HPC @ UL & new Trends in Europe and Abroad

Path to Exascale

Dr. Sebastien Varrette

Oct. 24th, 2017

University of Luxembourg (UL), Luxembourg

Keynote, 8th Intl. SuperComputing Camp (SC-Camp 2017), Cadiz, Spain

Summary

Summary

University of Luxembourg

http://www.uni.lu

- In Belval campus (South of the country) since 2015

A Model European Research University of the 21st Century

University of Luxembourg

http://www.uni.lu

Among Top 200 Universities Worldwide

- **Multilingual European research university**
  - ↪ since 2003, 6300 students, 120 nationalities
    - ✓ 12 Bachelor, 42 Ms. Degrees, 4 DS
- ~1700 employees (researchers + staff)
  - ↪ 1139 scientific publication in 2015
    - ✓ 570 in refereed journals
  - ↪ 100 exchange agreements/research coop.
- **Total Budget**: 215.4 million euros (2015)

- **International rankings (2017-2018)**
  - ↪ *Times Higher Education (THE) World University* **179th**
  - ↪ *Times Higher Education (THE)* ***Young*** *University* **11th**

Summary

Why High Performance Computing ?

"The country that out-computes will be the one that out-competes". *Council on Competitiveness*

- **Accelerates** research by accelerating **computation**

$\simeq$ 64 **G**Flops (Dual-core i5 2GHz)

206.772 **T**Flops (602 computing nodes,8452 cores)

- Increases **storage** capacity and velocity for Big Data processing

4 TB (1 disk, 250 MB/s)

7952.4TB (2015 disks, 10 GB/s)

- **Communicates faster** 1 GbE (1 Gb/s) vs Infiniband EDR (100 Gb/s)~~

HPC at the Heart of our Daily Life

- **Today**: Research, Industry, Local Collectivities

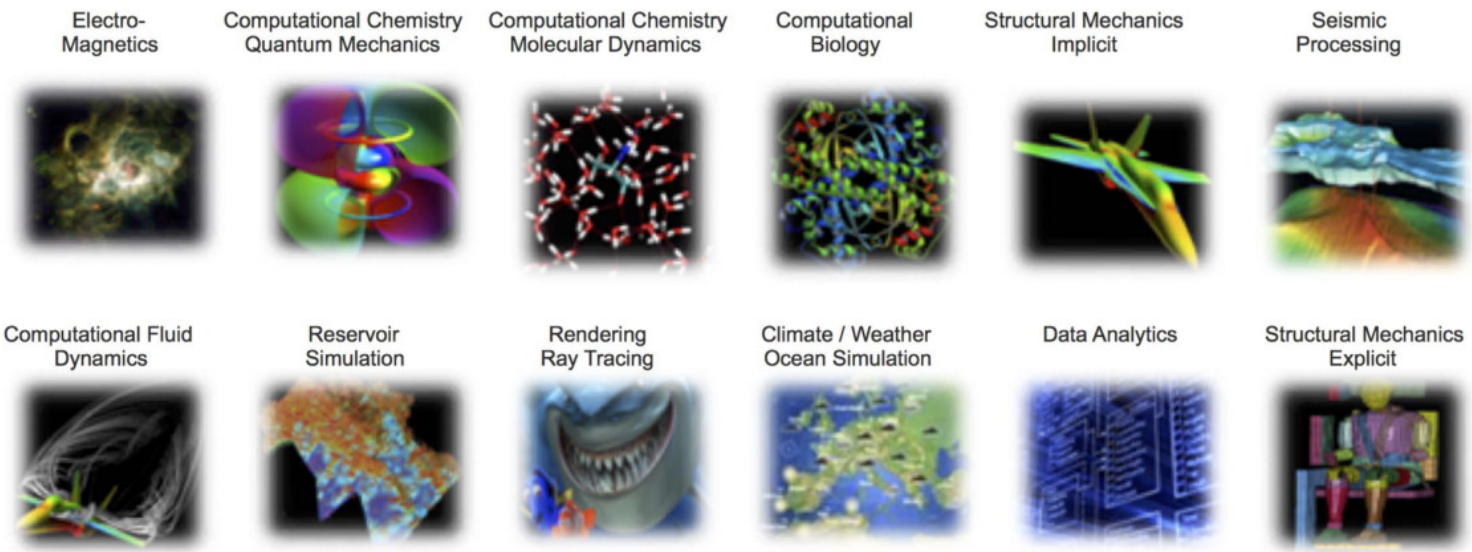

- . . . **Tomorrow**: applied research, digital health, nano/bio tech.

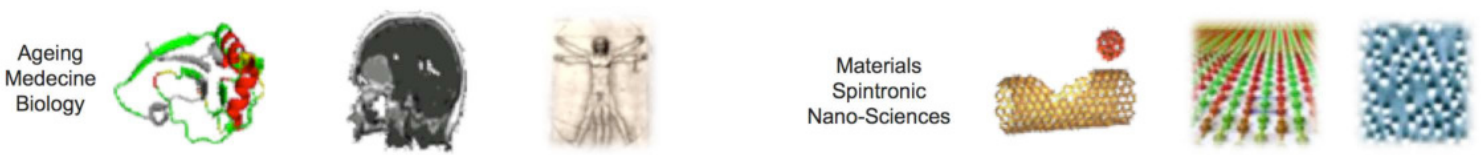

High Performance Computing @ UL

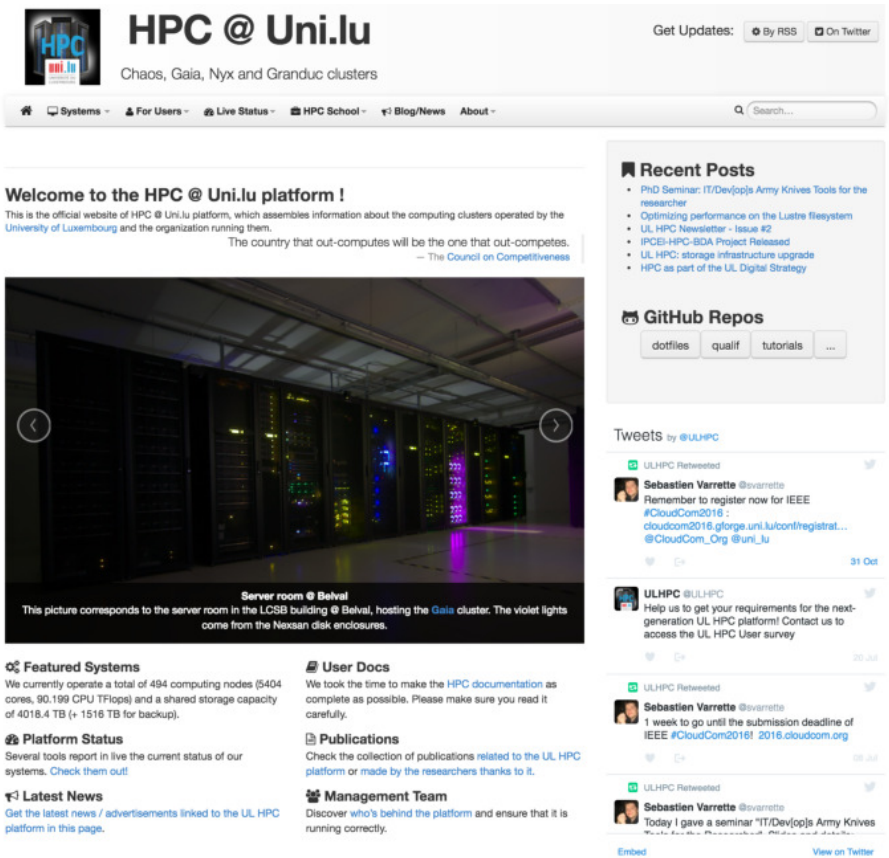

http://hpc.uni.lu

Key numbers

- 416 users
- 602 computing nodes
  - ↪ 8452 cores
  - ↪ **206.772 TFlops**
  - ↪ 50 accelerators (+ **76.22 TFlops**)
- **7952.4 TB**
- 130 (+ 71) servers
- 5 sysadmins
- 2 sites
  - ↪ Kirchberg
  - ↪ Belval

High Performance Computing @ UL

- **Enables & accelerates** scientific discovery and innovation
- **Largest facility** in Luxembourg (after GoodYear R&D Center)

			(CPU)	TFlops	TB (Shared)
Country	**Institute**	**#Nodes**	**#Cores**	**R_{peak}**	**Storage**
Luxembourg	**UL HPC** (Uni.lu)	**602**	**8452**	**206.772**	**7952.4**
	LIST	58	800	6.21	144
France	LORIA (G5K), Nancy	320	2520	26.98	82
	ROMEO, Reims	174	3136	49.26	245
Belgium	NIC4, University of Liège	128	2048	32.00	20
	Université Catholique de Louvain	112	1344	13.28	120
	UGent / VSC, Gent	440	8768	275.30	1122
Germany	bwGrid, Heidelberg	140	1120	12.38	32
	bwForCluster, Ulm	444	7104	266.40	400
	bwHPC MLS&WISO, Mannheim	604	9728	371.60	420

UL HPC User Base

- **416 Active** HPC Users

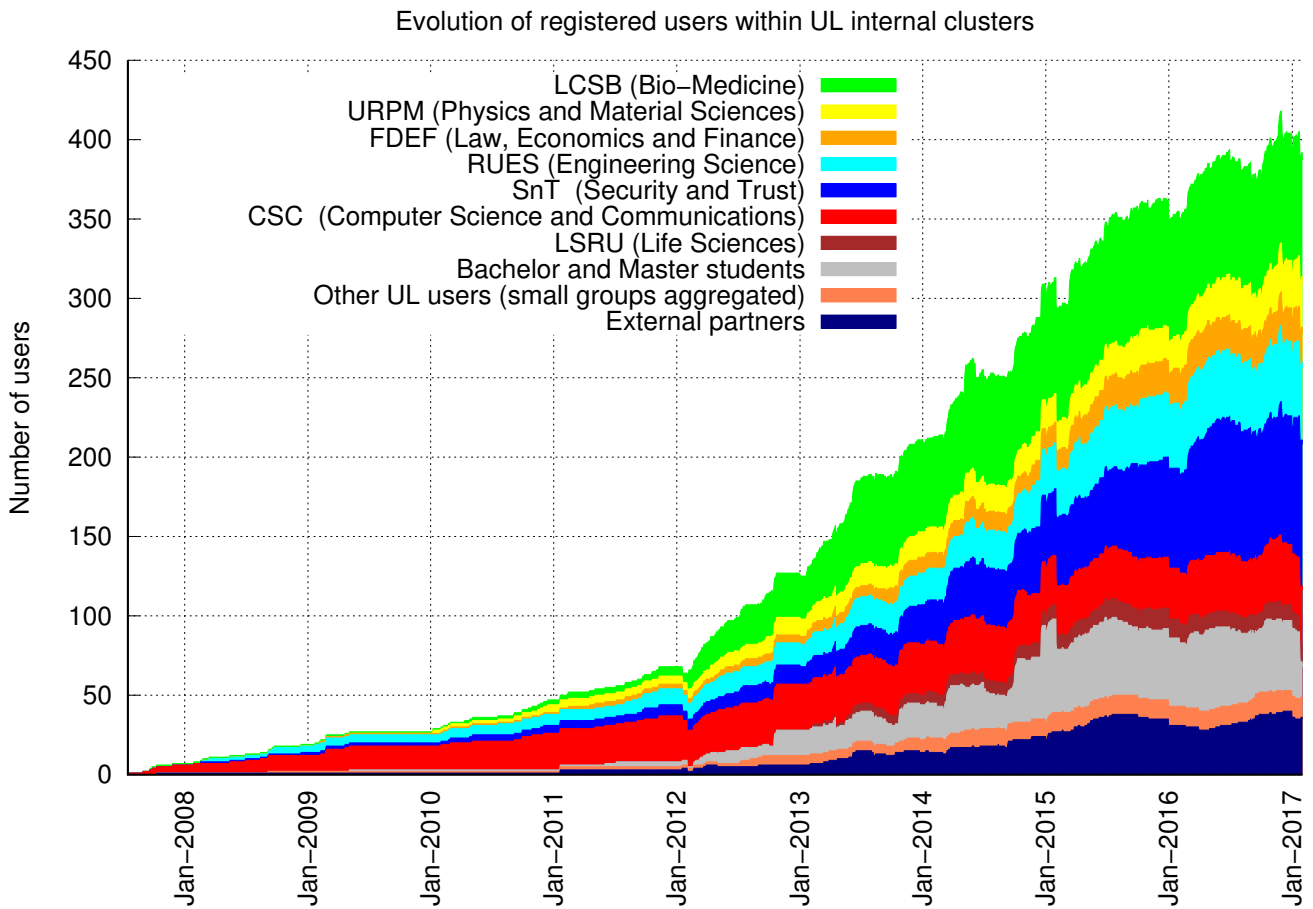

UL HPC Beneficiaries

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 feat. special systems targeting specific workloads:
  - ↪ **Machine Learning & AI**: GPU accelerators
    - ✓ 10 Tesla K40 + 16 Tesla K80 + 24 Tesla M20*: **76 GPU Tflops**
  - ↪ **BigData analytics & data driven science**: large memory systems
    - ✓ Large SMP systems with 1, 2, 3 & 4 TB RAM
  - ↪ **Scale-out workloads**: energy efficient systems
    - ✓ 90 HP Moonshot servers + 96 viridis ARM-based systems

Accelerating UL Research

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

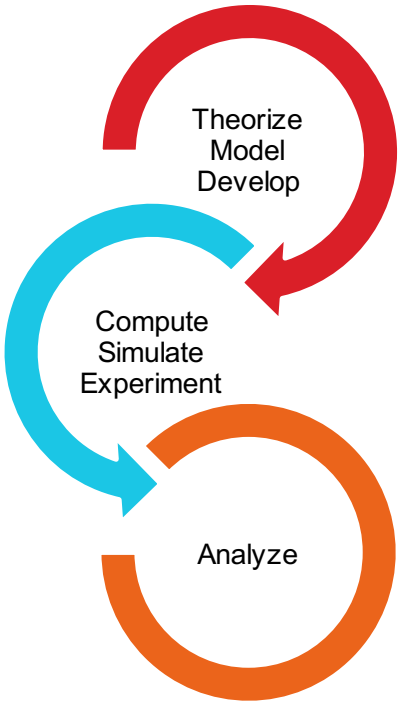

- **>140 software packages** available for researchers
  - ↪ **General purpose**, statistics, optimization:
    - ✓ Matlab, Mathematica, R, Stata, CPLEX, Gurobi Optimizer...
  - ↪ **Bioinformatics**
    - ✓ BioPython, STAR, TopHat, Bowtie, mpiHMMER...
  - ↪ **Computer aided engineering**:
    - ✓ ANSYS, ABAQUS, OpenFOAM...
  - ↪ **Molecular dynamics**:
    - ✓ NAMD, ABINIT, Q.ESPRESSO, GROMACS...
  - ↪ **Visualisation**: ParaView, VisIt, VMD, XCS portal
  - ↪ Compilers, libraries, performance modeling tools
  - ↪ [Parallel] debugging tools aiding development

