LUXEMBOURG

LET'S MAKE IT HAPPEN

Luxembourg HPC-BD Competence Center

Building up High Performance Computing to support the national commitments to EuroHPC and the Big Data strategy

uni.lu

UNIVERSITÉ DU LUXEMBOURG

Outline

1. Context
2. Modern and Next-Generation Computing at Scale: Current Status and Next Trends
3. Application and Computing Needs in Luxembourg: Case Studies
4. Luxembourg HPC-BD Solution
    Tailoring HPC and BD for Luxembourg
5. Conclusion and Next Steps

Context

UNIVERSITÉ DU LUXEMBOURG

Why HPC ?

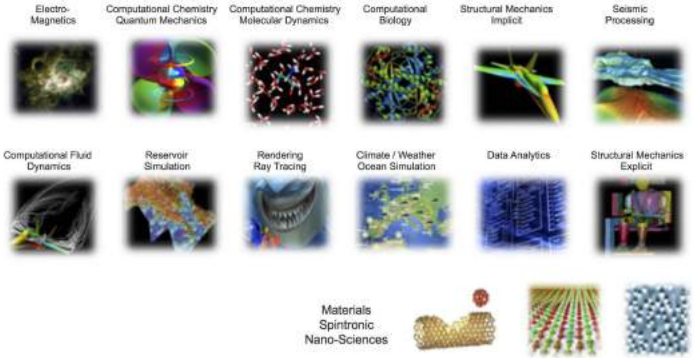

- **HPC**: **H**igh **P**erformance **C**omputing
- Essential tool for **Science**, **Society and Industry**

- 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, Russia, Brazil, India) around HPC

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.

Science & Technology Facilities Council

Prof Tony Hey CBE FREng, STFC Chief Data Scientist & Joint Chair, UK E-infrastructure Leadership Council

Excellent research requires excellent research infrastructure

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...
  - 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)

In this talk: Design and Implementation of the EuroHPC Center of Excellence in Luxembourg

Luxembourg's Smart Specialisation strategy

Luxembourg's strategy to succeed as a small country in a global competition

- **Be different**
  - by opening up **new avenues** (e.g. space mining, personalized medicine)
- **Be specialized**
  - by focusing on **solving problems** within Luxembourg Ecosystem needs
- **Be smart**
  - create favourable **legislative framework**
  - build on worldwide best-practices, **success stories** and experiences

Source: Shutterstock

The same principles should be applied to the HPC- Big Data (BD) strategy in Luxembourg

Potential goals for engagement in HPC - BD

- **Get reputation as leading computing site in Europe**
  - be at the forefront of adapting and pushing the next-generation technologies and hybrid architectures
  - … to solve problems that are not possible by other means
- **Gain leadership in big data computing**
  - keep the data in the country and further establish Luxembourg as data hub
- **Bring effective solutions for all computing and data problems of the country**
  - enable and assist in tackling (big) problems
- **Further improve research and development for national actors**
  - facilitate next-gen R&D in line with national priorities (Digital Luxembourg, Industry 4.0, Space)
- **Attract companies to the country**
  - **supported by all of the above**

Digital Luxembourg

Leading the way towards a smart nation

Path to the HPC-BD Center of Competence

Source: macafeeinstitite.com

- Build on worldwide **best-practices and success stories**
  - on modern and next-generation computing at scale
- Create a **flexible and modular** computing design
  - **offer efficient solutions** to all identified workloads
  - **embed foreseen trends** changing the HPC landscape
- Enable **successful partnership** with the industry
  - **joint projects** between research institutions and enterprises
  - **attract SMEs** by unveiling new opportunities
  - **link best students** with the industry
- Diffuse HPC-BD knowledge and experiences
  - **training events with hands-on activities**
  - **tailored courses on-site,** offer **technology consulting**

Modern and Next-Generation Computing at Scale

Current Status and Next Trends

Computing Hardware

Base

- **CPU** (Central Processing Unit)
  - highest software flexibility
  - high performance across all computational domains

CPU: 2017 Intel Xeon Skylake-SP

GPU accelerator: 2017 Nvidia Tesla V100

Accelerators

- **Intel MIC** (Many Integrated Core) **Accelerator**
  - less software flexibility
  - higher performance for acceleration enabled software
- **ASIC (**Application-Specific Integrated Circuits)
  - least software flexibility
  - highest performance for specialized problems (AI)
- **GPU** (Graphics Processing Unit) **Accelerator**
  - less software flexibility
  - higher performance on acceleration enabled software (physics, machine/deep learning...)
- **FPGA** (Field Programmable Gate Array)
  - least software flexibility
  - highest performance for specialized problems

Future trends:

Increased heterogeneity in hybrid platforms + Deep learning enabled accelerators (GPU, FPGA, ASIC)

Communication

What

- Internal network(s): **key differentiator** between HPC/supercomputers and regular datacenters
  - bandwidth and latency characteristics **much more performant** than in traditional datacenters

Why

- Enabler of **fast processing** and **fast data movement at scale**
  - including internal transfers of large datasets in hybrid architectures
  - essential for **streaming workloads** (e.g. real-time analytics)

How

- Different **topologies** for different needs
  - **design is critical** for accelerating particular workloads
  - **multiple networks** that separate traffic flows within a unified system

