Large-scale Research Data Management @ UL HPC

Road to GDPR compliance

Prof. Pascal Bouvry, Dr. Sebastien Varrette

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

Belval Campus, April 25th, 2018

University of Luxembourg (UL), Luxembourg

Summary

Why HPC and BD ?

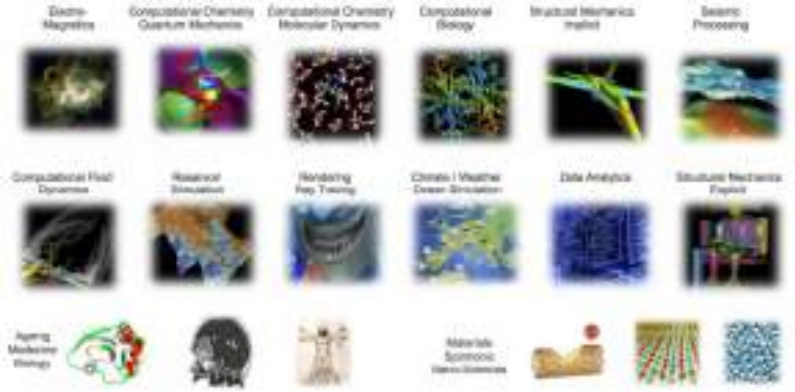

HPC: **H**igh **P**erformance **C**omputing
BD: **B**ig **D**ata

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 ?

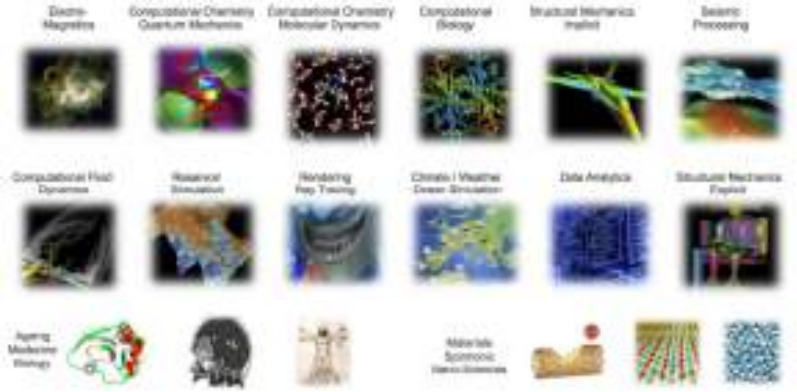

HPC: **H**igh **P**erformance **C**omputing
BD: **B**ig **D**ata

- Essential tools for **Science, Society and Industry**
  - ↪ All scientific disciplines are becoming computational today
    - ✓ requires very high computing power, handles **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 ?

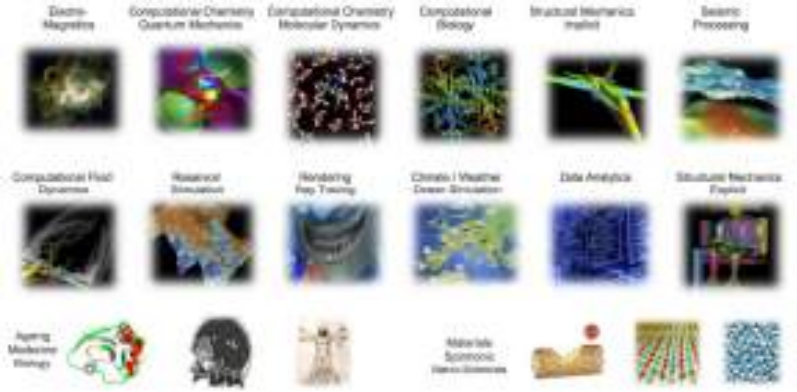

HPC: **H**igh **P**erformance **C**omputing
BD: **B**ig **D**ata

- Essential tools for **Science, Society and Industry**
  - ↪ All scientific disciplines are becoming computational today
    - ✓ requires very high computing power, handles **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:
  - ↪ 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

