



High Performance
Computing &
Big Data Services



hpc.uni.lu



hpc@uni.lu



[@ULHPC](https://twitter.com/ULHPC)

HPC & BD Services @ Uni.lu

Building up High Performance Computing & Big Data Competence Center to support national priorities

Dr. Sébastien Varrette

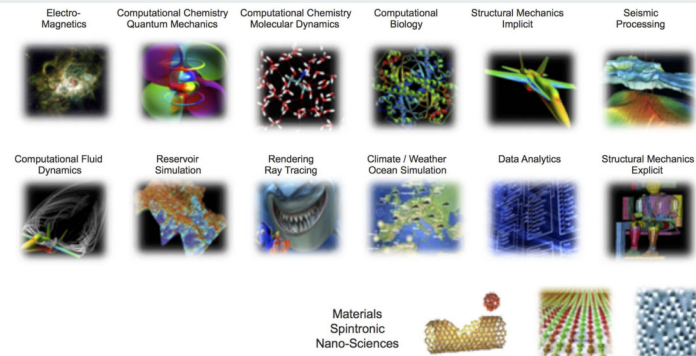
*International Supercomputing Conference (ISC'18)
2nd Workshop on HPC Collaboration Between Europe and Latin America*

Thursday, June 28th, 2018



Why HPC & BD ?

- **HPC: High Performance Computing**
 - **BD: Big Data**
 - **Essential tool for Science, Society and Industry**
-
- All scientific disciplines are becoming **computational** today
 - modern scientific discovery requires very high computing power and handles huge volumes of data
 - cf. J. Rifkin report: “[3rd Industrial Revolution Strategy for the Grand Duchy of Luxembourg](#)”
 - **Industry and SMEs** are increasingly relying on the power of supercomputers...
 - ... to invent innovative solutions while reducing cost and decreasing time to market
 - HPC is part of a **global race** (recognized as a strategic priority) - [EU is taking up the challenge](#)
 - Ambitious plans from many countries (USA, China, Japan, **Brazil & South america**, India...) around HPC



European HPC strategy and its implementation

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

- Latest advances:

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

- PRACE (Partnership for Advanced Computing in Europe) transitioning to PRACE2.
 - Luxembourg: 25th country to join in Oct. 17th, 2017
 - Official Delegate/Advisor (P. Bouvry/S. Varrette) from Uni.lu

◦ EU Member States sign the EuroHPC initiative and prepare its implementation (Mar. 2017)

- A common effort to create and grow the European supercomputing ecosystem
- Federation of national and regional HPC centers (see also PRACE/PRACE2)
 - Funding next-generation Peta-scale / Pre-exascale / Exascale systems
- EuroHPC Joint Undertaking (JU) effective starting Jan. 2019

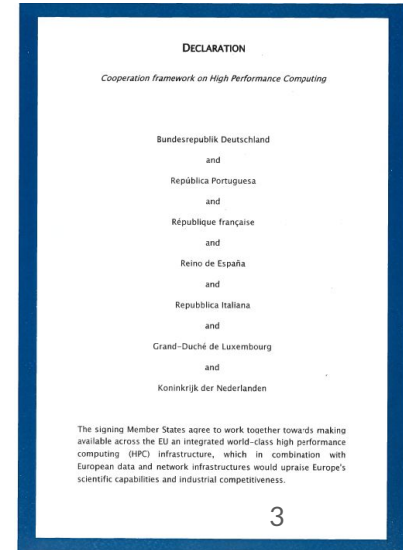
HPC - High on the European agenda



European Commission President
Jean-Claude Juncker

"Our ambition is for Europe to become one of the top 3 world leaders in high-performance computing by 2020"

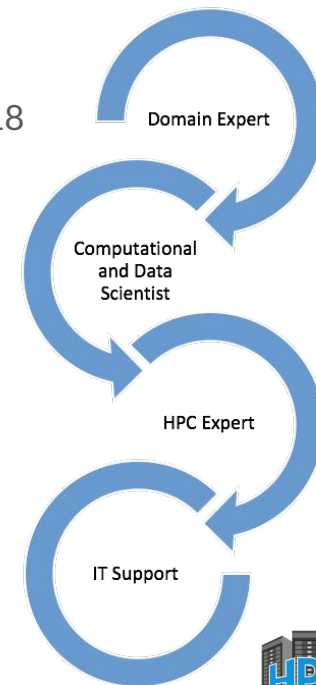
Paris, 27 October 2015



University of Luxembourg as a Jewel



- **Created in 2003, moved to Belval in 2015**
 - An international university serving its country
- **Ranked 12** among the young universities for the Times Higher Education Ranking in 2018
- *With regards to HPC, University of Luxembourg offers:*
 - **People**
 - Domain experts
 - Computational and data scientists
 - Specialists in parallel algorithms
 - **Services**
 - HPC clusters and management team
 - IT team (IT department)
 - Infrastructure team (IT department) in collaboration with Fonds Belval
 - **Infrastructure**
 - Data center and a set high-end clusters
 - **Education & Training**