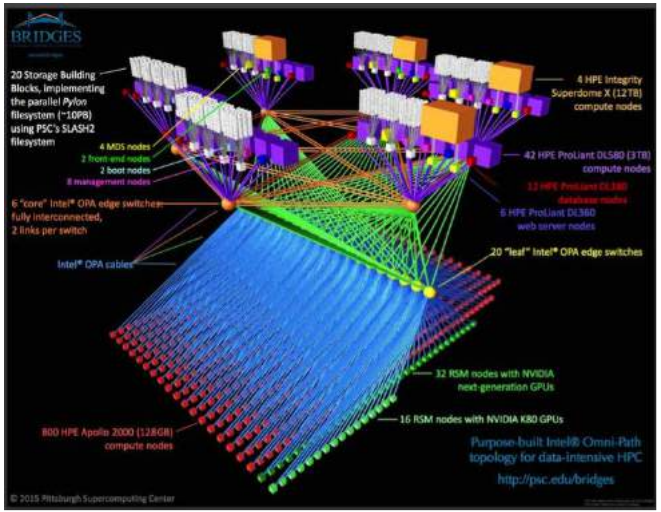

Omni-Path interconnect at Pittsburgh SC

Storage

1.44PB SpectrumScale storage at UL

20PB SpectrumScale storage at Jülich SC

- Computing at scale requires **scalable and performant storage**
  - HPC-oriented storage design is very different than in regular datacenters
  - **Necessity to co-locate Data repositories within the HPC facility**

- **Multi-layered storage** is always needed
  - Different performance characteristics for different needs
  - Storage tiers with varying data retention policies

- Highly specific software/hardware elements for large scale
  - **Throughput up to TB/s** over Distributed File Systems
  - Data migration logic; Data safety and security