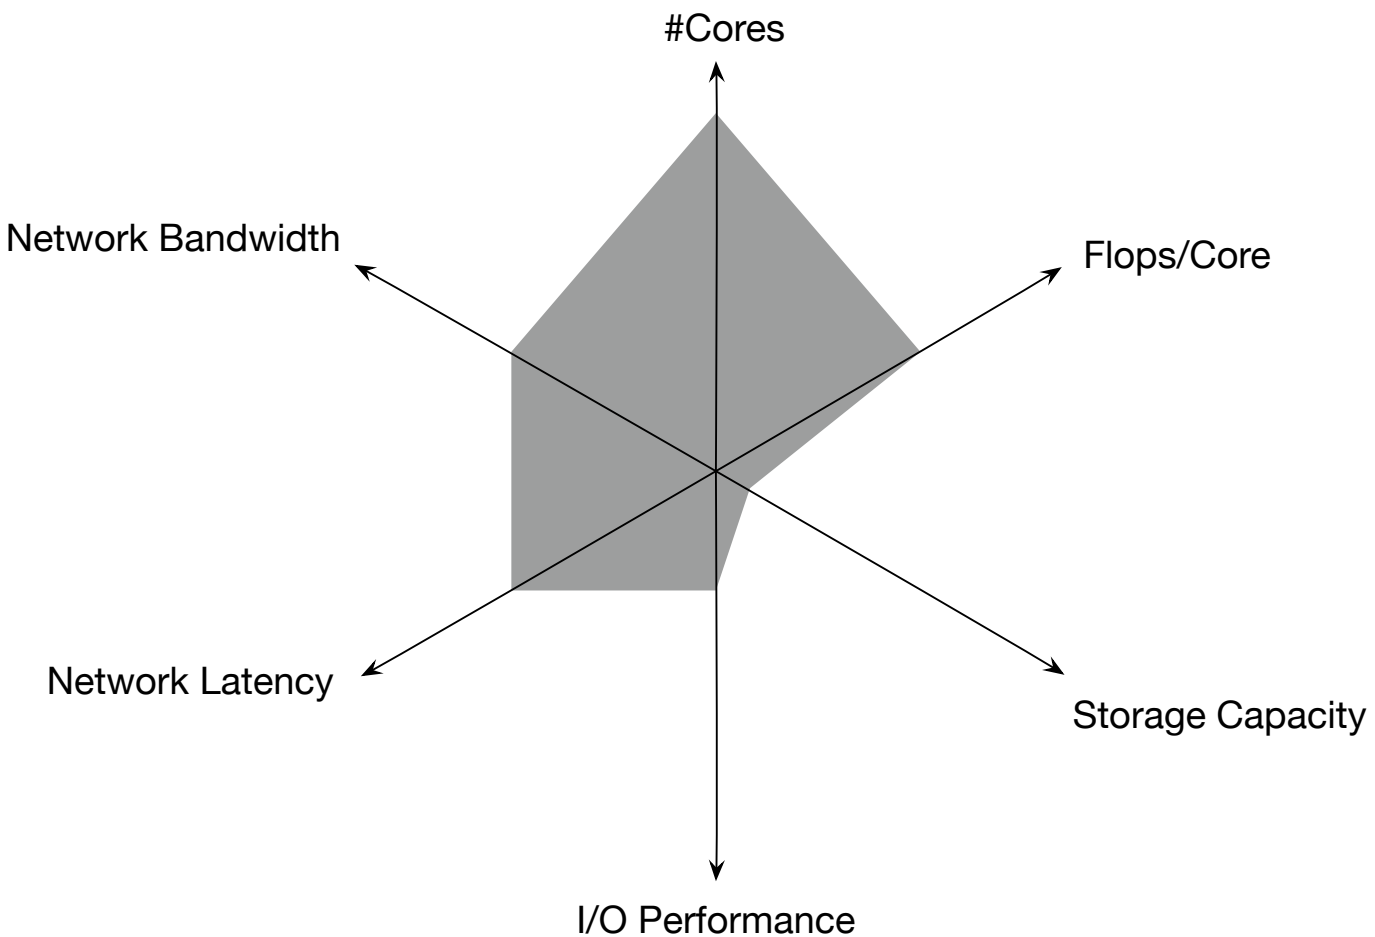

Different HPC Needs per Domains

Biomedical Industry / Life Sciences

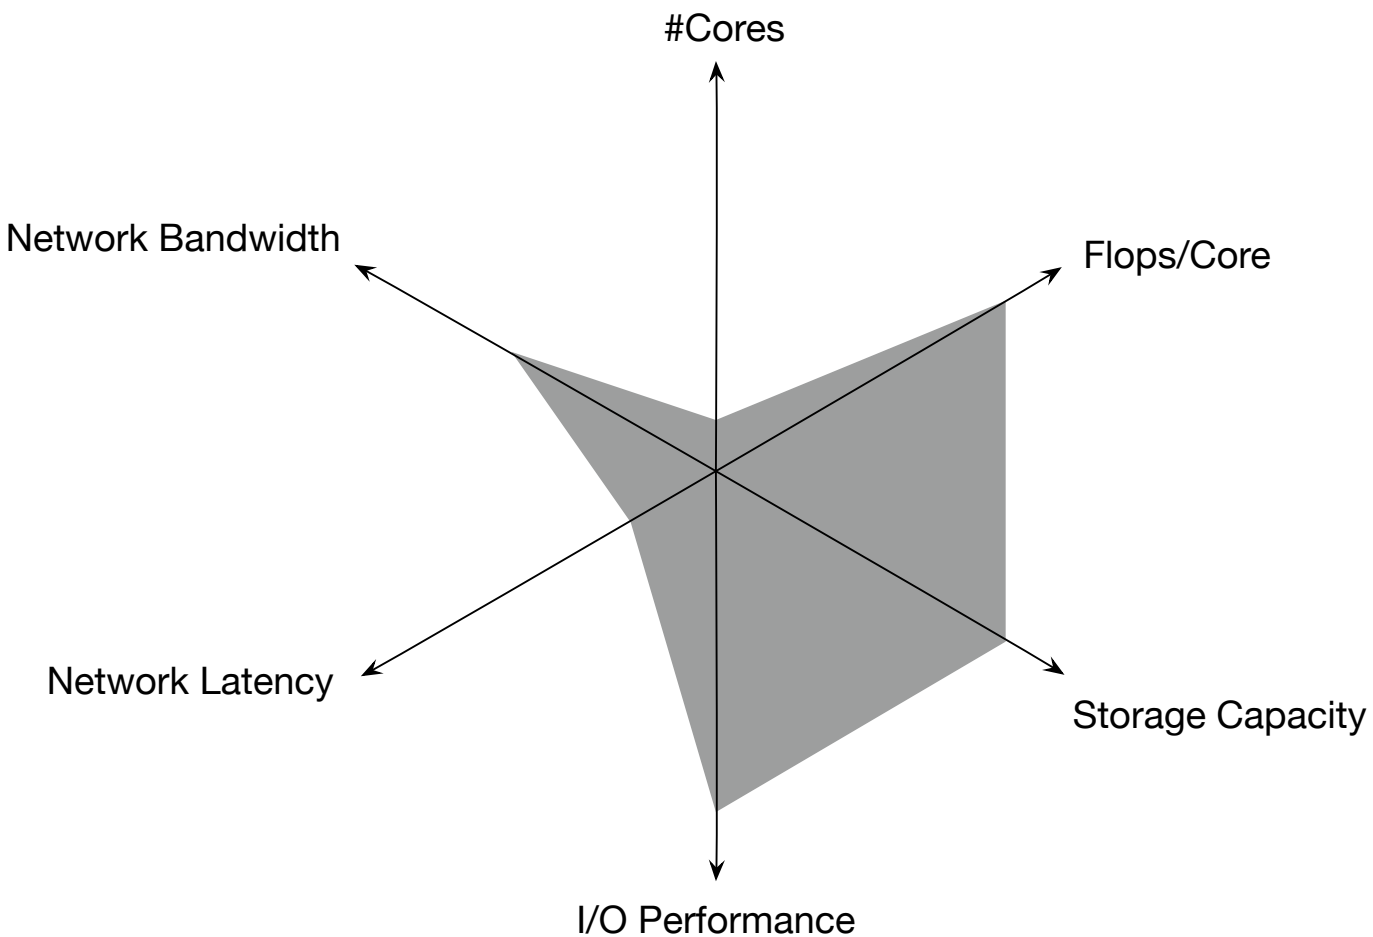

Different HPC Needs per Domains

Deep Learning / Cognitive Computing

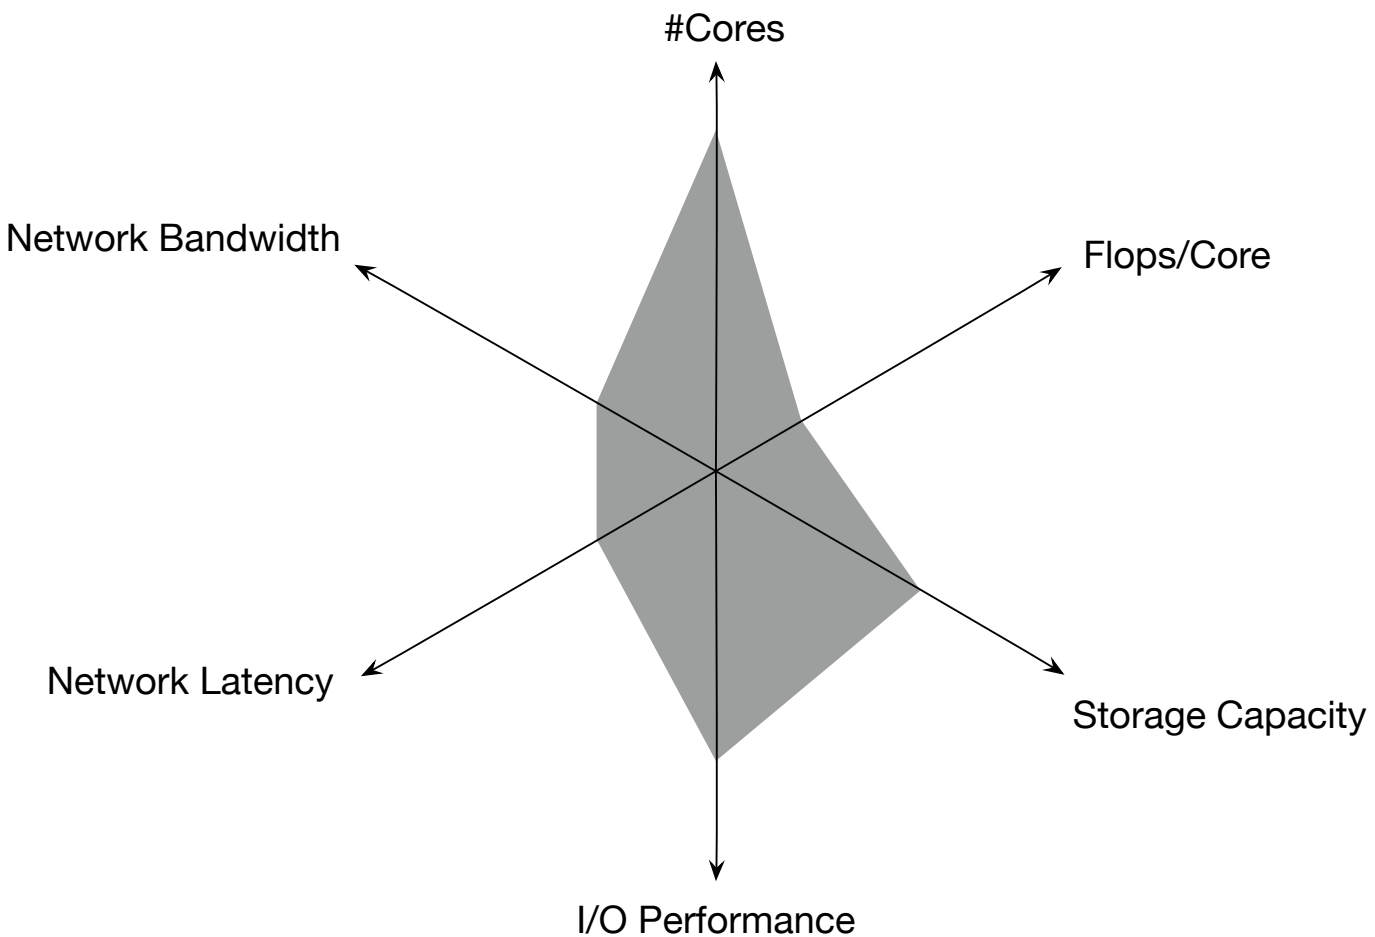

Different HPC Needs per Domains

IoT, FinTech

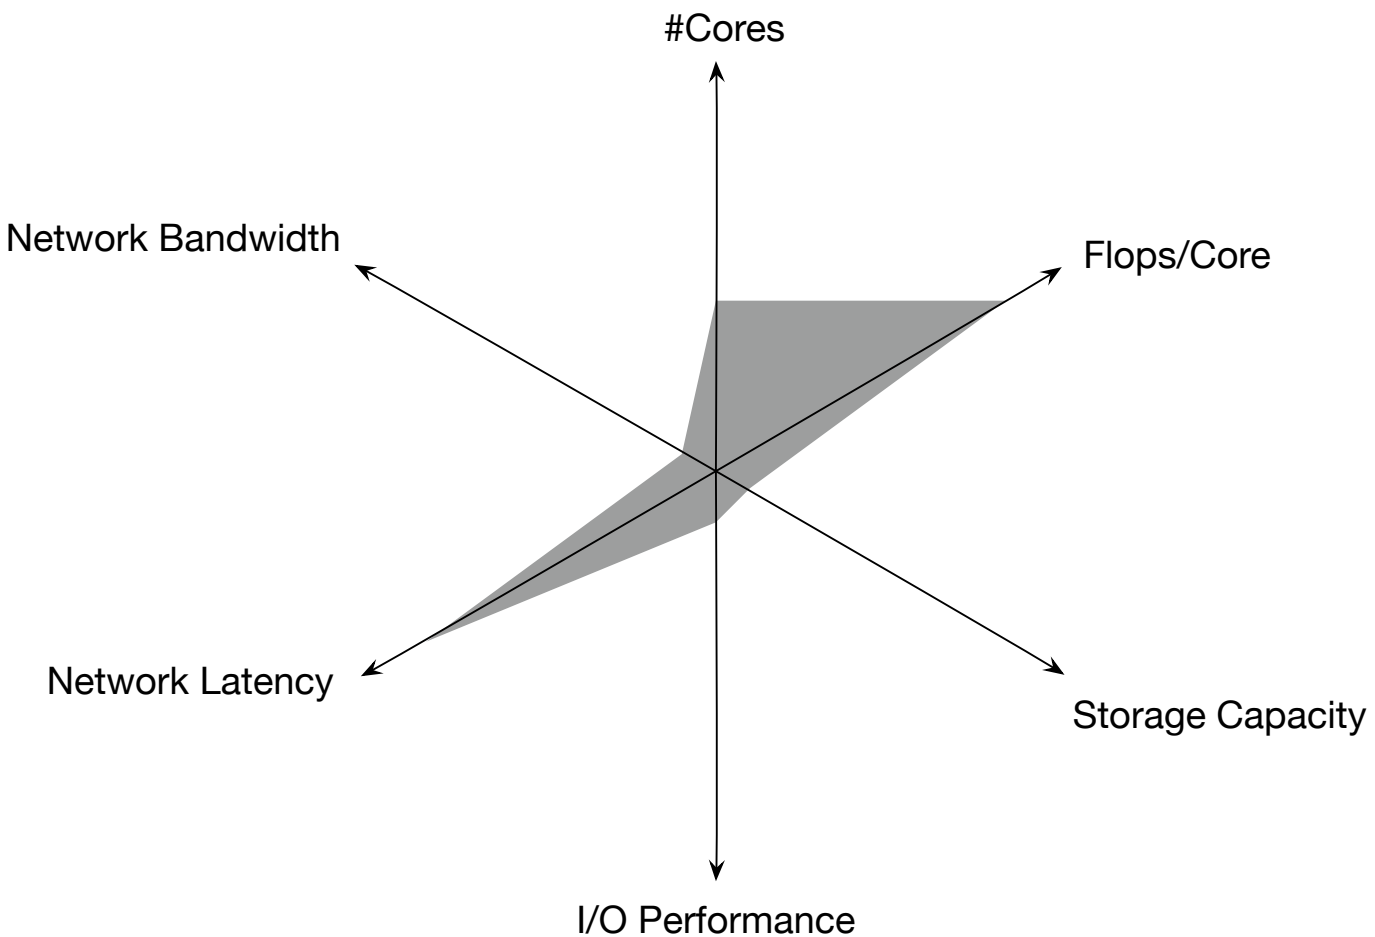

