



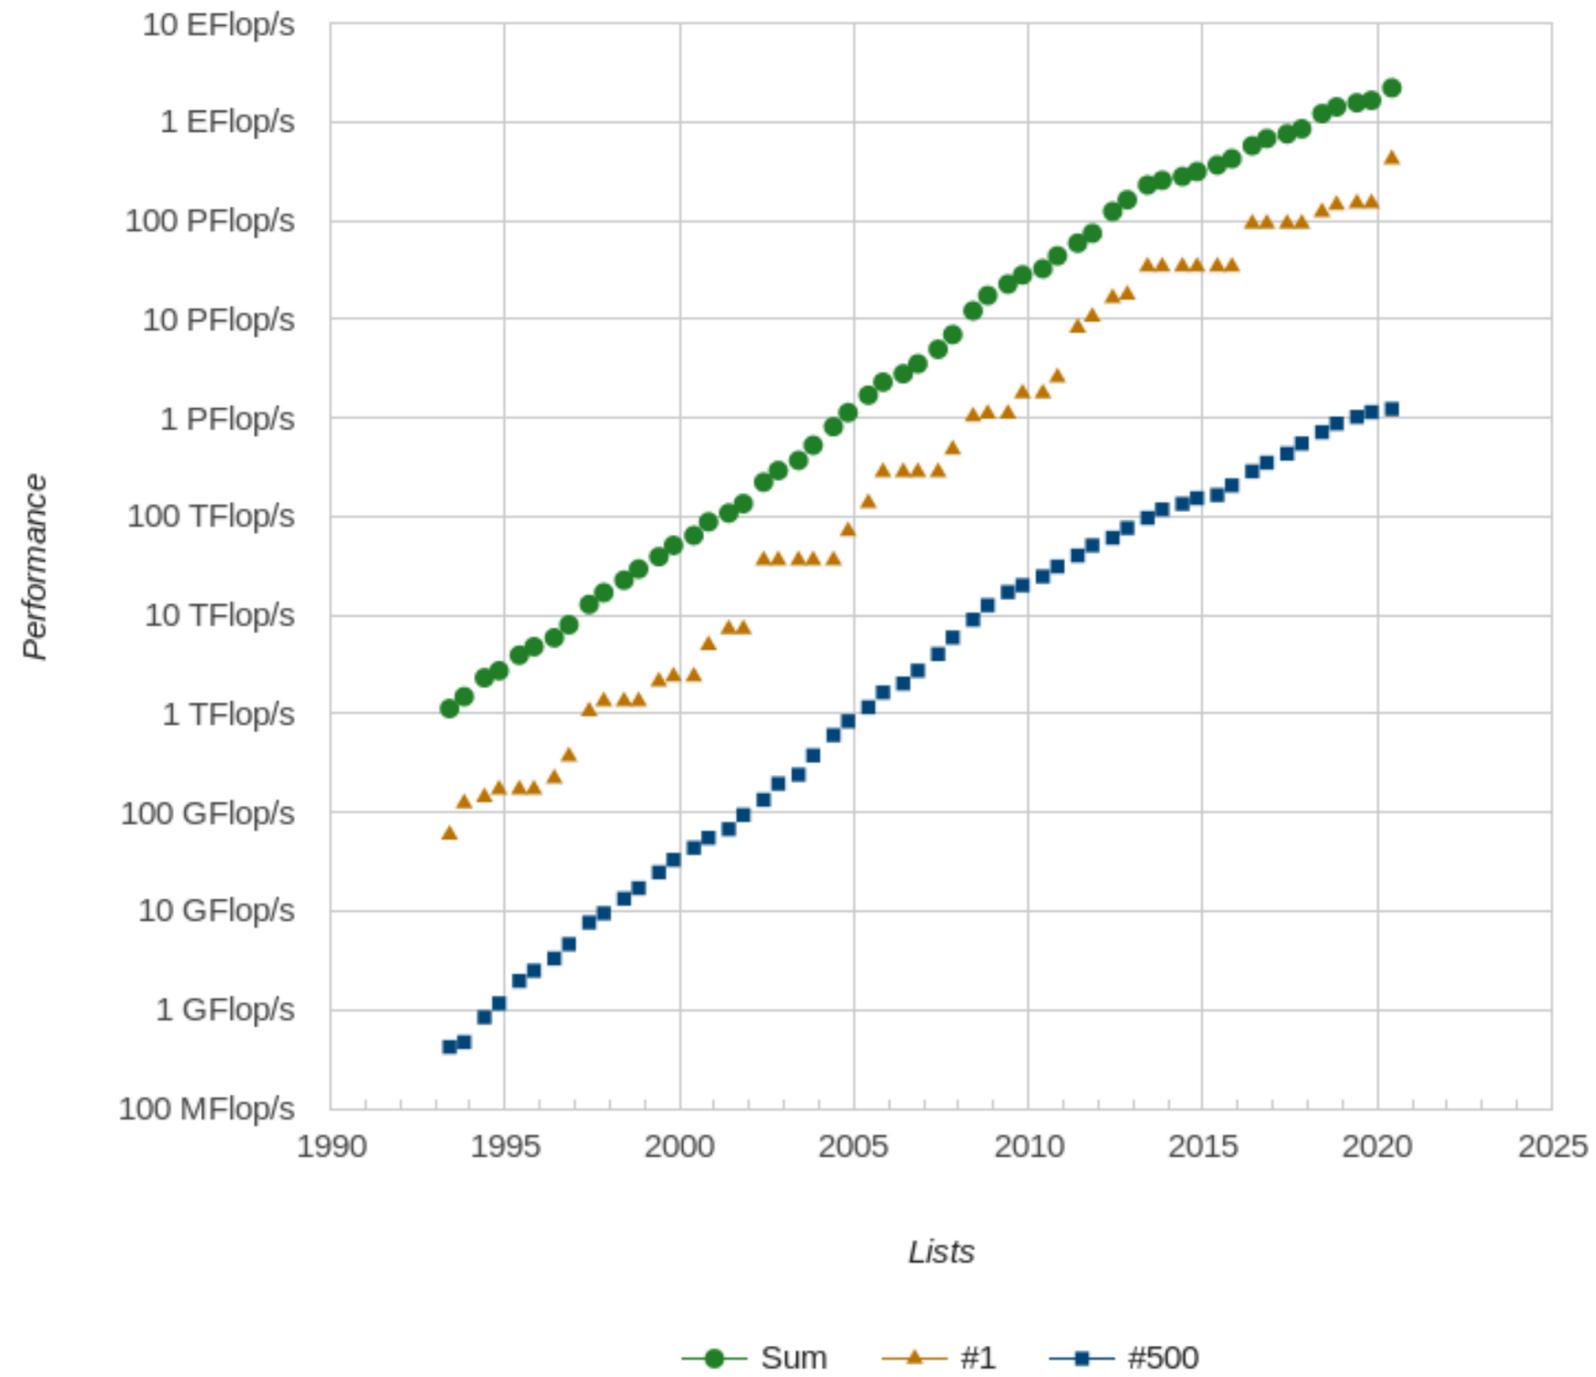
HIGH-PERFORMANCE COMPUTING & DATA ANALYTICS

MeluXina, a Competitive Advantage for Luxembourg

Prof. Dr. Pascal Bouvry
Université du Luxembourg
pascal.bouvry@uni.lu

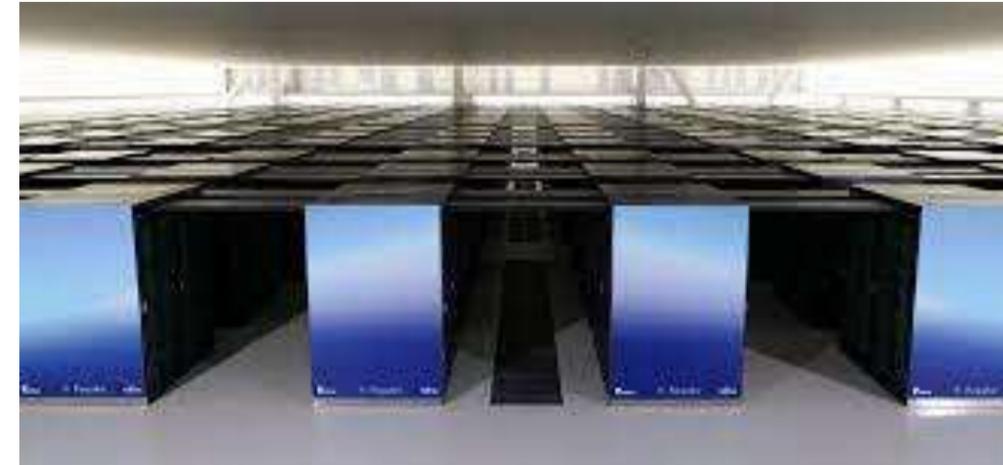
Dr. Matthieu Lefebvre
Head of Supercomputing & Data Solutions
matthieu.lefebvre@lxp.lu

Performance Development

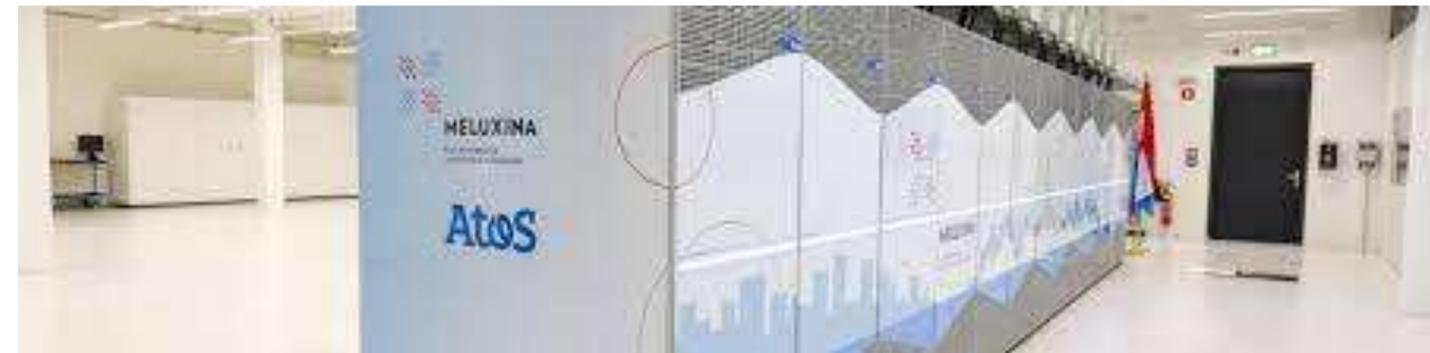




1. Frontier (Oak Ridge, USA) 1,1 Exaflops



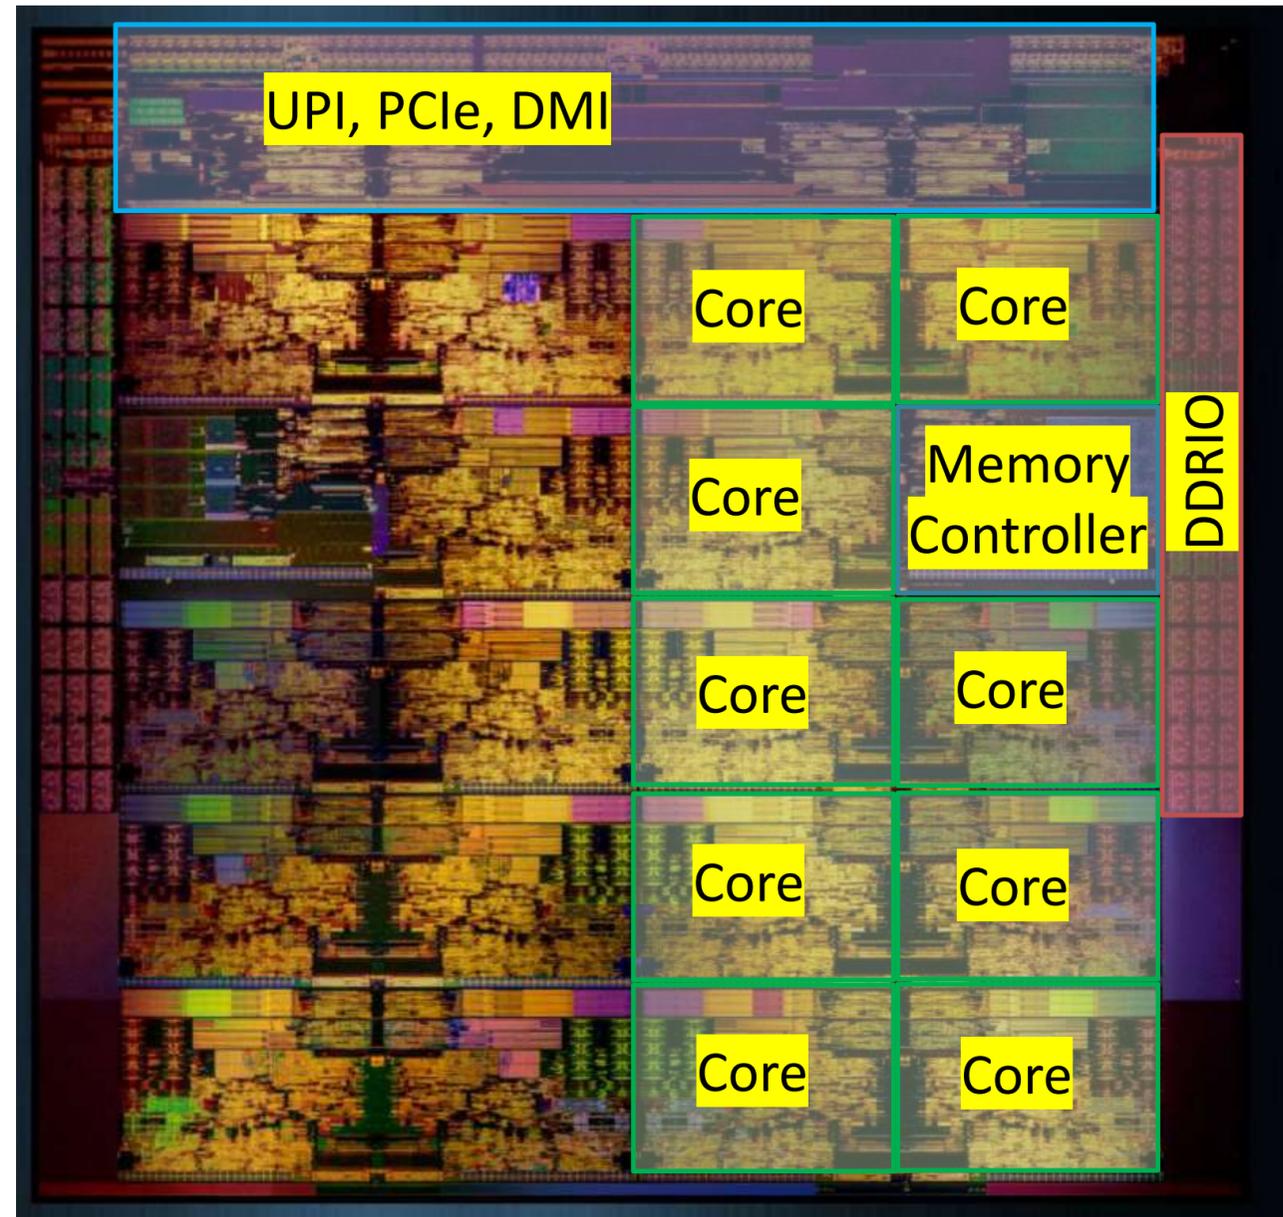
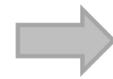
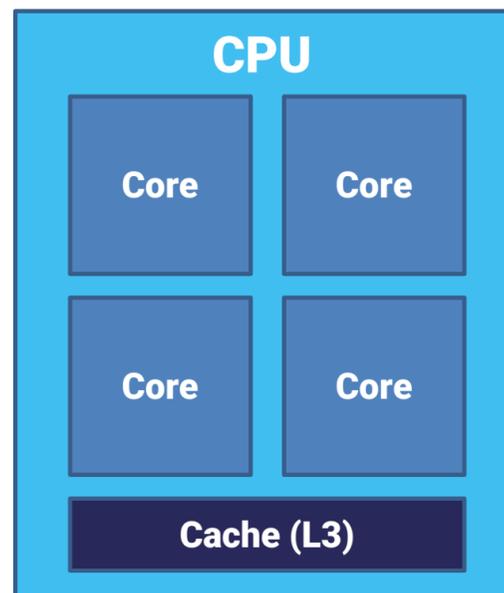
2. Fugaku (Riken, Japan) 442 PFlops



48. MeluXina (Bissen, Luxembourg) 10 PFlops

Ranking Top500 June 2022

CENTRAL PROCESSING UNIT (CPU)



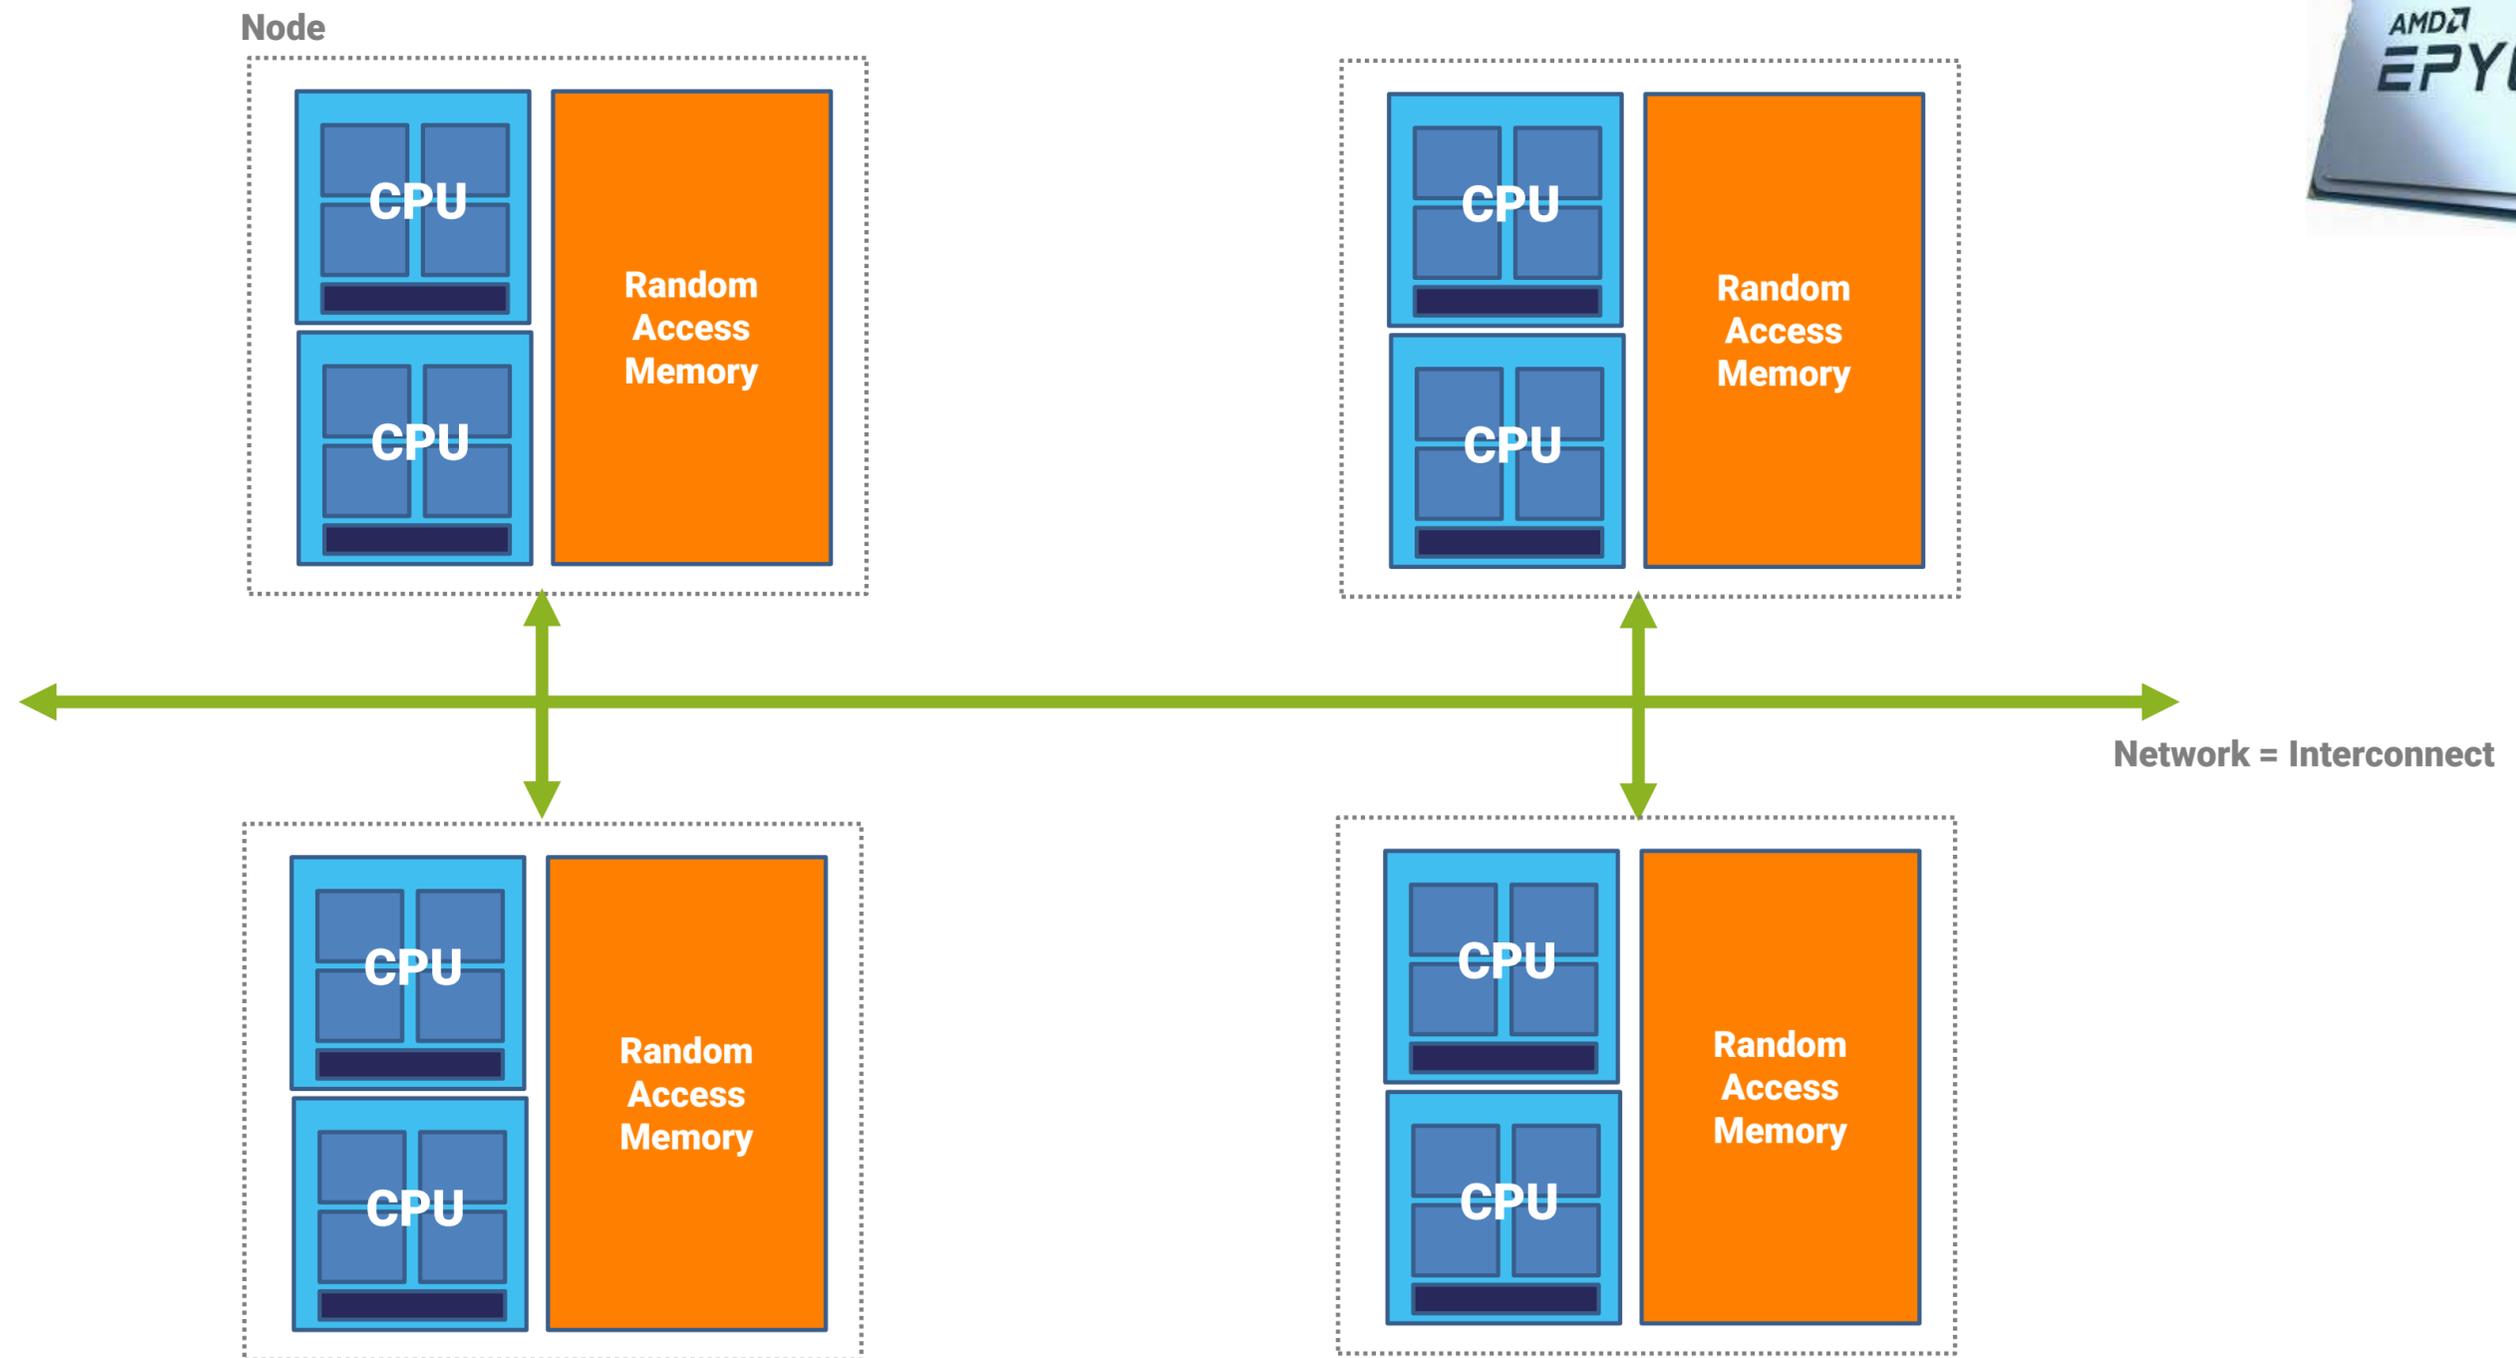
Intel Skylake (Courtesy of Intel)