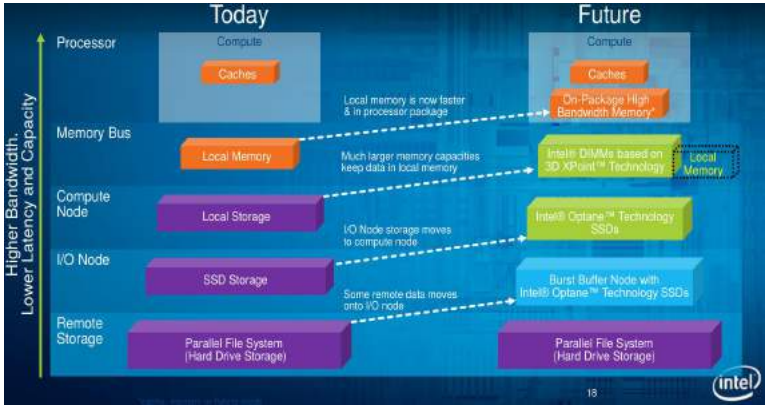

2016 Intel vision for storage

Supercomputing in the World

Tsubame 3, Japan, 2017

Piz Daint, Switzerland, 2016

ECMWF Anemos, UK, 2016

TaihuLight, China, 2016

Stampede 2, US, 2017

MareNostrum 4, Spain, 2017

Titan, US, 2012

Hazel Hen, Germany, 2015

Marconi, Italy, 2017

Different Needs for Different Domains

Material Science & Engineering

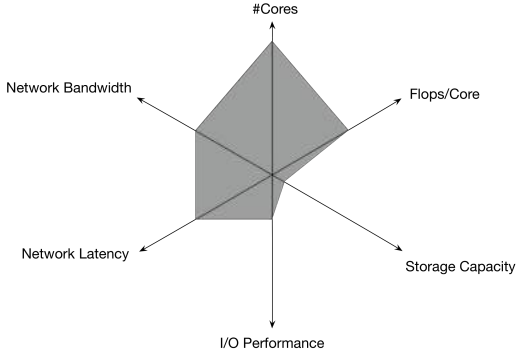

Biomedical Industry / Life Sciences

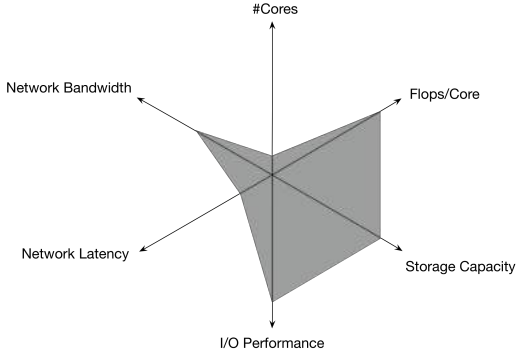

Deep Learning / Cognitive Computing

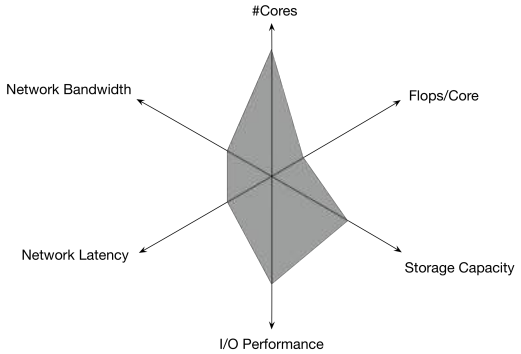

IoT, FinTech

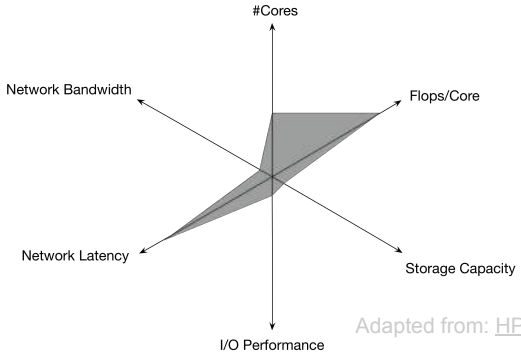

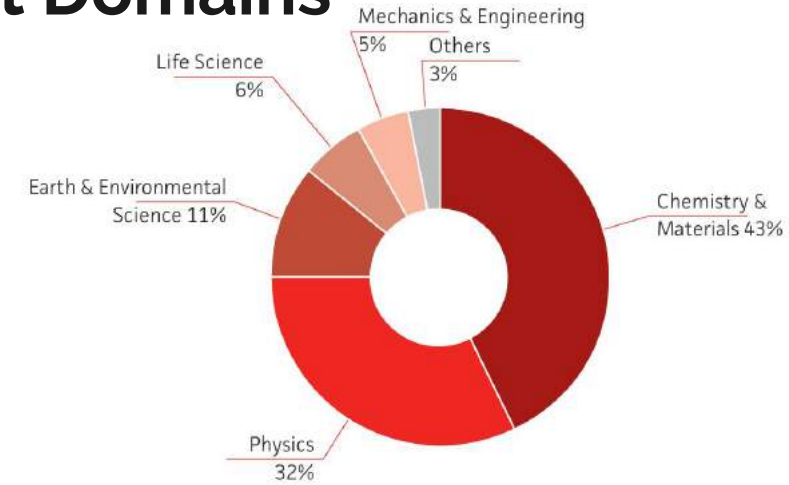

Source: Swiss National Supercomputing Centre (CSCS) - 2016 Activity Report
Note: CSCS Piz Daint - largest supercomputer in Europe, 3rd worldwide (Sept.2017)

Additional Opportunities:
Space, Smart City, IoT, FinTech etc.

Adapted from: HPC and HPDA for the Cognitive Journey with OpenPOWER, Srini Chari, May 2016

New Trends in HPC

Source : IDC report: RIKEN (2016)

- **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

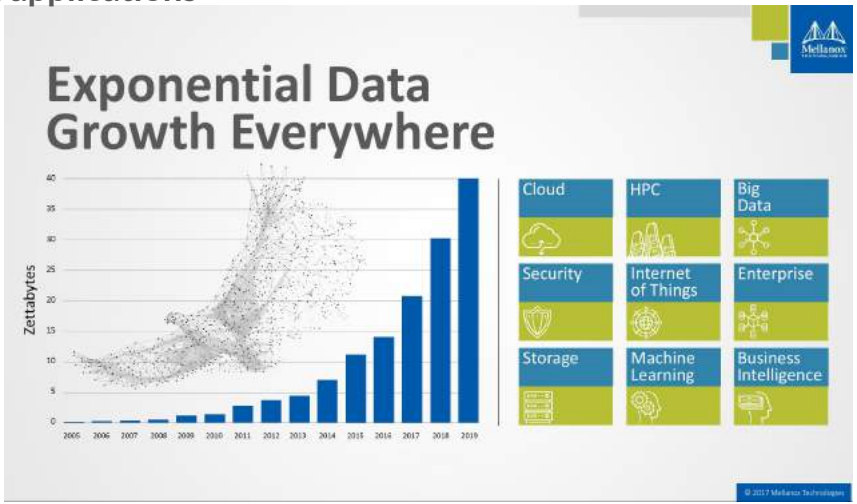

Source : 2017 Mellanox Corporate Deck

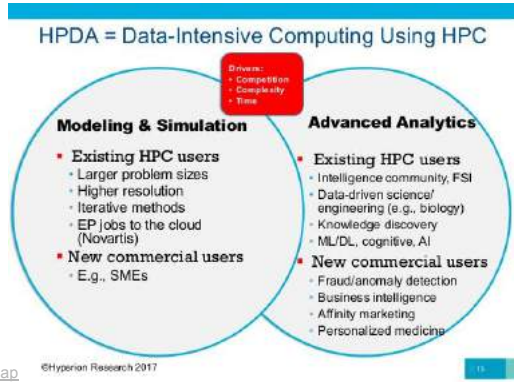

Source: EU Eurolab-4-HPC Long-Term roadmap (to 2030) for High Performance Computing (2017)

These trends should be reflected in the design of the HPC-BD solution for Luxembourg

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

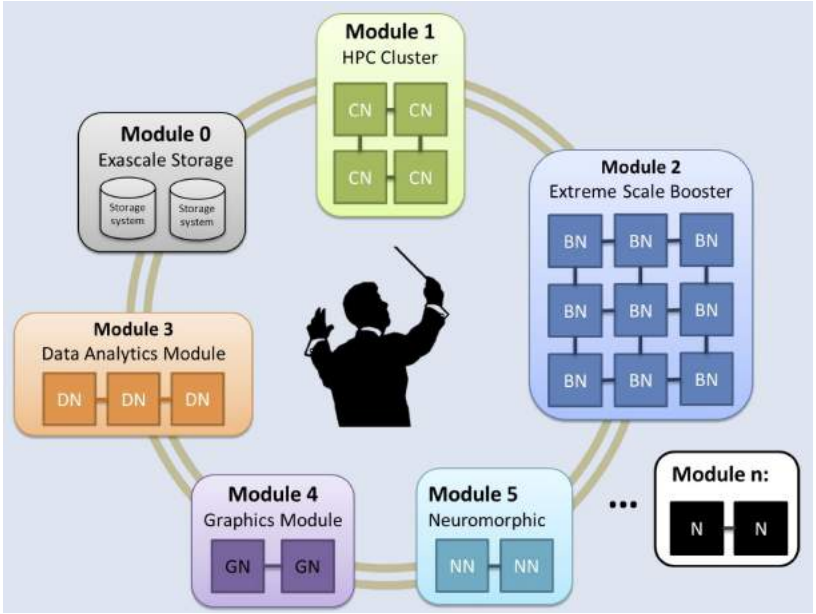

Source: "Towards Modular Supercomputing: The DEEP and DEEP-ER take on Heterogeneous Cluster Architectures", Norbert Eicker, SC'16 (Nov. 2016)