Different HPC Needs per Domains

ALL Research Computing Domains

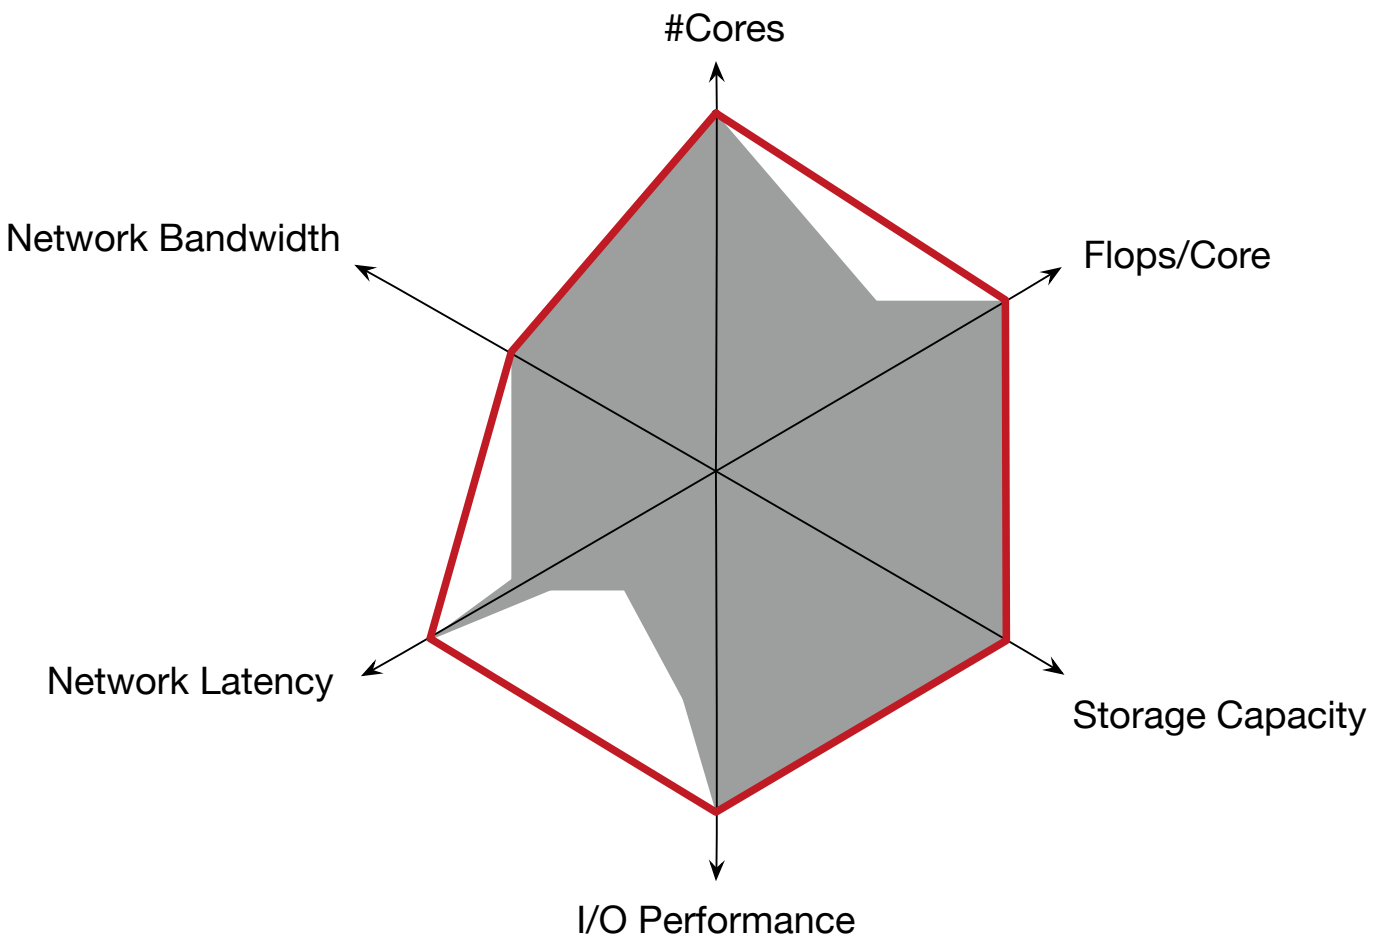

High Performance Computing @ UL

- **Started in 2007**, under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ expert UL HPC team *(S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot)*
  - ↪ **8,173,747€** cumulative investment in hardware

http://hpc.uni.lu

Key numbers

- 469 users
- 662 computing nodes
  - ↪ 10132 cores, **346.652 TFlops**
  - ↪ 50 accelerators (+ **76.22 TFlops**)
- **9232.4 TB** storage
- 130 (+ 71) servers
- 5 sysadmins
- 2 sites: Kirchberg / Belval

Sites / Data centers

Kirchberg

CS.43, AS. 28

Belval

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 Computing capacity

5 clusters
346.652 TFlops
662 nodes
10132 cores
34512GPU cores

INFINIBAND TRADE ASSOCIATION intel

UL HPC Storage capacity

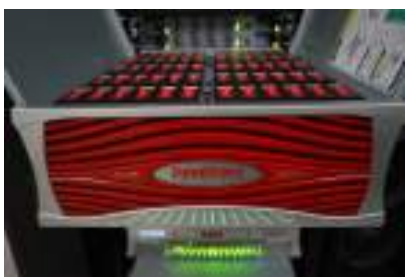

4 distributed/parallel FS
2183 disks
9232.4 TB

(incl. 2116TB for Backup)

[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.