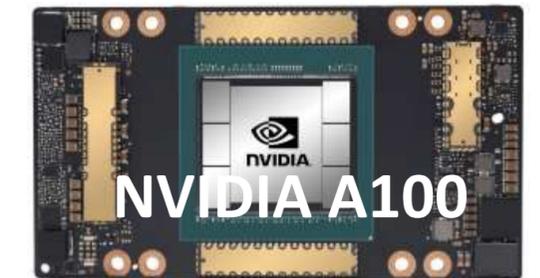
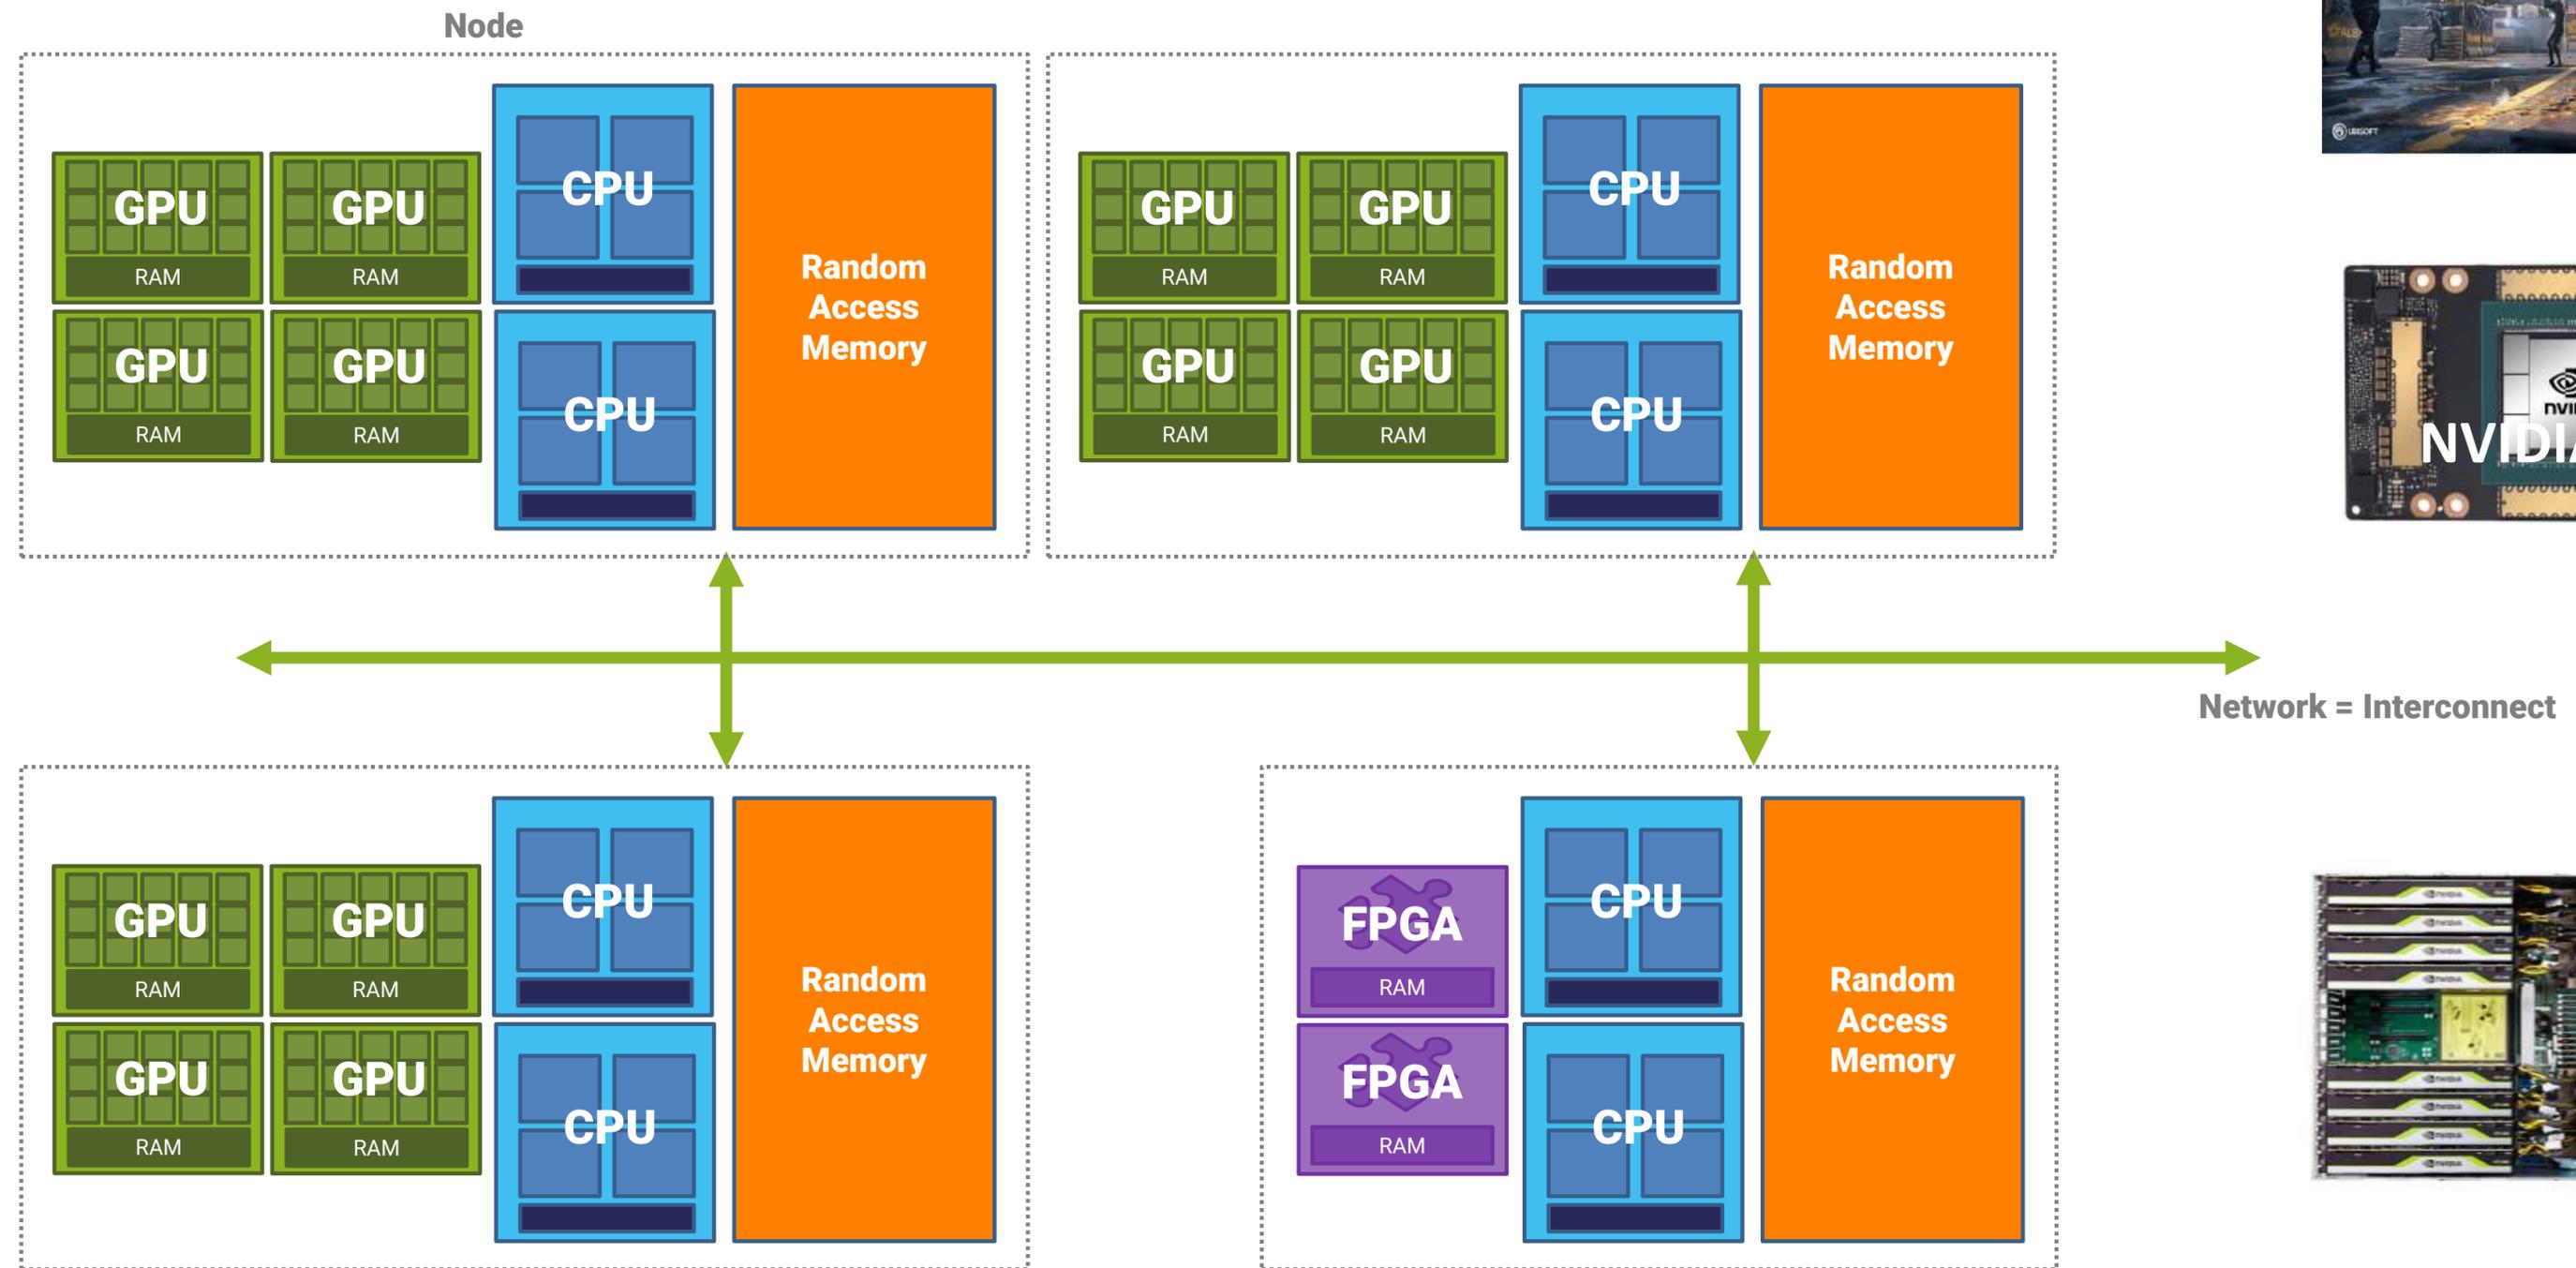
Core Definition (Jarp):

“A complete ensemble of execution logic, and cache storage as well as register files plus instruction counter (IC) for executing a software process or thread.”

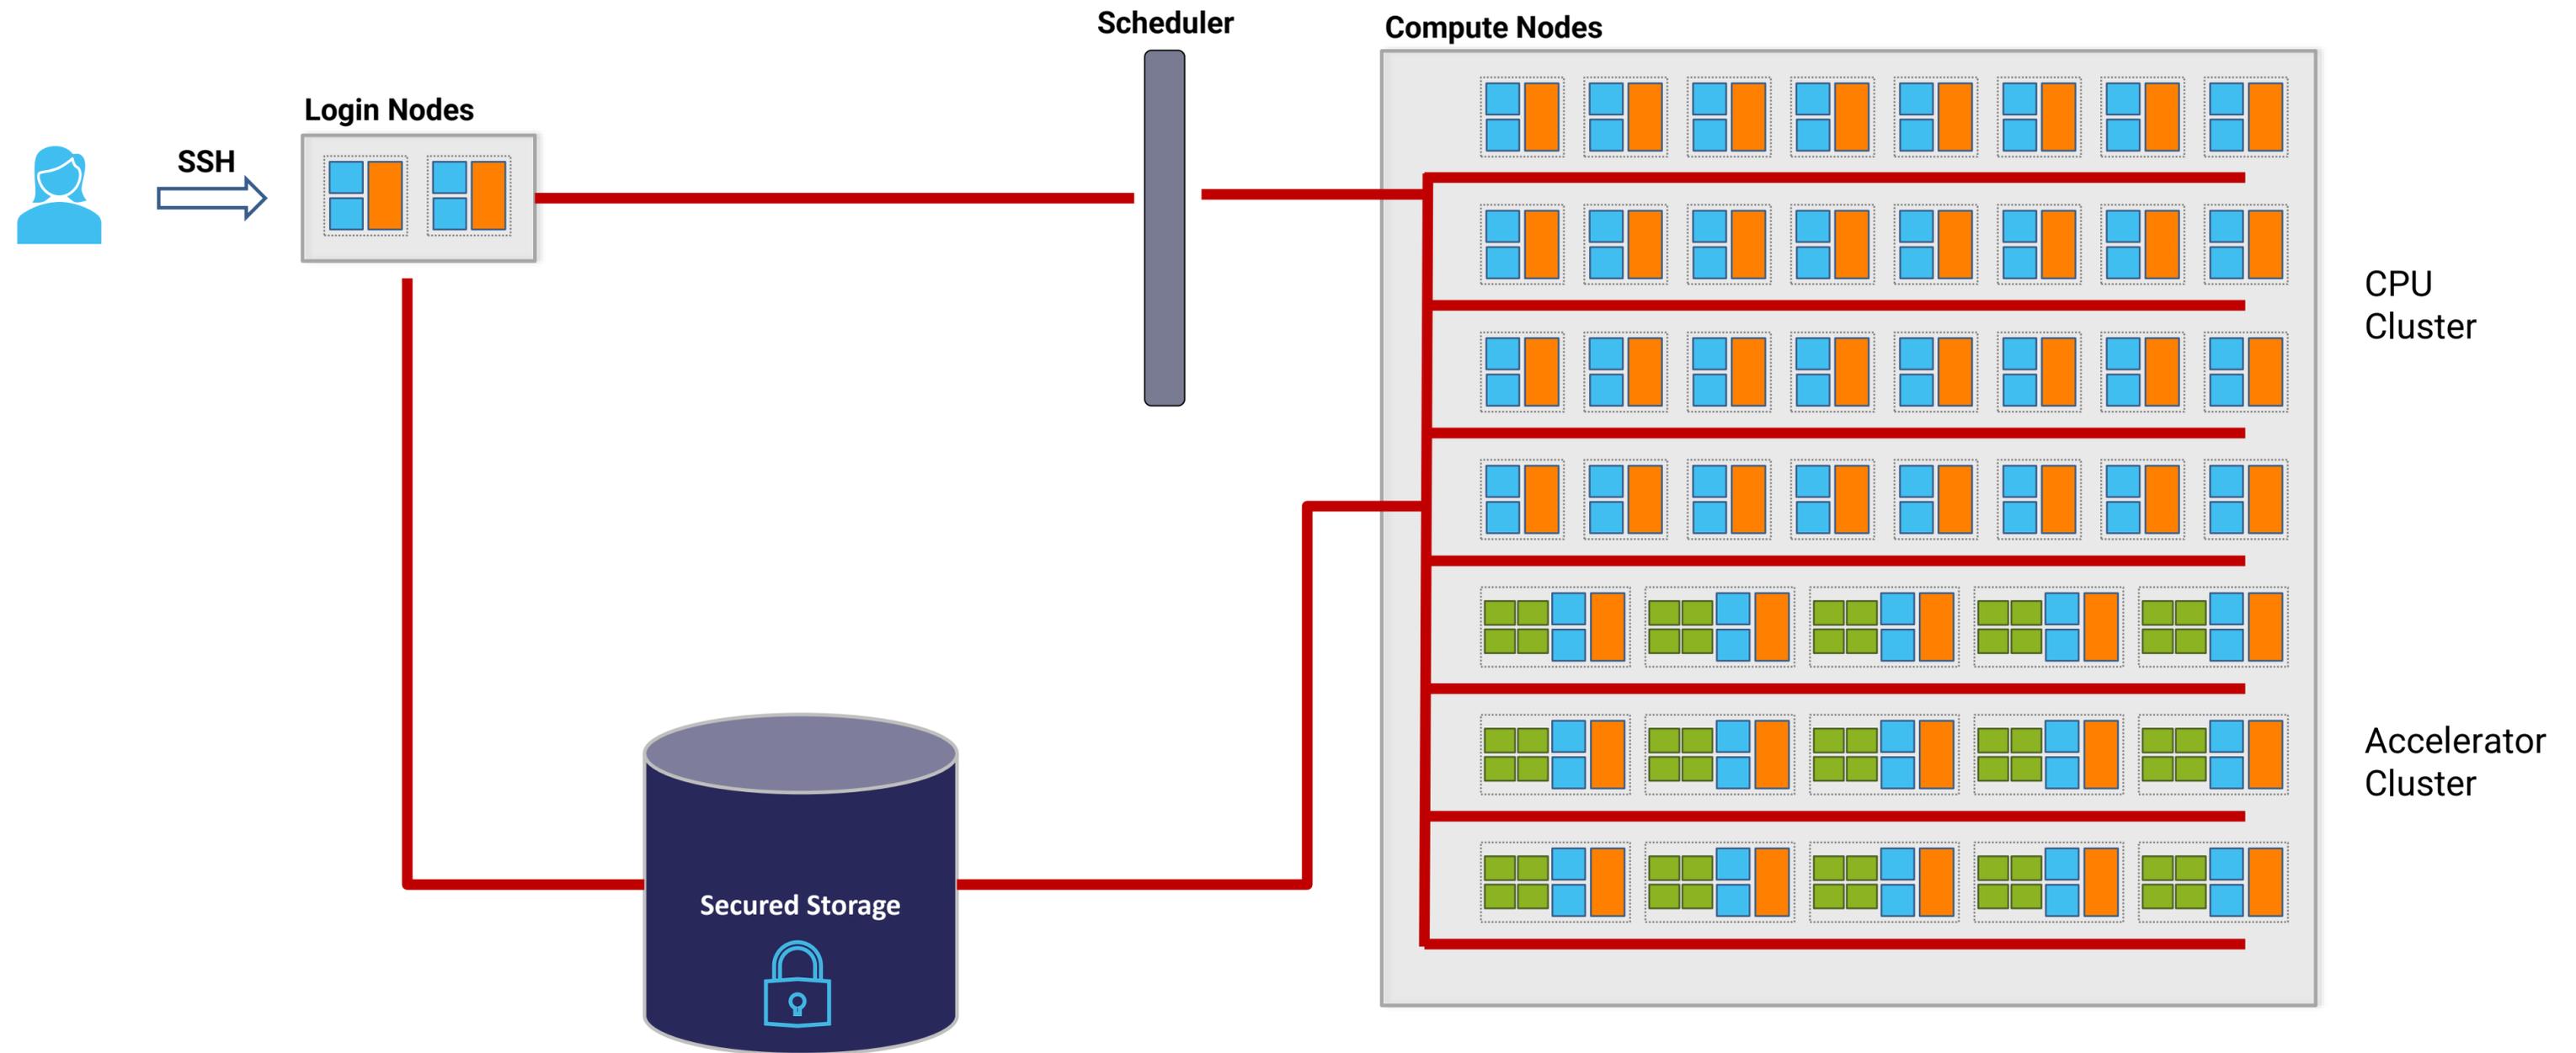
CPU "CLUSTER"



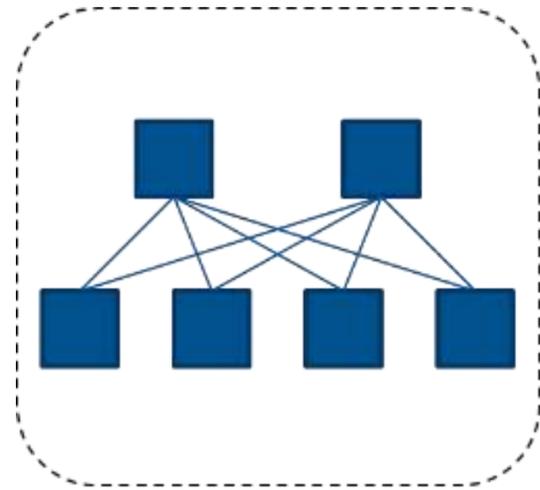
ACCELERATOR "CLUSTER"