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

* Source: Addison Snell, Intersect360

Application and Computing Needs in Luxembourg

Case Studies

UNIVERSITÉ DU LUXEMBOURG

Main Categories of Computing Needs

- **Traditional High Performance Computing (HPC)**
  - High compute power with modest data in/out (PetaByte range)
    - e.g. numerical simulation, engineering, artificial intelligence, deep learning, FinTech
  - *Lightweight* virtual containers (enabling HPC interconnect) supporting rapid application deployment
- **High throughput computing (HTC)**
  - Fast input / output of large data sets (10s of PetaBytes) with high compute demands *e.g.* genome research
- **Cloud Computing (CC)**
  - Web services executing pre-configured applications (user brings data)
  - On demand virtual instances (VMs), low level customization of environment
- **Interactive and Cognitive Computing**
  - Real-time computing of many individual data points
    - e.g. driverless cars (imaging, satellites), Smart City and IoT (sensors)

Future HPC-BD should cover all these fields

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)

Tyre modelisation, GoodYear

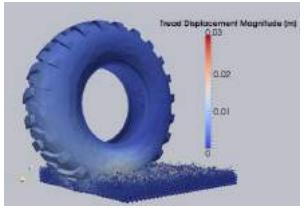

Snow-tyre interaction, LuxDem

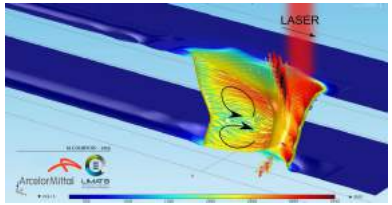

Calculated molten pool shape, temperatures and fluid streamlines, ArcelorMittal

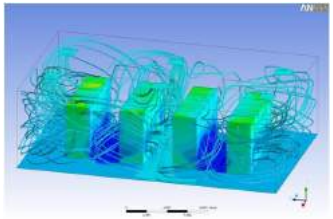

CFD Analysis of Data Center Cooling

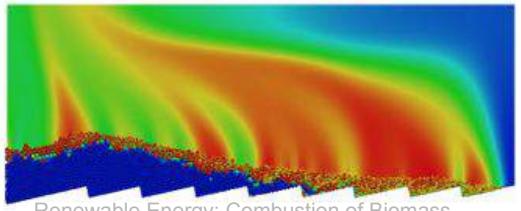

Renewable Energy: Combustion of Biomass

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

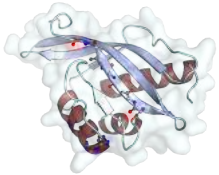

Source: CSC/LCSB (UL)

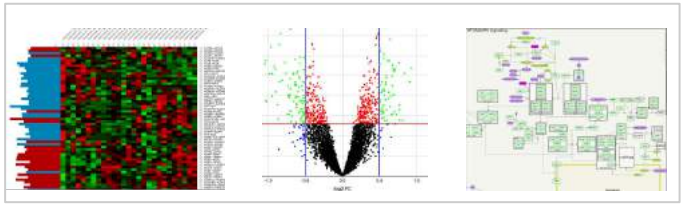

Source: LCSB (UL)

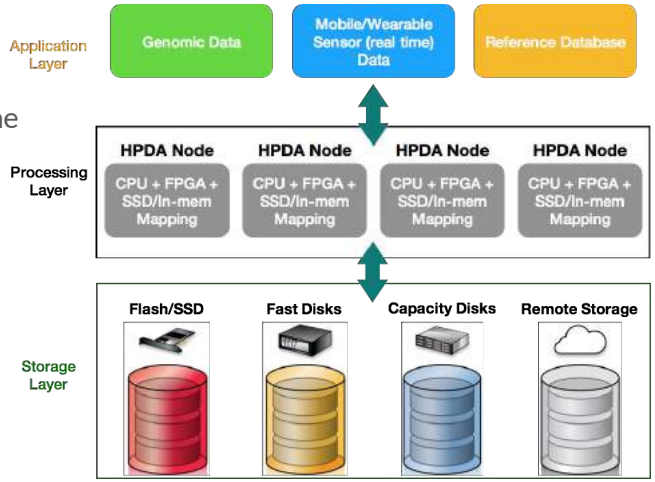

Case Study 3: Deep Learning - Cognitive Computing

- **Companies & Research centers in Luxembourg**
  - Churchill Frank, SES, Aiva Technologies, 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