Summary

Data Intensive Computing

- Data volumes increasing massively
  - ↪ Clusters, storage capacity increasing massively
- Disk speeds are not keeping pace.
- Seek speeds even worse than read/write

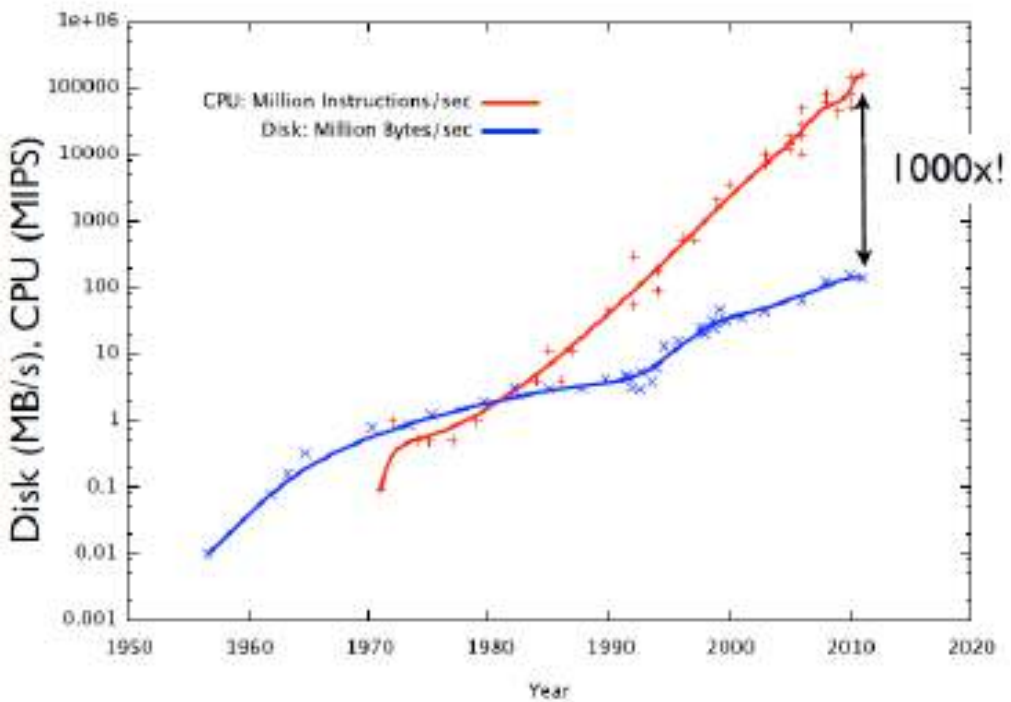

Speed Expectation on Data Transfer

http://fasterdata.es.net/

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

Speed Expectation on Data Transfer

http://fasterdata.es.net/

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

ULHPC Storage Performances: GPFS

- Self Encrypting Disks (SED)-based storage

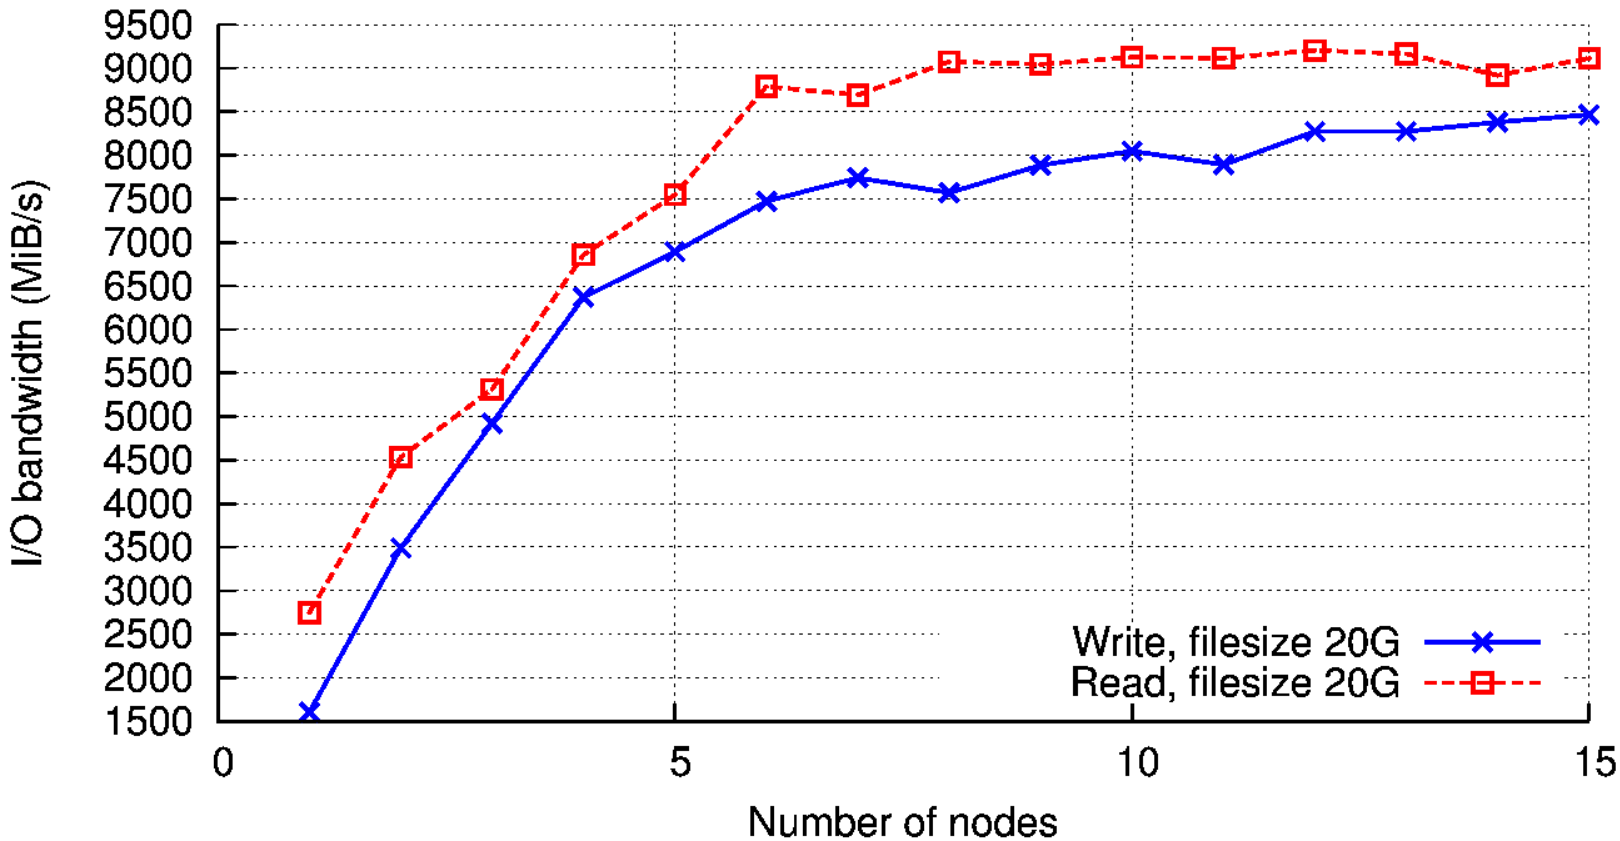

ULHPC Storage Performances: Lustre

- Self Encrypting Disks (SED)-based storage

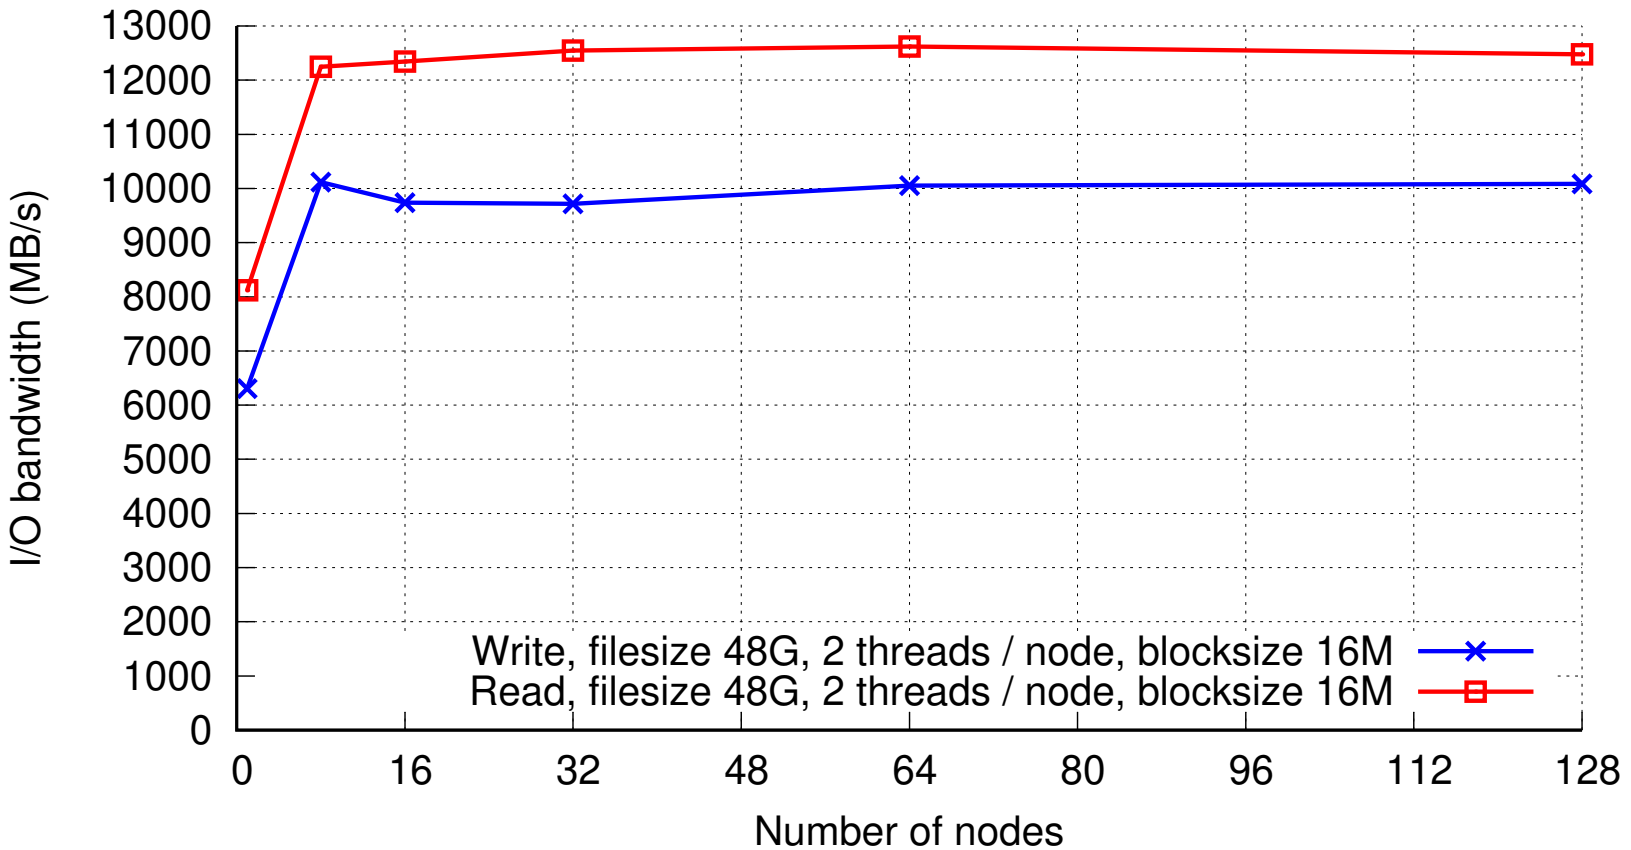

GDPR and UL HPC

www.eugdpr.org

- EU General Data Protection Regulation (**GDPR**)
  - ↪ replaces the Data Protection Directive 95/46/EC
  - ↪ legislation comes into effect May 25th 2018.

- **The UL HPC facility handles both**:
  - ↪ **data about people** (facility users identification details)
    - ✓ ULHPC Identity Management (IdM) system
    - ✓ on Google Drive (account request form results) (**bad**)
  - ↪ **large scale data** that may contain Personally Identifiable Info
    - ✓ stored by facility users in networked, parallel & distributed filesystems used across the HPC infrastructure
    - ✓ can be considered as falling under GDPR regulations.

