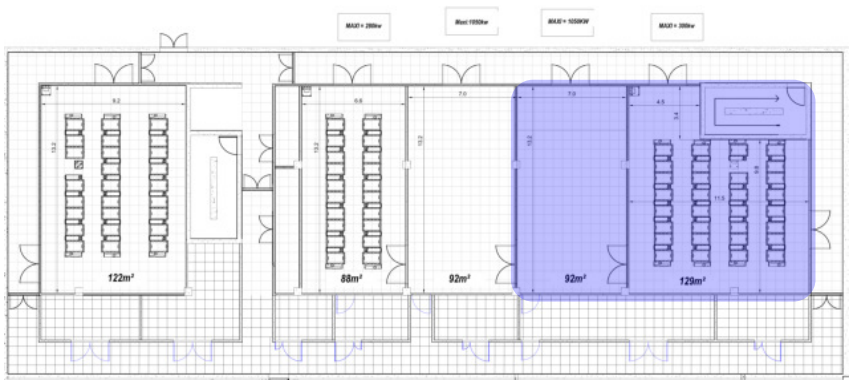

- $\simeq$ **1050kW** per **HPC** room
  - ↪ Direct Liquid Cooling (DLC)
- $\simeq$ **300kW** per **storage** room
  - ↪ rooms 1, 2 & 5
  - ↪ Air-flow storage / HPC

- **Short term actions (by end 2019)**
  - ↪ **RFP 190027** `aion`(?) cluster deployment in DLC room
    - ✓ First phase Q4 2019 - 2021
    - ✓ Further extensions depends on collaboration agreement with Meluxina
  - ↪ `gaia/chaos` **decommissioning end 2019**
  - ↪ **GDPR**: WIP, extending work initiated @ LCSB

Incoming Milestones in MSA CDC-02

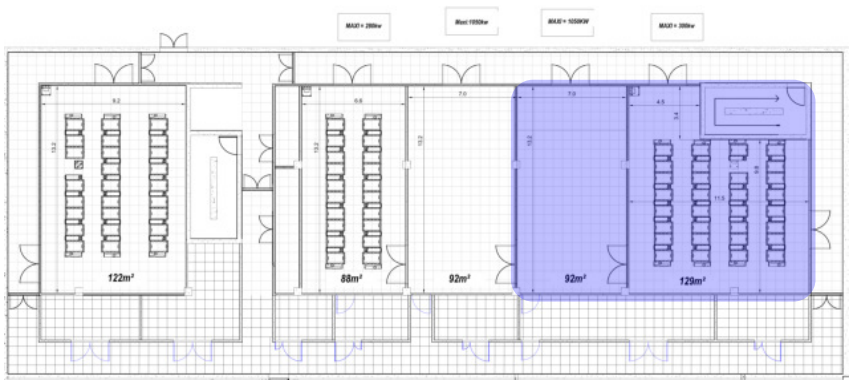

- $\simeq$ **1050kW** per **HPC** room
  - ↪ Direct Liquid Cooling (DLC)
- $\simeq$ **300kW** per **storage** room
  - ↪ rooms 1, 2 & 5
  - ↪ Air-flow storage / HPC

- **Short term actions (by end 2019)**
  - ↪ **RFP 190027** `aion`(?) cluster deployment in DLC room
    - ✓ First phase Q4 2019 - 2021
    - ✓ Further extensions depends on collaboration agreement with Meluxina
  - ↪ `gaia/chaos` **decommissioning end 2019**
  - ↪ **GDPR**: WIP, extending work initiated @ LCSB
  - ↪ **Submission portal 2.0** (draft), **Identity Infrastructure 2.0** (with incoming maintenance)

Incoming Milestones in MSA CDC-02

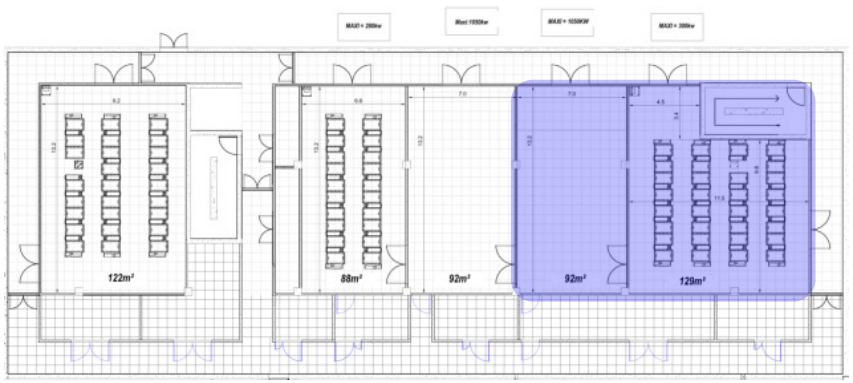

- $\simeq$ **1050kW** per **HPC** room
  - ↪ Direct Liquid Cooling (DLC)
- $\simeq$ **300kW** per **storage** room
  - ↪ rooms 1, 2 & 5
  - ↪ Air-flow storage / HPC

- **Short term actions (by end 2019)**
  - ↪ **RFP 190027** `aion`(?) cluster deployment in DLC room
    - ✓ First phase Q4 2019 - 2021
    - ✓ Further extensions depends on collaboration agreement with Meluxina
  - ↪ `gaia/chaos` **decommissioning end 2019**
  - ↪ **GDPR**: WIP, extending work initiated @ LCSB
  - ↪ **Submission portal 2.0** (draft), **Identity Infrastructure 2.0** (with incoming maintenance)
  - ↪ Reference platform article update (now that all nodes are integrated)

Incoming Milestones

- **Mid term actions (by end 2020)**
  - ↪ **Uni.lu Digital strategy**
    - ✓ to be released Sept. 2019,
    - ✓ implementation Q4-2019 / 2020
  - ↪ **NVidia Joint AI Lab** continuous developments
  - ↪ Meluxina / **Luxembourg HPC-BD Competence Center**
    - ✓ Expected in production Q3-Q4 2020
    - ✓ MoU LuxConnect / Uni.lu (pending details)
    - ✓ Etienne Schneider: "*Access will be free for Uni.lu and LIST researchers*"

Questions?

http://hpc.uni.lu

High Performance Computing @ uni.lu

Prof. Pascal Bouvry
Dr. Sebastien Varrette
Valentin Plugaru
Sarah Peter
Hyacinthe Cartiaux
Clement Parisot
Dr. Frederic Pinel
Dr. Emmanuel Kieffer

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

Questions?

http://hpc.uni.lu

High Performance Computing @ uni.lu

Prof. Pascal Bouvry
Dr. Sebastien Varrette
Valentin Plugaru
Sarah Peter
Hyacinthe Cartiaux
Clement Parisot
Dr. Frederic Pinel
Dr. Emmanuel Kieffer

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

Picture Time!!!!