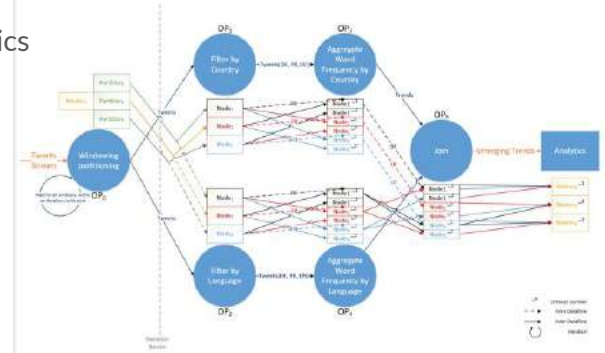

Real-Time Big-Data Analytics: Emerging trends in Twitter streams, UL

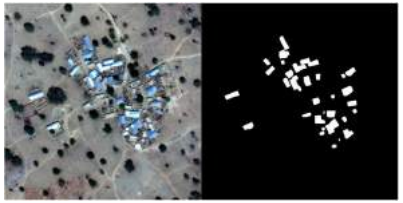

Satellite imagery via image segmentation from Deep Learning, Deepsense.ai, 2017

Space Mining, http://www.spaceresources.public.lu/

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

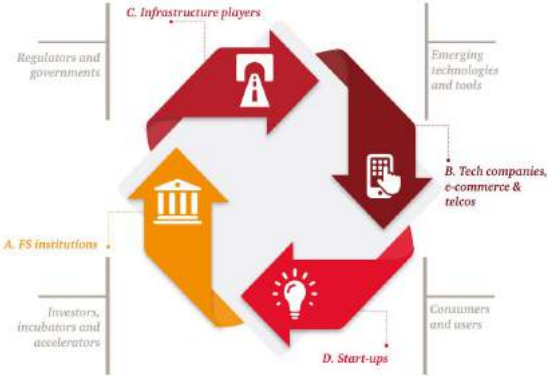

FinTech ecosystem, from 2016 PwC report "How FinTech is shaping Financial Services in Luxembourg"

Luxembourg HPC-BD Solution

Tailoring High Performance Computing and Big Data for Luxembourg

UNIVERSITÉ DU LUXEMBOURG

How to be Smart in HPC

- **Start from your goal**
  - What do you want to achieve?
  - Who are your users?
  - What are your applications and domains?

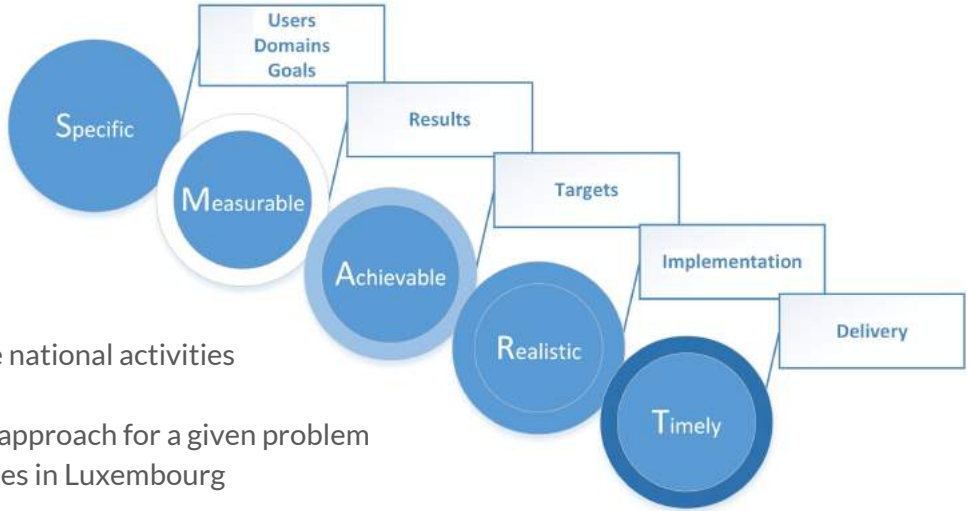

- **Target your customer**
  - Create a modular infrastructure targeted at the national activities
  - Bring HPC to the companies
  - Provide knowledge to find the right computing approach for a given problem
  - Offer computing and data expertise to companies in Luxembourg

“Wirtschaftsförderung durch Kompetenzförderung”

- Attract companies to Luxembourg by attractive computing infrastructure and support in its use

Design of a National Platform

Source: Getty Images/iStockphoto

- **Capacity**
  - To deliver **accelerated time-to-solution**
  - While sustaining a **multitude of clients**
  - ***Through*** PetaScale computing, PetaByte storage
- **Capability**
  - To drive current and **next-gen processing**
  - **Flexible designs** for many work patterns
  - ***Through*** CPU + Accelerators (GPU, FPGA), Tiered storage (with burst-buffers)
- **Efficiency & Expertise**
  - Need to leverage capacity & capability in an optimal way
  - Consolidate the pool of HPC-BD experts in Luxembourg, attract international talent
  - ***Through*** cutting-edge infrastructure from top international vendors & the best people