ULHPC Governance

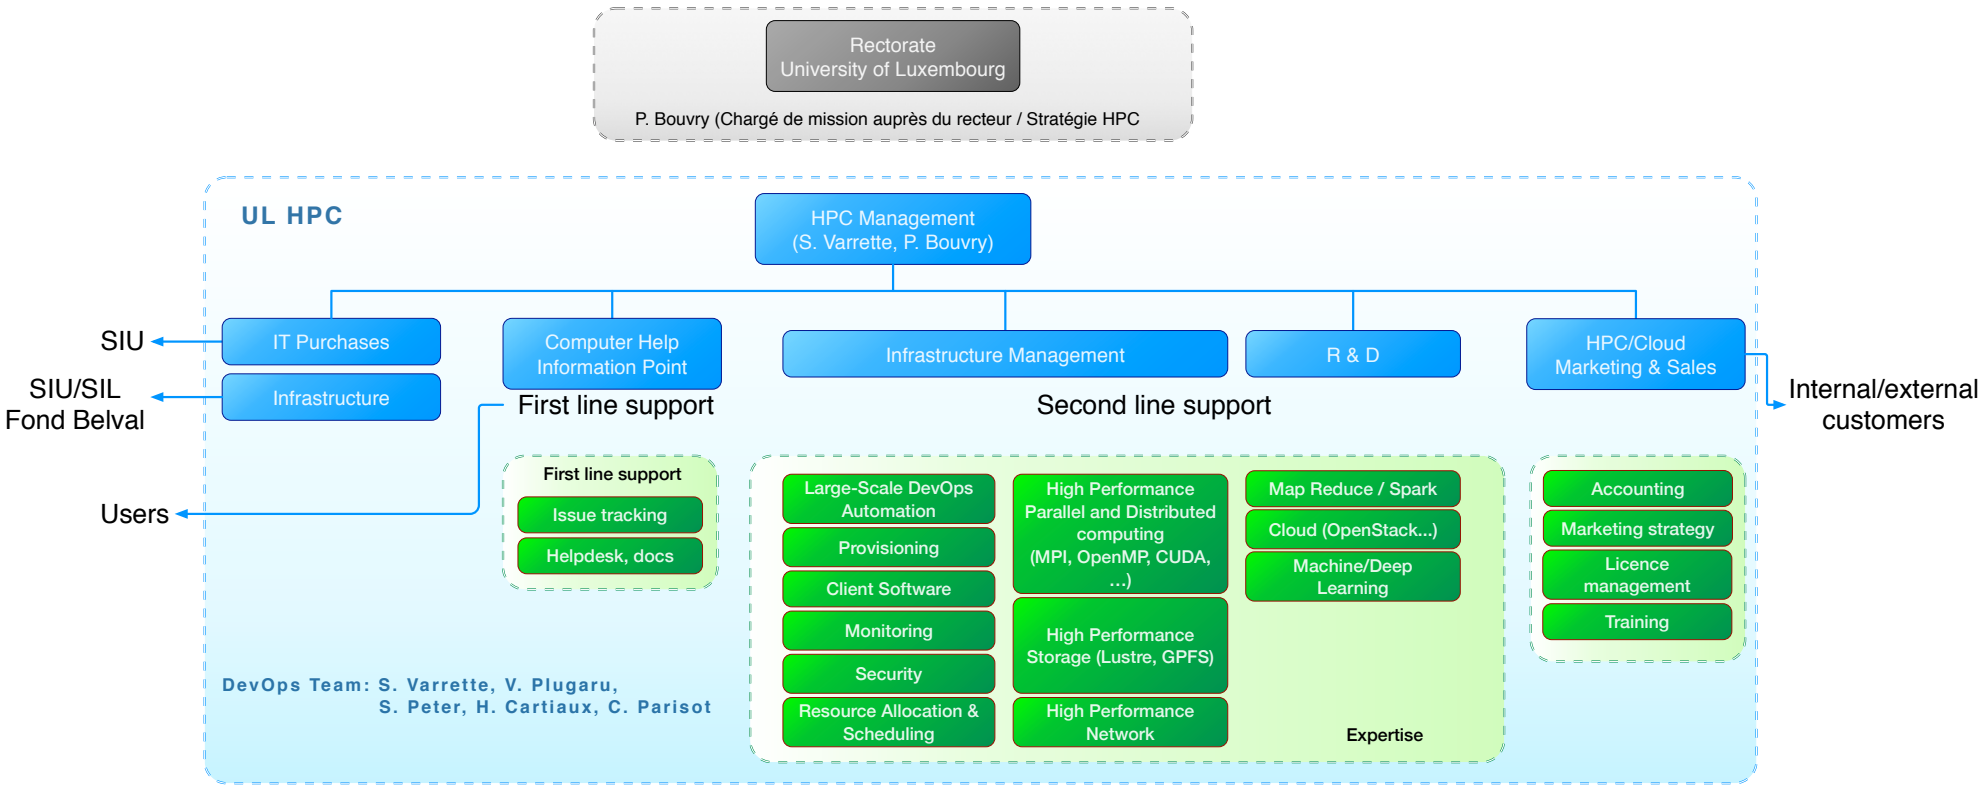

UL HPC Team

Prof. Pascal Bouvry
Director of DS-CSCE, Leader of PCO Group
Senior advisor for the president as regards the HPC strategy

Sébastien Varrette, PhD
CDI, Research Scientist (CSC, FSTC)

Valentin Plugaru, MSc.
CDI, Research Associate (CSC, FSTC)

Sarah Peter, MSc.
CDD, Research Associate (LCSB)

Hyacinthe Cartiaux
CDI, Support (SIU)

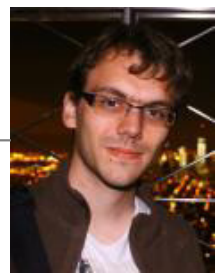

Clément Parisot
CDI, Support (CSC, FSTC)

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: General cluster organization