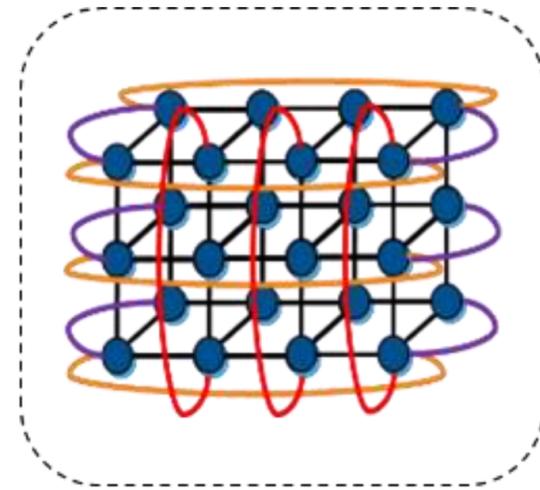
BASIC CLUSTER ACCESS LAYOUT



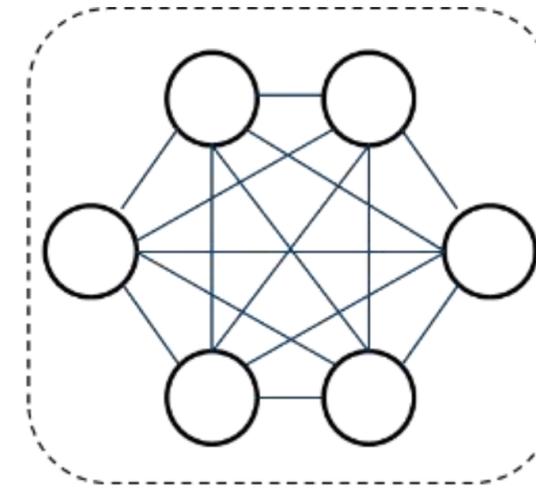
Network Topologies



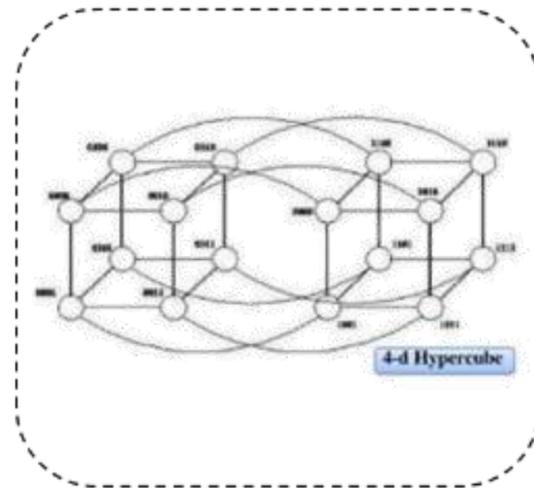
Fat Tree



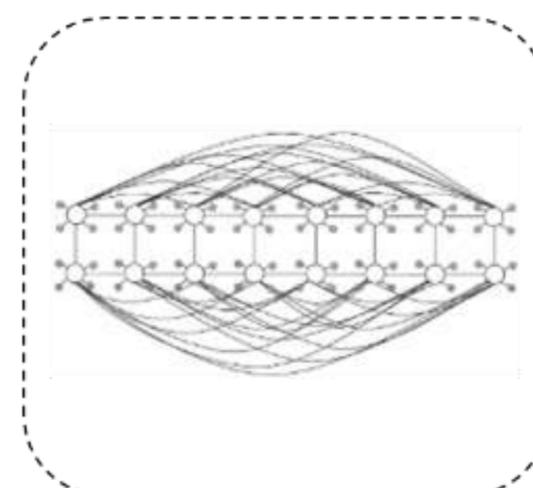
Torus



Dragonfly

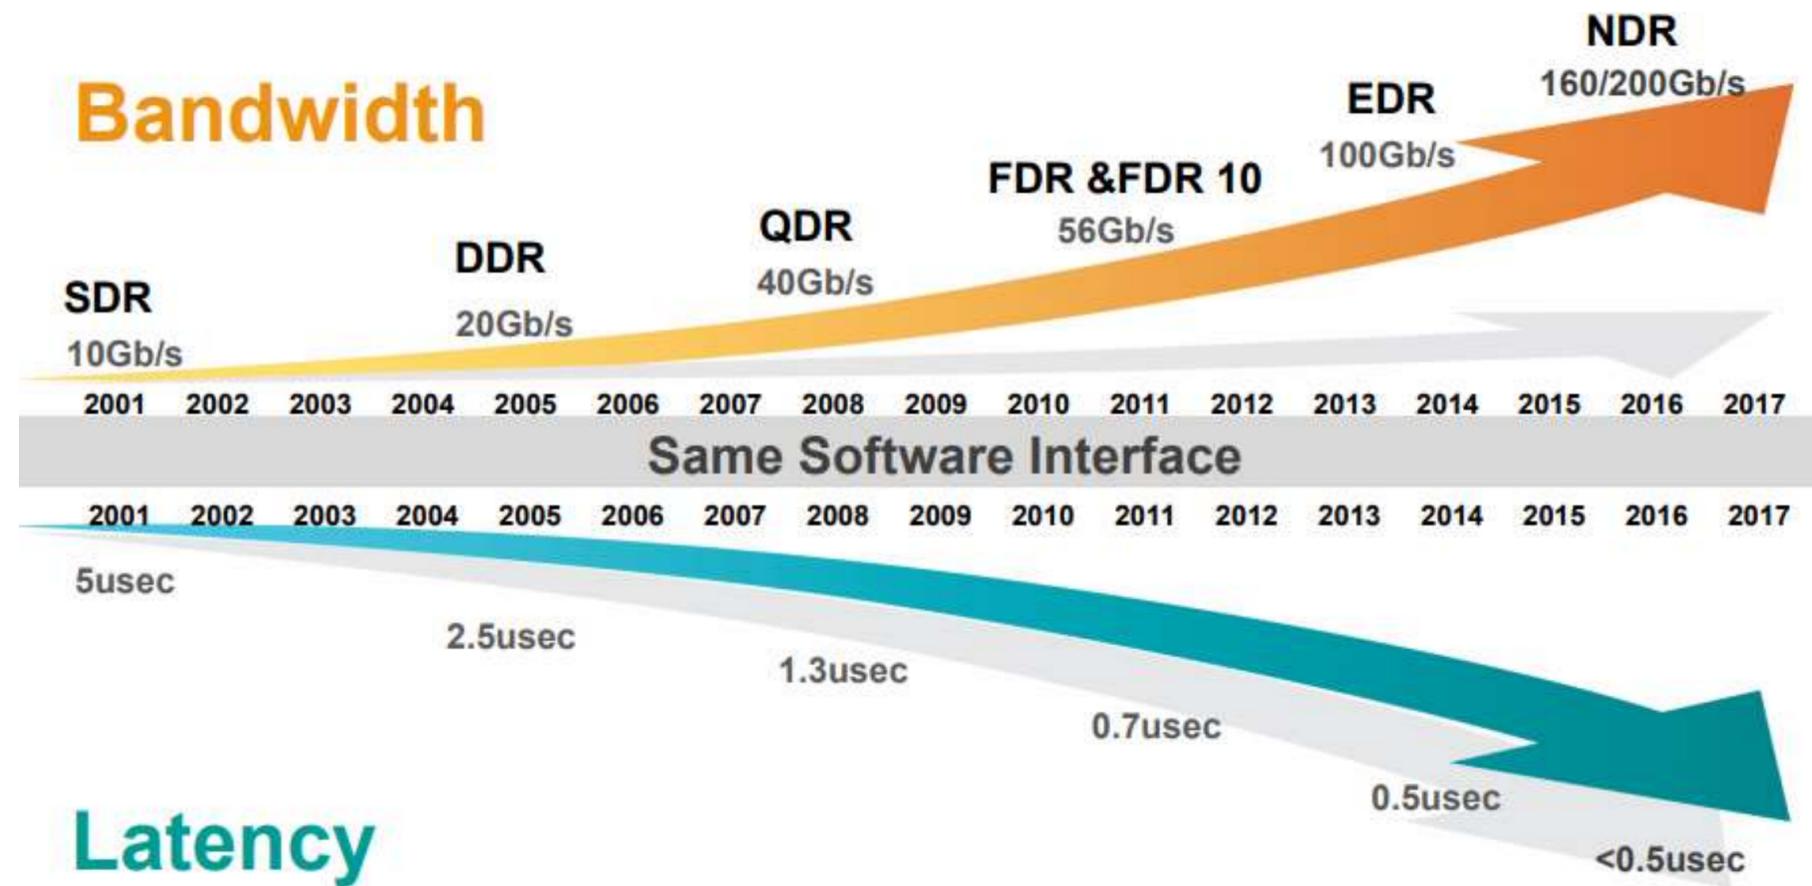


Hypercube

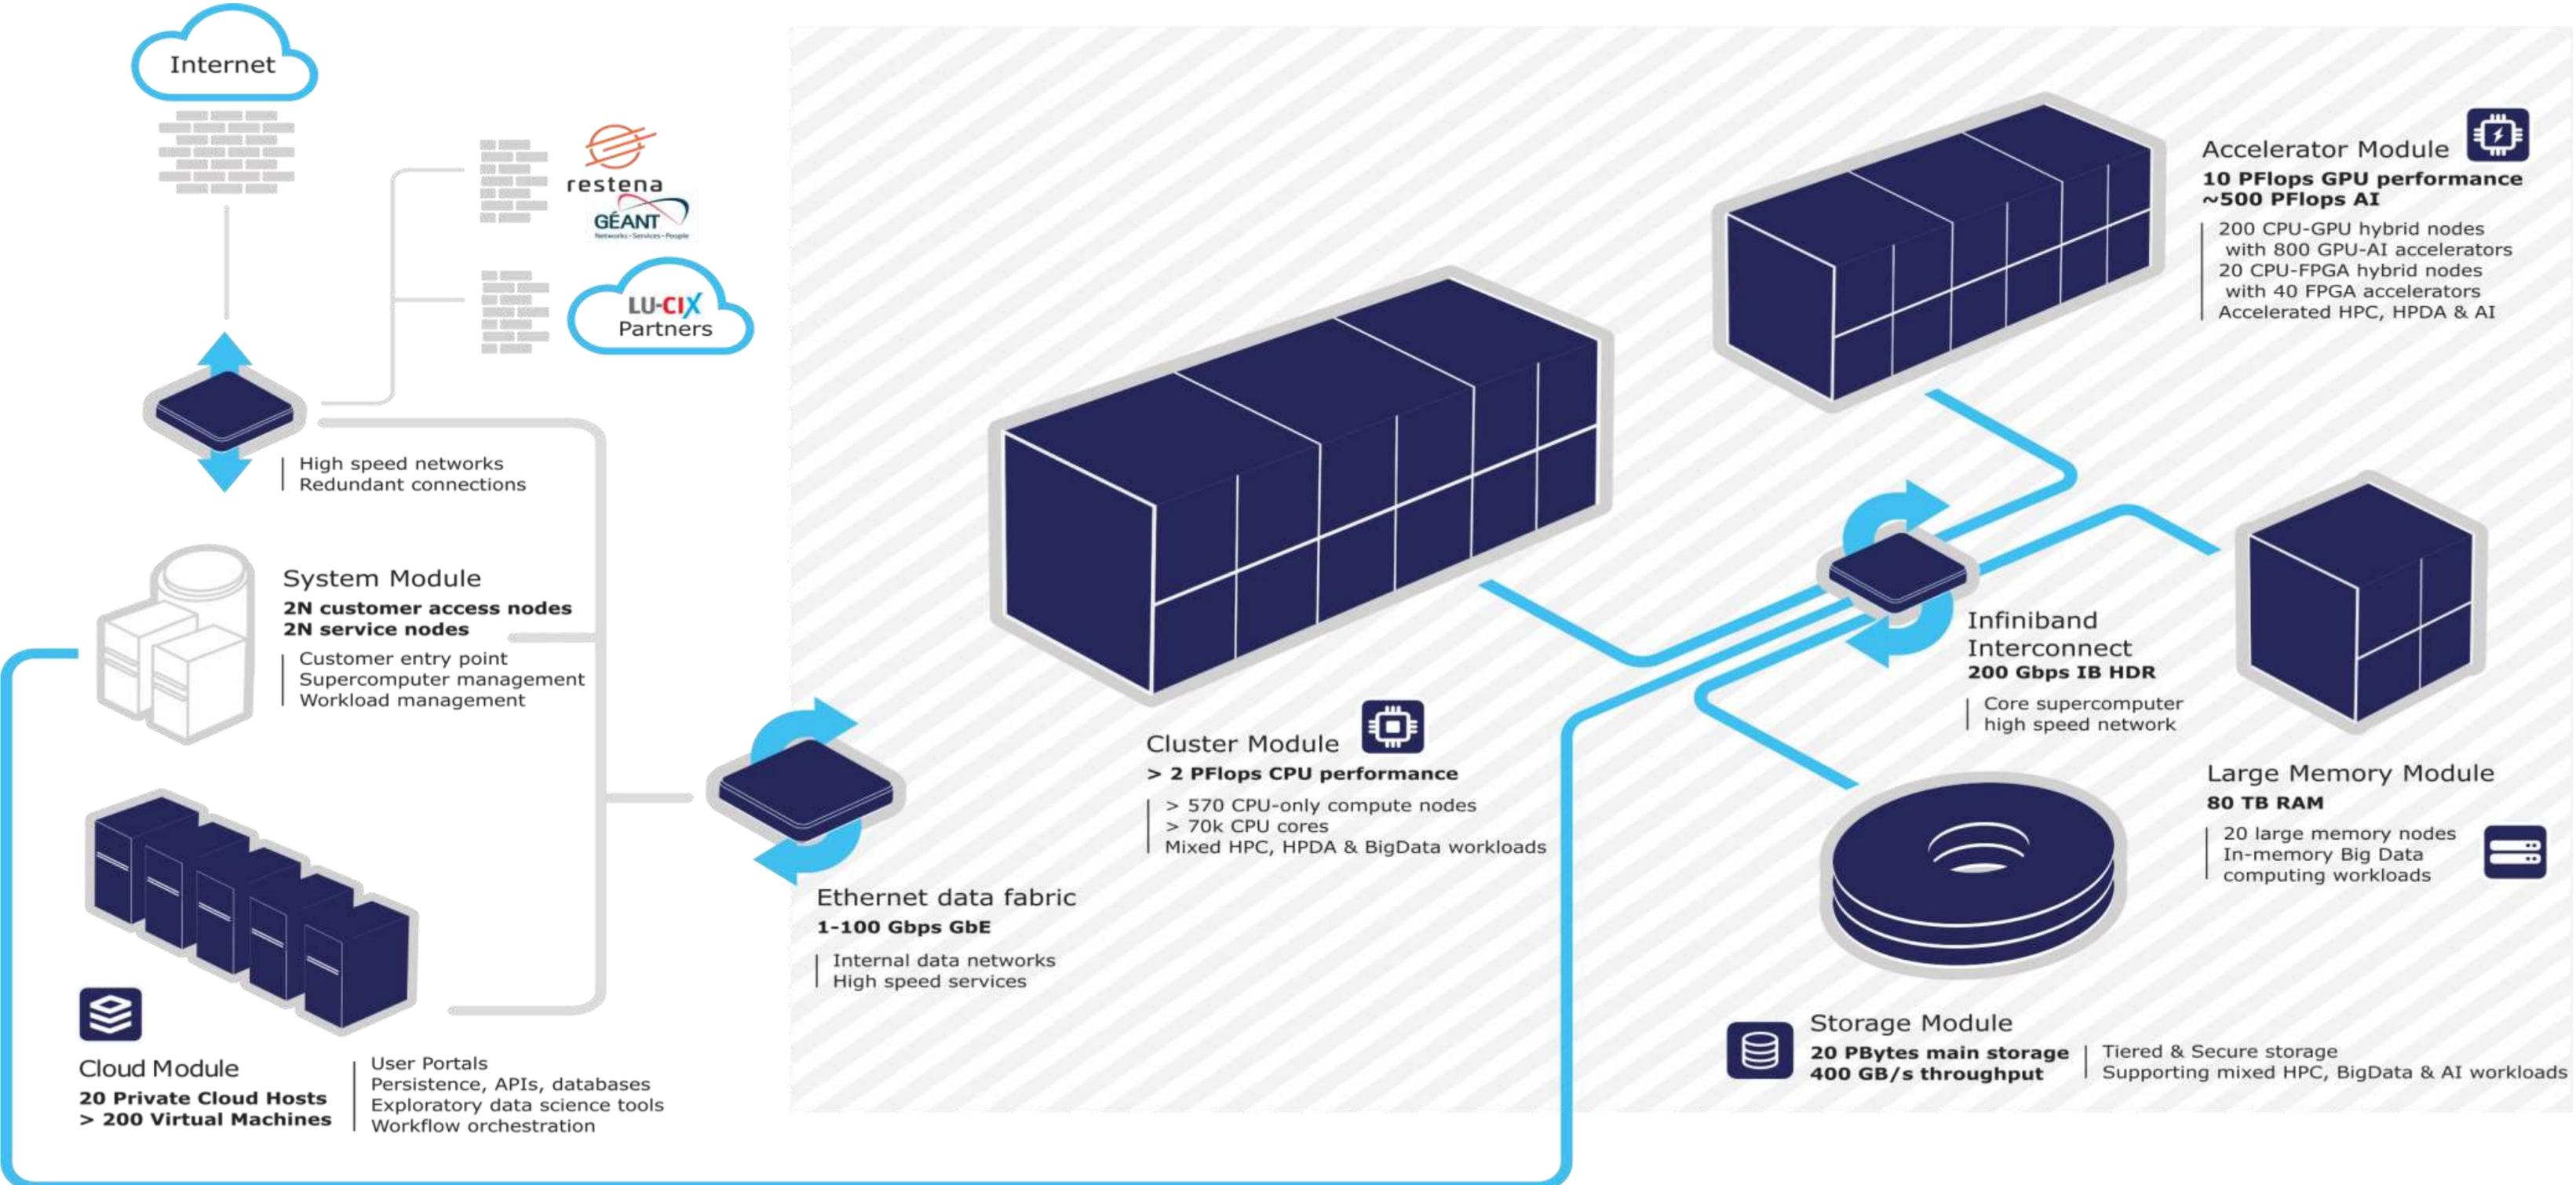


HyperX

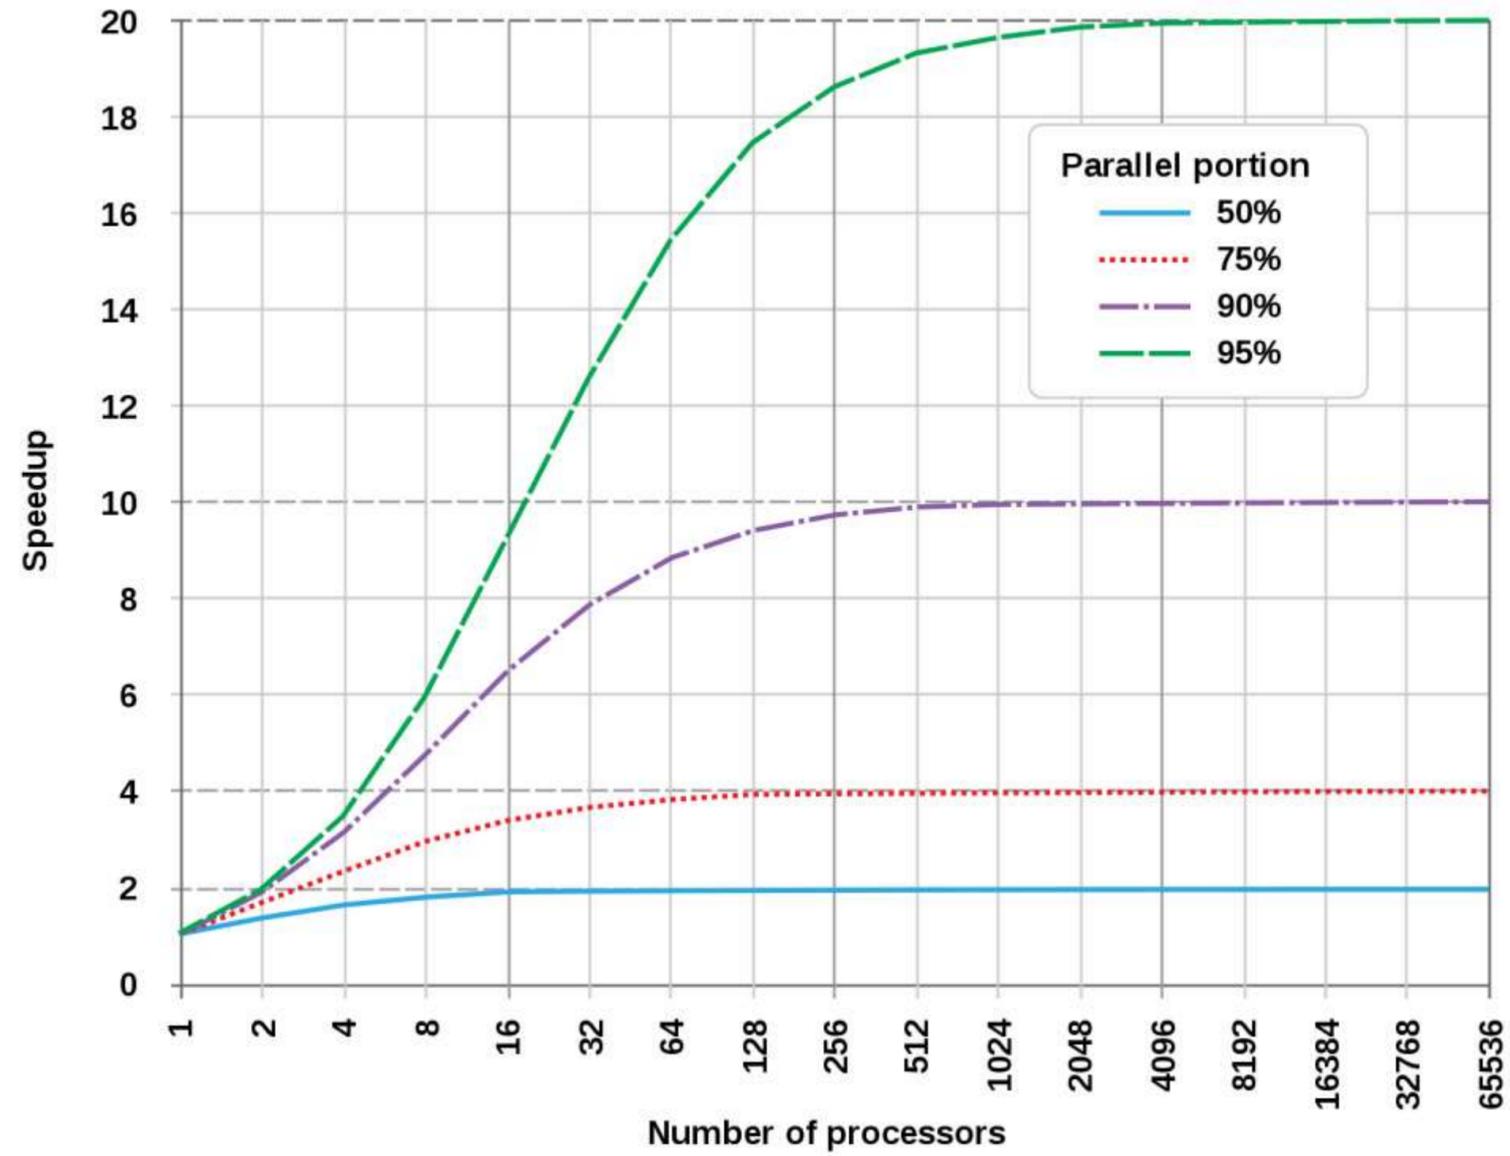
Infiniband



MELUXINA



Amdahl's Law



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

- Federation of national and regional HPC centers (see also PRACE/PRACE2) [EU HPC strategy](#) initiated in 2012 implementation within H2020 program
- EU Member States sign the [EuroHPC initiative](#) and prepare its implementation (Mar. 2017)
- Creation of EuroHPC JU (Jan. 2019) based in Luxembourg with 1 Billion Euros funding for Peta, Pre-Exa and Exascale solutions
- A common effort to create and grow the [European supercomputing ecosystem](#)