Special Study

Analysis of the Characteristics and Development Trends of the Next-Generation of Supercomputers in Foreign Countries

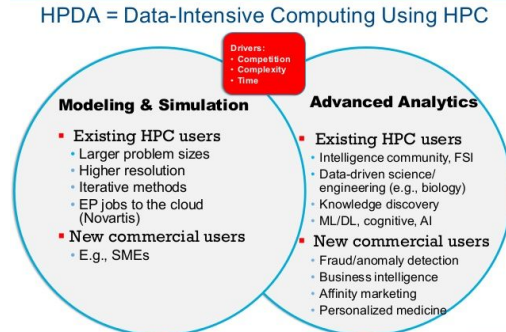
Earl C. Joseph, Ph.D.
Steve Conway

Robert Sorensen
Kevin Monroe

Source : [IDC report: RIKEN](#) (2016)


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



©Hyperion Research 2017


15



Source : 2017 Mellanox [Corporate Deck](#)

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
- 88%* of stakeholders will have multiple architectures
- Creates new adopters, targets also SME market



Source: "Towards Modular Supercomputing: The DEEP and DEEP-ER take on Heterogeneous Cluster Architectures", Norbert Eicker, SC'16 (Nov. 2016)

**Next-generation HPC-BD platforms expected to increase modularity
... and thus flexibility**

HPC Facility @ Uni.lu

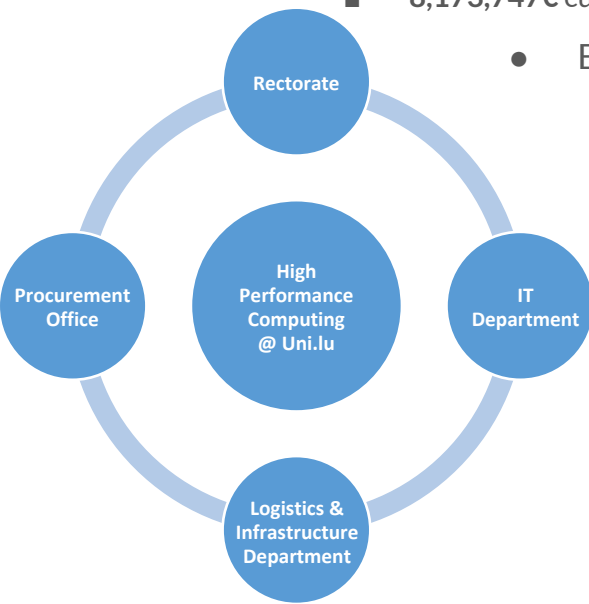
- Started in 2007, under responsibility of Prof. P. Bouvry & Dr. S. Varrette

- expert UL HPC team

- S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot, among others
 - 8,173,747€ cumulative investment hardware (excl. Server rooms)

- Enables & accelerates scientific discovery & innovation

- Largest HPC facility in Luxembourg w. GoodYear



HPC/Computing Capacity


423 TFlops → 1035,8 TFlops by eoy

HTC/Storage Capacity

9,8 PB shared




High Performance
Computing &
Big Data Services



HPC Facility @ Uni.lu / Computing



- 5 clusters, 2 sites
- **Total computing capacity: 423 TFlops**
 - 10,130 CPU cores + 120,704 GPU cores
 - Planned extension to **1035,8 TFlops** by Q4 2018 (RFP 180027)
 - Fast interconnect based on Infiniband
 - ... typically over a non-blocking Fat-tree network topology




HPC Facility @ Uni.lu / Storage



Total shared storage capacity: 9.8 PB


- 3.24 PB on GPFS/SpectrumScale
- 1.94 PB on Lustre
- 4.67 PB on other FS (OneFS...)
 - incl. backups