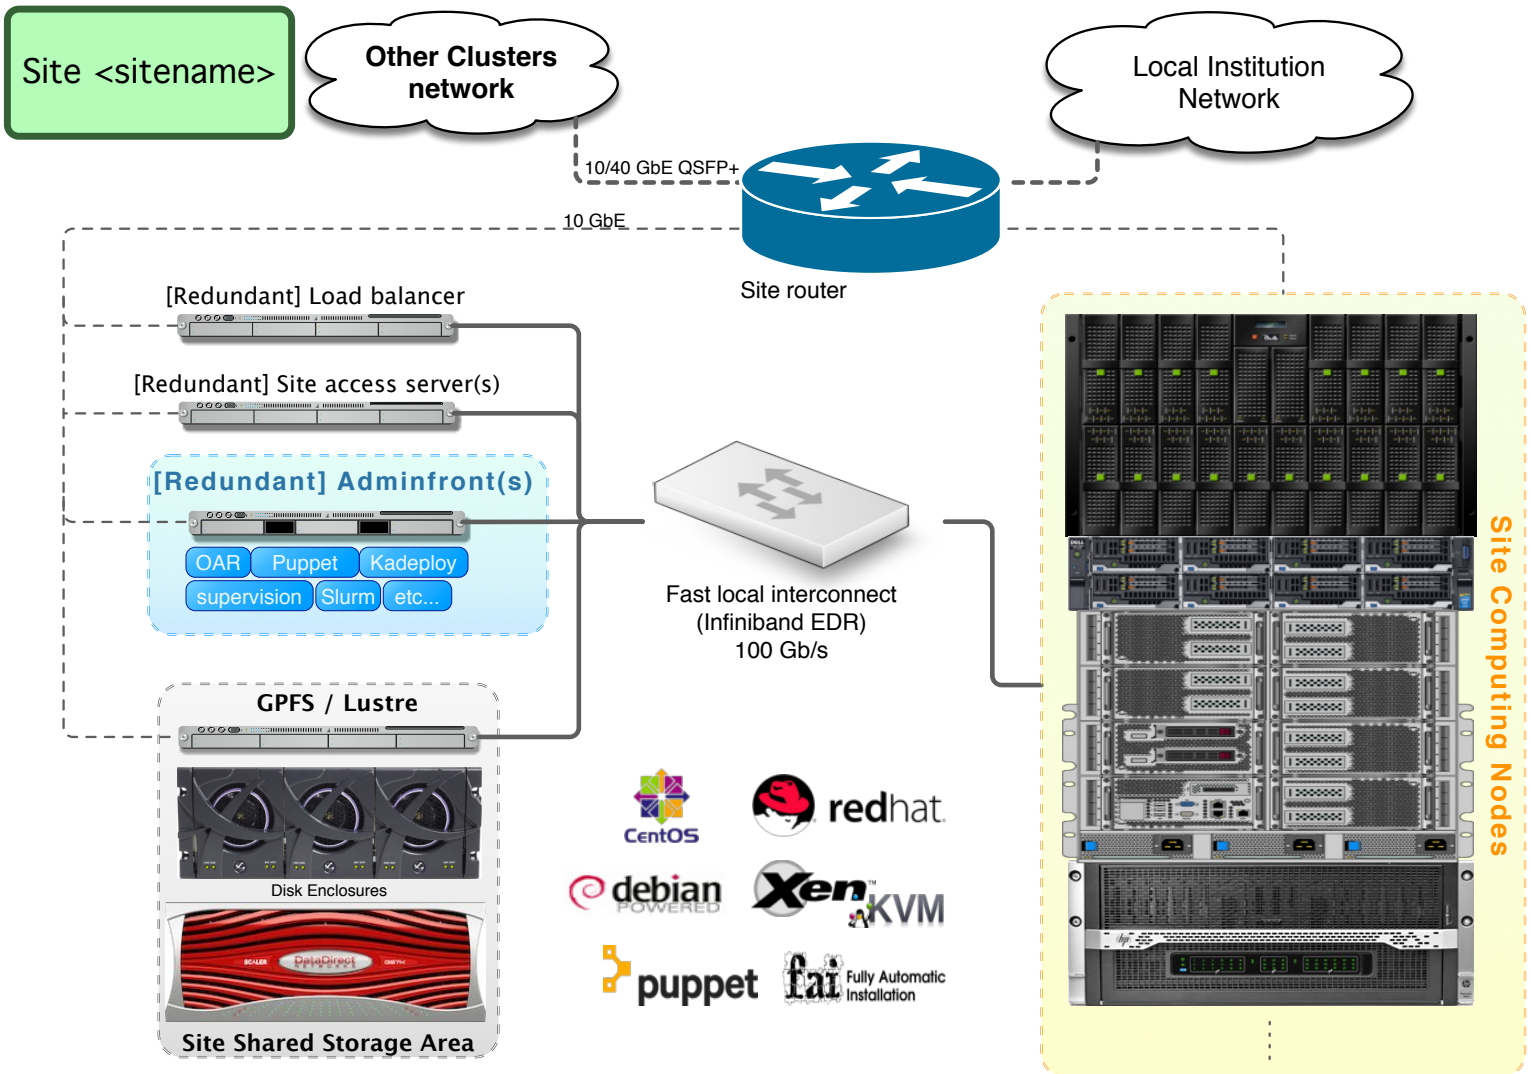

Ex: The gaia cluster

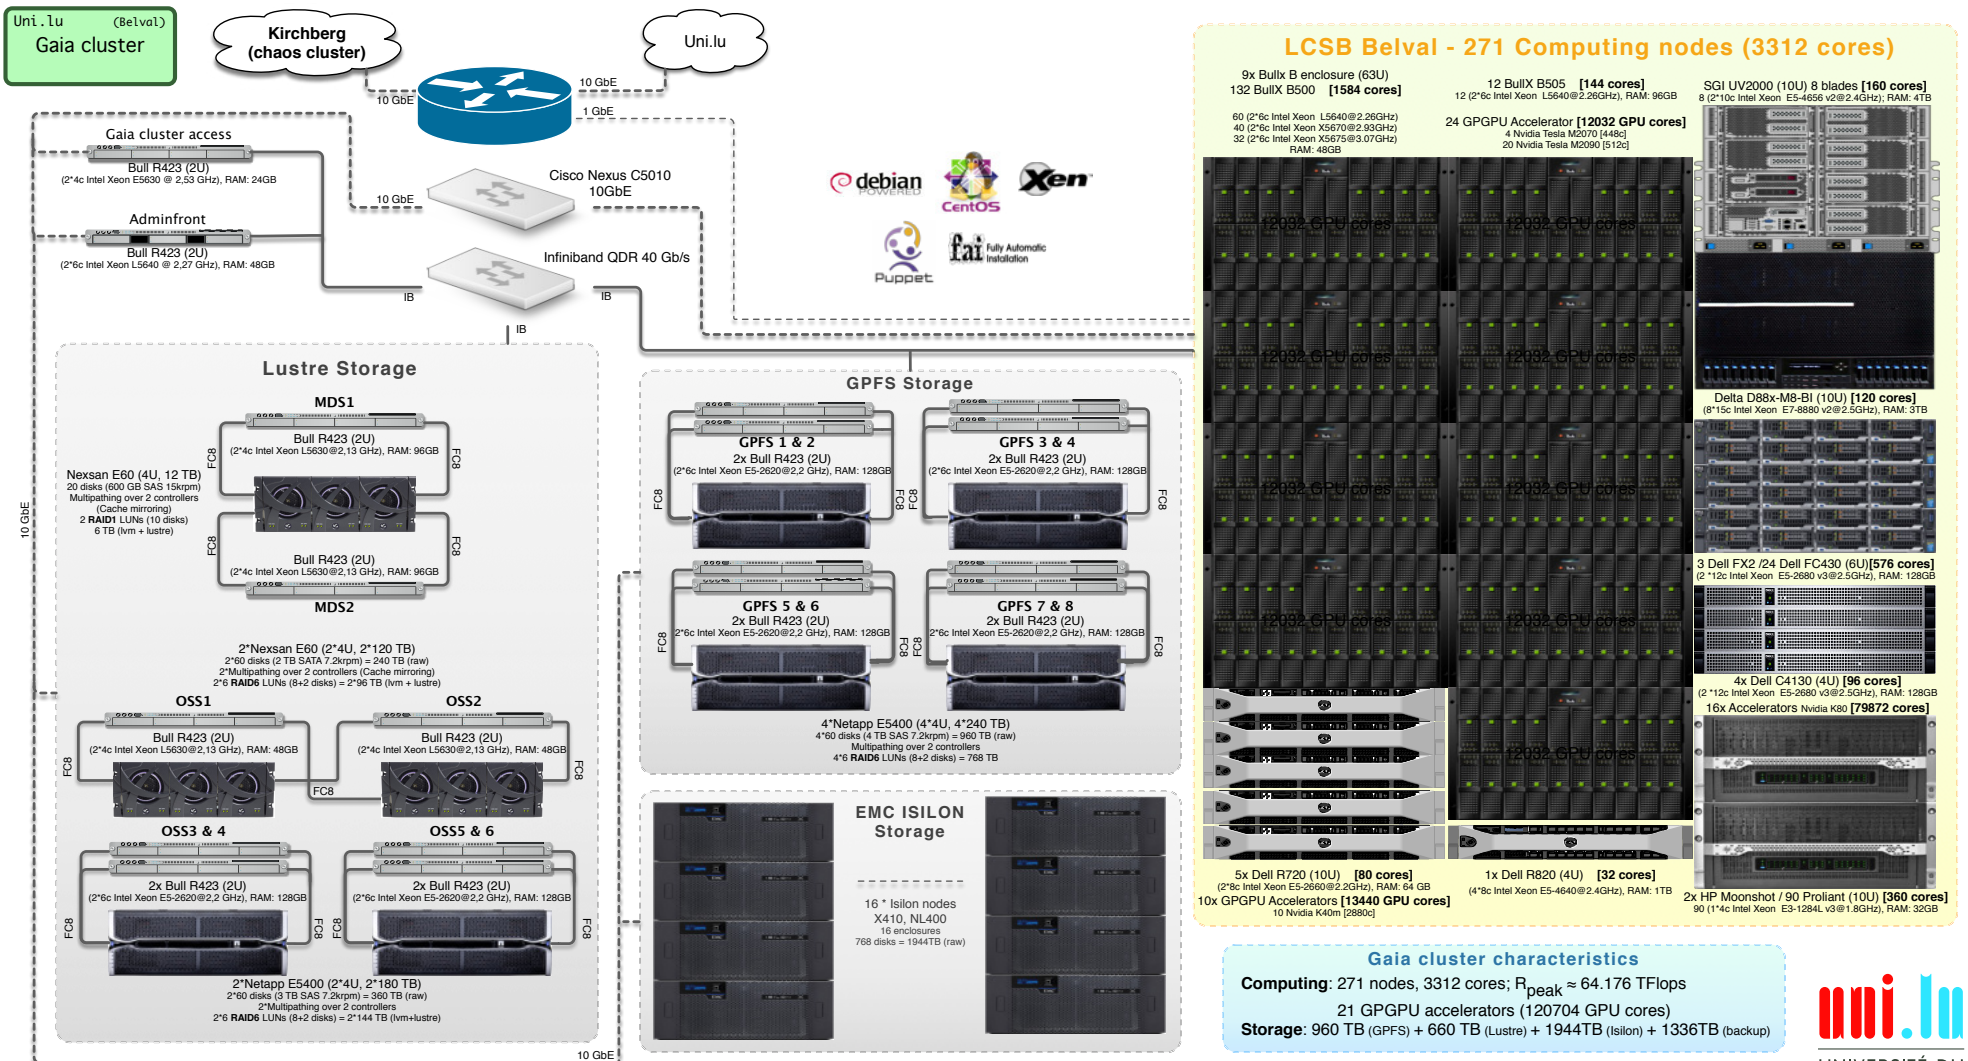

Gaia cluster characteristics

Computing: 271 nodes, 3312 cores; $R_{peak} \approx 64.176$ TFlops
21 GPGPU accelerators (120704 GPU cores)
Storage: 960 TB (GPFS) + 660 TB (Lustre) + 1944TB (Isilon) + 1336TB (backup)

The new iris cluster

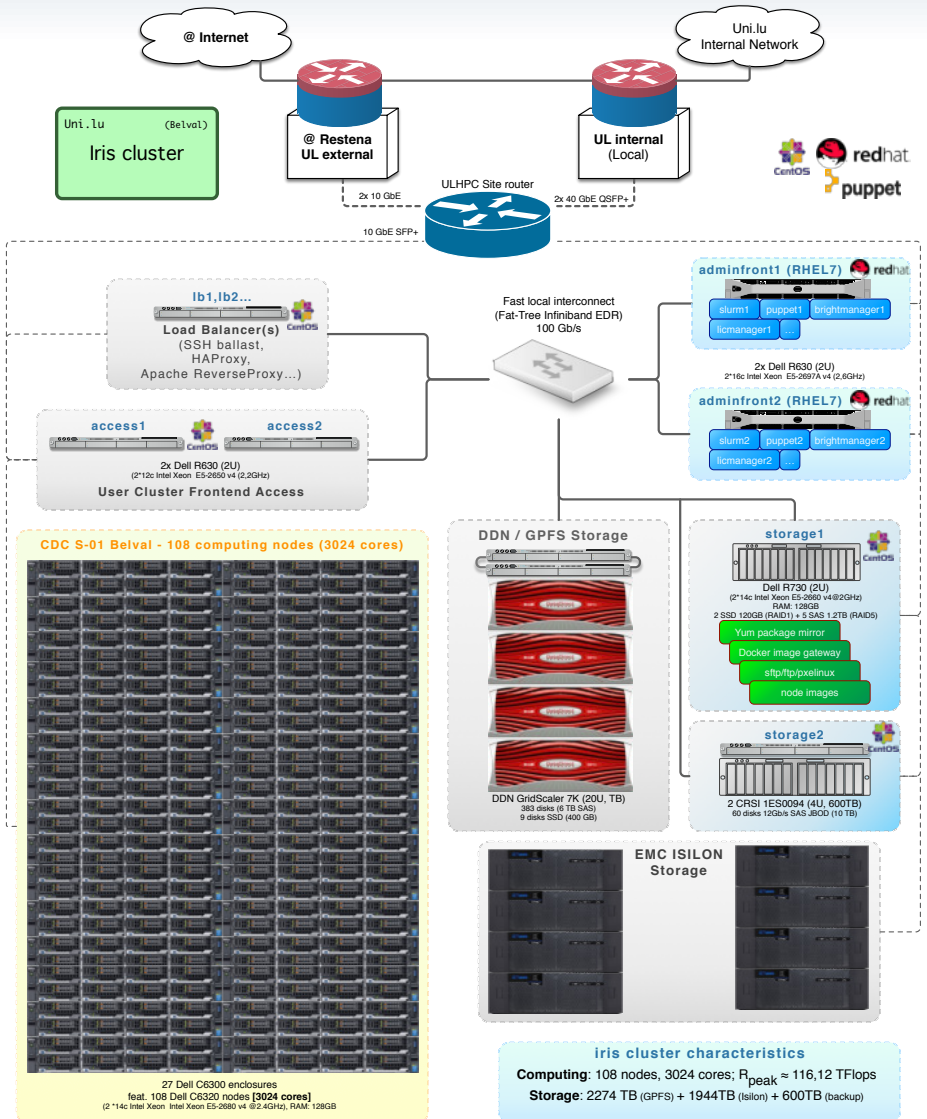

UL HPC Computing capacity

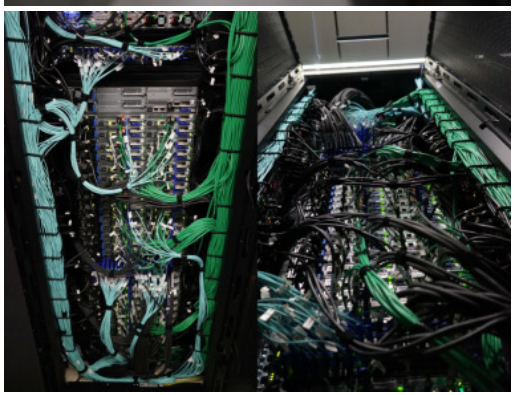

5 clusters

206.772 TFlops

602 nodes

8452 cores

34512GPU cores

UL HPC Computing Clusters

Cluster	Location	**#N**	**#C**	**Rpeak**	GPU Rpeak
`iris`	CDC S-01	108	3024	116.12	0
`gaia`	BT1	273	3440	69.296	76
`chaos`	Kirchberg	81	1120	14.495	0
`g5k`	Kirchberg	38	368	4.48	0
`nyx` (experimental)	BT1	102	500	2.381	0
	TOTAL:	**602**	**8452**	**206.772**	**+ 76 TFlops**

UL HPC – Detailed Computing Nodes

	Date	Vendor	Proc. Description		#N	#C	R_{peak}
iris	2017	Dell	Intel Xeon E5-2680 v4@2.4GHz	2 × 14C,128GB	100	2800	107,52 TFlops
				`iris` TOTAL:	**108**	**3024**	**116.12 TFlops**
gaia	2011	Bull	Intel Xeon L5640@2.26GHz	2 × 6C,48GB	72	864	7.811 TFlops
	2012	Dell	Intel Xeon E5-4640@2.4GHz	4 × 8C, 1TB	1	32	0.614 TFlops
	2012	Bull	Intel Xeon E7-4850@2GHz	16 × 10C,1TB	1	160	1.280 TFLops
	2013	Dell	Intel Xeon E5-2660@2.2GHz	2 × 8C,64GB	5	80	1.408 TFlops
	2013	Bull	Intel Xeon X5670@2.93GHz	2 × 6C,48GB	40	480	5.626 TFlops
	2013	Bull	Intel Xeon X5675@3.07GHz	2 × 6C,48GB	32	384	4.746 TFlops
	2014	Delta	Intel Xeon E7-8880@2.5 GHz	8 × 15C,1TB	1	120	2.4 TFlops
	2014	SGi	Intel Xeon E5-4650@2.4 GHz	16 × 10C,4TB	1	160	3.072 TFlops
	2015	Dell	Intel Xeon E5-2680@2.5 GHz	2 × 12C,128GB	28	672	26.88 TFlops
	2015	HP	Intel E3-1284Lv3, 1.8GHz	1 × 4C,32GB	90	360	10.368 TFlops
	2016	Dell	Intel Xeon E7-8867@2.5 GHz	4 × 16C,2TB	2	128	5.12 TFlops
				`gaia` TOTAL:	**273**	**3440**	**69.296 TFlops**
chaos	2010	HP	Intel Xeon L5640@2.26GHz	2 × 6C,24GB	32	384	3.472 TFlops
	2011	Dell	Intel Xeon L5640@2.26GHz	2 × 6C,24GB	16	192	1.736 TFlops
	2012	Dell	Intel Xeon X7560@2,26GHz	4 × 6C, 1TB	1	32	0.289 TFlops
	2012	Dell	Intel Xeon E5-2660@2.2GHz	2 × 8C,32GB	16	256	4.506 TFlops
	2012	HP	Intel Xeon E5-2660@2.2GHz	2 × 8C,32GB	16	256	4.506 TFlops
				`chaos` TOTAL:	**81**	**1120**	**14.495 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**

UL HPC Storage capacity

4 distributed/parallel FS

2015 disks

7952.4 TB

(incl. 2116TB for Backup)

UL HPC Shared Storage Capacities

Cluster	GPFS	Lustre	Other (NFS...)	Backup	TOTAL
iris	1440	0	6	600	**2046** TB
gaia	960	480	0	1336	**2776** TB
chaos	0	0	180	180	**360** TB
g5k	0	0	32.4	0	**32.4** TB
nyx (experimental)	0	0	242	0	**242** TB
TOTAL:	**2400**	**480**	**2956.4**	**2116**	**7952.4 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.

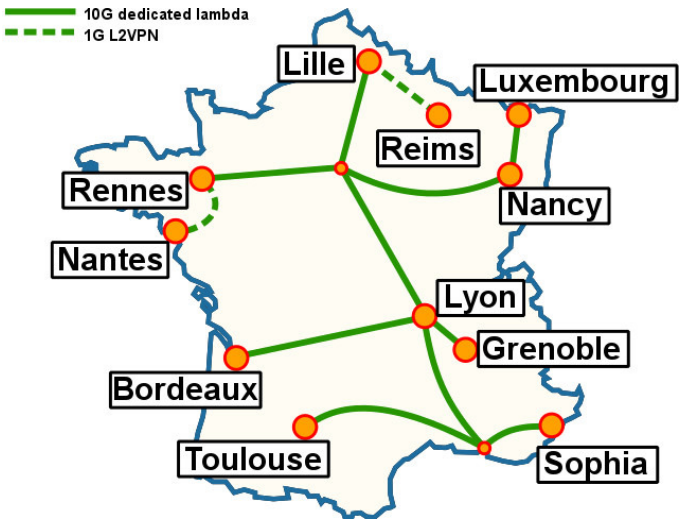

- 10 sites in France
  - ↪ **Abroad**: Luxembourg, Porto Allegre
  - ↪ Total: **7782** cores over **26** clusters
- 1-10GbE / Myrinet / Infiniband
  - ↪ **10Gb/s dedicated** between all sites
- Unique software stack
  - ↪ **kadeploy, kavlan, storage5k**

The case of Grid'5000

http://www.grid5000.fr

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

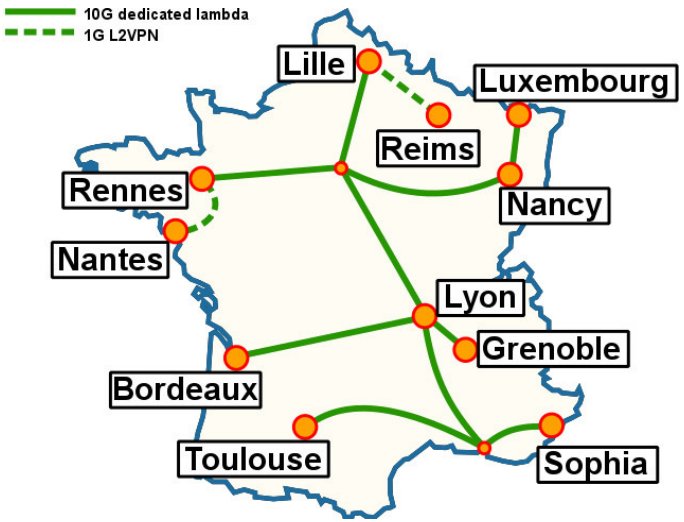

- 10 sites in France
  - ↪ **Abroad**: Luxembourg, Porto Allegre
  - ↪ Total: **7782** cores over **26** clusters
- 1-10GbE / Myrinet / Infiniband
  - ↪ **10Gb/s dedicated** between all sites
- Unique software stack
  - ↪ **kadeploy, kavlan, storage5k**