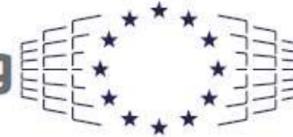
EuroHPC JU



EuroHPC
Joint Undertaking

2019-2020: PUTTING EUROPE IN THE LEAD

The EuroHPC Joint Undertaking



EuroHPC
Joint Undertaking

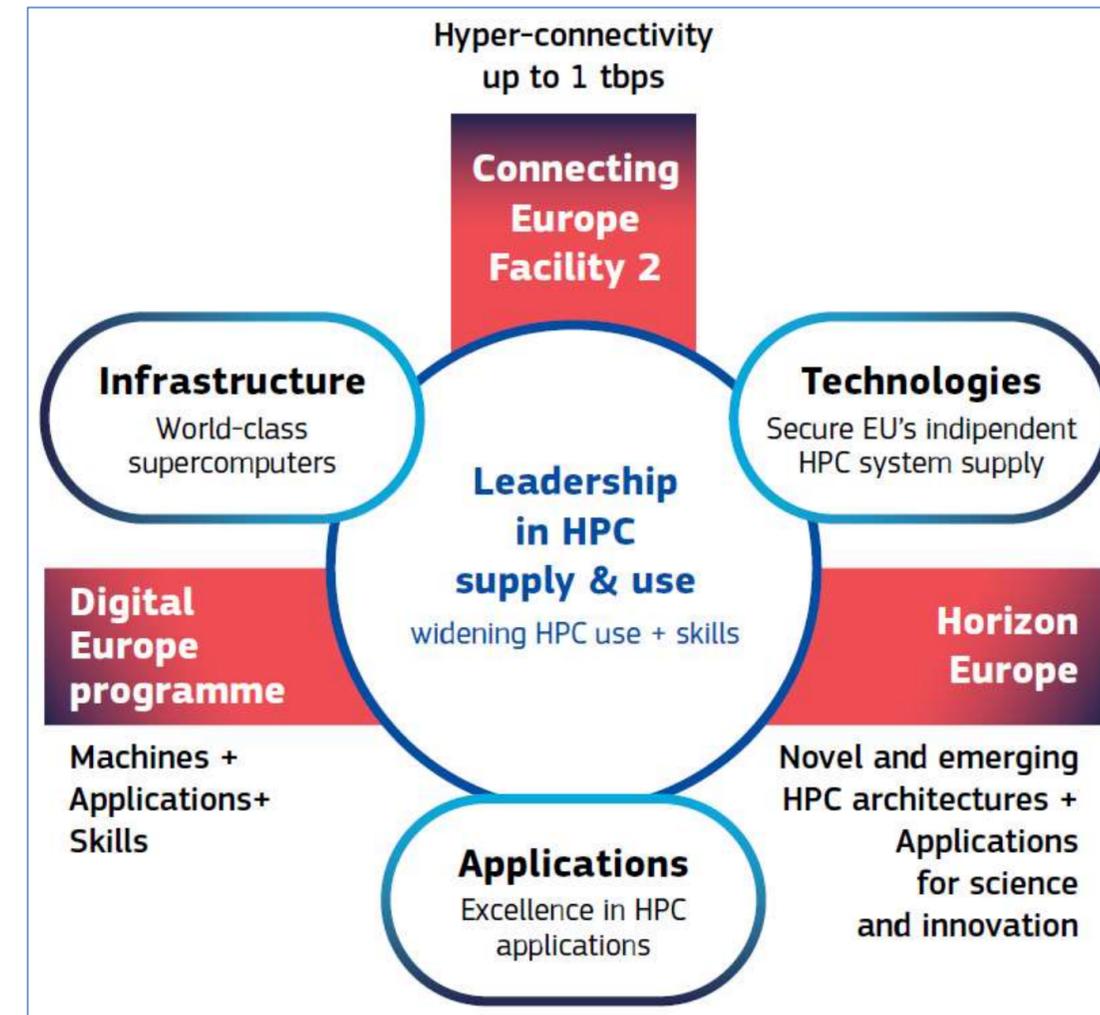
Legal and funding body created by the Council in September 2018



Mission: Pool European resources to establish an integrated world-class supercomputing & data infrastructure and support a highly competitive and innovative HPC and Big Data ecosystem.



- ➔ 29 Participating States + EU
- ➔ Site: Luxembourg
- ➔ Budget (2019-2020): €1.5 Billion
- ➔ Operational: 11/2018 to 2026



ec.europa.eu/digital-single-market/en/policies/high-performance-computing

2021

DIGIBYTE | Publication 28 May 2021

Commission welcomes Council's support for a €7 billion investment in Europe's supercomputing infrastructure

The European Commission welcomes today's General Approach in the COMPET Council by Research Ministers on the European High Performance Computing Joint Undertaking Regulation.

EuroCC & CASTIEL Projects



MORE THAN INFRASTRUCTURE...

- EuroCC project about kickstarting **National (HPC, HPDA & AI) Competence Centers**
- 1 NCC per country, 33 participating countries
- **Luxembourg NCC:** LuxInnovation, the University of Luxembourg and LuxProvide

... ALL ABOUT EXPERTISE

- NCCs will provide **broad service portfolio** tailored to national needs
- ... of industry, academia & public administration
- LuxProvide leading work on **Technology Transfer, Business Development & Collaboration with Industry**
- ... will also organize trainings, workshops

... AND COORDINATION

- CASTIEL project will promote interaction & exchanges across NCCs

Find out more about EuroCC Luxembourg: eurocc-luxembourg.lu



[linkedin.com/company/eurocc-luxembourg](https://www.linkedin.com/company/eurocc-luxembourg)



twitter.com/EuroCC_Lux



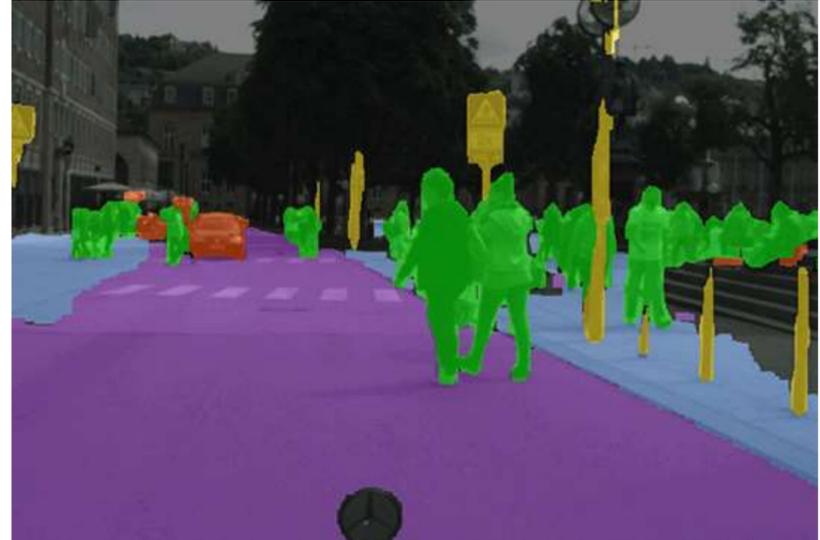
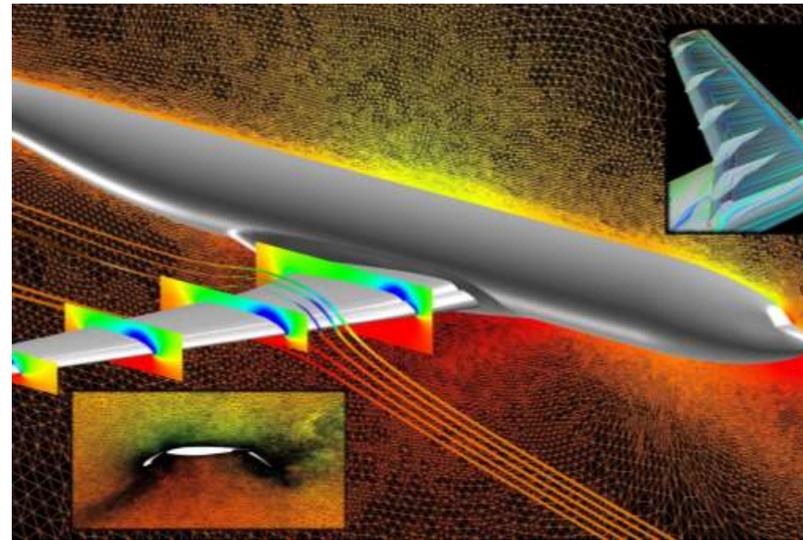
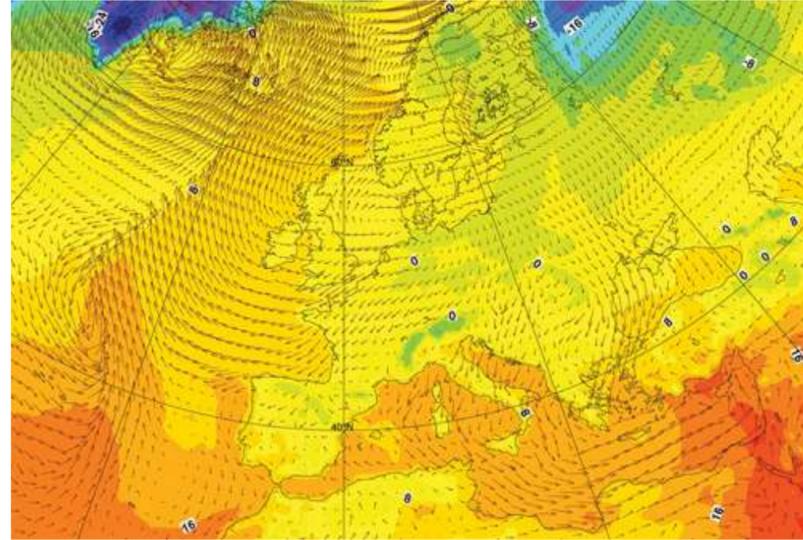
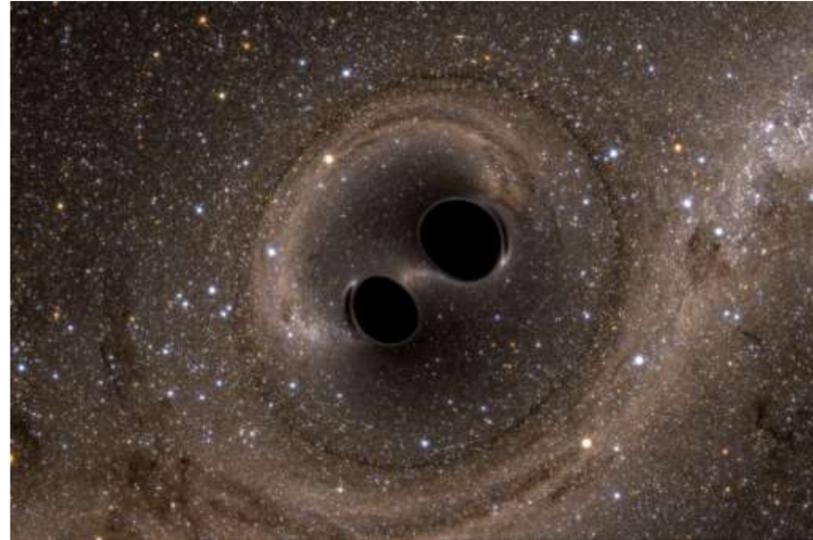


The Scientific Case for Computing in Europe 2018-2026

The PRACE Scientific Steering Committee

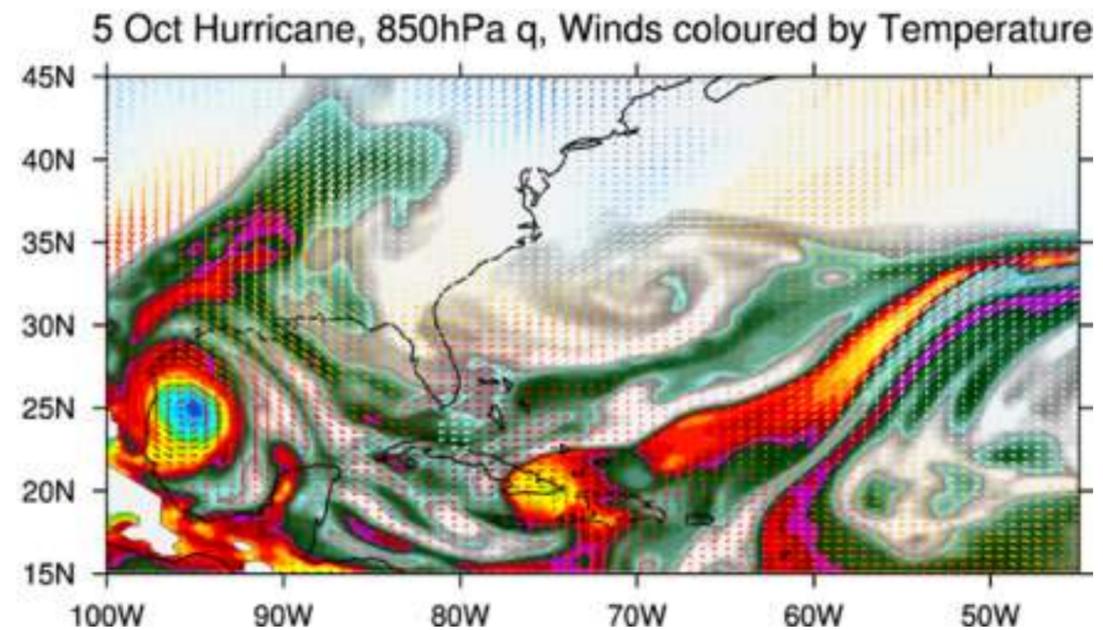
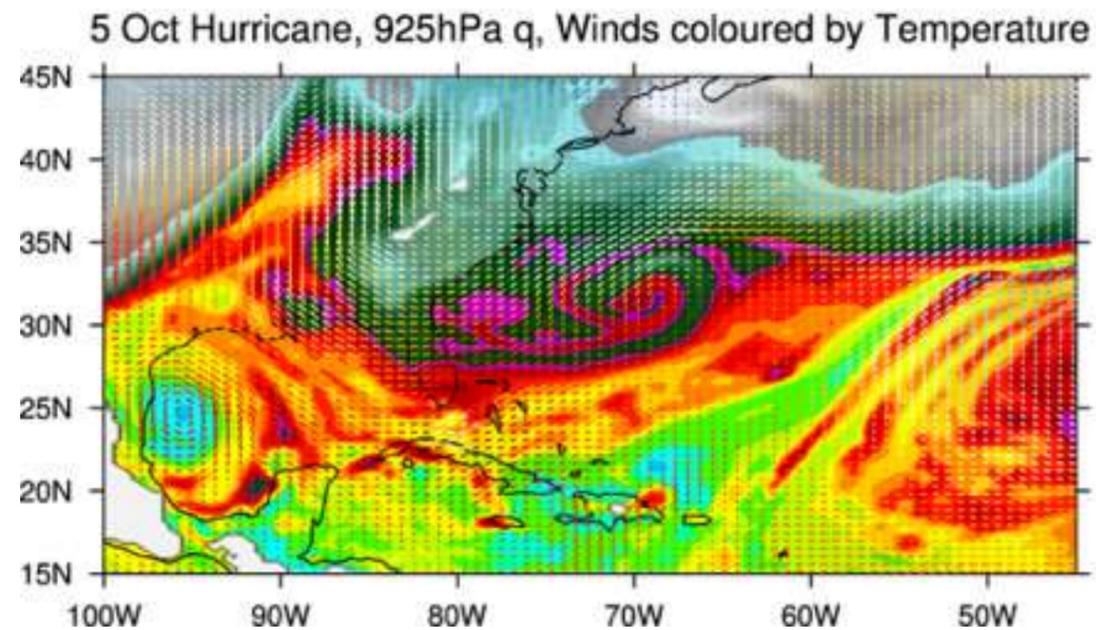
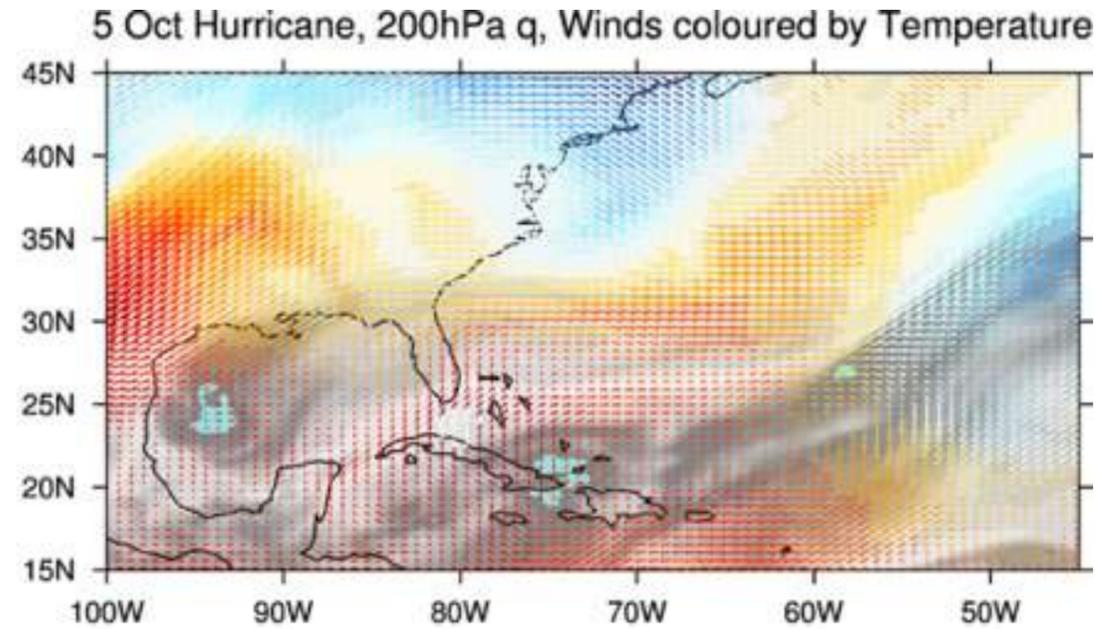
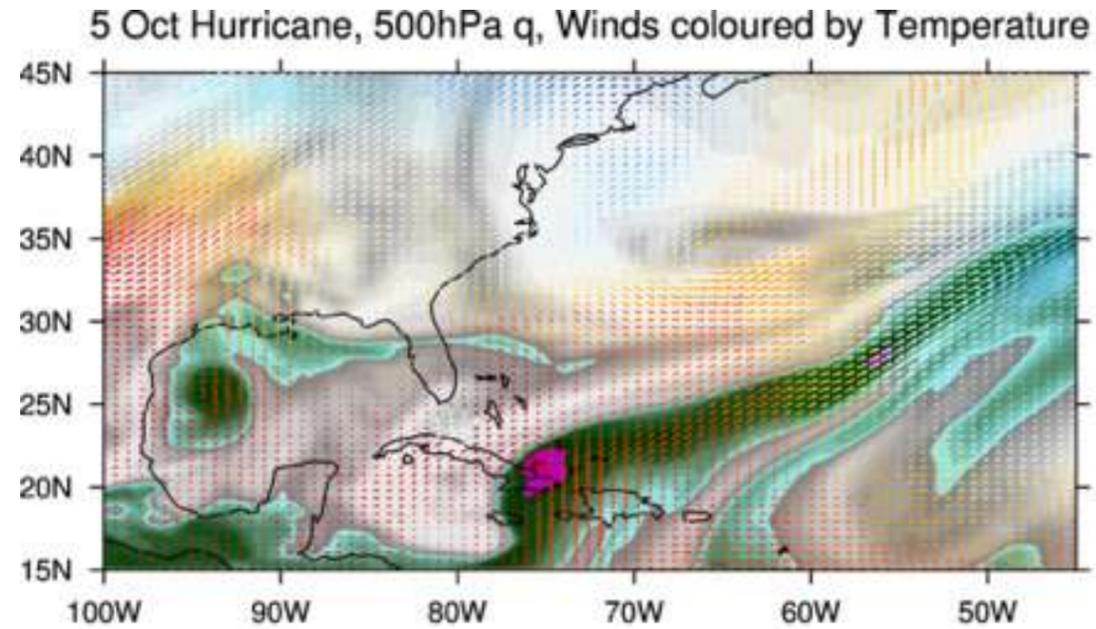
Erik Lindahl & Sinéad Ryan

European Computing Solves Societal Challenges

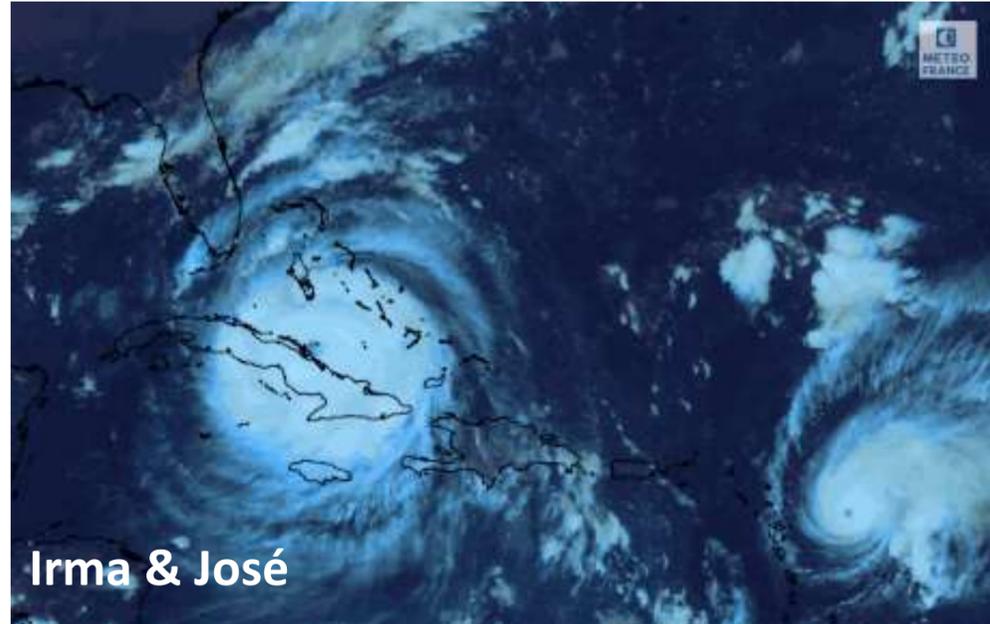


PRACE's goal is to help solve these challenges. The days when scientists did not have to care about the hardware are over, and so are the days when compute centers did not have to worry about the scientific application!