GDPR and UL HPC

- **Personal data is/may be visible, accessible or handled**:
  - ↪ directly on the HPC clusters
  - ↪ through *Resource and Job Management System* (RJMS) tools
    - ✓ glue for a parallel computer to execute parallel jobs
    - ✓ **Goal**: satisfy users demands for computation
    - ✓ comes with web interfaces, eventually public — Monika, Ganttchart
  - ↪ through service portals — hpc-tracker, XCS, Galaxy
  - ↪ on code management portals — GitLab, GitHub
  - ↪ on secondary storage systems — DropIT, OwnCloud

Toward a ULHPC QoS Master Plan

Objectives

- **Formalizing the way we tackle security hardening**
  - ↪ Work in progress with **continuous improvement**
  - ↪ Completes other initiatives at SIU, LCSB, SnT etc.
  - ↪ Ongoing adaptation to match GDPR compliance
  - ↪ In line with UL guidelines expected this year
    - ✓ public release expected max Q4 2018

Toward a ULHPC QoS Master Plan

- Covers specific protection operations, either **in general**:
  - ↪ default protections at the level of network (VLANs, firewall), OS...
  - ↪ secure access over SSH etc.
- ... or in **special SLAs when dealing with sensitive project**:
  - ↪ physical security protection: data center/rack access, BIOS...
  - ↪ data protection:
    - ✓ umask, Self Encrypting Disks (SED)-based storage...
    - ✓ GDPR data only stored on SED capable systems