The Appropriate Architecture for the Appropriate Workload

Luxembourg's HPC-BD Strategy

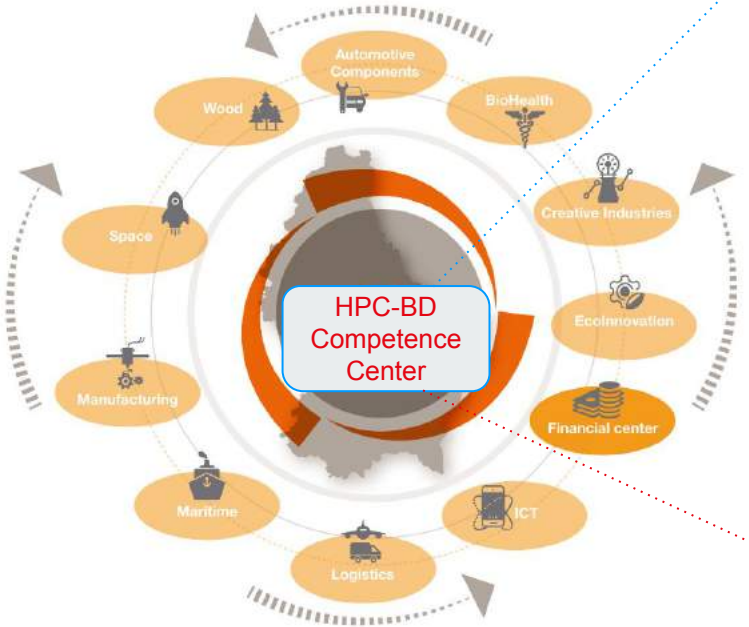

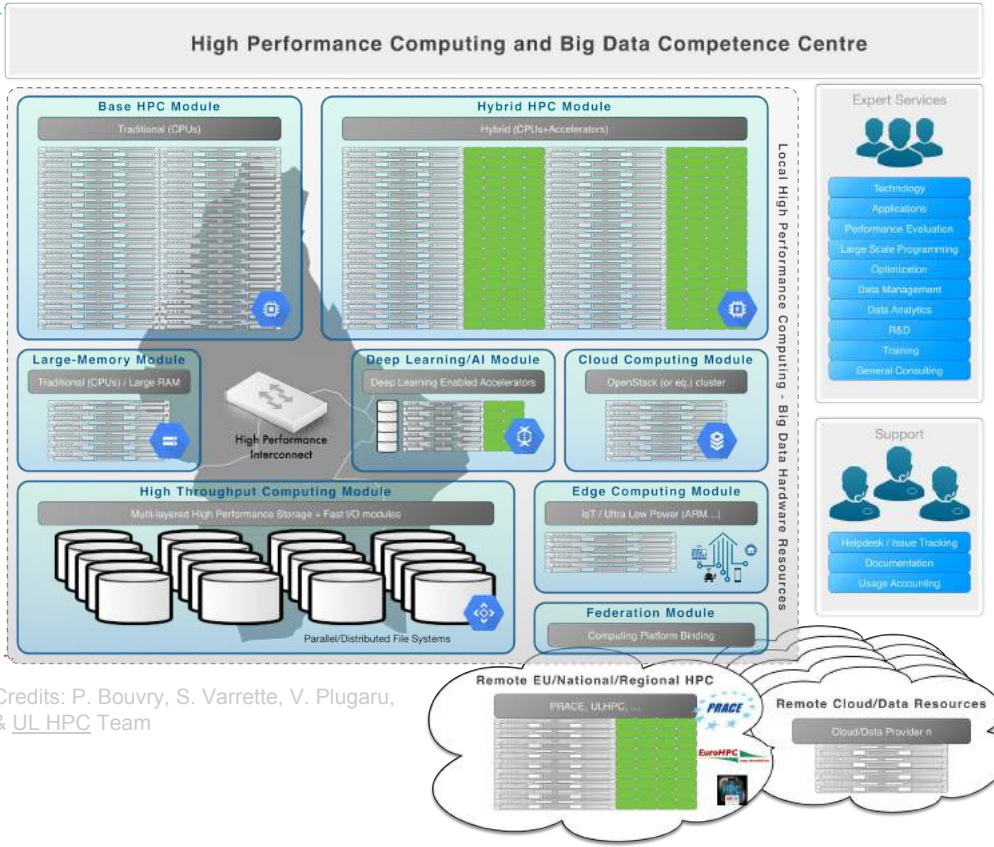

Credits: P. Bouvry, S. Varrette, V. Plugaru, & UL HPC Team

Source: Synergies between the Luxembourg Clusters
2017 PwC study "Inside the mind of the CEO - A Luxembourg perspective"

The Competence/Excellence Center

- Implementation and deployment of national HPC in Luxembourg
- Built by ministerial, academic, industrial stakeholders
- Comprehensive centre:
  - High Performance Computing 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

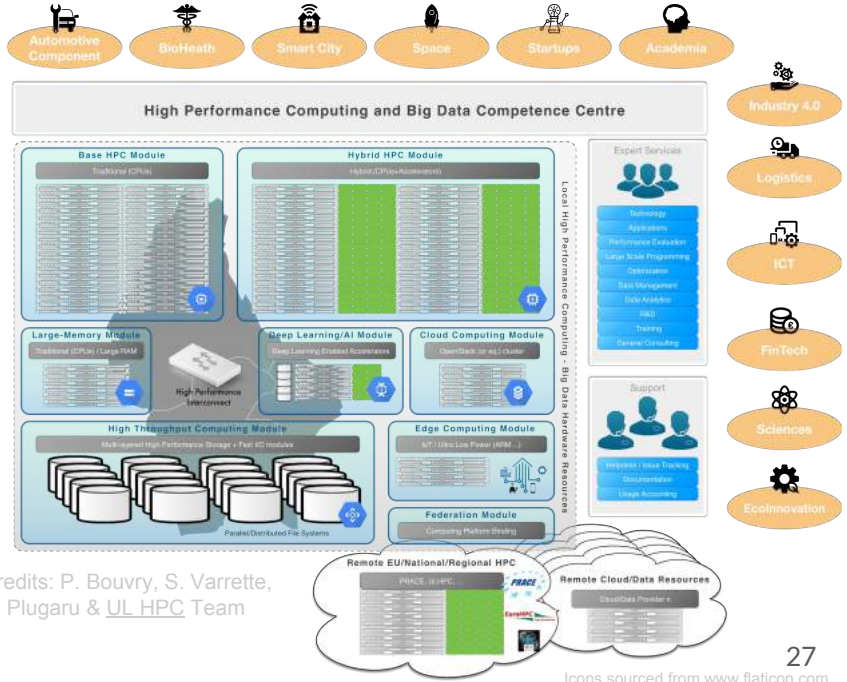

Credits: P. Bouvry, S. Varrette, V. Plugaru & UL HPC Team

Icons sourced from www.flaticon.com

Enabling Successful Partnership with Industry

- **Adjust mission statements** to add the focus on industry engagement
  - Choose HPC systems with industry as well as science in mind
- **Focus on domains the HPC center knows well**
  - Hire a person with business experience to lead the industrial partnership program
- **Assign experienced people to work with industrial users**
- **Involve industrial users in the peer review process**
- **Streamline the process for Intellectual Property and Contract Agreements**