Climate, Weather & Earth Ecosystems



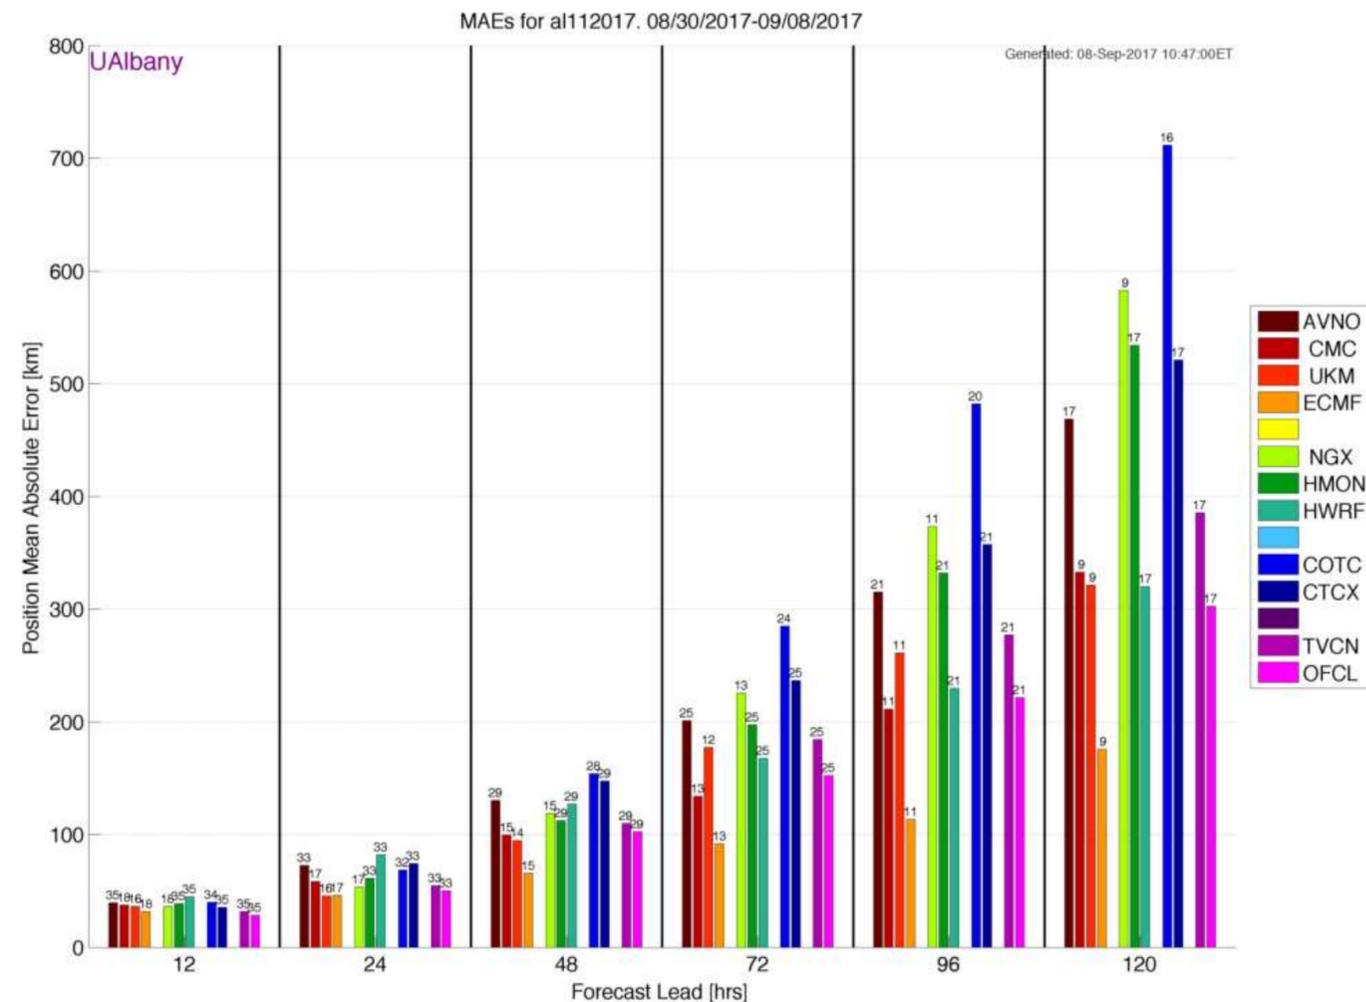
Europe leads international code development



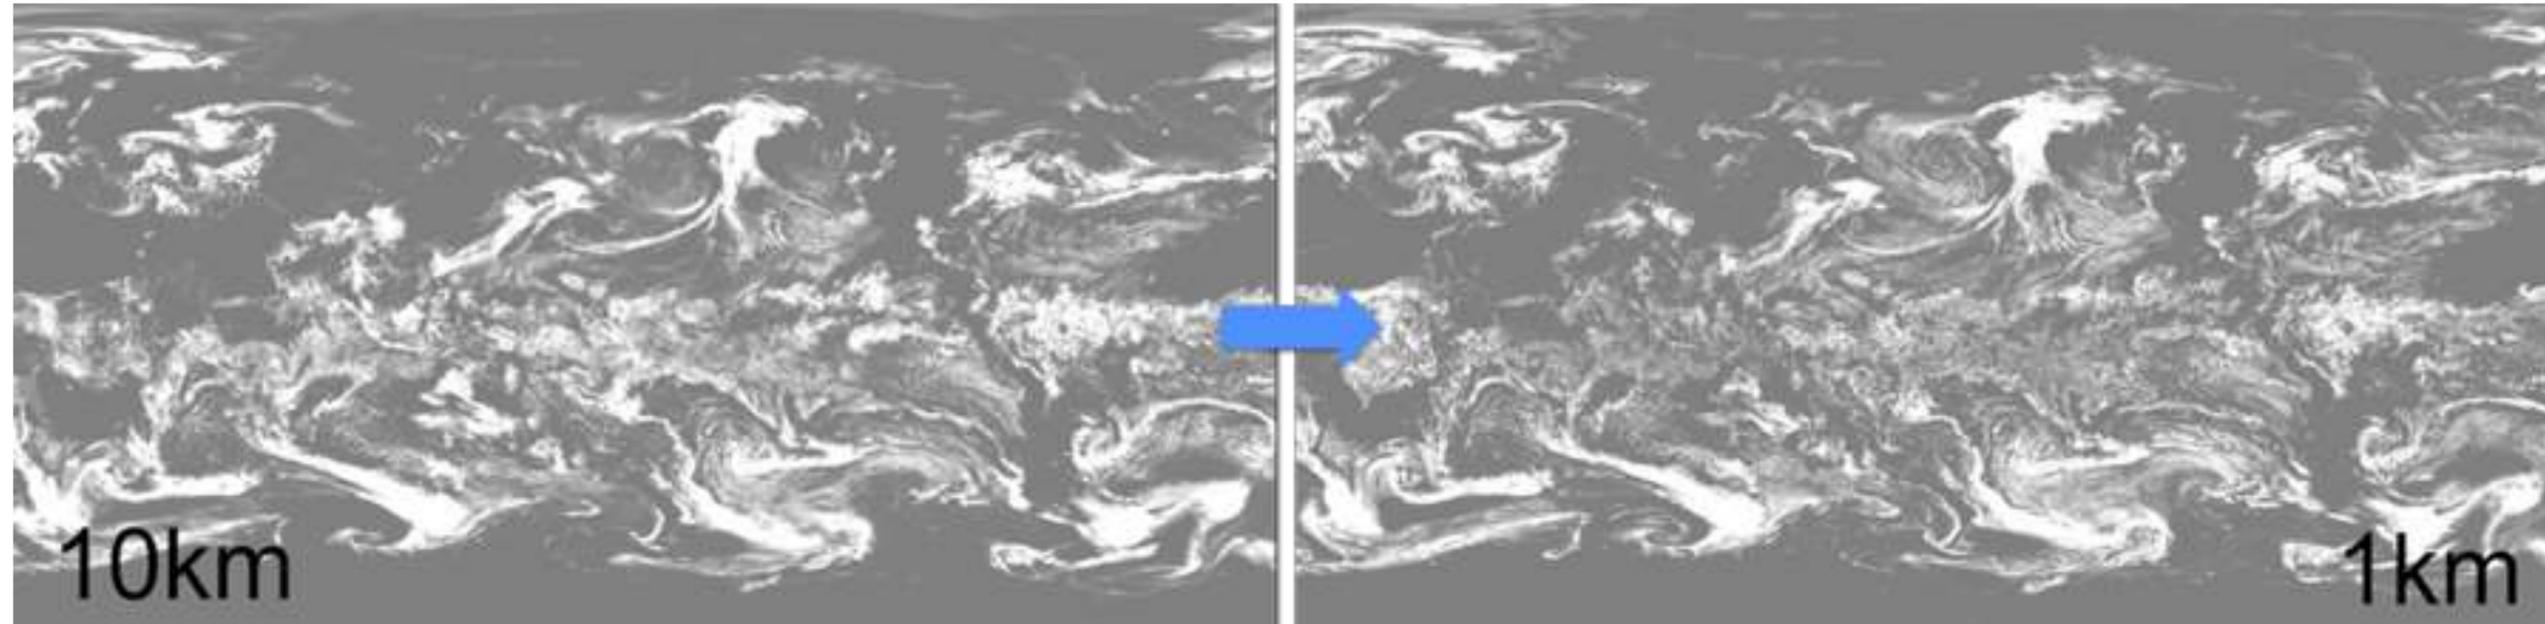
Predicting evacuation needs is a life/death matter – but avoiding it saves €250M

Maintaining our European lead on software should be one of our investment priorities

The best forecast codes are European
Arguably, the US would have done better by investing less in machines and more in software



The Centers of Excellence are Working

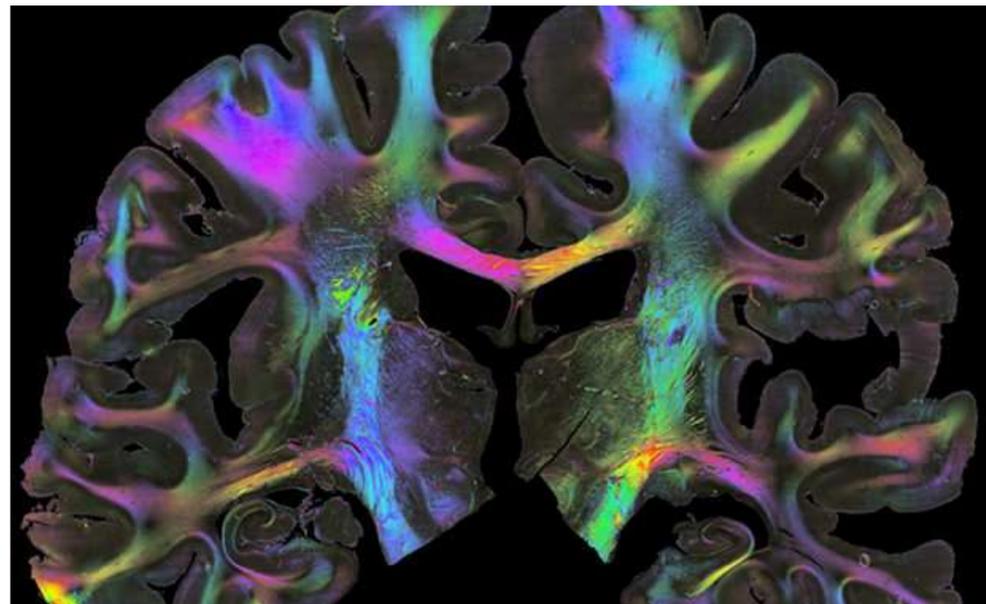


The Centers of Excellence are Working

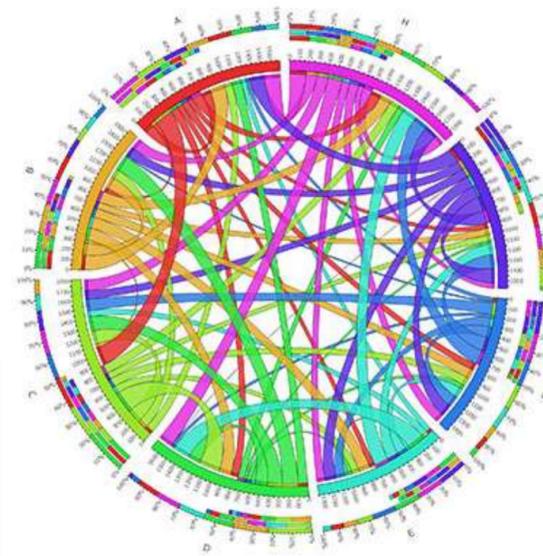
Accurate 20 day weather forecasts would translate to higher revenues in agriculture, tourism and production

- Excellent to outstanding impact in a number of research & code areas
- Important not to be short-sighted: Changing codes is a *longer* process than building an Exascale machine, but with greater long-term European impact. Give them room to focus on the real goal.
- Substantial impact on European competitiveness through workforce education & degrees
- But: Beware that the present CoEs do not cover all areas of science!

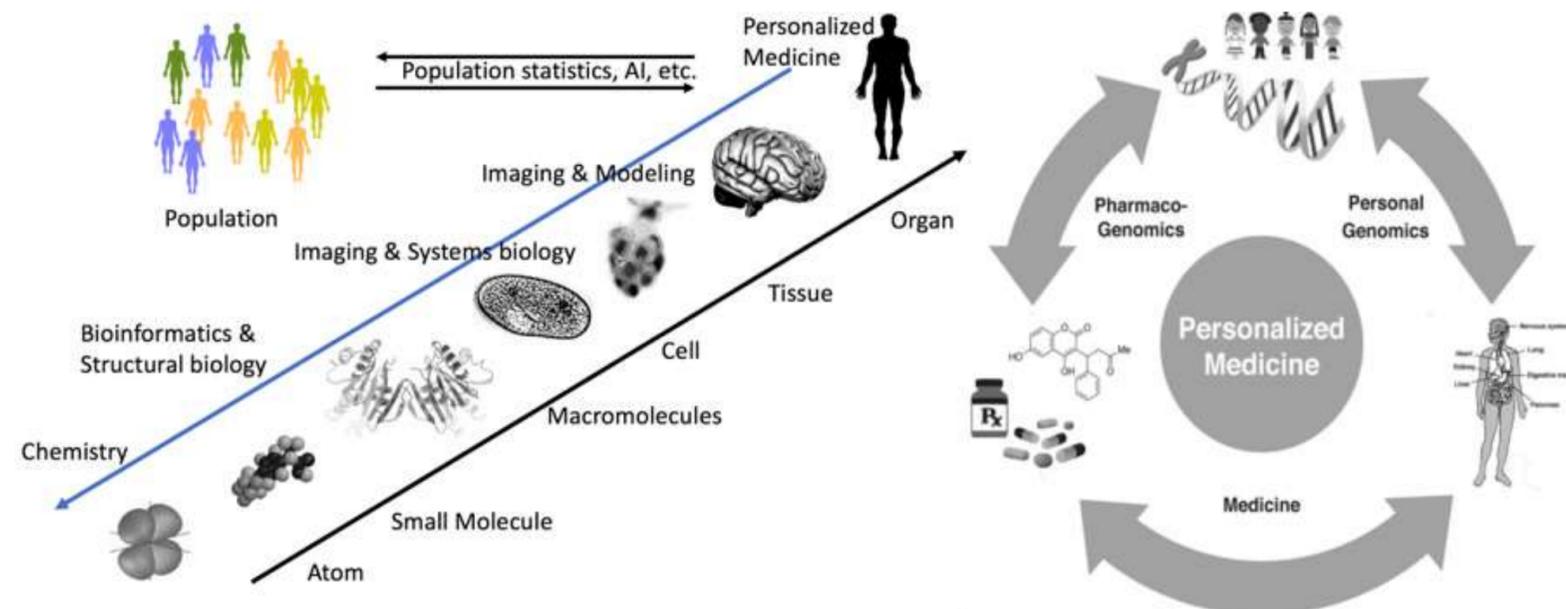
Improving Human Health



The Human Brain EU Flagship project

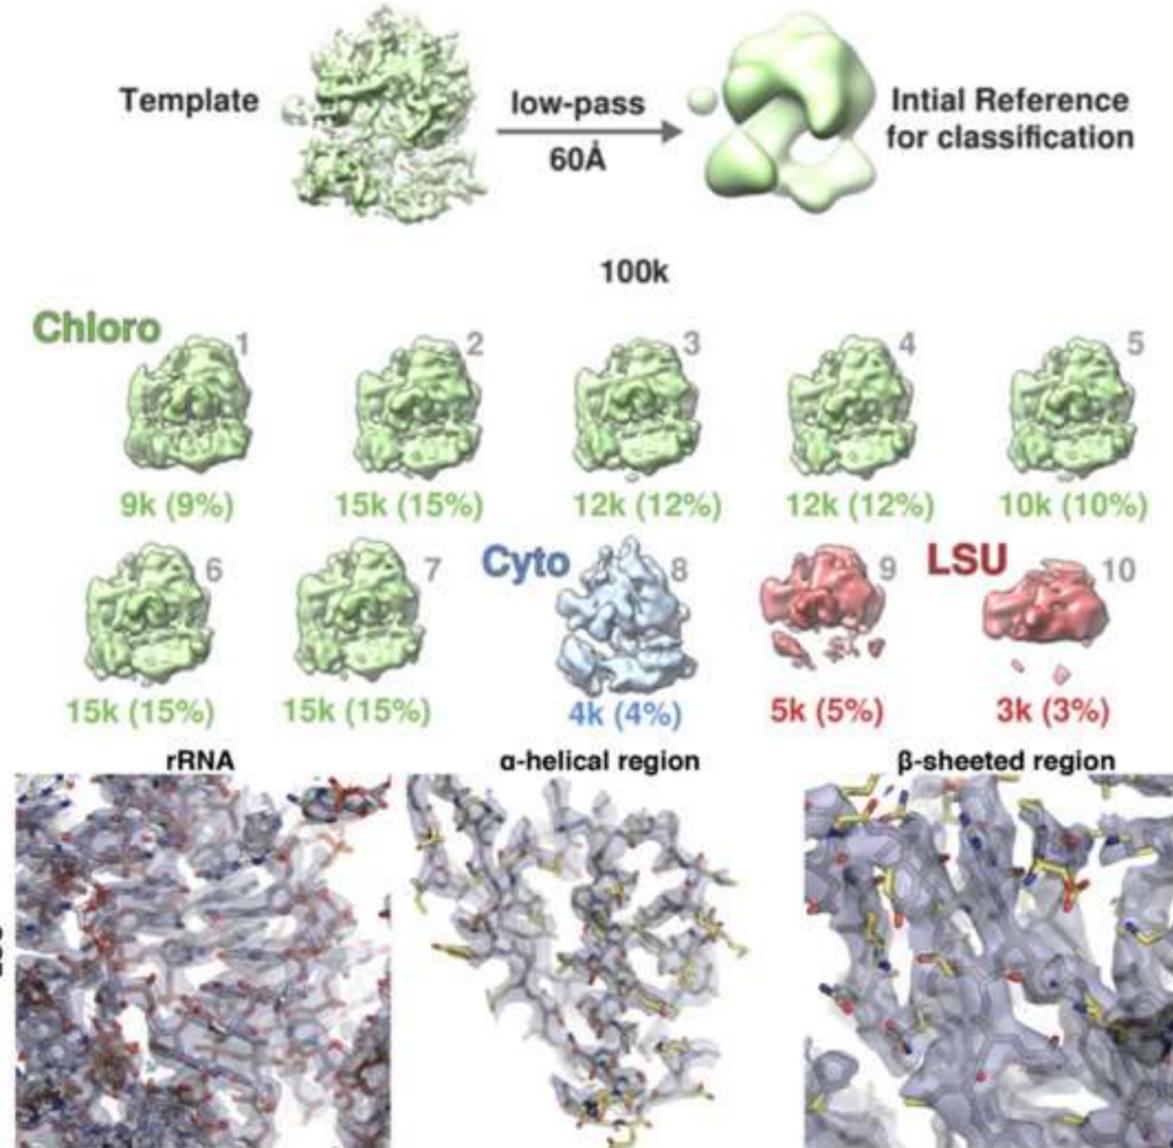
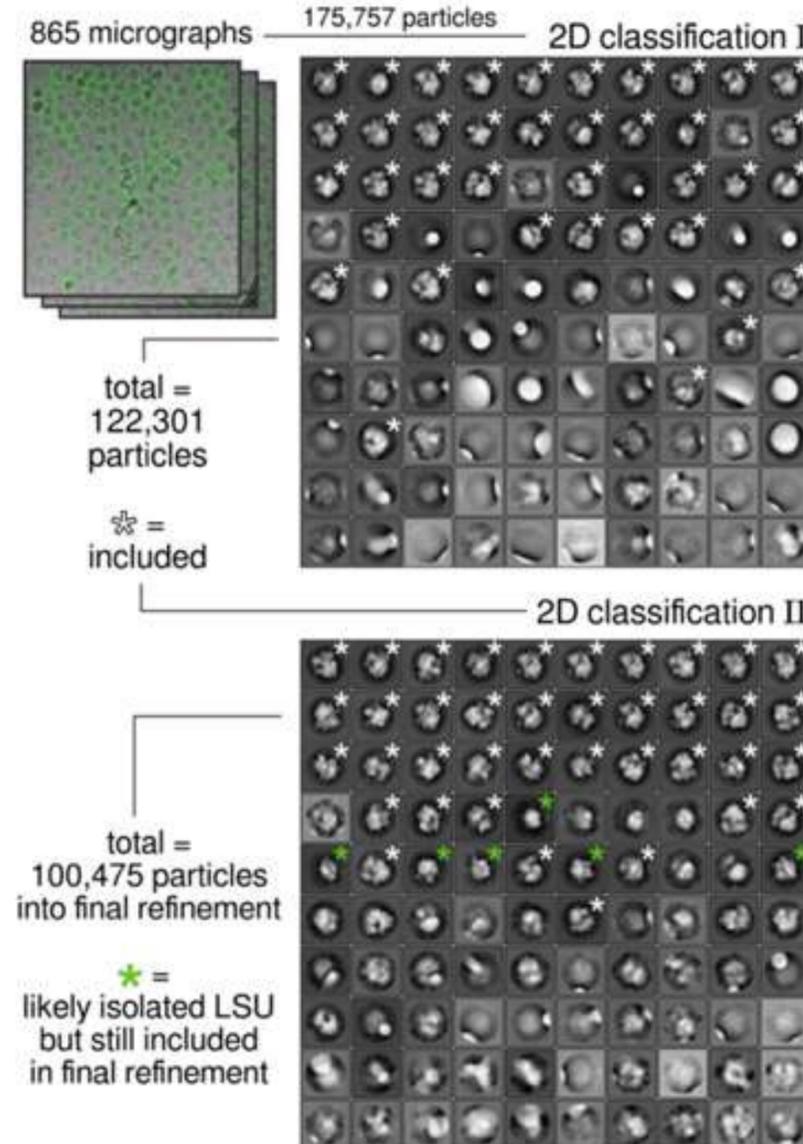
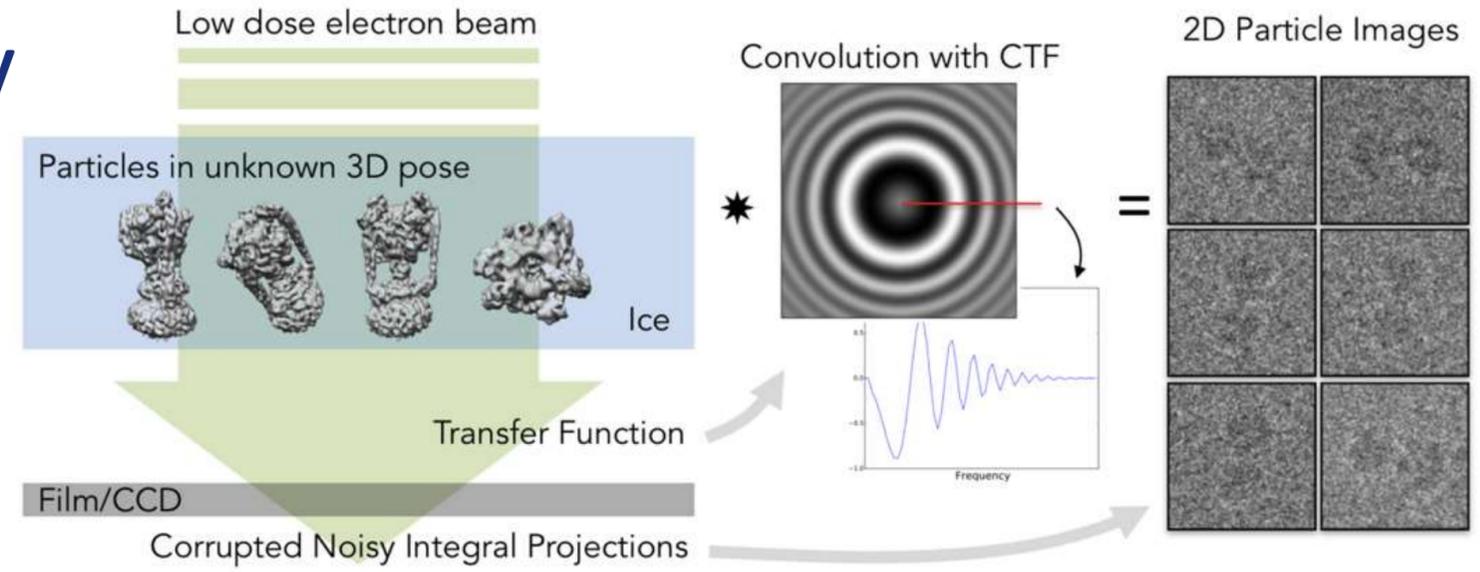


Genomics data grows even faster than computer speed. Future computational resources will enable real-time diagnostics, and predict disease before it happens – preventive care based on broad range of data.