    - ✓ GDPR data is only processed in memory?
    - ✓ private data encrypted (EncFS?) with per-job de-encryption?
    - ✓ RJMS scheduling policy from core-level to node-level
  - ↪ data transfer
  - ↪ scheduling aspects (exclusive mode),
  - ↪ job epilog/prolog re-formatting the nodes...

Summary

Conclusion & Perspectives

- **Luxembourg government priority on HPC & Big Data**
  - ↪ sustained by University of Luxembourg HPC developments
    - ✓ started in 2007, under resp. of Prof P. Bouvry & Dr. S. Varrette
    - ✓ expert UL HPC team *(S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot)*
  - ↪ UL HPC (as of 2018): **346.652 TFlops, 9232.4TB (shared)**
  - ↪ consolidate and extend Europe efforts on HPC/Big Data

- **EU GDPR compliance expected**
  - ↪ (especially) with **large-scale research data**
  - ↪ **Incoming**: **ULHPC QoS Master plan** in line with UL guidelines
- Elements to take into account at **all levels** of UL [HPC] services:
  - ↪ Access restrictions, Data minimisation, Encryption
  - ↪ Access control, Data integrity, Backups
  - ↪ Reviews & testing
  - ↪ Training & awareness

Questions?

http://hpc.uni.lu

Prof. Pascal Bouvry
Dr. Sebastien Varrette & The UL HPC Team
(V. Plugaru, S. Peter, H. Cartiaux & C. Parisot)
University of Luxembourg, Belval Campus:
Maison du Nombre, 4th floor
2, avenue de l'Université
L-4365 Esch-sur-Alzette
mail: hpc@uni.lu