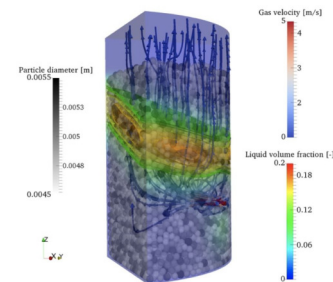
Lustre IOR performance on the iris cluster, 2018

Case Study 1: Material Science & Engineering



- **Companies & Research centers in Luxembourg**
 - GoodYear, IEE, Delphi, ArcelorMittal, ProNewTech, Prosciens, CrmClouder.com, UL, LIST...
- **Application domains**
 - Physics and Chemistry (materials design, new insights), Finite Element Analysis (FEA), Computational Fluid Dynamics (CFD), Optimization, Visualisation...
- **Computing infrastructure answering these needs**
 - Traditional (CPU only) or Hybrid (CPU + Accelerators)




Snow-tire interaction, [LuxDem](#) (UL)



Hydrodynamic study within counter-current packed bed reactor, [LuxDem](#) (UL)



Molecular crystal structure prediction with FHI-aims, PMSRU (UL)

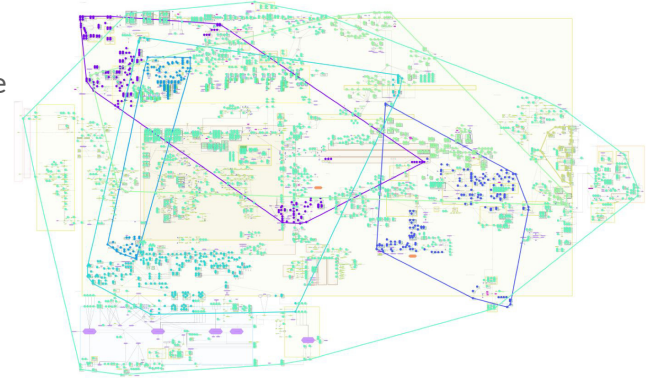



Case Study 2: Biomedical Industry




Source: Illumina


- **Companies & Research centers in Luxembourg**
 - ITTM, IBBL, UL, LIH...
- **Potential to attract external companies**
 - Edico Genome (US), Fabric Genomics (US), Swarm64 (DE)...
- **Application domains**
 - System Bio-medicine, BD Analytics, Pharmacology, Personalised Medicine
- **Computing infrastructure answering these needs**
 - Traditional (CPU only), HTC, Hybrid (CPU + Accelerators)
 - [High Performance] Data Storage



Visualization of biological knowledge: Bi-level optimization/ Parkinson disease map. CSC/LCSB (UL)



Inverse Protein Folding
CSC/LCSB (UL)



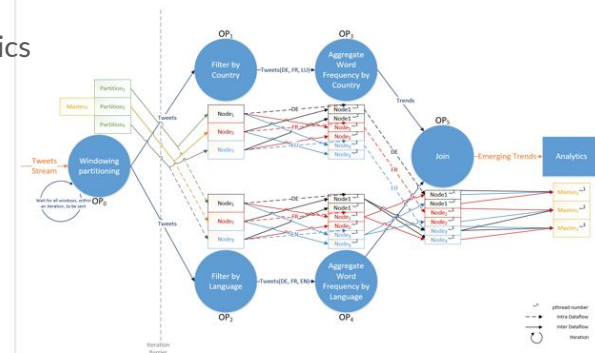
Source: LCSB (UL)

Case Study 3: Deep Learning - Cognitive Computing

- **Companies & Research centers in Luxembourg**
 - Churchill Frank, SES, Aiva Technologies, IEE, UL...
- **Potential to attract external companies**
 - Amazon (US), Google (US), Uber (US), Tesla (US), Deepsense.AI (US)...
- **Application domains**
 - Data Mining, Self-Driving cars, Satellite & Communications, Big Data Analytics
- **Computing infrastructure answering these needs**
 - Traditional (CPU only), HTC, Hybrid HPC (CPUs + Accelerators)



Source: "Luxembourg strikes deal to create 'driverless car' test zone", Luxembourg Wort, 2017




Real-Time Big-Data Analytics: Emerging trends
Twitter streams, UL

Case Study 4: Data science, IoT and FinTech


- Companies & Research centers in Luxembourg
 - Big Four (E&Y, Deloitte, PwC, KPMG), ExaMotive, AXA, BIL, BCEE, UL...
- Potential to attract external companies
 - Amazon (US), NEST (US), Somfy (FR), Google (US)...
- Application domains
 - Risk & Asset Management, Data intensive IoT, Smart City
- Computing infrastructure answering these needs
 - Cloud Computing, HTC, HPC