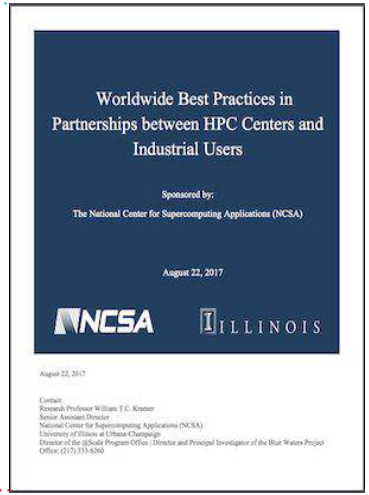

Source: Best Practices for Effective Partnerships between Public HPC Centers and Industrial Users, NCSA, Aug. 2017

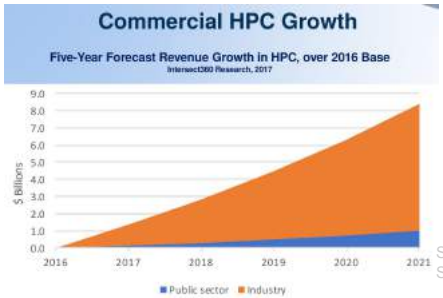

Source: HPC Market Update
Sept. 2017, Intersect360 Research

Growing Luxembourg's HPC & BigData Expertise and Knowledge Base

Governance

- **Initial creation as GIE**
  - Groupement d'Intérêt Économique
- **Creation of additional legal entities**
  - better suited for commercial operations
- **Involvement of stakeholders**
  - Funders
  - Technology and knowledge providers
  - Users

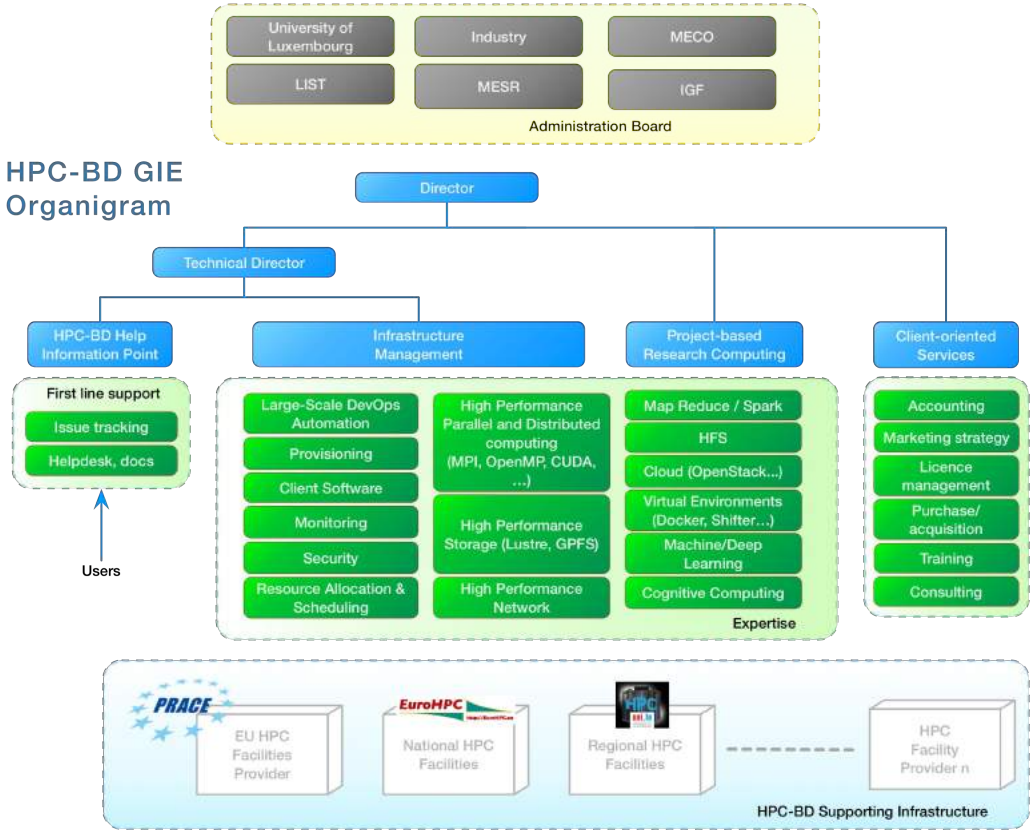

GIE needs Executive Power to decide on Development of the Centre

Success Criteria: Implementation

2017 **SES-10** launch
Credit: SpaceX

- **Time-scale: 2 years**
  - **Creation of the GIE,** participation of the major Luxembourg actors from private & public sectors
  - **Launch of the national HPC infrastructure**
    - Hybrid supercomputer *i.e.* >1PFlops using CPUs/GPUs/FPGAs, >10PB BD storage
  - Existence of a **core expert team** (20 people)
  - Definition of a **set of use cases** with quick wins (industrial simulation, satellite imagery processing, etc.)
- **Time-scale: 3.5 years**
  - Connected infrastructure with at least 3 other EU countries (PRACE2 members a prime target)
  - Proof-of-concept of use cases
  - Administrative / organisational structure in place
- **Time-scale: 5 years**
  - Full integration with EU HPC facilities
  - Needs driven upscale of infrastructure
  - HPC expert team growth to 50 people, plus a team of technical support and administration