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

CPU-year usage since 2010

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

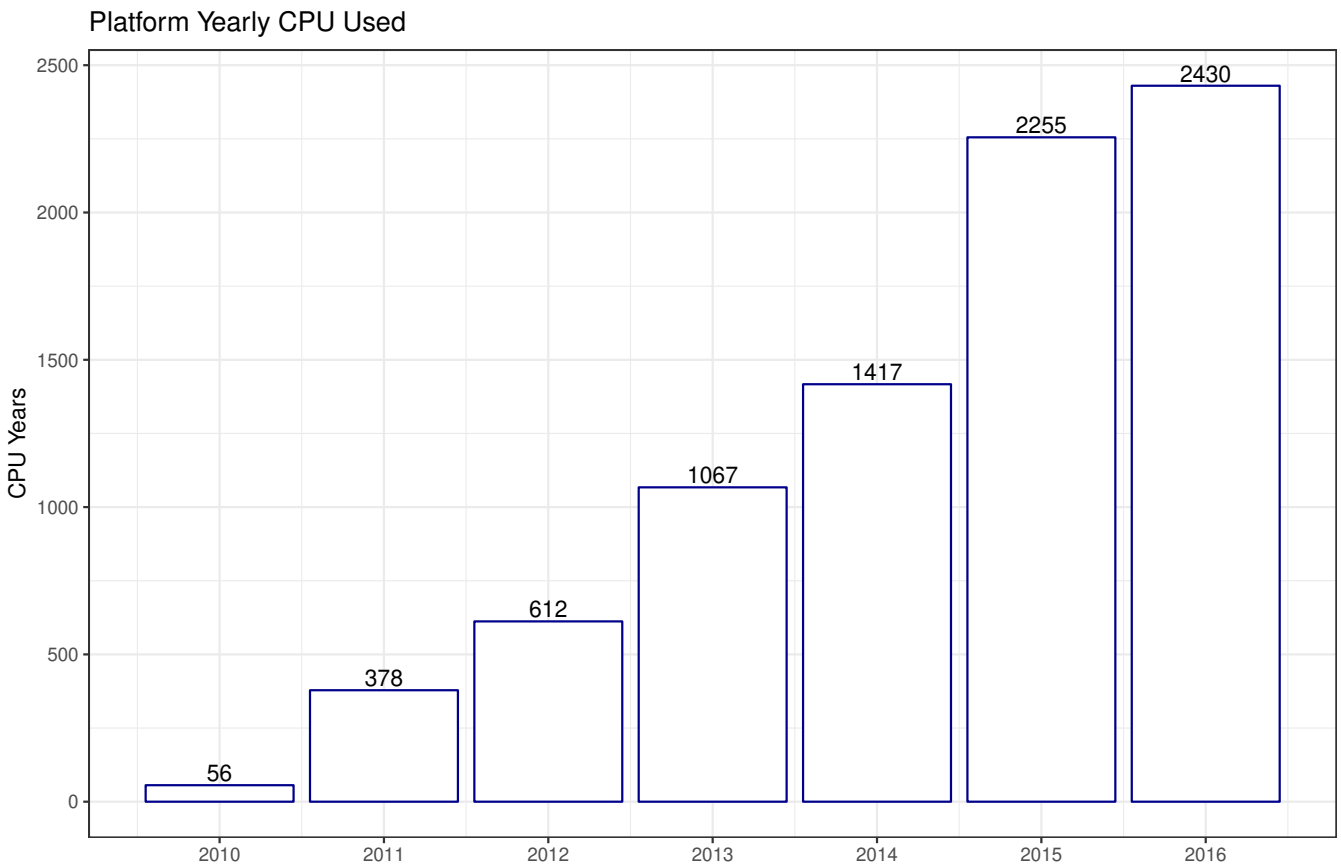

Computing nodes Management

Node deployment by FAI/Bright Manager

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

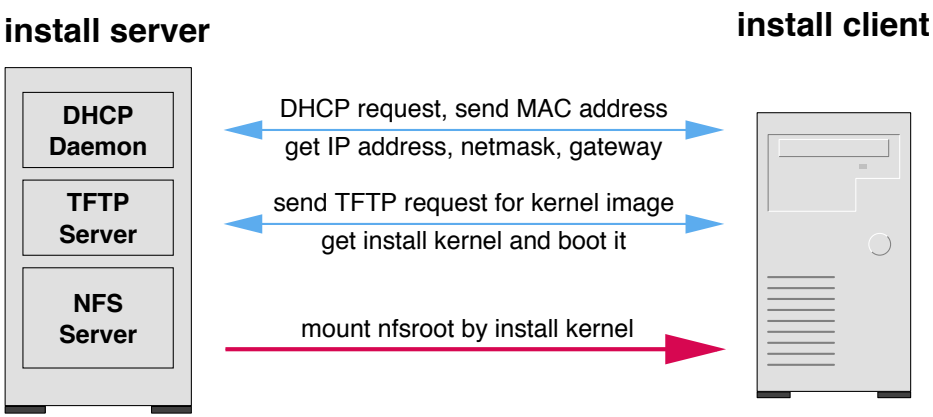

Computing nodes Management

Node deployment by FAI/Bright Manager

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

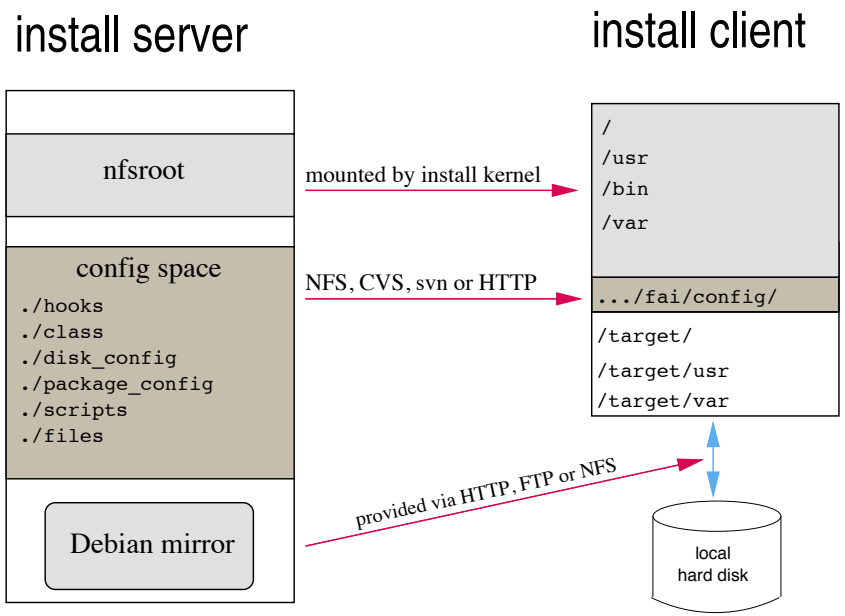

Computing nodes Management

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 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 command
  - ↪ configure OS and additional software
  - ↪ save log files to install server, then reboot new system

Average reinstallation time: $\simeq$ 500s

IT Serv[er|ice] Management: Puppet

Server/Service configuration by Puppet

puppet

http://puppetlabs.com

- **IT Automation** for configuration management
  - ↪ idempotent
  - ↪ agent/master OR stand-alone architecture
  - ↪ cross-platform through Puppet's Resource Abstraction Layer (RAL)
  - ↪ Git-based workflow
  - ↪ PKI-based security (X.509)
- DevOps tool of choice for configuration management
  - ↪ Declarative Domain Specific Language (DSL)

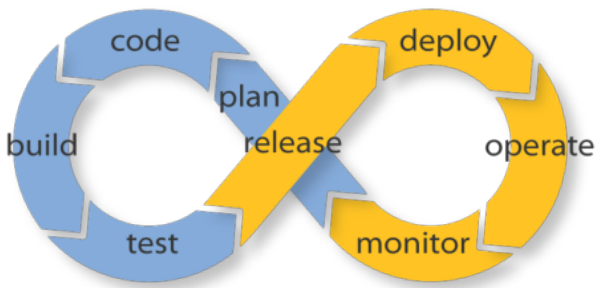

Endless Possibilities: DevOps can create an infinite loop of release and feedback for all your code and deployment targets.

IT Serv[er|ice] Management: Puppet

Server/Service configuration by Puppet

http://puppetlabs.com

- **IT Automation** for configuration management
  - ↪ idempotent
  - ↪ agent/master OR stand-alone architecture
  - ↪ cross-platform through Puppet's Resource Abstraction Layer (RAL)
  - ↪ Git-based workflow
  - ↪ PKI-based security (X.509)
- DevOps tool of choice for configuration management
  - ↪ Declarative Domain Specific Language (DSL)

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

General Puppet Infrastructure

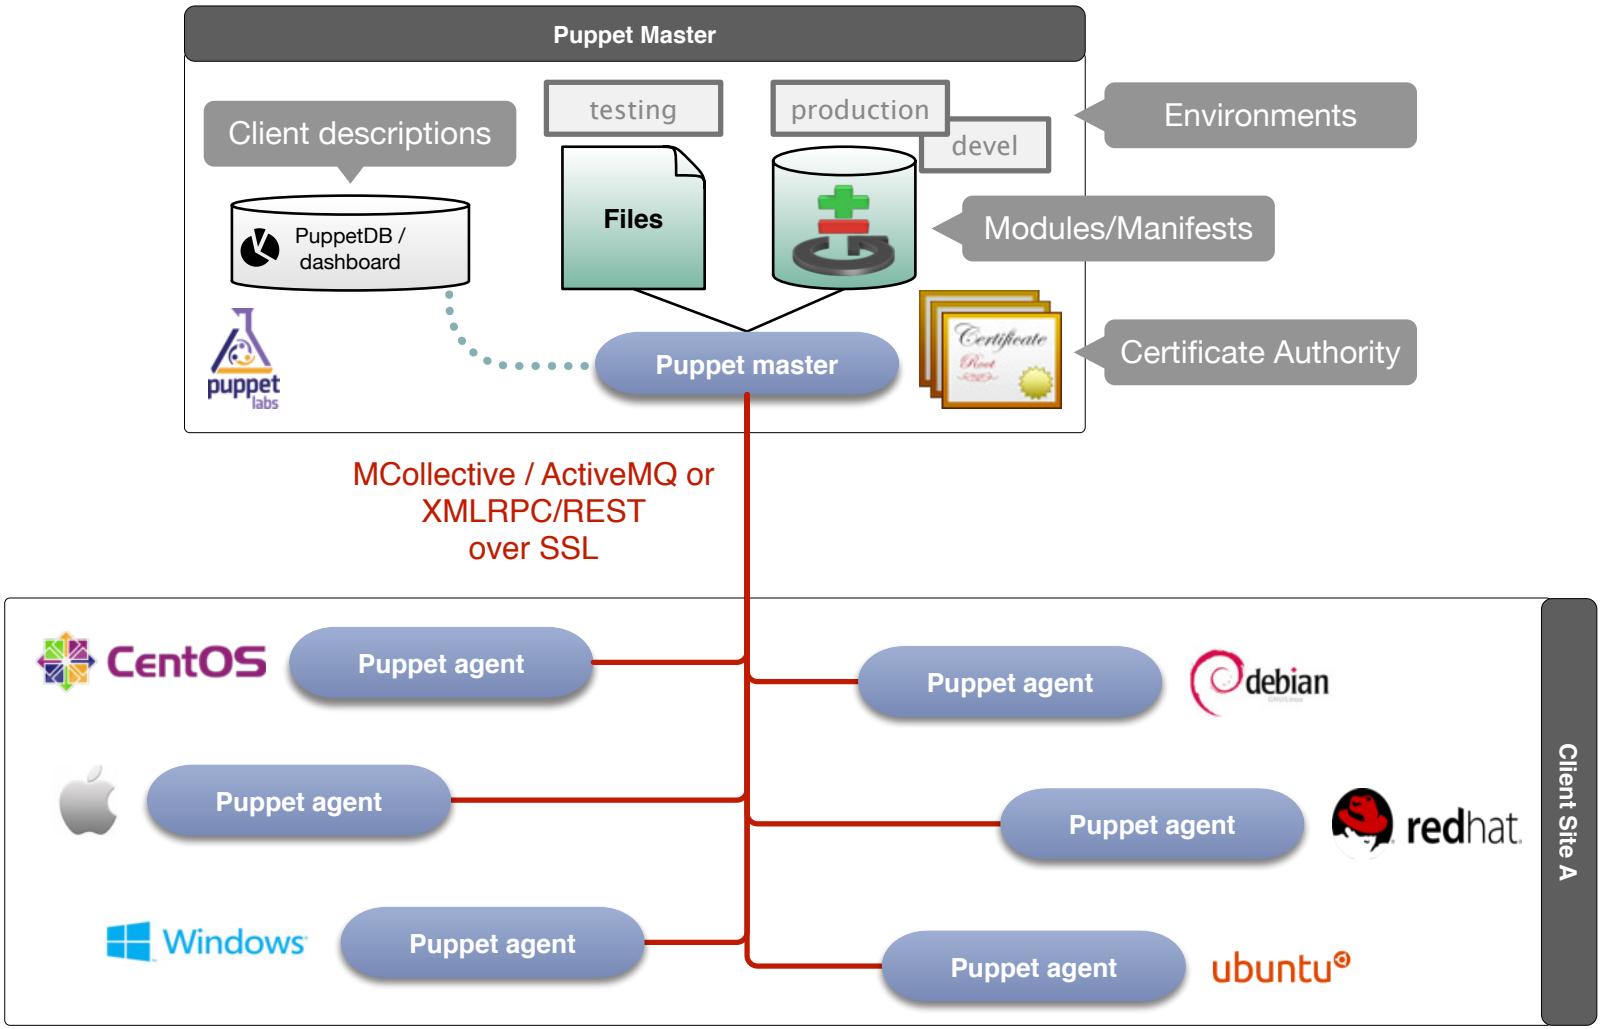

Software/Modules Management

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:
  - ↪ > **140 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