A Smart Day in Luxembourg, Luxinnovation




FinTech ecosystem, from 2016 PwC report
"How FinTech is shaping Financial Services in
Luxembourg"




CACOC+CA with MPC+CA
Chaos-enhanced mobility models for
multilevel swarms of UAVs (PCOG)


HPC Facility @ Uni.lu usage




Statistics extracted on HPC active users
(submitted at least one job in the last year)




HPC Facility @ Uni.lu usage




Statistics extracted on HPC active users
(submitted at least one job in the last year)



HPC Facility @ Uni.lu usage



Statistics extracted on HPC active users
(submitted at least one job in the last year)




Uni.lu CDC (Centre de Calcul):

Toward Energy-Efficient HPC enabling DLC

- 2x500 m² deployed since 2015
 - Electric energy , ventilation and chilled water produced by Fonds Belval power plant (located basement MNO)
 - CDC managed by SIU, one *floor dedicated to HPC* developments
- 4 server rooms **ready and operational** in 2018
 - One room (CDC S-02-005) holds our (current) flagship HPC facility (*iris*)

Data center ready for DLC with a dedicated high temperature network for energy efficiency

| Location | Cooling | Usage | Max Capacity [kW] |
|--------------|------------|--|-------------------------------------|
| CDC S-02-001 | Airflow | <i>Future extension</i> | 280 kW (120 m ²) |
| CDC S-02-002 | Airflow | Storage / Traditional HPC /Cloud/FPGA | 280 kW (88 m ²) |
| CDC S-02-003 | DLC | High Density/Energy efficient HPC | 1050 kW (90 m ²) |
| CDC S-02-004 | DLC | High Density/Energy efficient HPC | 1050 kW (92 m ²) |
| CDC S-02-005 | Airflow | Storage / Traditional HPC (<i>iris</i> cluster) | 300 kW (128 m ²) |



Education & Training

- **COST NESUS Winter School on Data Science and Heterogeneous Computing**
 - [Last edition](#) in January 2018 in Zagreb, Croatia
- **PRACE**
 - [Last edition](#) PRACE Days in May 2018 : European HPC Summit Week 2018 in Ljubljana
 - H2020-INFRAEDI-2018-2020 - PRACE-6IP
- **Bachelor degree BICS/BINFO: AI & Middleware**
- **Master degree MICS: Parallel and Grid Computing**
 - **Master in Information and Computer Sciences** - see misc.uni.lu
 - Master 2 courses, lectured since 2008
 - Evaluation based on projects evaluated on the UL HPC platform. Ex:
 - Fault Tolerance and Performance of MPI toolchains over HPC Applications
 - Measurement and Optimization of Energy Consumption in HPC Applications
 - Parallelization of the Barnes-Hut N-Body Simulation Algorithm
 - Evaluation of Big Data Framework (Hadoop, Spark...)
 - Evaluation of [Parallel/Distributed] Deep/Machine Learning Framework (Pytorch, Tensorflow, caffe...)
- **Lifelong-learning Smart-ICT for Business Innovation certificate: Cloud Computing/Big Data/IoT**





Education & Training



- **SC-Camp 2017 (Cadiz)**
- **Bi-annual HPC School @ Uni.lu** (Part of the doctoral program)
 - 7th edition June 12th - 13th, 2018 in Belval
 - Material publically available: ulhpc-tutorials.rtf.dio
 - Last edition joint event: **UL-NSTDA workshop (EU-ASEAN E-READI)**



| Day 1 | Main Track (MSA 4.520) | Speaker |
|-------|--|-------------|
| 9h00 | PS1a: Getting Started (part I: SSH) | C. Pariset |
| 10h00 | Coffee break | |
| 10h30 | Overview & Challenges of the UL HPC Facility at the EuroHPC Horizon | S. Varrette |
| 11h45 | PS1b: Getting Started (part II) | H. Cartiaux |
| 12h30 | LUNCH | |
| 13h30 | PS2: HPC workflow with sequential jobs | H. Cartiaux |
| 14h30 | PS4a: UL HPC Monitoring in practice | H. Cartiaux |
| 15h30 | Coffee break | |
| 16h00 | PS5: Parallel computations with OpenMP/MPI | S. Varrette |
| 17h30 | PS6: User environment and storage data management | S. Peter |