Success Criteria: Business and Finance

- **Within 3-4 years**
  - **Proof-of-concept projects with SMEs**
    - on adaptation, optimisation or new implementation of software / hardware on HPC/HTC systems
  - First set of SMEs benefit from the Competence Centre
    - through hands-on training, technology and application consulting
    - **product development** and **optimisation** using the Centre's infrastructure
- **Beyond 4 years**
  - Set of key companies either created or attracted to Luxembourg
    - have substantial funding or other economic activities based on the Competence Centre
  - Local companies using the Competence Centre that have never used HPC before
  - Improved business for existing companies, first on local then global market

View of Kirchberg plateau
Source: Luxembourg for Finance

Other Success Criteria (5 years)

- **As per J. Rifkin Report** "3rd Industrial Revolution Strategy for the Grand Duchy of Luxembourg"
  - Digitalisation improved in >3 sectors in Luxembourg
- **Reputation**
  - International computing and data hub:
    - Place to go to for hosting and processing your (sensitive) data
  - Very strong dynamics in growth of clients, both from inside and outside the country
  - > 3 international highly visible projects, where HPC-facility with the expertise of the analytics platform was the attractor
- **Publications**
  - > 5 Top papers that would not be there without the IPCEI HPC-BD
- **Support and Training**
  - Master in "High-Performance Computing" introduced by the University of Luxembourg
  - HPC-facility developed into a premier "Computing Data scientist"-training centre

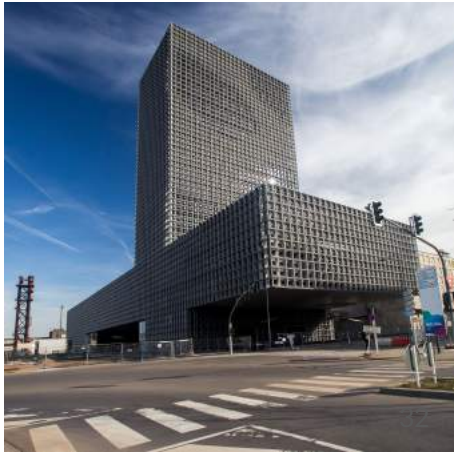

University of Luxembourg - House of Knowledge
Source: Inspiring Luxembourg

Conclusions and Next Steps

UNIVERSITÉ DU LUXEMBOURG

Luxembourg HPC-BD Competence Center

The MACHINE + the TEAM

=

Multiple Approaches in Computing with High INtegration and Efficiency

+

Technical Expertise and Application Management

Next Steps

- Nation-wide agreement on strategy
- Creation of expert team
- Creation of HPC-BD GIE/Foundation
- Development of implementation strategy
- Establishment Foresight Think Tank to identify early future business developments
- Identification of partners
- Identification of quick wins
- Implementation of demonstrators and exploration of longer term use cases

Disclaimer: These slides hold the contributions from many UL researchers.
Editors:

- P. Bouvry, S. Varrette, V. Plugaru & UL HPC Team
- R. Becker, W. Gu (LCSB)

Backup slides

UNIVERSITÉ DU LUXEMBOURG

Supercomputing in the World (2016 & beyond)

Supercomputer [latest update]	Computing capability	Storage capability	Core computational domains
Piz Daint (Switzerland) [2016]	361 760 cores ~4500 GPU accelerators (NVIDIA P100)	8.7 PB (max. 138 GB/s)	Chemistry & Materials, Physics, Earth, Life Science, Engineering, Data Science
Mare Nostrum 4 (Spain) [2017]	155 000 cores 2 types of accelerators to be added later (GPU+Xeon Phi)	14 TB	Computer Sciences, Earth Sciences, Life sciences, Software development for large scale (biomedical, geophysics, atmospheric, energy, social and economic simulations)
Tsubame 3 (Japan) [2017]	7560 cores 2160 GPU accelerators (NVIDIA P100)	15.9 PB (max. 150 GB/s)	AI (Deep Learning) applications in Healthcare, Energy and Transportation
Aurora (US) [2018+]	> 50 000 nodes Xeon Phi accelerators	> 150 PB (> 1TB/s)	Climate Science, Accelerator Design, Biological Science, Materials Science, Transportation Efficiency, Chemistry, Cosmology, Energy Storage


TaihuLight, China


Titan, US


Piz Daint, Switzerland


MareNostrum 4, Spain


Hazel Hen, Germany

HPC in the World

Physics, Chemistry and Materials Science
Earth and Environmental Sciences
Life Sciences, Engineering

AI applied in:
Healthcare
Energy
Transportation

Tsubame 3, Japan, 2017

Piz Daint, Switzerland, 2016

ECMWF Anemos, UK, 2016

Weather forecast
Environmental services
Meteorological data archival

Oil prospecting
Life sciences
Weather forecast
Industrial design
Pharmaceutical research

TaihuLight, China, 2016

Stampede 2, US, 2017

MareNostrum 4, Spain, 2017

Computer Science
Material Science
Physics, Chemistry
Earth Sciences
Life Sciences

Physics, Chemistry
Materials Science
Biological Sciences
Earth Sciences
Computer Science
Engineering

Titan, US, 2012

Hazel Hen, Germany, 2015

Marconi, Italy, 2017

Physics and Chemistry
Computational Engineering
Earth and Climate Science
Life Sciences

Materials Science, Particle Physics
Astrophysics, Environment and Energy
Computational and Scientific Engineering

Mathematics & Physical Sciences
Biological, Behavioral & Social Sciences
Geosciences, Engineering, Computer & Information Sciences