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

- Easybuild: open-source framework to (automatically) build scientific software
- **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 /path/to/easybuild
$> module load tools/EasyBuild
$> eb -S HPL    # Search for recipes for HPL software
$> eb HPL-2.2-intel-2017a.eb # Install HPL 2.2 w. Intel toolchain
```

Software/Modules Management

http://resif.readthedocs.io/en/latest/

- **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
  - ↪ "Easy to use" yet **pending workflow rework**

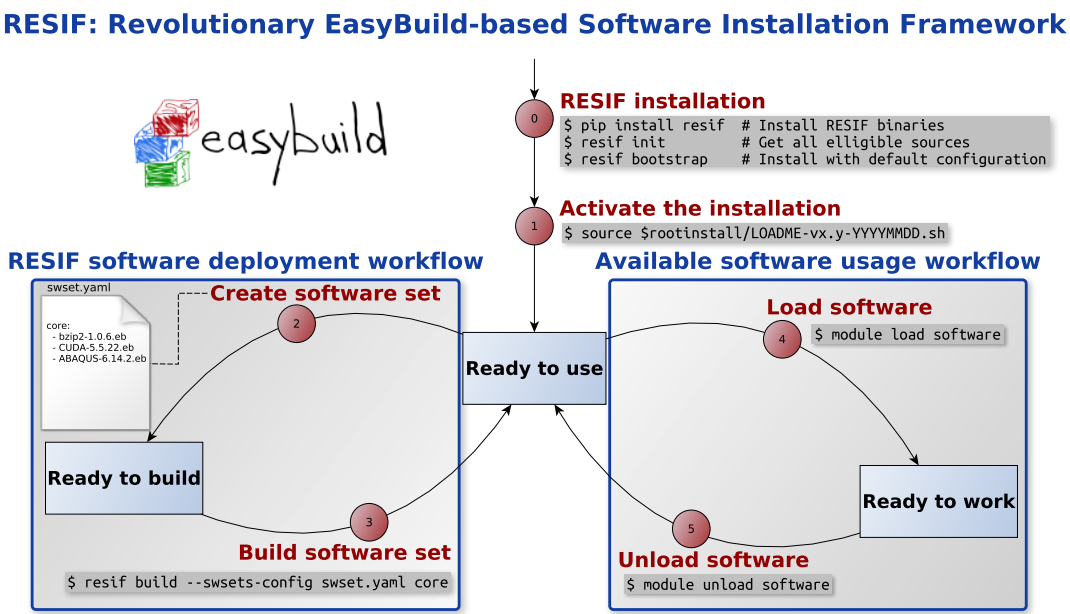

Platform Monitoring

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

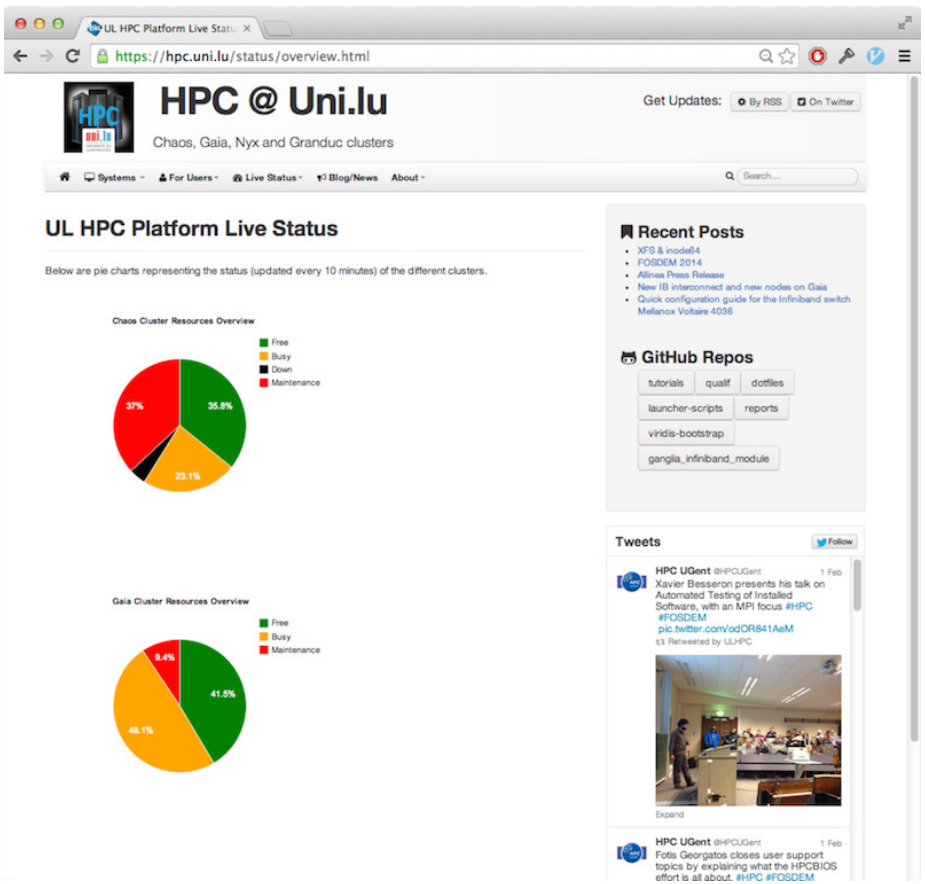

Platform Monitoring

- **Monika** `http://hpc.uni.lu/{gaia,chaos,g5k}/monika`

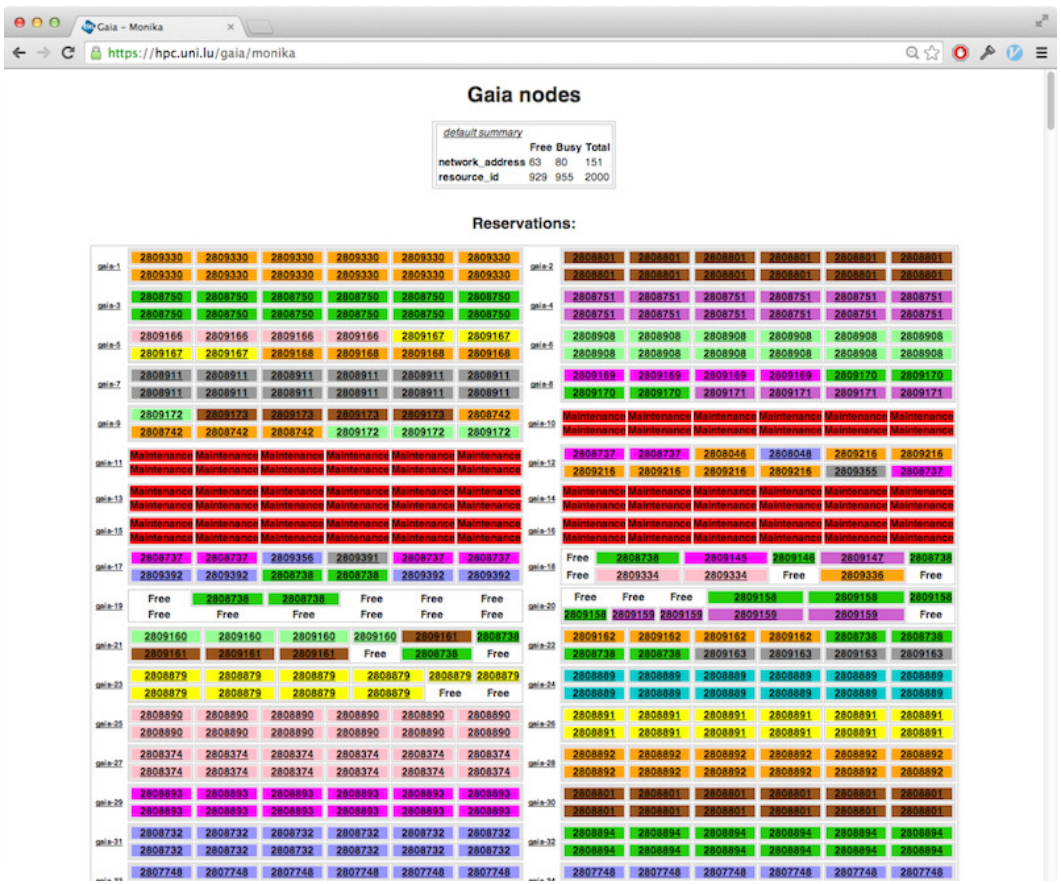

Platform Monitoring

- **Drawgantt** `http://hpc.uni.lu/{gaia,chaos,g5k}/drawgantt`

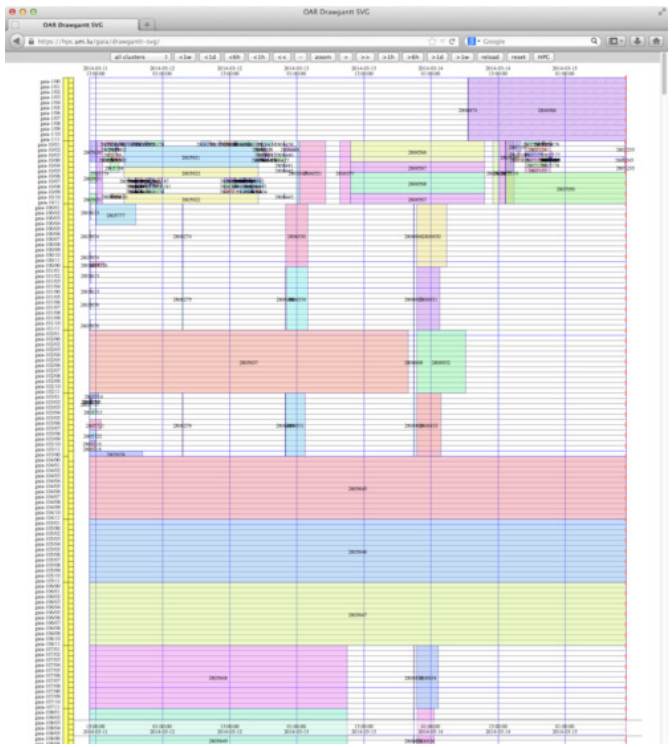

Platform Monitoring

- **Ganglia**

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

Platform Monitoring

- **SLURM-Web** `http://hpc.uni.lu/iris/slurm/` (soon)

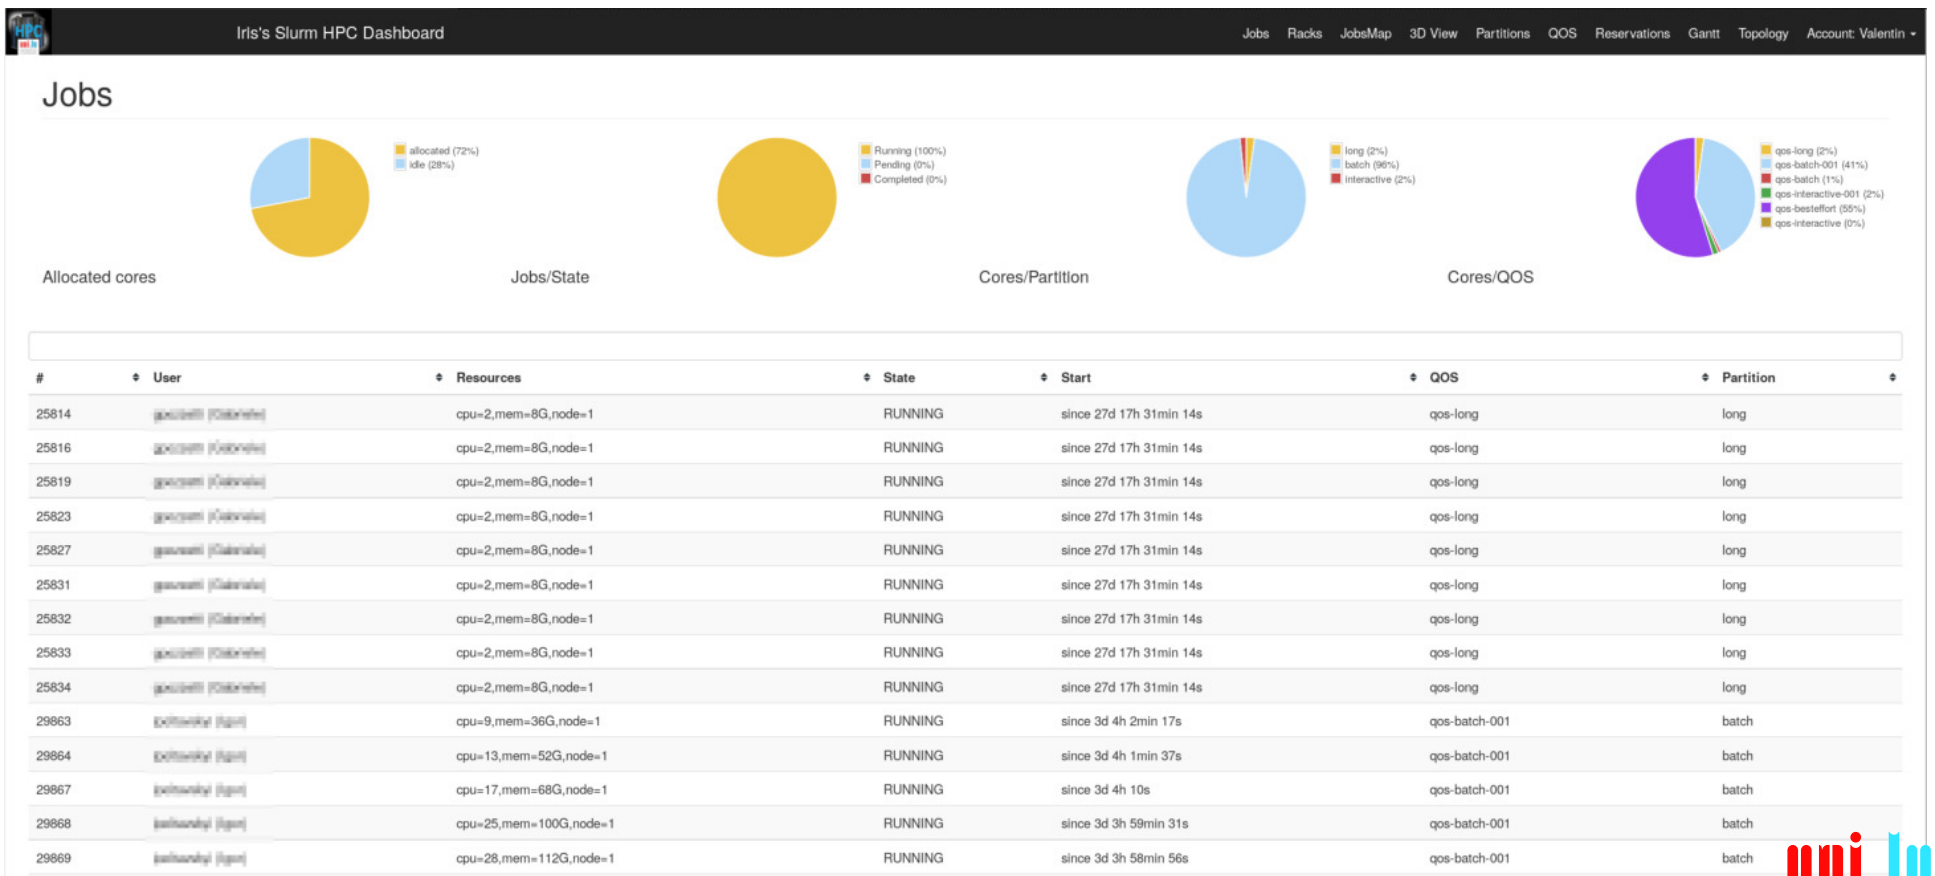

#	User	Resources	State	Start	QOS	Partition
25814	[redacted]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25816	[redacted]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25819	[redacted]	cpu=2,mem=8G,node=1	RUNNING	since 27d 17h 31min 14s	qos-long	long
25823	[redacted]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

Platform Monitoring

- **CDash** http://cdash.uni.lu/

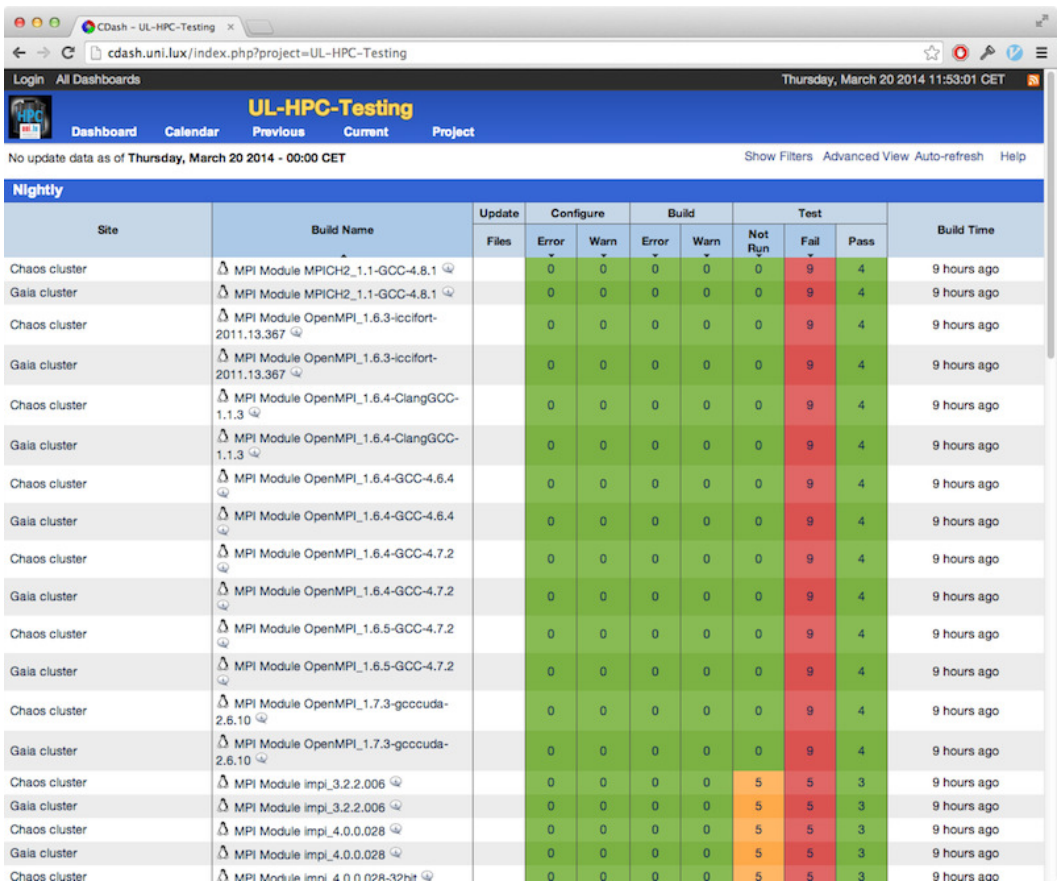

Site	Build Name	Update	Configure		Build		Test			Build Time
		Files	Error	Warn	Error	Warn	Not Run	Fail	Pass	
Chaos cluster	MPI Module MPICH2_1.1-GCC-4.8.1		0	0	0	0	0	9	4	9 hours ago
Gaia cluster	MPI Module MPICH2_1.1-GCC-4.8.1		0	0	0	0	0	9	4	9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.3-iccifort-2011.13.367		0	0	0	0	0	9	4	9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.3-iccifort-2011.13.367		0	0	0	0	0	9	4	9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.4-ClangGCC-1.1.3		0	0	0	0	0	9	4	9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.4-ClangGCC-1.1.3		0	0	0	0	0	9	4	9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.4-GCC-4.6.4		0	0	0	0	0	9	4	9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.4-GCC-4.6.4		0	0	0	0	0	9	4	9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.4-GCC-4.7.2		0	0	0	0	0	9	4	9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.4-GCC-4.7.2		0	0	0	0	0	9	4	9 hours ago
Chaos cluster	MPI Module OpenMPI_1.6.5-GCC-4.7.2		0	0	0	0	0	9	4	9 hours ago
Gaia cluster	MPI Module OpenMPI_1.6.5-GCC-4.7.2		0	0	0	0	0	9	4	9 hours ago