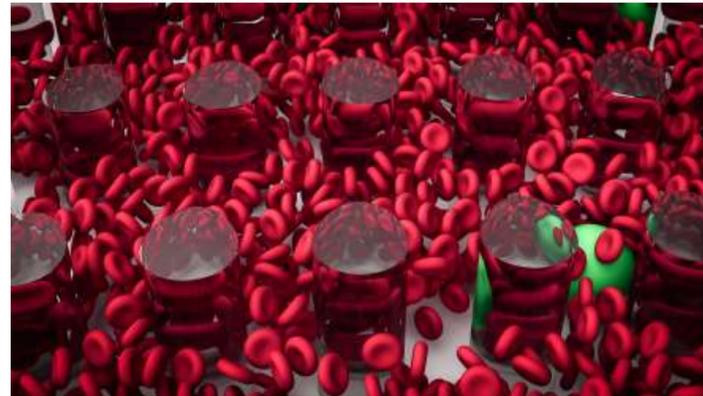


Future Structural Biology is Computational

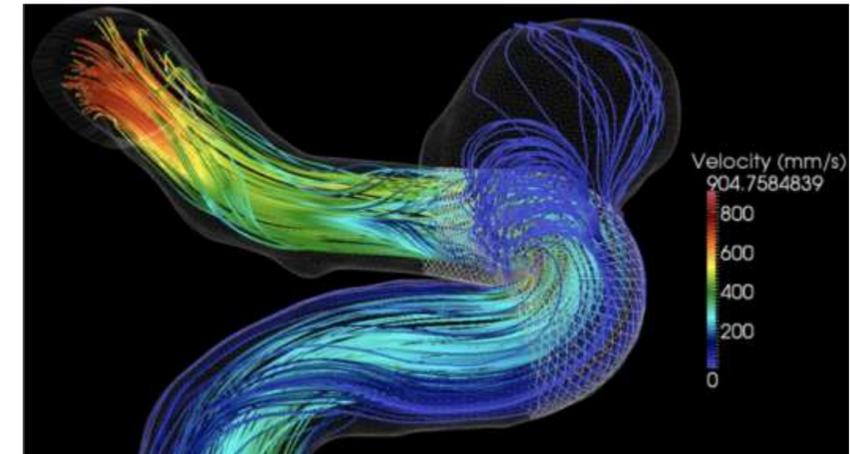
Cryo-electron microscopy (Nobel 2017):
Better computers & algorithms make it possible
to understand how molecules
move, and design new types of drugs



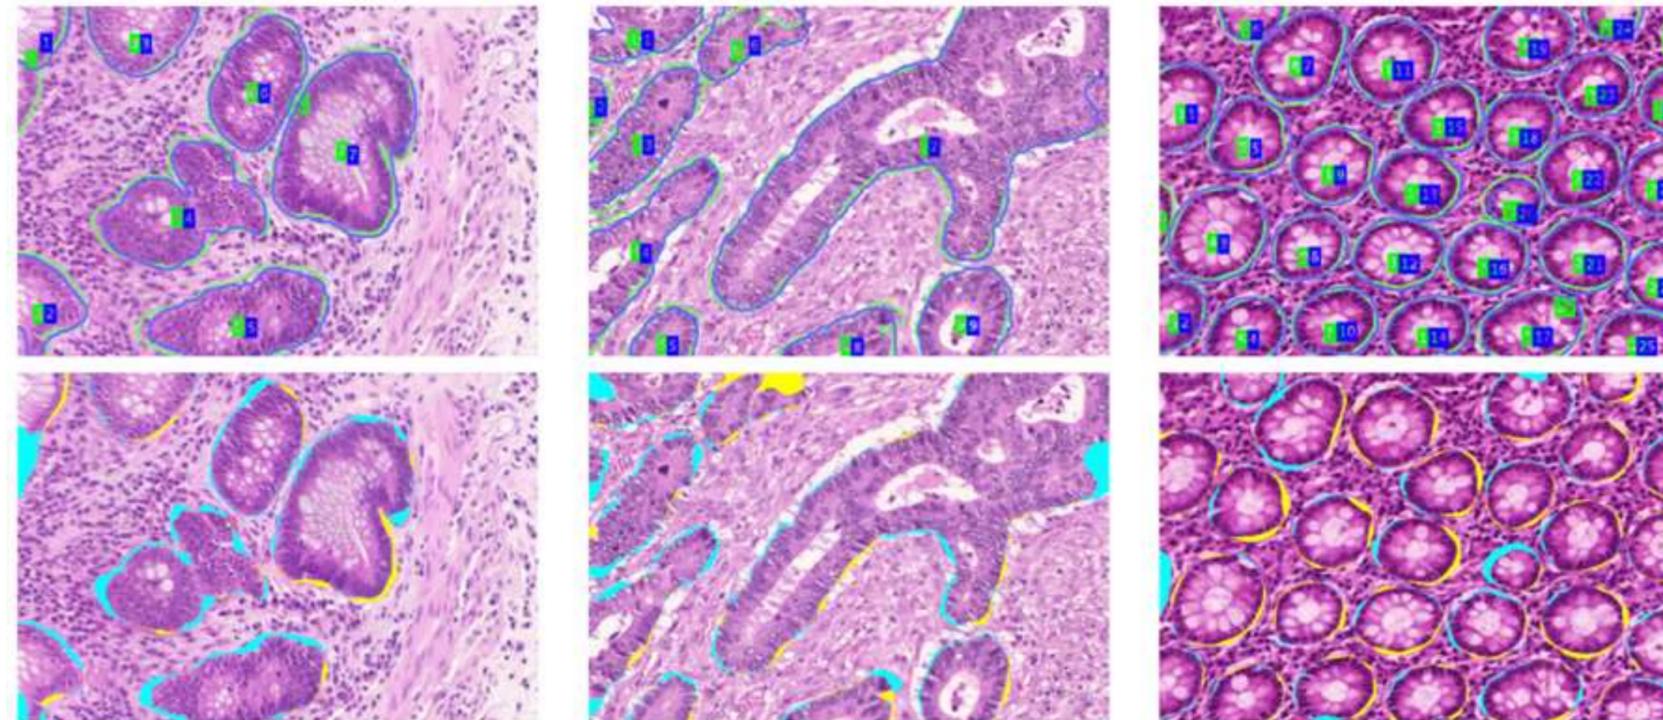
Human Health Computing is Translational



It will become possible to design microfluidics & nanobiotechnology devices to target e.g. specific cancer cells by modeling flow/interactions



Lattice-Boltzmann enable simulations of blood flow, e.g. effects of using a stent to treat brain aneurysm. On-demand computing!



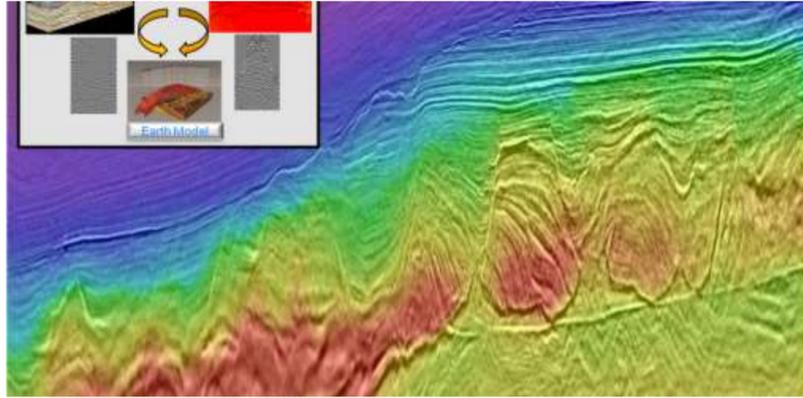
(a) benign

(b) malignant

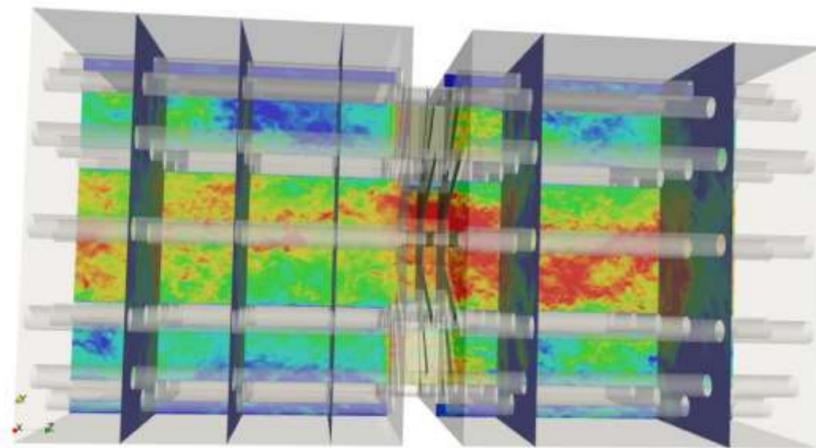
(c) benign

Tumor classification with deep learning neural networks now beats the best pathologists

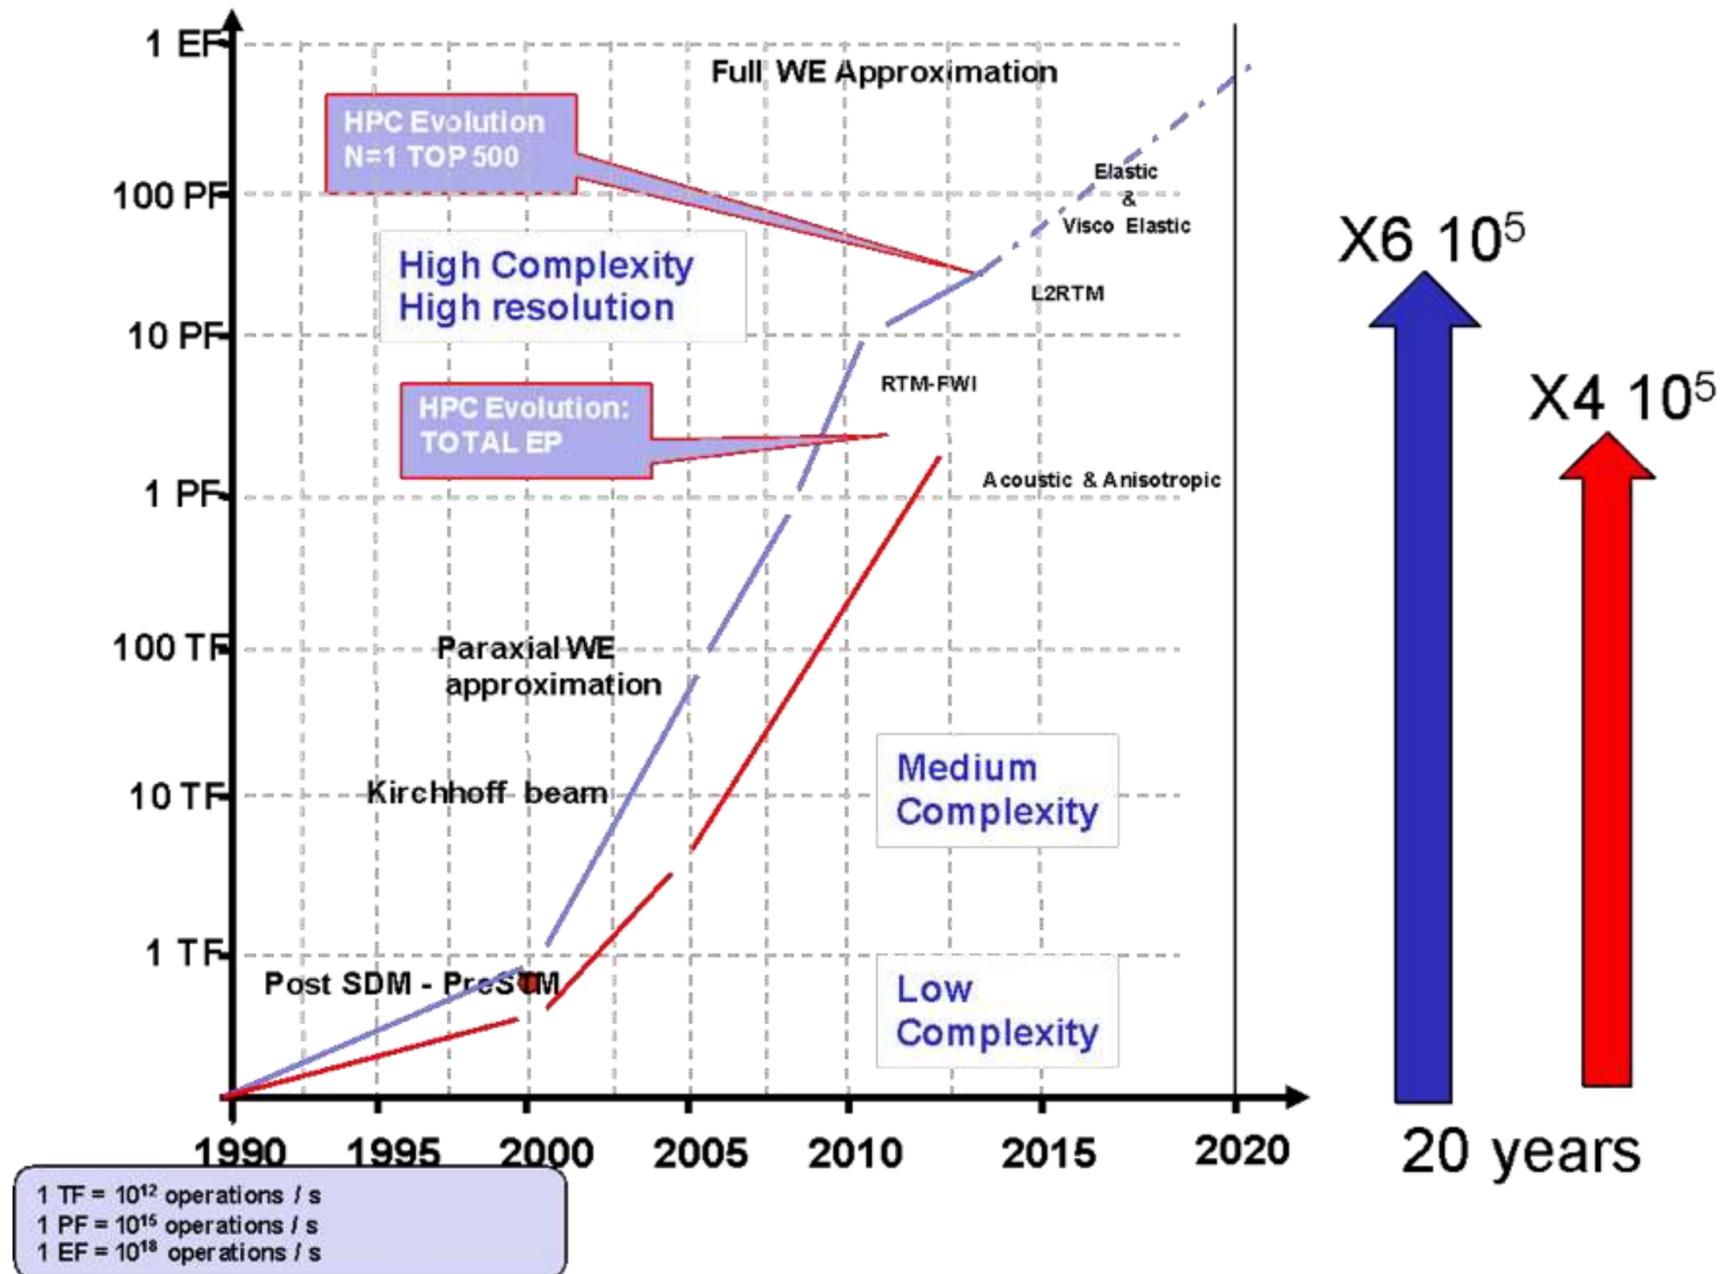
Energy applications need Exascale



Oil & gas is one of the world's largest users of HPC – because computing saves time & money



Fluid dynamics is used to understand heat & flow in nuclear reactors



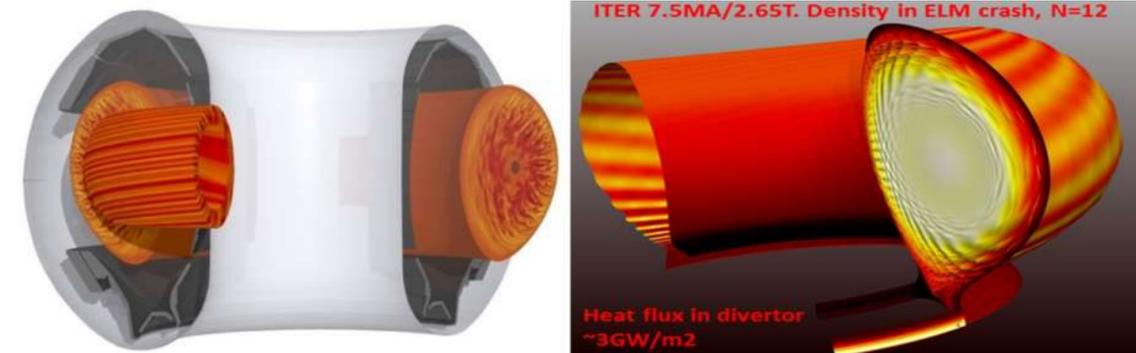
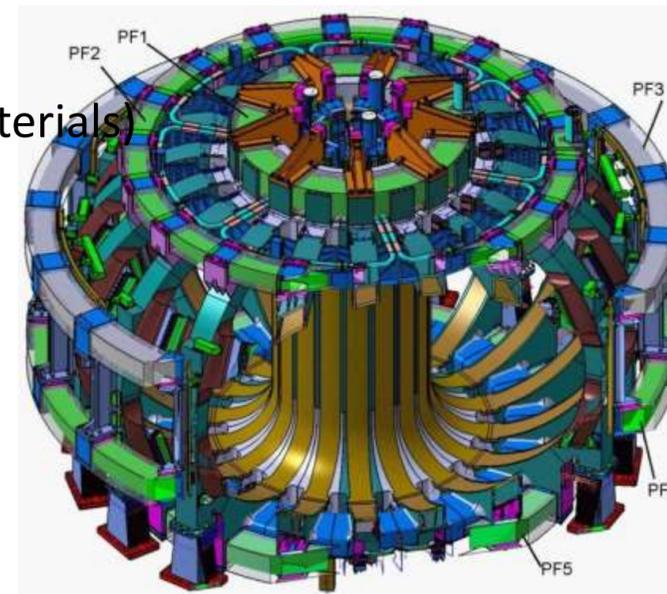
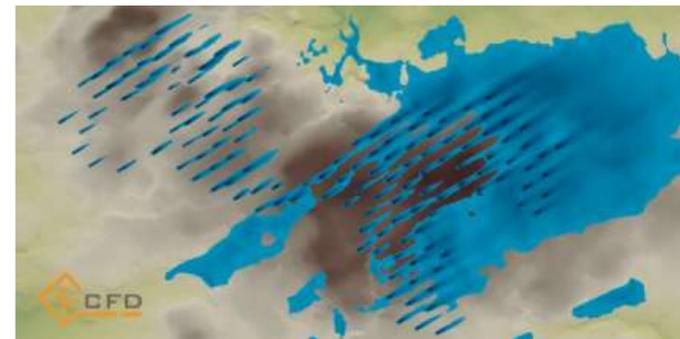
Computing is Driving Renewable Energy



HPC is creating new generations of insulators that enable higher voltage cables, which reduces losses