| Day 2 | Main Track (MSA 4.520) | Speaker |
|-------|--|-------------|
| 9h00 | PS7: Multi-Physics workflows: test cases on CFD/MD/Chemistry applications | V. Plugaru |
| 10h30 | Coffee break | |
| 11h00 | Users' session: UL HPC experiences | |
| 12h40 | LUNCH | |
| 13h30 | PS9: [Basic + Advanced] Prototyping with Python | C. Pariset |
| 15h30 | Coffee break | |
| 15h45 | PS11: Big Data Applications | S. Varrette |
| 17h15 | PS13: Machine / Deep learning (Pytorch, Tensorflow, Caffe2) | S. Varrette |
| 18h15 | Closing Keynote: Take Away Messages | S. Varrette |

| Day 1 | Parallel Track (MSA 4.410) | Speaker |
|-------|--|------------|
| 10h30 | Keynote in 4.520 | |
| 13h30 | PS3a: Advanced scheduling (SLURM, OAR) | V. Plugaru |
| 14h30 | PS3b: Software environment management using Easybuild | S. Peter |
| 15h30 | Coffee break (in 4.520) | |
| 16h00 | PS4b: Performance engineering - HPC debugging and profiling | V. Plugaru |


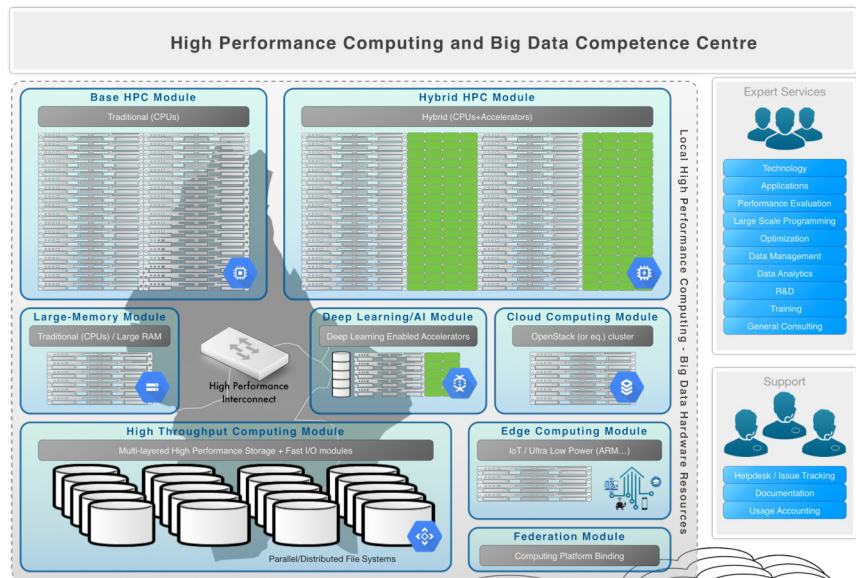
| Day 2 | Parallel Track (MSA 4.510) | Speaker |
|-------|--|-------------|
| 9h00 | PS8: Bio-informatics workflows and applications | S. Peter |
| 10h30 | Coffee break (in 4.520) | |
| 11h00 | Users' session: UL HPC experiences (in 4.520) | |
| 12h40 | LUNCH | |
| 13h30 | PS10: Scientific computing using MATLAB | V. Plugaru |
| 15h30 | Coffee break | |
| 15h45 | PS12: R - statistical computing | A. Ginolhac |
| 17h15 | PS14: HPC Containers: Singularity | V. Plugaru |
| 18h15 | Closing Keynote in 4.520 | S. Varrette |

Next steps: National HPC-BD Competence Center




Built by ministerial, academic, industrial stakeholders

- **Comprehensive centre:**
 - HPC *and* data infrastructure
 - Expertise in technology
 - Domain knowledge in applications
- **More than just computing services:**
 - “Bring HPC and BD to the users”
- **Creates twofold innovation:**
 - Innovation in applications and
 - Innovation in technology