Chaos cluster	MPI Module OpenMPI_1.7.3-gcccuda-2.6.10		0	0	0	0	0	9	4	9 hours ago
Gaia cluster	MPI Module OpenMPI_1.7.3-gcccuda-2.6.10		0	0	0	0	0	9	4	9 hours ago
Chaos cluster	MPI Module impi_3.2.2.006		0	0	0	0	5	5	3	9 hours ago
Gaia cluster	MPI Module impi_3.2.2.006		0	0	0	0	5	5	3	9 hours ago
Chaos cluster	MPI Module impi_4.0.0.028		0	0	0	0	5	5	3	9 hours ago
Gaia cluster	MPI Module impi_4.0.0.028		0	0	0	0	5	5	3	9 hours ago
Chaos cluster	MPI Module impi_4.0.0.028-32bit		0	0	0	0	5	5	3	9 hours ago

Platform Monitoring

- **Internal Monitoring**

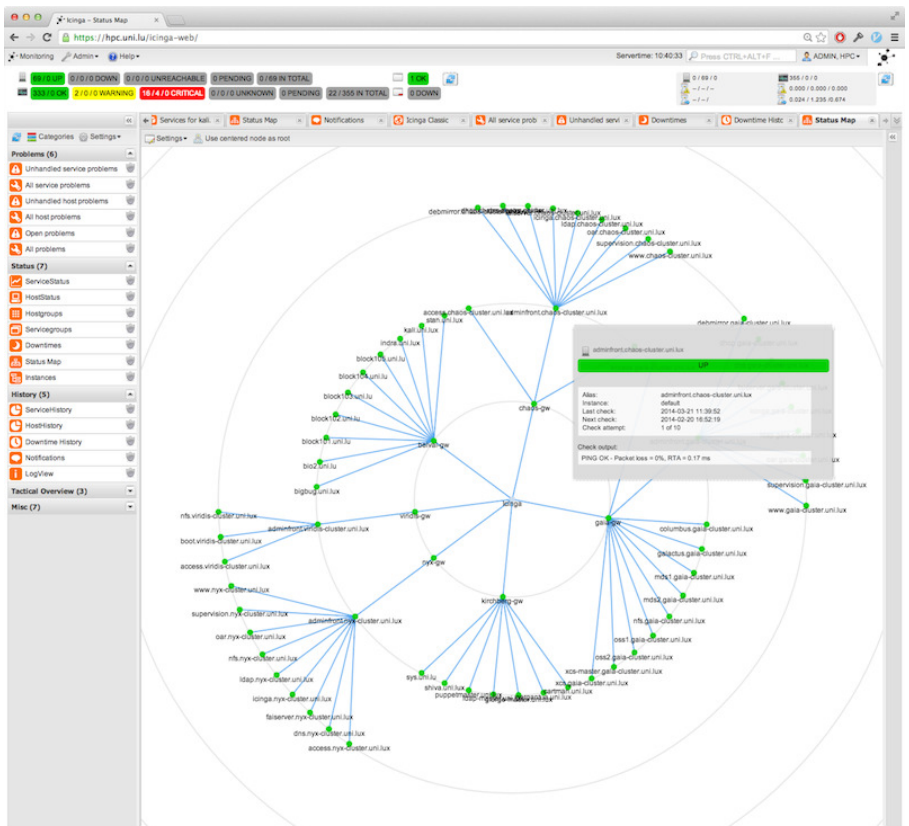

Icinga / Puppet / NetXMS (networking)

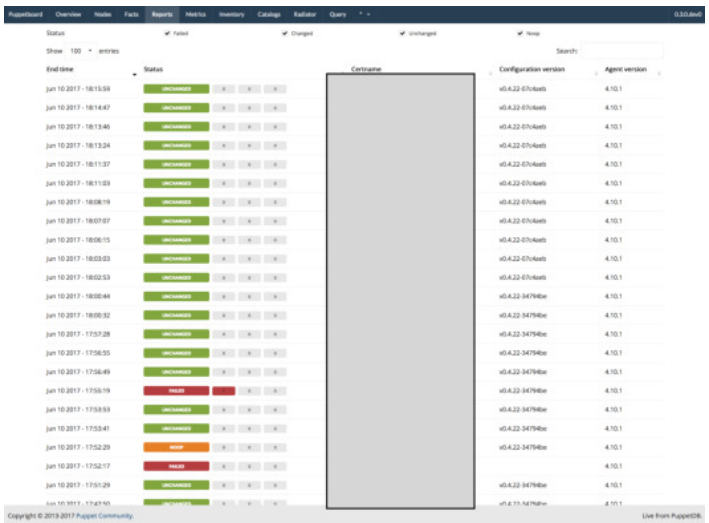

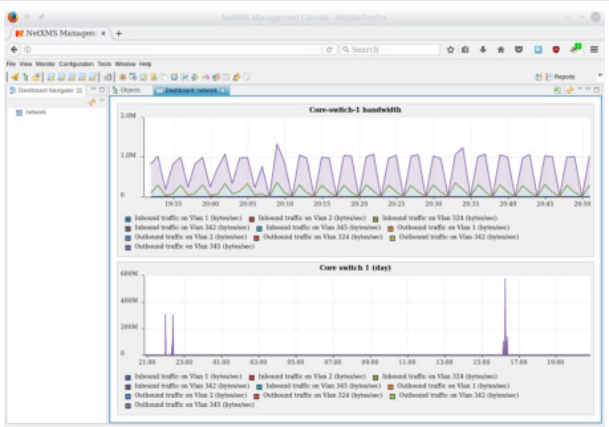

Platform Monitoring

- **Internal Monitoring**

[Disk] Enclosure status

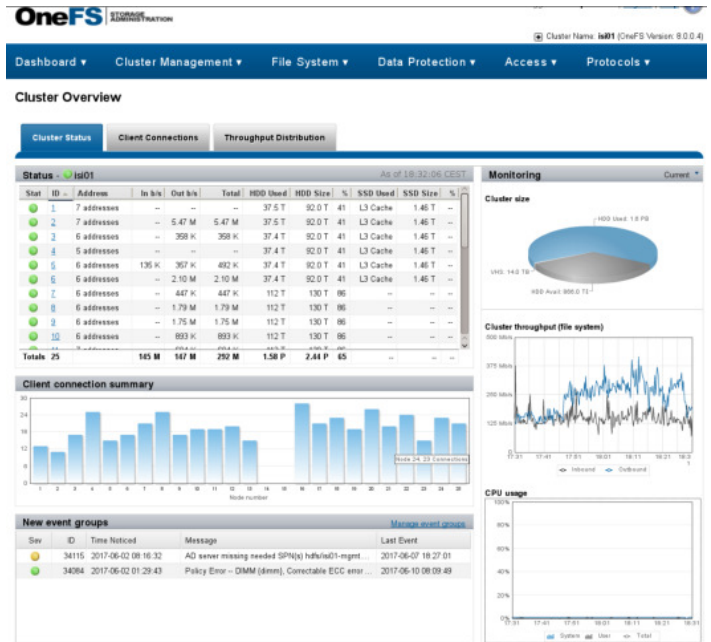

Summary

International State of Affairs

- **Exascale Race/Technology**

IDC-Projected Exascale Investment Levels (In Addition to System Purchases)

U.S.

- $1 to $2 billion a year in R&D (including NRE)
- Investments by both governments & vendors
- Plans are to purchase multiple exascale systems

EU

- About 5 billion euros in total
- Investments in multiple exascale and pre-exascale systems
- Investments mostly by country governments with a little from the EU

China

- Over $1billion a year in R&D
- Investments by both governments & vendors
- Plans are to purchases multiple exascale systems each year
- Already investing in 3 pre-exascale systems by 2017/18

Japan

- Planned investment of just over $1billion* (over 5 years) for both the R&D and purchase of 1 exascale system
- To be followed by a number of smaller systems ~$100M to $150M each
- Creating a new processor and a new software environment

©Hyperion Research 2017

** Note that this includes both the system and R&D*

International State of Affairs

- **Exascale Race/Technology**

IDC-Projected Exascale Dates and Suppliers

U.S.
- Sustained ES: 2023
- Peak ES: 2021
- Vendors: U.S.
- Processors: U.S.
- Initiatives: NSCI/ECP
- Cost: $300-500M per system, plus heavy R&D investments

EU
- Sustained ES: 2023-24
- Peak ES: 2021
- Vendors: U.S., Europe
- Processors: U.S., ARM
- Initiatives: PRACE, ETP4HPC
- Cost: $300-$350 per system, plus heavy R&D investments

China
- Sustained ES: 2023
- Peak ES: ~~2020~~ 2019…
- Vendors: Chinese
- Processors: Chinese (plus ~~U.S.~~?)
- 13th 5-Year Plan
- Cost: $350-500M per system, plus heavy R&D

Japan
- Sustained ES: 2023-24
- Peak ES: Not planned
- Vendors: Japanese
- Processors: Japanese
- Cost: $600-850M, this includes both 1 system and the R&D costs…will also do many smaller size systems

Exascale Feasibility

We Can Build an Exascale System Today?

Connect together 10 Sunway TaihuLight systems