Optimisation of wing shapes
placement of wind farms

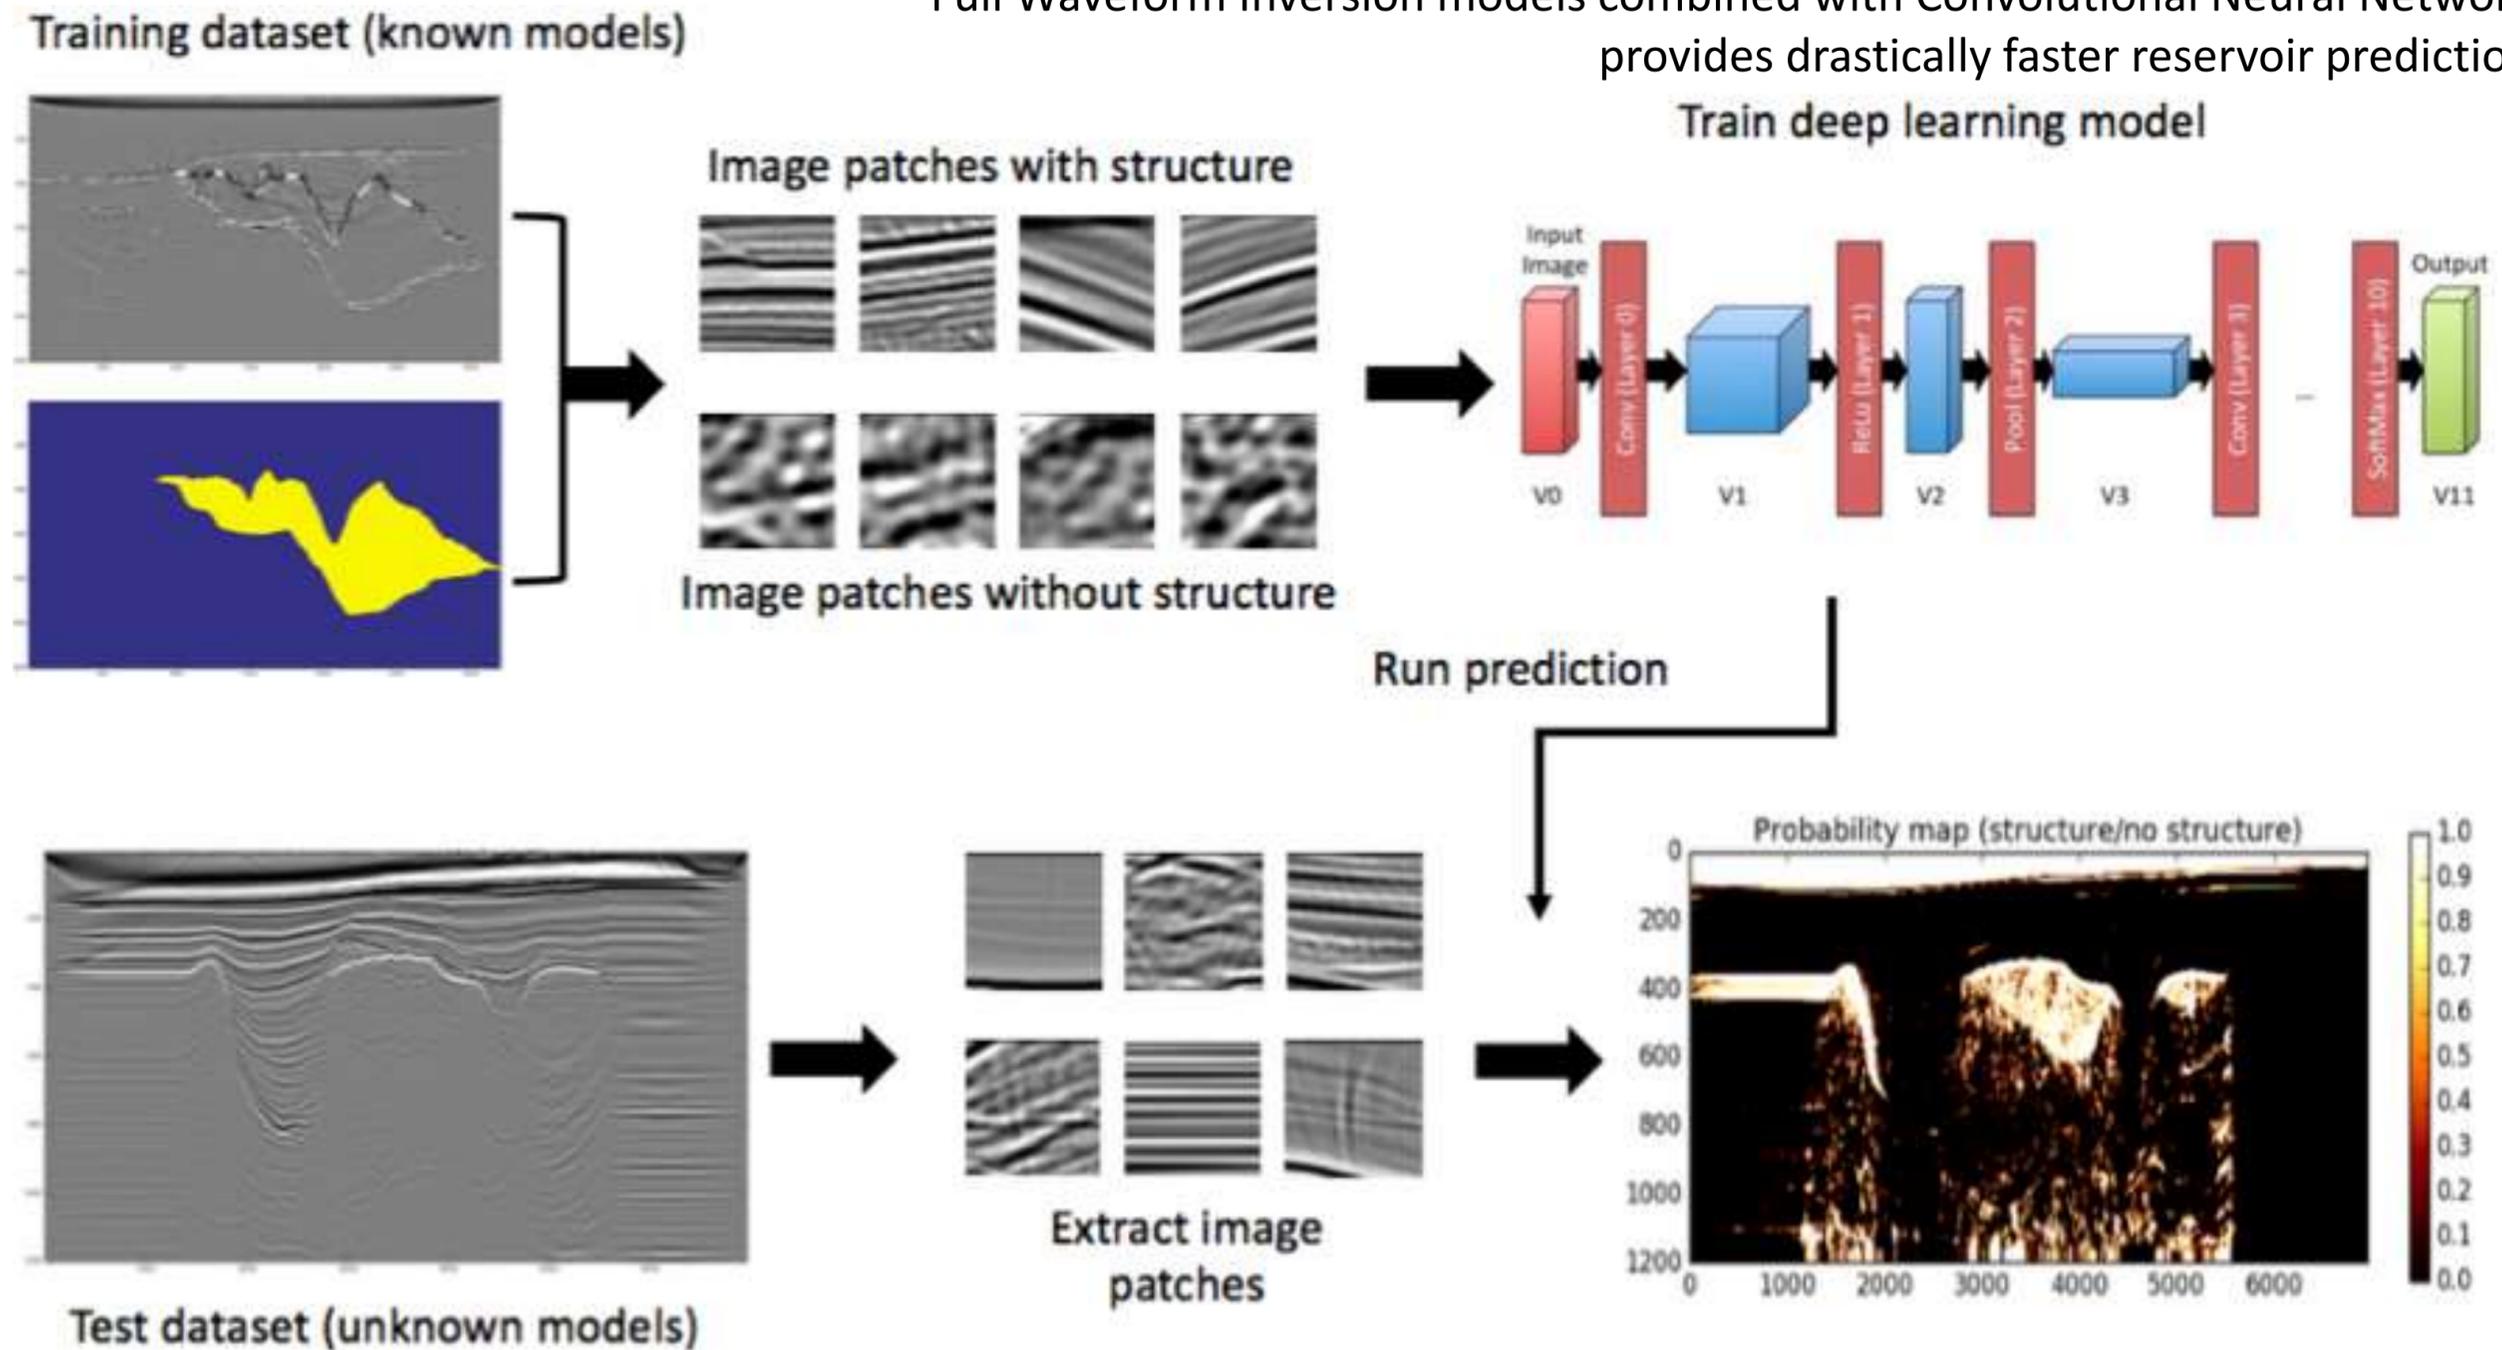
New generations of solar cells (materials)



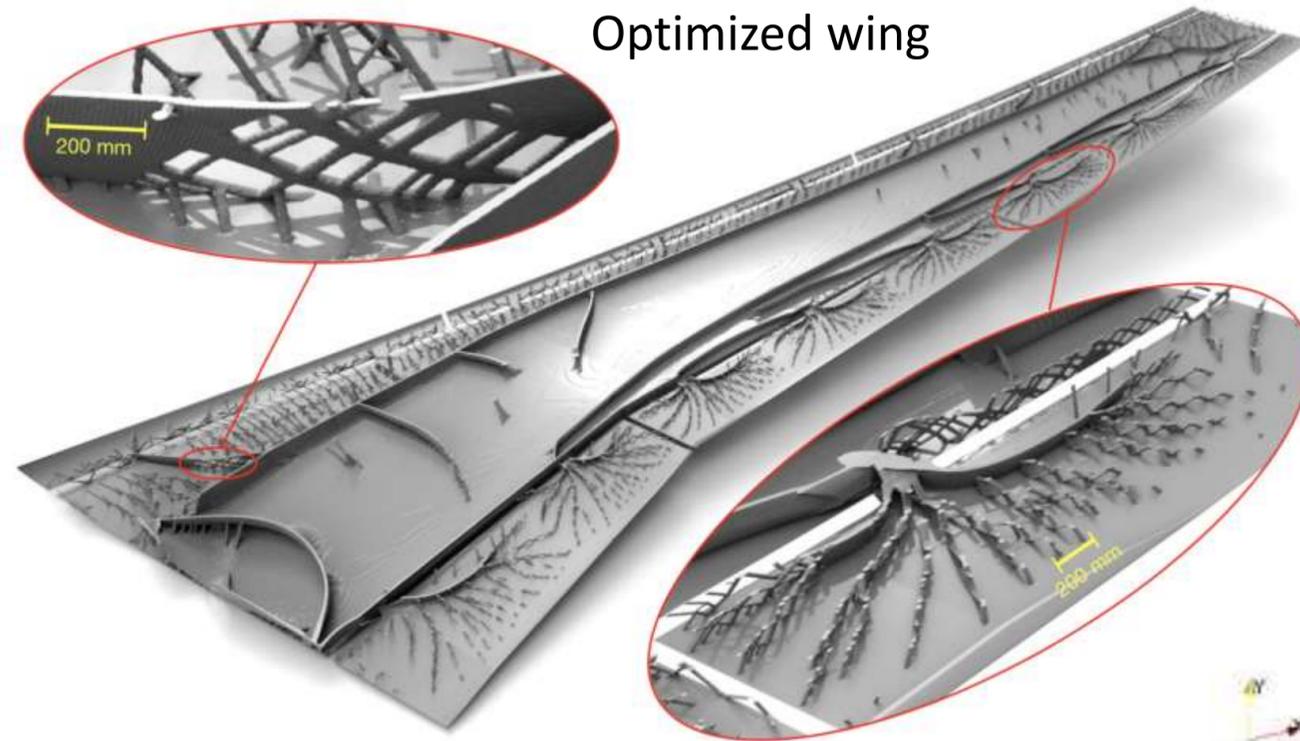
Advanced MHD simulations are critical to make the plasma stable enough to make Tokamaks like ITER useful for energy production through fusion

Exascale is causing even traditional fields to turn to AI

Full Waveform Inversion models combined with Convolutional Neural Networks provides drastically faster reservoir predictions



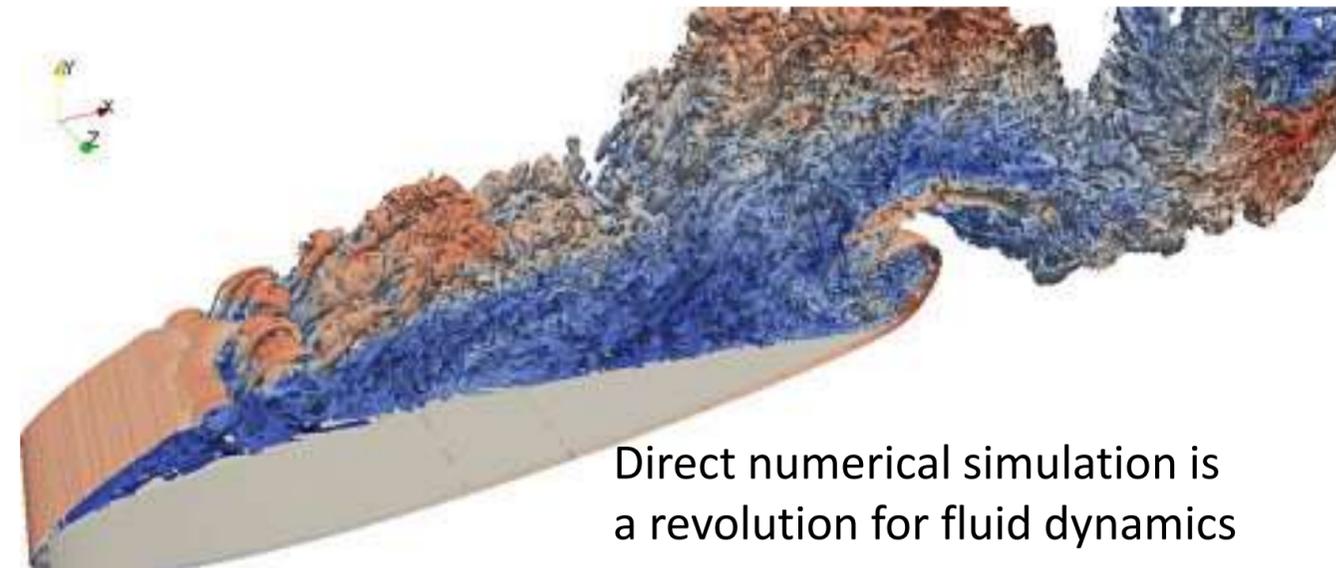
Infrastructure & Manufacturing for Humankind



Aage et al., Nature 550, 84 (2017)
PRACE industrial project

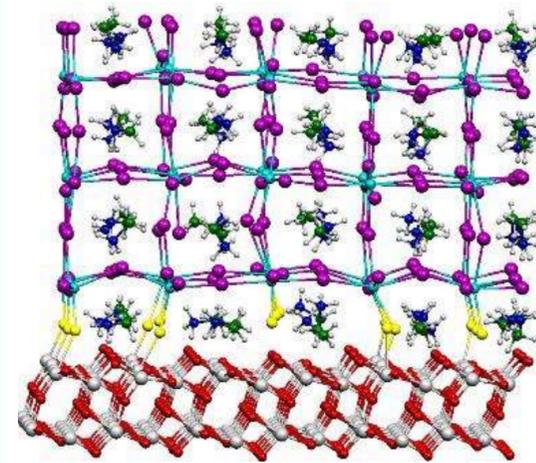
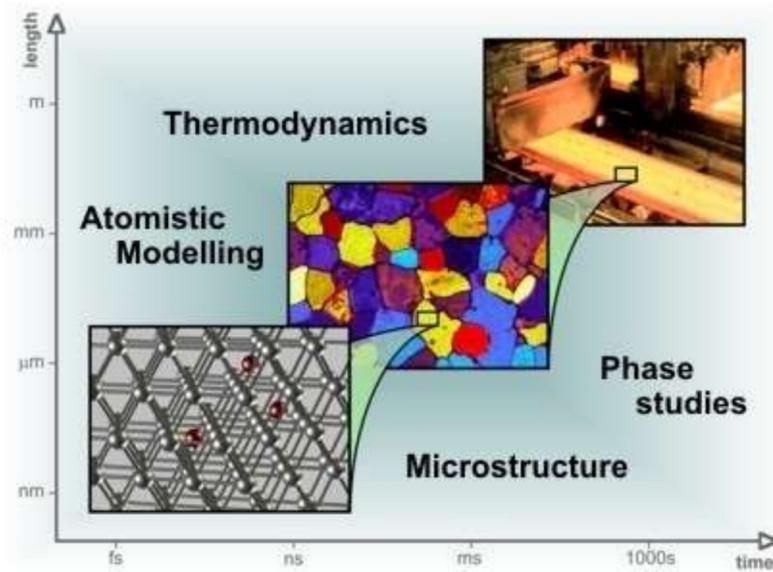


Reach design goals such as:
Stronger, Lighter, Faster
Use less raw material
Cheaper, Safer
Reduce time-to-market

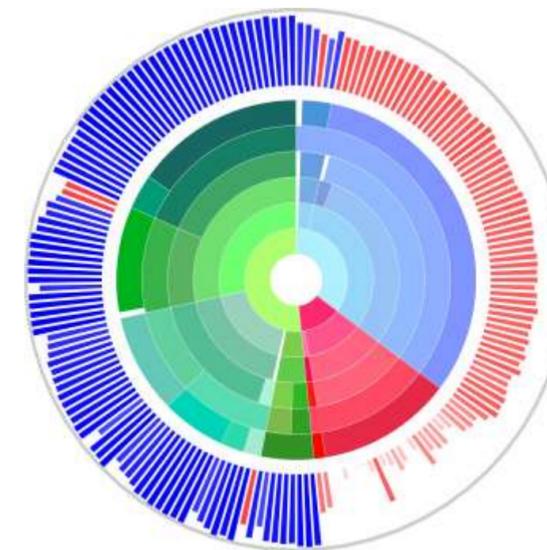
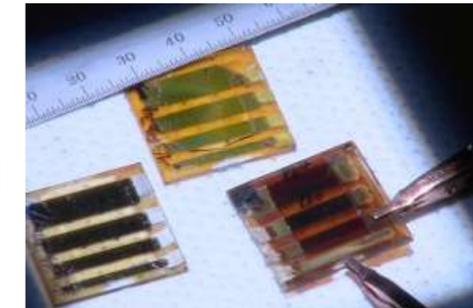


Direct numerical simulation is
a revolution for fluid dynamics

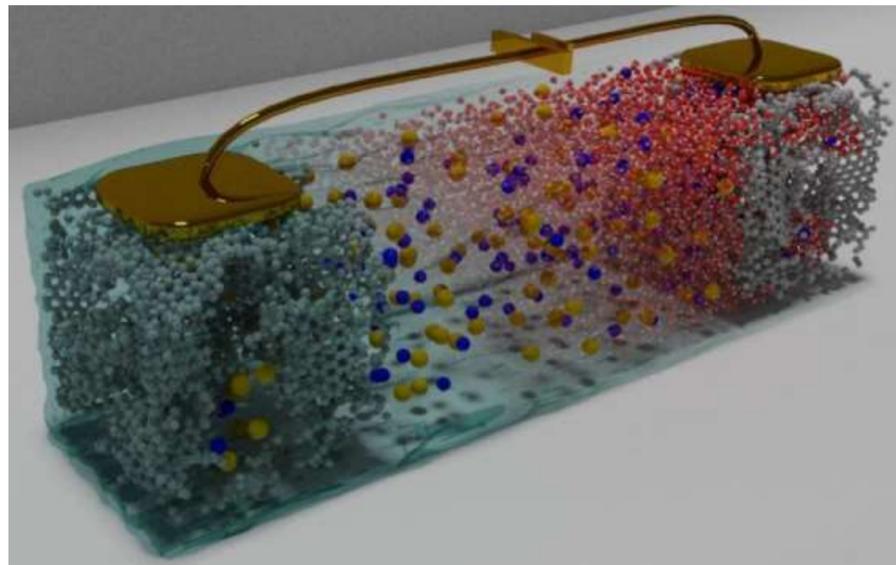
Future Materials: From Molecules to Machines



Perovskite solar cells have outstanding efficiency, but contain lead. Simulations used to develop lead-free versions

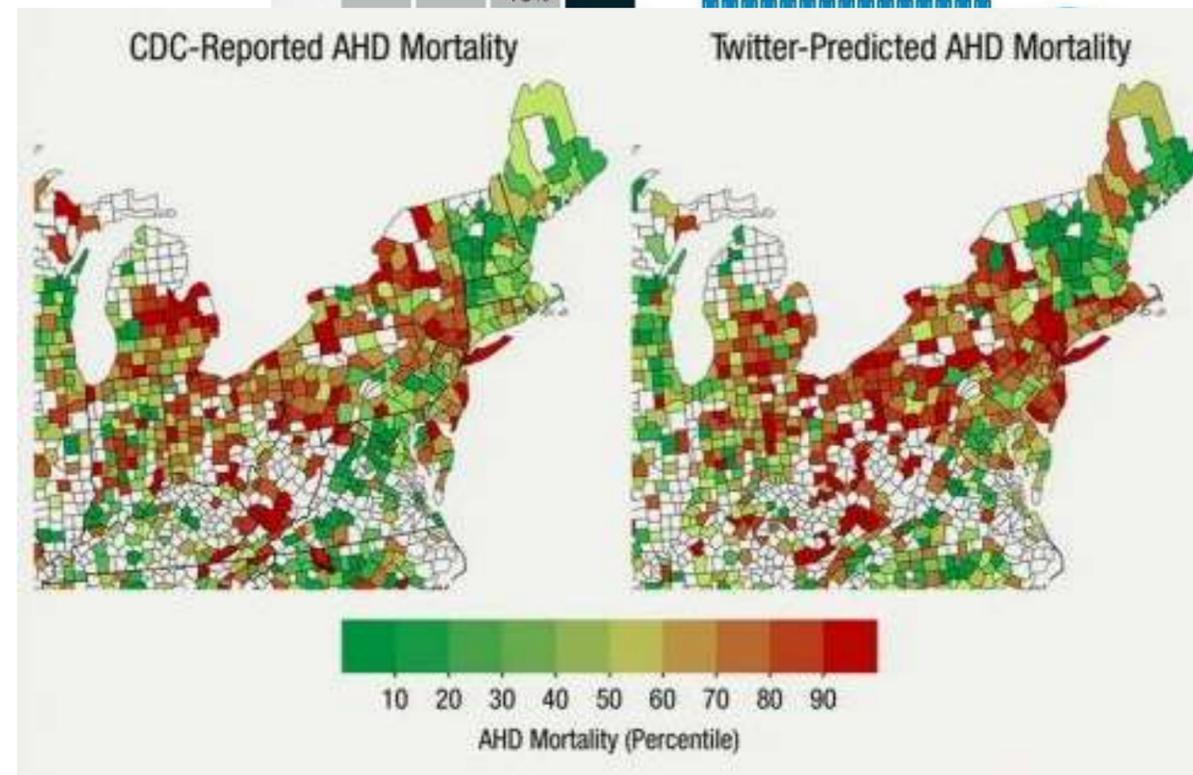
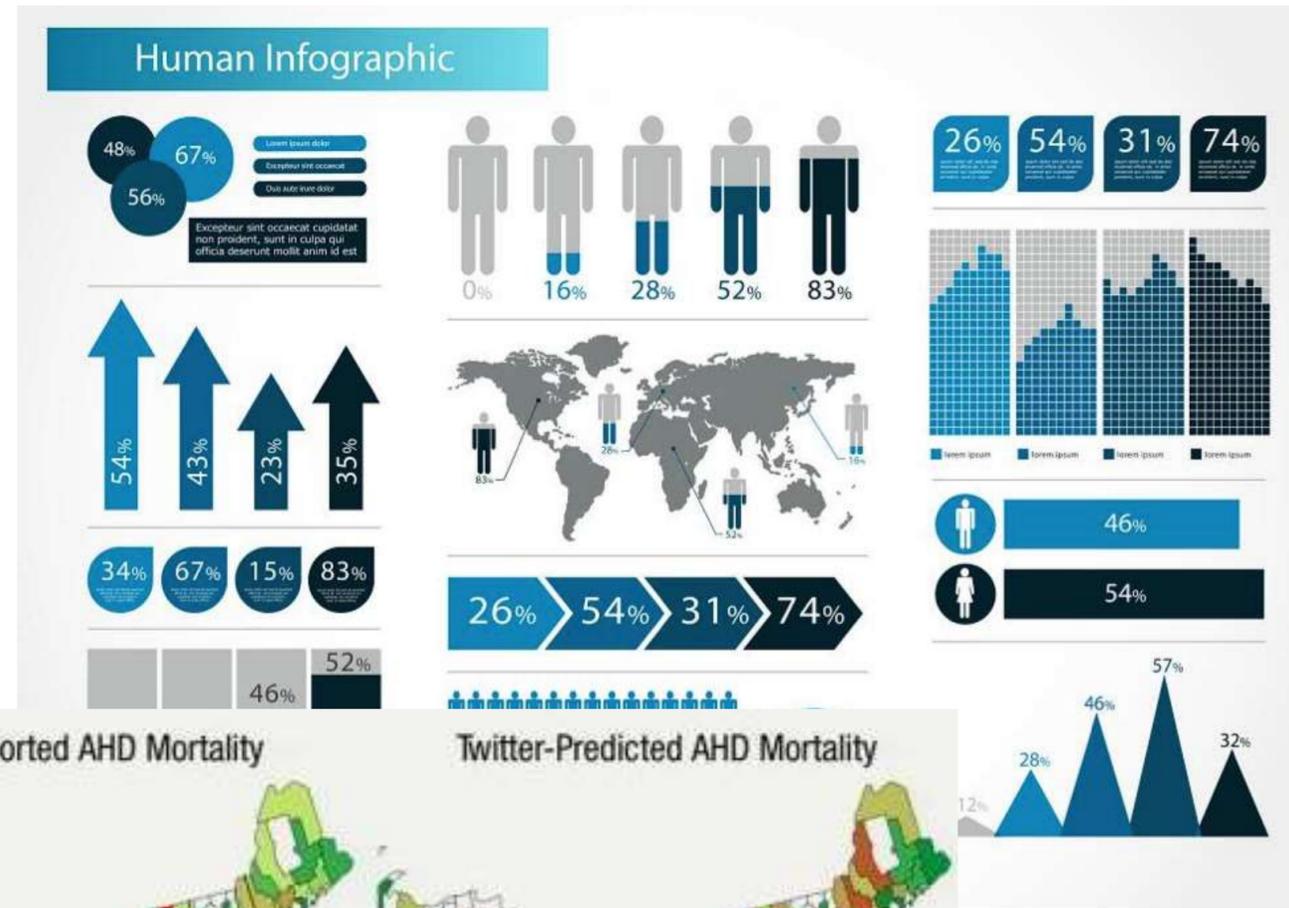
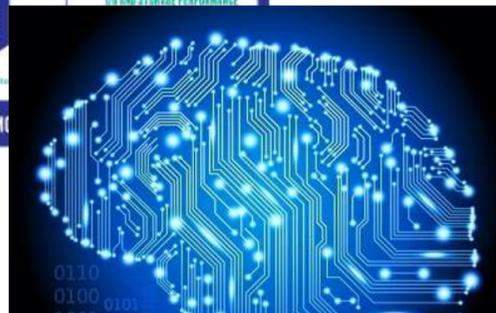
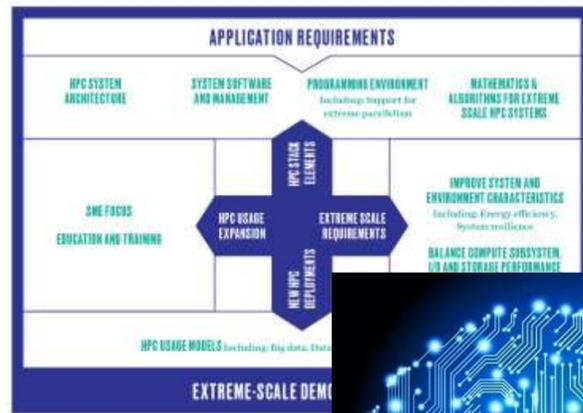


Modern materials design is increasingly data-driven, using AI/deep learning methods to rapidly predict properties without full DFT calculations



Supercapacitors are used to store energy e.g. in braking kinetic-energy recovery systems in cars/buses. Molecular simulations are used to understand the charge build-up to improve their capacity and efficiency

Complexity & Data



Centres of Excellence in HPC



ChEESE

Center of Excellence for Exascale in Solid Earth



esiwace

CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

FocusCoE



HiDALGO



Key Infrastructure Requirements

- Urgent need for more compute cycles - but also high demands for memory bandwidth & I/O
- Scientific Applications are not only ready to use Exascale, but have specific high-impact research goals that can only be realized with access to next-generation resources
- However: Exascale will require new approaches to scaling by using ensembles, deep learning, and statistical models. High-end computing will dominate, but tomorrow's solutions will be different than yesterday's
- Exascale systems will need to be able to handle tens of thousands of active jobs and large I/O requests
- Software & algorithms take longer to change than hardware - first-generation systems must be based on present concrete needs, rather than hopes about being able to co-design
- Linux rules the modern computing world: All systems will need to support a full Linux stack of development tools
- Centers & their staff need to engage directly in software development together with scientific communities

Aside from the HPC efforts, the European Commission is increasing its annual investments in AI by 70% under the research and innovation programme Horizon 2020. It will reach EUR 1.5 billion for the period 2018-2020.