Credits: P. Bouvry, S. Varrette,
V. Plugaru & UL HPC Team




Icons sourced from www.flaticon.com

Appendix



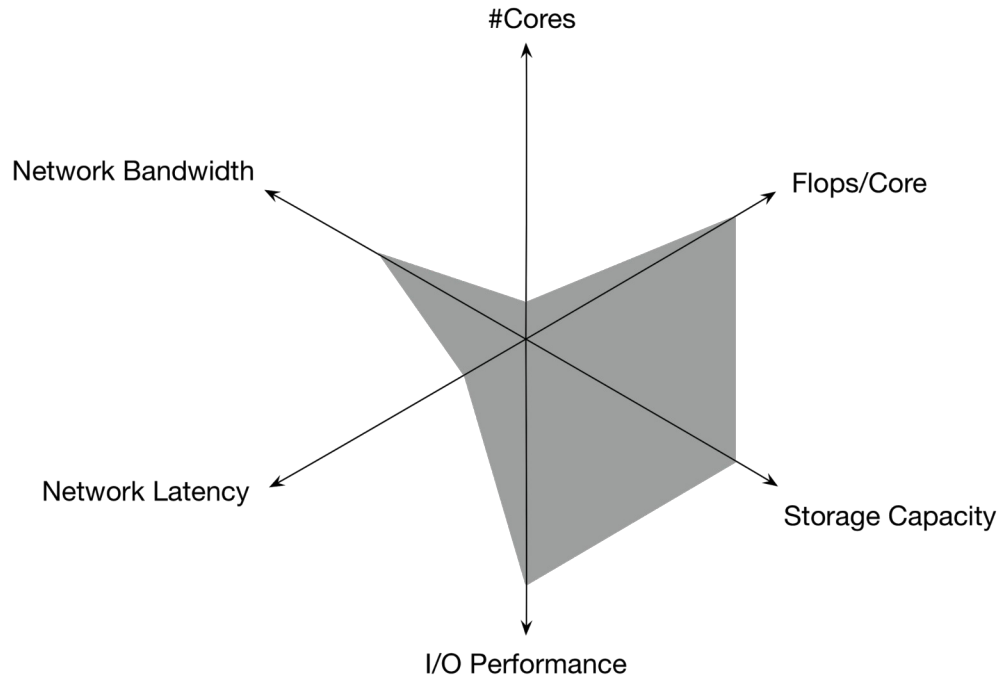
Different Needs for Different Domains

Material Science & Engineering




Different Needs for Different Domains

Biomedical Industry / Life Sciences




Different Needs for Different Domains

Deep Learning / Cognitive Computing

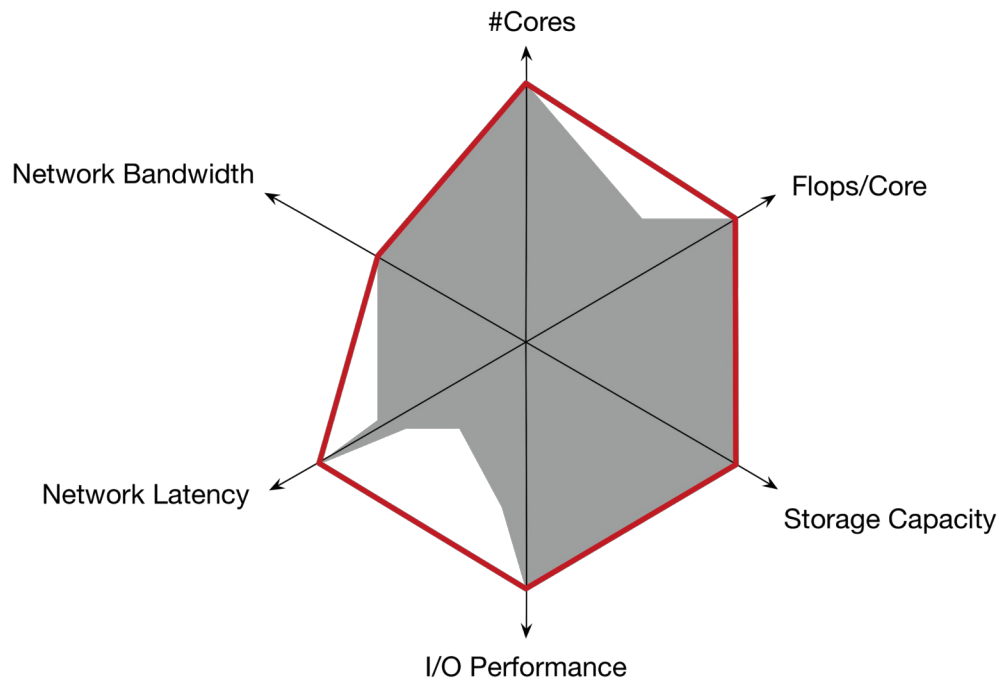


Different Needs for Different Domains



Different Needs for Different Domains

ALL Research Computing Domains



High Performance
Computing &
Big Data Services