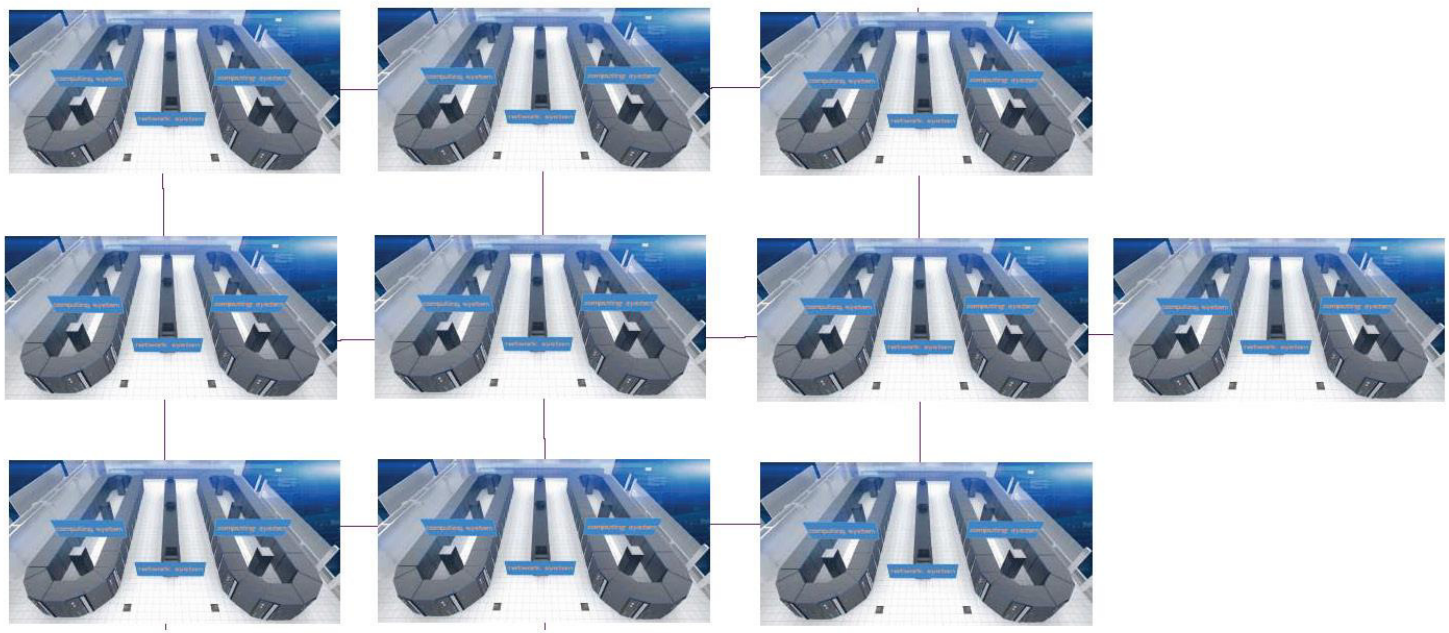

Require 150 MW of power, programming for 100 M threads, and $2.7B price tag

22

European HPC strategy

- EU HPC strategy initiated in 2012
  - ↪ implementation within H2020 program

European HPC strategy

- EU HPC strategy initiated in 2012
  - ↪ implementation within H2020 program

- More recently:
  - ↪ 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

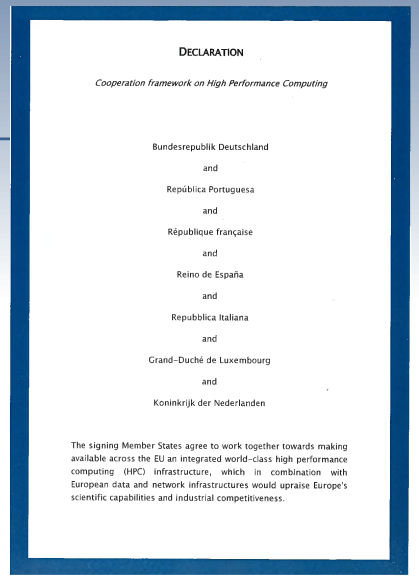
DECLARATION

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
- More recently:
  - ↪ 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...
- **Latest advances**:
  - ↪ EU Member States sign EuroHPC (Mar. 2017)
    - ✓ common effort to create/grow European supercomputing ecosystem
    - ✓ Federation of national/regional HPC centers (see also PRACE2)
  - ↪ EU Objective with EuroHPC:
    - ✓ 2-3 **Pre-exascale** systems by 2019, **2 exascale** systems by 2021

European HPC strategy

- EU HPC strategy initiated in 2012
  - ↪ implementation within H2020 program

- More recently:
  - ↪ 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...

- **Latest advances**:
  - ↪ EU Member States sign EuroHPC (Mar. 2017)
    - ✓ common effort to create/grow European supercomputing ecosystem
    - ✓ Federation of national/regional HPC centers (see also PRACE2)
  - ↪ EU Objective with EuroHPC:
    - ✓ 2-3 **Pre-exascale** systems by 2019, **2 exascale** systems by 2021

⇒ which architecture/approach to sustain these developments?

Different Needs for Different Domains

Material Science & Engineering

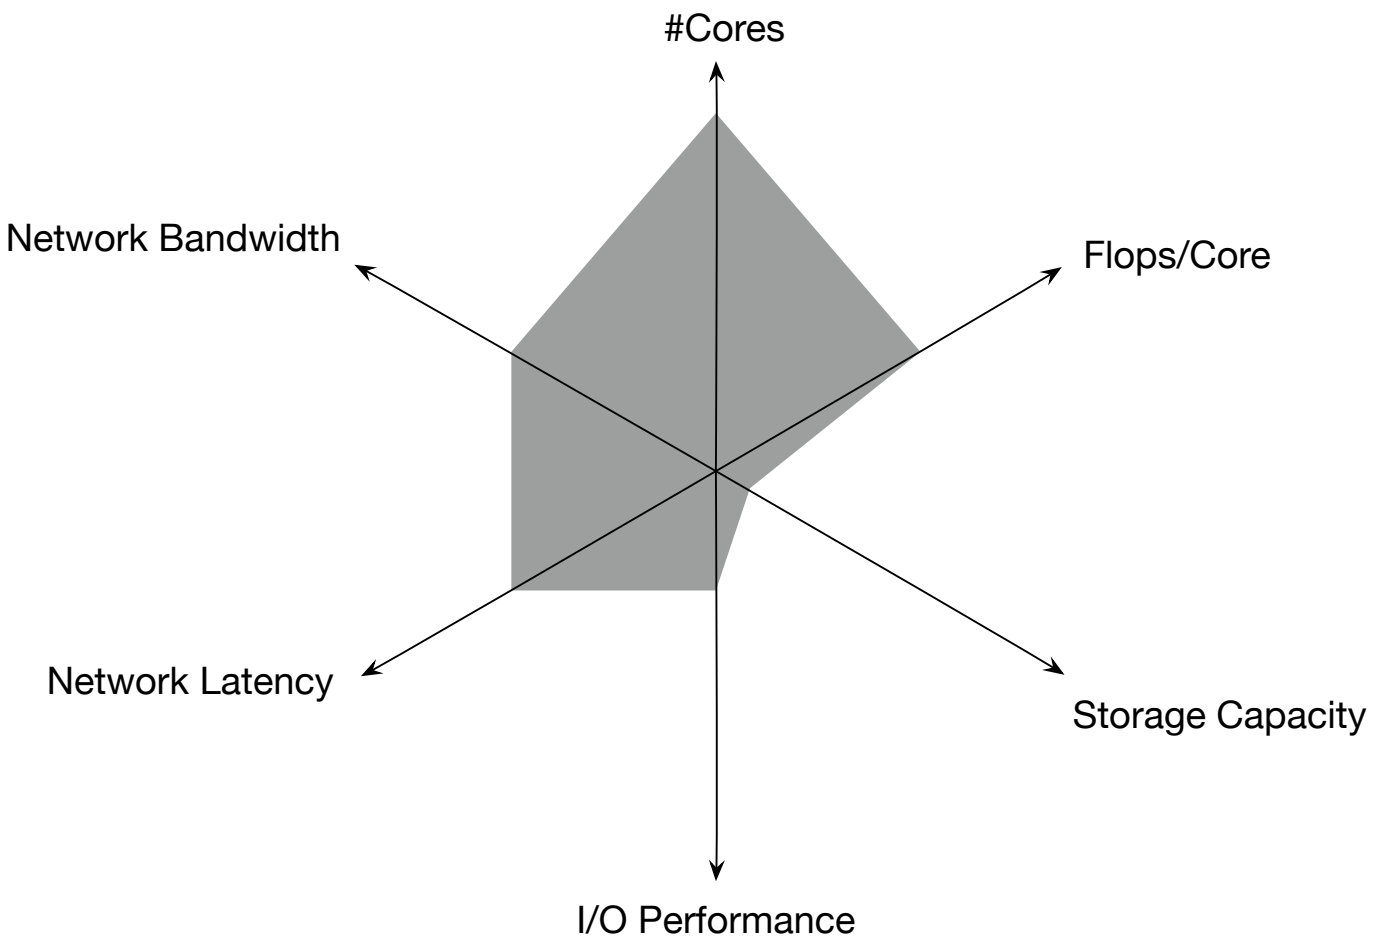

Different Needs for Different Domains

Biomedical Industry / Life Sciences

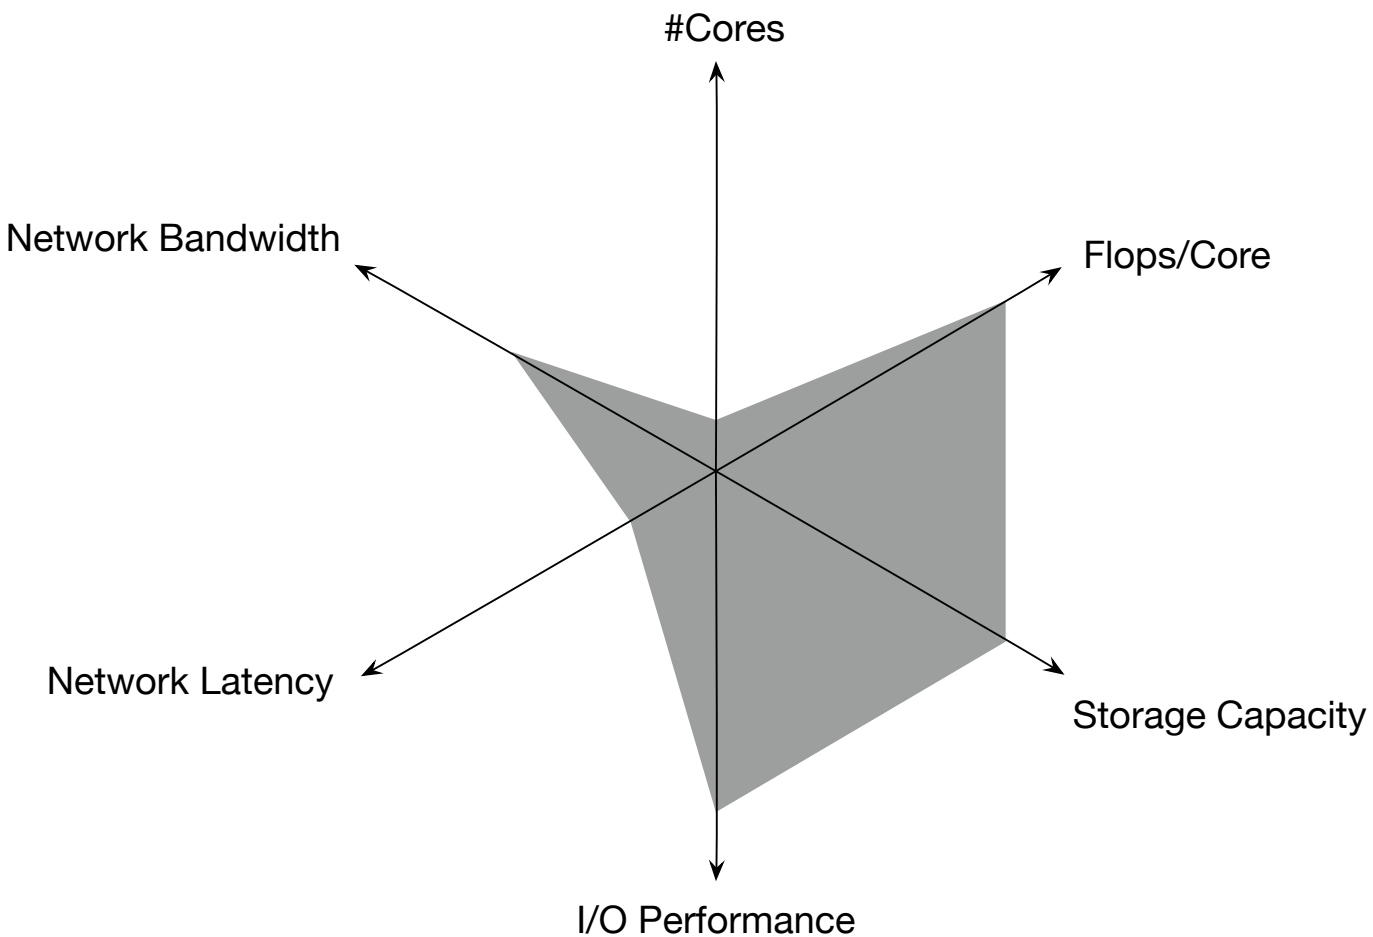

Different Needs for Different Domains

Deep Learning / Cognitive Computing

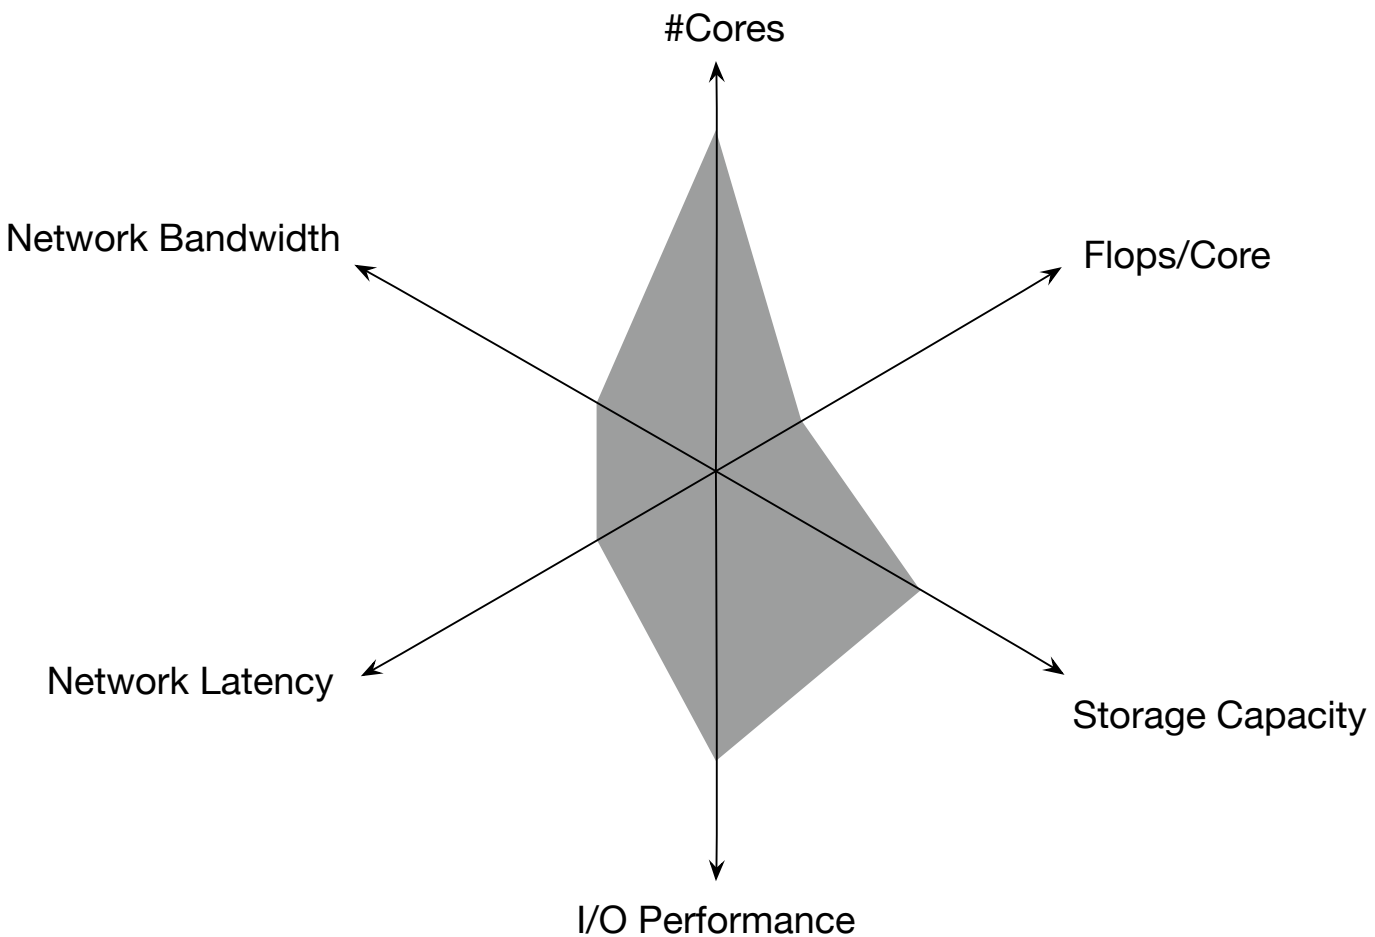

Different Needs for Different Domains

IoT, FinTech

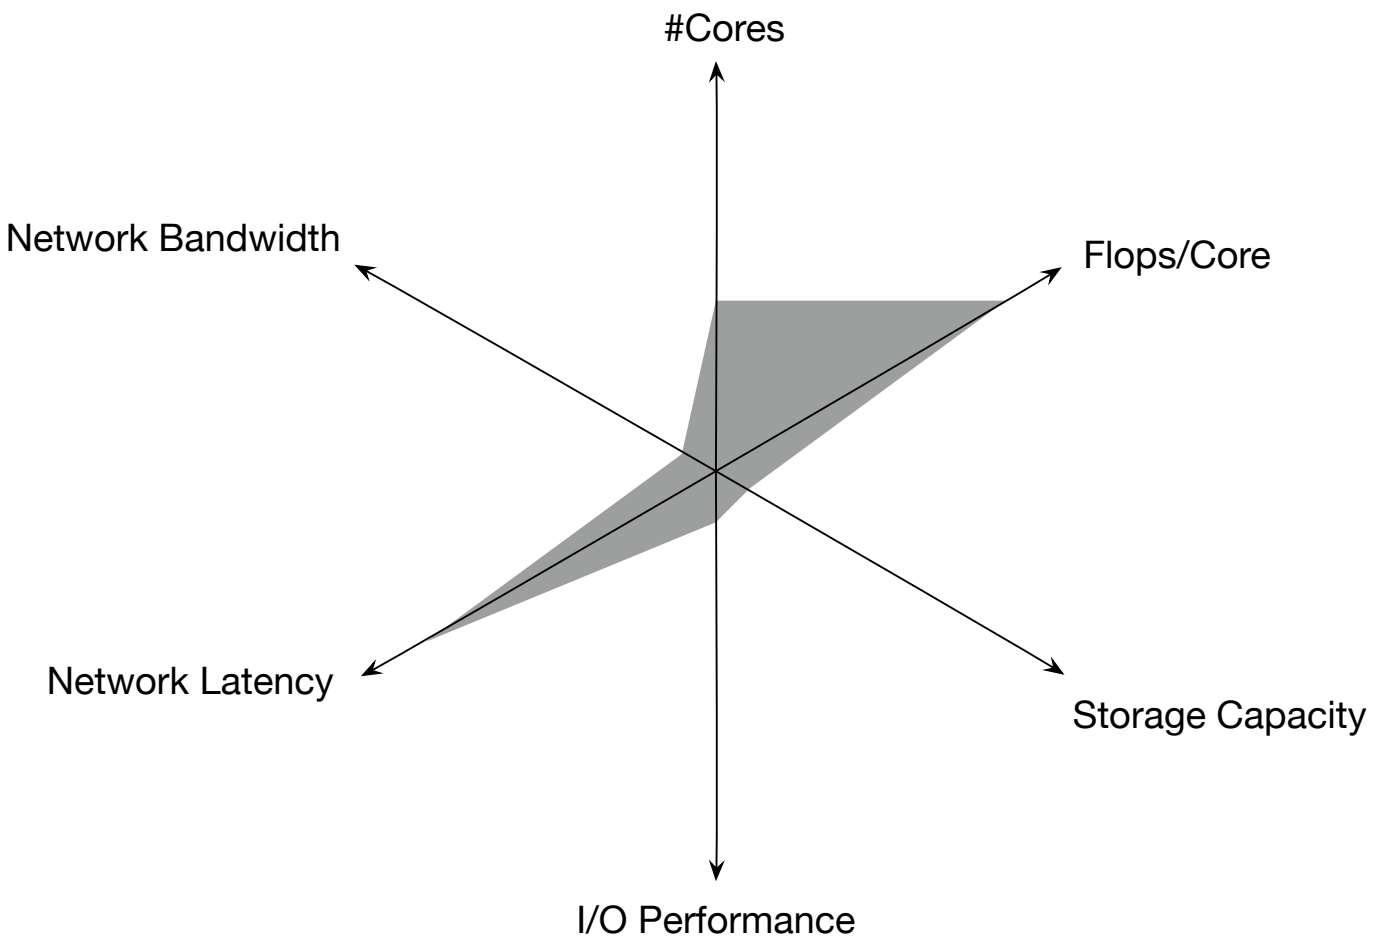

Different Needs for Different Domains

ALL Research Computing Domains

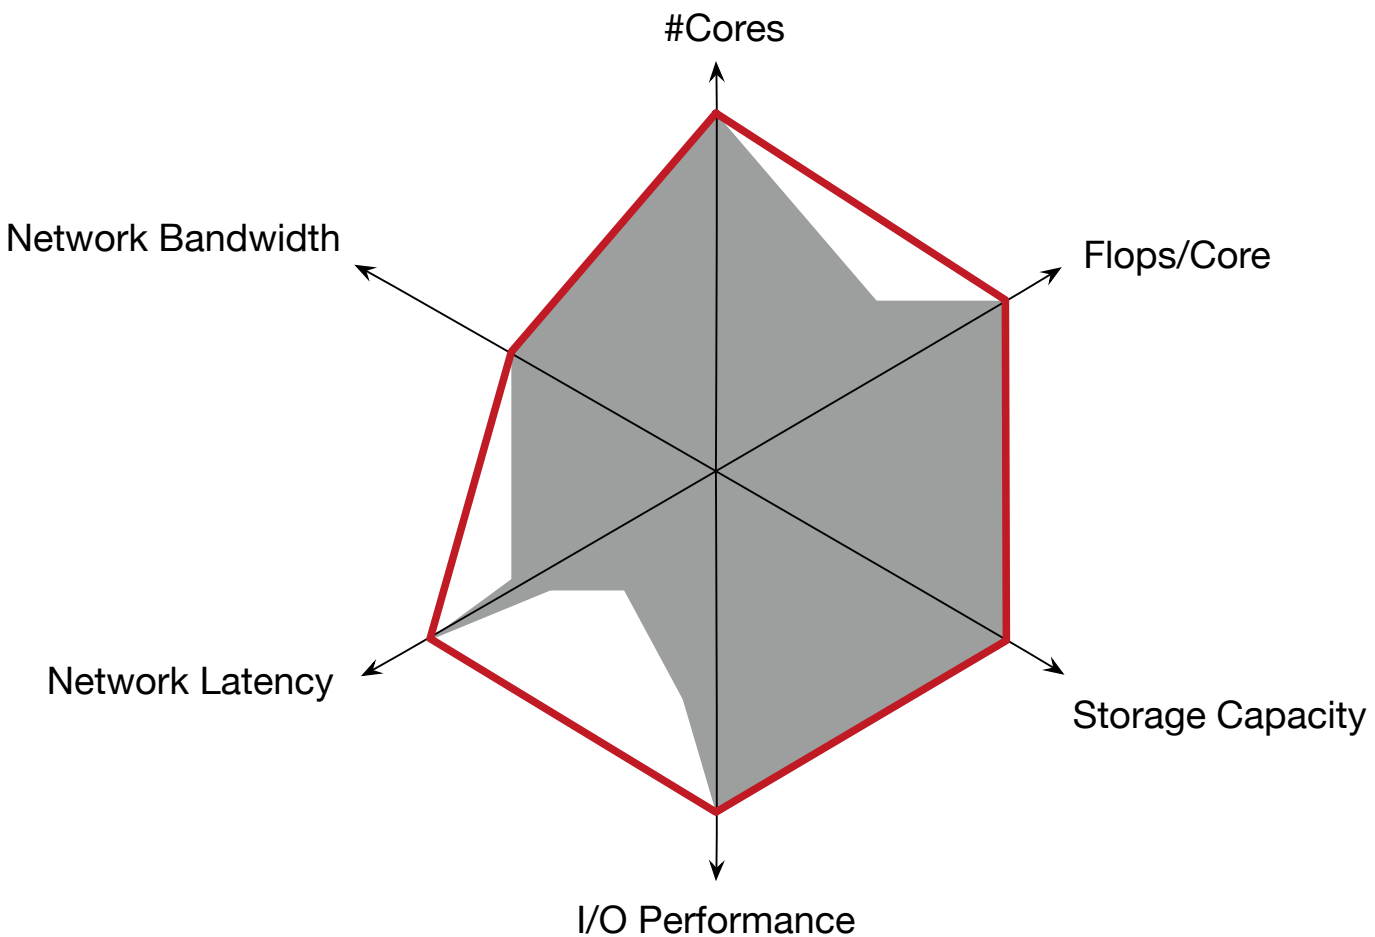

New Trends in HPC

- **Continued scaling** of scientific, industrial & financial applications
  - ↪ . . . well beyond Exascale
- New trends changing the landscape for HPC
  - ↪ Emergence of **Big Data analytics**
  - ↪ Emergence of (**Hyperscale**) **Cloud Computing**
  - ↪ **Data intensive Internet of Things (IoT)** applications
  - ↪ **Deep learning & cognitive computing** paradigms

[Source : EuroLab-4-HPC]

[Source : IDC RIKEN report, 2016]

Toward Modular Computing

- Aiming at **scalable**, **flexible HPC infrastructures**
  - ↪ Primary processing on CPUs and accelerators
    - ✓ HPC & Extreme Scale Booster modules
  - ↪ Specialized modules for:
    - ✓ HTC & I/O intensive workloads; Data Analytics and AI

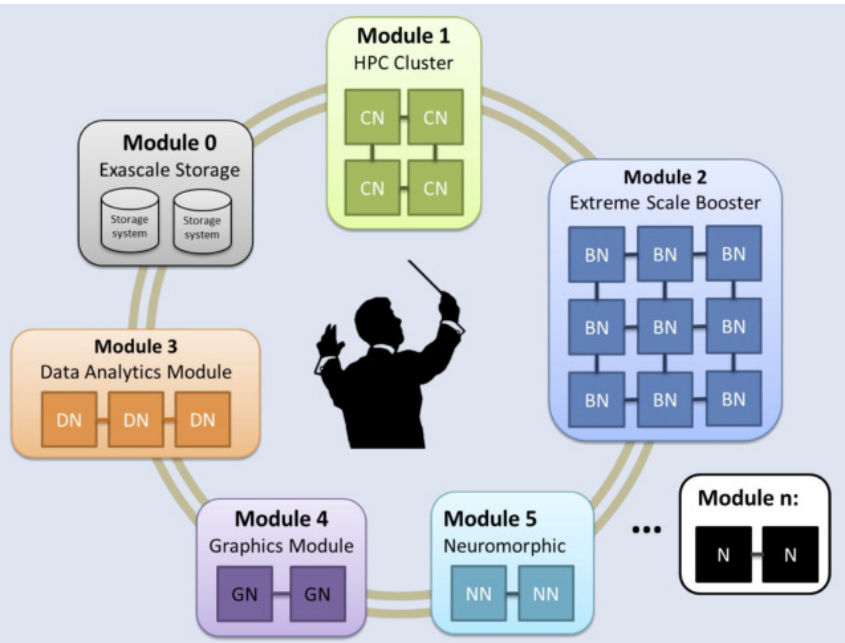

[Source : "Towards Modular Supercomputing: The DEEP and DEEP-ER projects", 2016]

Implementation Frameworks in Europe

- **European Technology Platform (ETP) for HPC**
  - ↪ Industry-led forum
  - ↪ founded by stakeholders of HPC technology
  - ↪ 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 COST Actions
  - ↪ Ex: NESUS: Network for Sustainable Ultrascale Computing

Implementation Frameworks in Europe

- **European Technology Platform (ETP) for HPC**
  - ↪ Industry-led forum
  - ↪ founded by stakeholders of HPC technology
  - ↪ 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 COST Actions
  - ↪ Ex: NESUS: Network for Sustainable Ultrascale Computing
- **PRACE** – Partnership for Advanced Computing in Europe
  - ↪ Non-profit association, 25 member countries
  - ↪ Providing access to **Five EU Tier-0** compute & data resources

Implementation Frameworks in Europe

- **European Technology Platform (ETP) for HPC**
  - ↪ Industry-led forum
  - ↪ founded by stakeholders of HPC technology
  - ↪ 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 COST Actions
  - ↪ Ex: NESUS: Network for Sustainable Ultrascale Computing
- **PRACE** – Partnership for Advanced Computing in Europe
  - ↪ Non-profit association, 25 member countries
  - ↪ Providing access to **Five EU Tier-0** compute & data resources

- Luxembourg officially entered PRACE on **Oct. 17th, 2017**
  - ↪ Official Delegate/Advisor (P. Bouvry/S. Varrette) from UL

Summary

Conclusion & Perspective

- **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 *(S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot)*
  - ↪ UL HPC (as of 2017): **206.772 TFlops, 7952.4TB (shared)**
  - ↪ consolidate and extend Europe efforts on HPC/Big Data

Conclusion & Perspective

- **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 *(S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot)*
  - ↪ UL HPC (as of 2017): **206.772 TFlops, 7952.4TB (shared)**
  - ↪ consolidate and extend Europe efforts on HPC/Big Data

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
  - ↪ EuroHPC / IPCEI on HPC and Big Data (BD) Applications

Conclusion & Perspective

- **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 *(S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot)*
  - ↪ UL HPC (as of 2017): **206.772 TFlops, 7952.4TB (shared)**
  - ↪ consolidate and extend Europe efforts on HPC/Big Data

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
  - ↪ EuroHPC / IPCEI on HPC and Big Data (BD) Applications

⇒ Still a long way to go to fill the gap with US,China,Japan...

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