Now How to Transfer Knowledge to SMEs?

National HPC Competence Centres
Education and Training
Next generation portals and domain workflows
Easier access to Computing cycles

MELUXINA: A WORLD CLASS HPC SYSTEM

Compute

18 Petaflops (10^{15})

Peak aggregated performance
Heterogeneous supercomputing

Compute intensive workloads Excellent for
simulation, data analytics & AI

Faster than
17,000 x 

Storage

25 Petabytes (10^{15})

Tiered HPC storage, excellent for HPDA & AI
+ Long Term Archival capabilities

Massive data sets
Perfect for innovative data exploration

More than
20,000 x 

Networking

400 Gigabits/s (10^9)

Infiniband HDR links on every MeluXina node +
High speed external connectivity

Large scale workloads
Lightning-fast data processing

Better than
768,000 x 

It is like running a full marathon in 0.4 seconds or flying from Paris to London in 0.1 seconds.

The *MUSE*: MeluXina User Software Environment

- Operating System: **Linux**
- Rich software environment with Open Source & commercial tools
- Ease of provisioning and use of provided software
- User can install their own software
- ... or use containers

Compilers, Languages & Performance Eng.	Parallelization tools, MPI suites & acceleration libraries	Numerical & data libraries	Frameworks, runtime & platform tools	End user applications
<ul style="list-style-type: none"> • AOCC • GCC • Intel • NVIDIA HPC SDK <ul style="list-style-type: none"> ✓ incl. PGI • Support for various programming languages <ul style="list-style-type: none"> ✓ Python ✓ R ✓ Julia ✓ Go ✓ Rust • Performance and debugging tools <ul style="list-style-type: none"> ✓ Intel ✓ NVIDIA ✓ ARM (Forge) ✓ Scalasca ✓ SCORE-P ✓ Extrae ✓ PAPI ✓ Valgrind ✓ GDB ✓ AMD-uProf ✓ Nsight-Systems ✓ Nsight-Compute ✓ Vtune ✓ gperf ✓ extrap ✓ Inspector • Many build & support tools (Autotools, CMake, ...) 	<ul style="list-style-type: none"> • OpenMPI • Intel MPI • ParaStationMPI • NVHPC • TBB • PETSc • KOKKOS • cuBLAS • cuFFT • cuDNN • NCCL • TensorRT 	<ul style="list-style-type: none"> • BLIS • Intel MKL • FFTW • OpenBLAS • ScalaPACK • Boost • Eigen • ARPACK • HDF5 • netCDF • OpenCV • CDO 	<ul style="list-style-type: none"> • PyTorch • Torch Text+Vision • TensorFlow + Hub • Horovod • Keras • Theano • Jupyter Lab • Apache Spark • Matlab Runtime • dotNET Core + SDK • Dakota • Quantum Computing/AI <ul style="list-style-type: none"> • Cirq • QsimCirq 	<ul style="list-style-type: none"> • GROMACS • OpenFOAM • FOAM-Extend • CP2K • QuantumESPRESSO • NAMD • QMCPACK • NWChem • HOOMD-blue • Freud-analysis • DualSPHysics • POV-Ray • Blender • MDAnalysis • BioPython • QUDA • Visualisation <ul style="list-style-type: none"> ✓ ParaView ✓ VMD ✓ OVITO ✓ NCO ✓ NCView

Over 300 tools pre-installed on MeluXina today



MELUXINA: DEMOCRATIZING HPC & HPDA

```

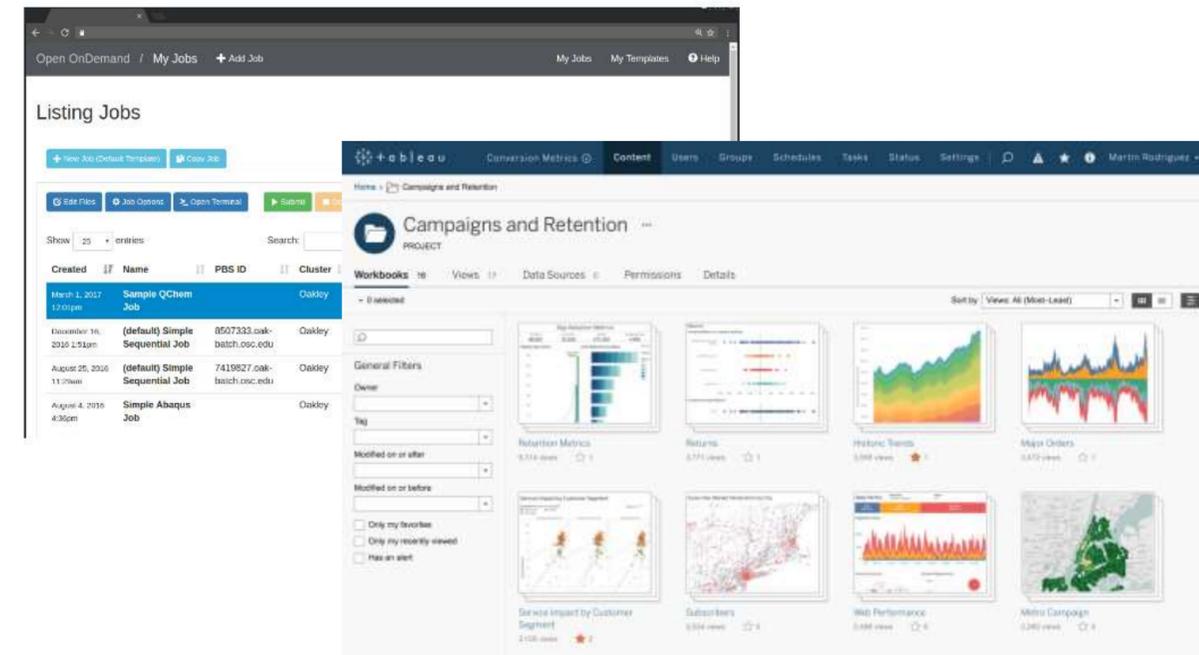
Welcome the Luxembourg - EuroHPC supercomputer

Meluxina

-----
You are on a MeluXina login node

-----
System information: Compute
-----
Nodes | CPU | RAM | Accelerator | Disk
-----|-----|-----|-----|-----
573N | 2x AMD 7H12: 128c @2.6G | 512GB | - | -
200N | 2x AMD 7452: 64c @2.3G | 512GB | 4x NVIDIA A100-40 | 1.92G
20N | 2x AMD 7452: 64c @2.3G | 512GB | 2x Intel Stratix 10MX-16 | 1.92G
20N | 2x AMD 7H12: 128c @2.6G | 4096GB | - | 1.92G
-----
System information: Data
-----
Tier | Capacity | Speed | Type | Location on compute/login
-----|-----|-----|-----|-----
Scratch | 0.5PB | 400GB/s | NVMe | /project/scratch
Home/Project | 12.5PB | 180GB/s | HDD | /home/users, /project/home
Backup | 7.5PB | 30GB/s | HDD | -
-----
System information: Interconnect
-----
Fabric: Infiniband HDR, 200Gbps, DragonFly+ topology
Links : 1x on CPU nodes, 2x on GPU, FPGA & LargeMemory nodes
-----
System information: Software
-----
Production software stack: 2021.2
Modules system: LMod, use `module av` to discover the environment
-----
Center information
-----
News & Events : luxprovide.lu
Documentation : docs.lxp.lu
Support : servicedesk.lxp.lu, servicedesk@lxp.lu
-----
LinkedIn & Twitter : @luxprovide #meluxina @EuroHPC_JU

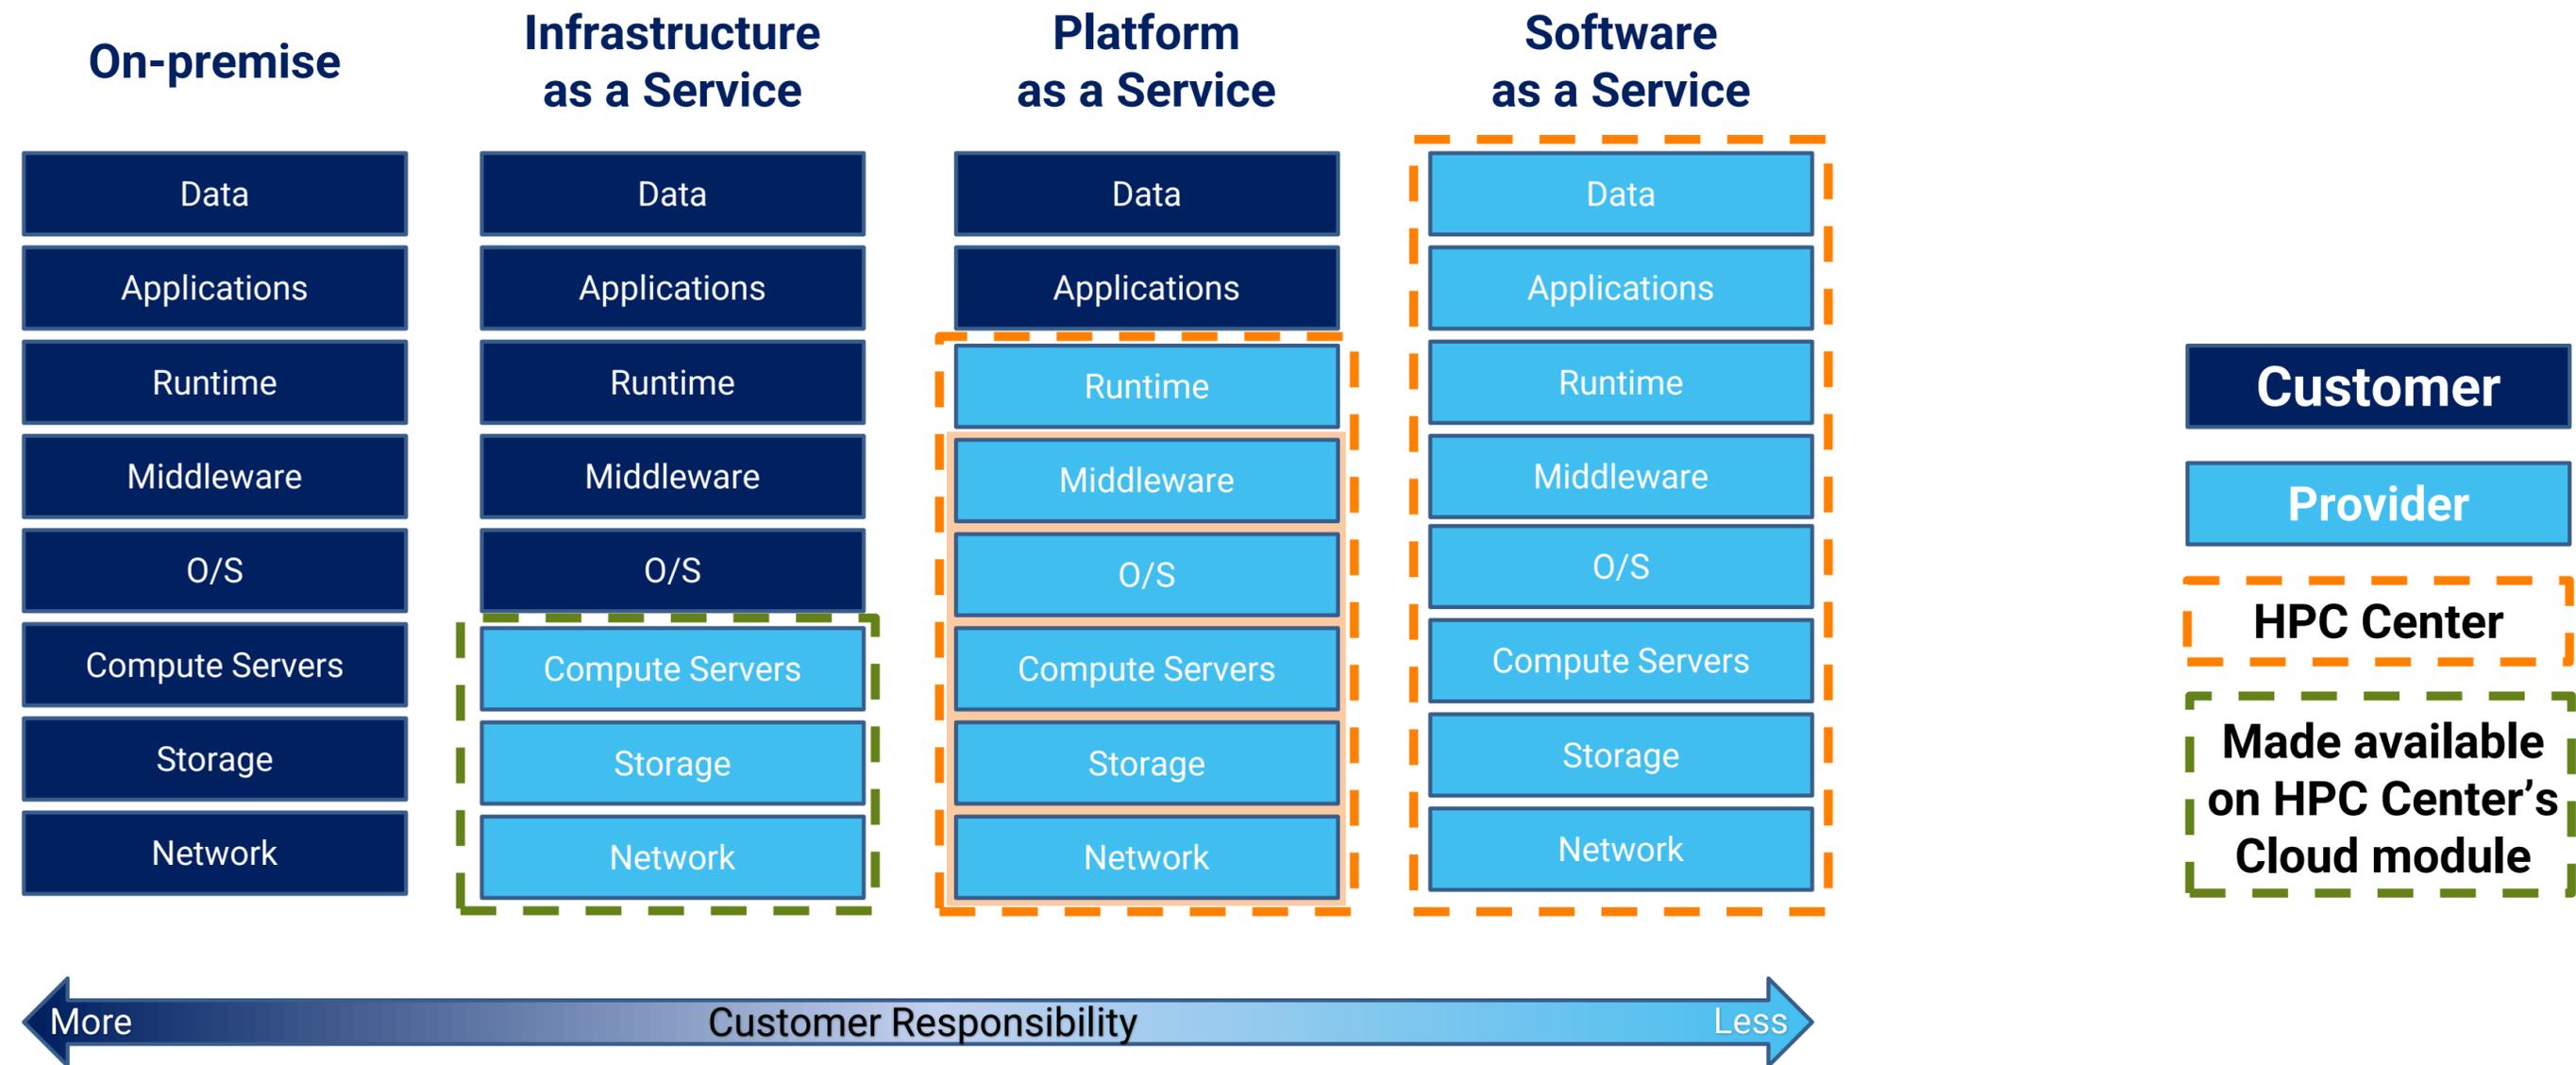
```



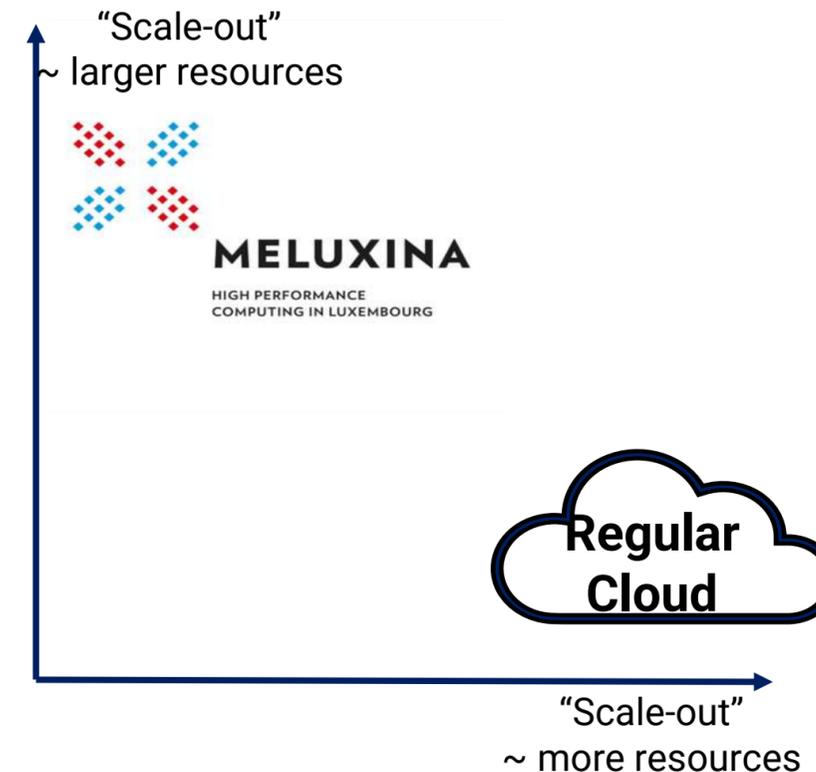
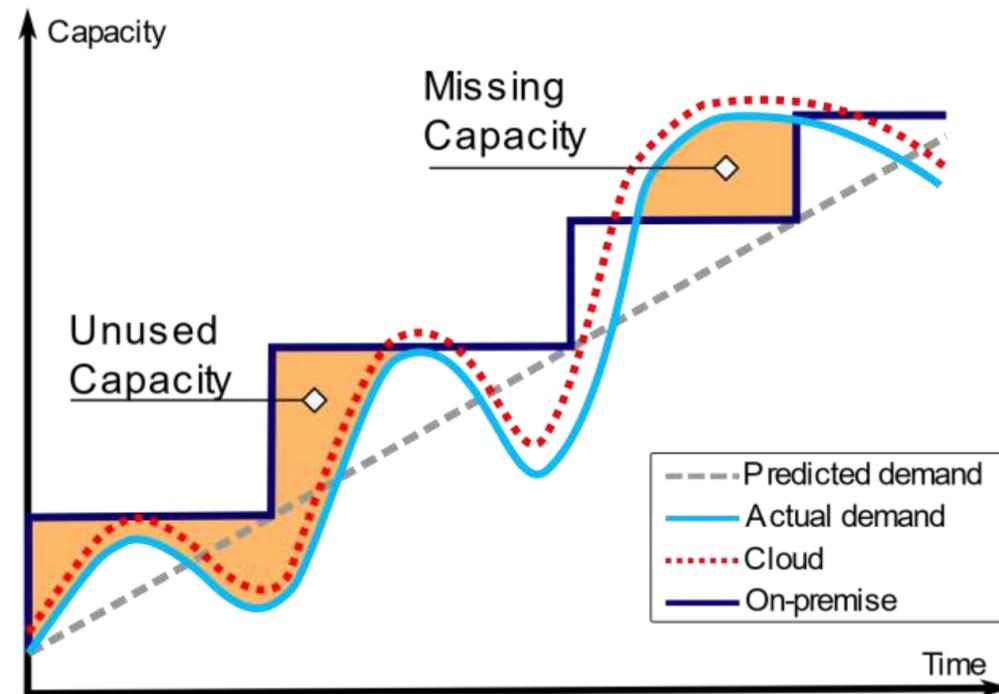
MeluXina's Cloud Module:

- User Portals
- Dedicated VMs for special workloads
- Data-Scientist friendly AI & ML platforms
- Data streaming for IoT
- Persistent databases
- ...

CLOUD RESPONSIBILITY MODES & HPC



CLOUD & HPC



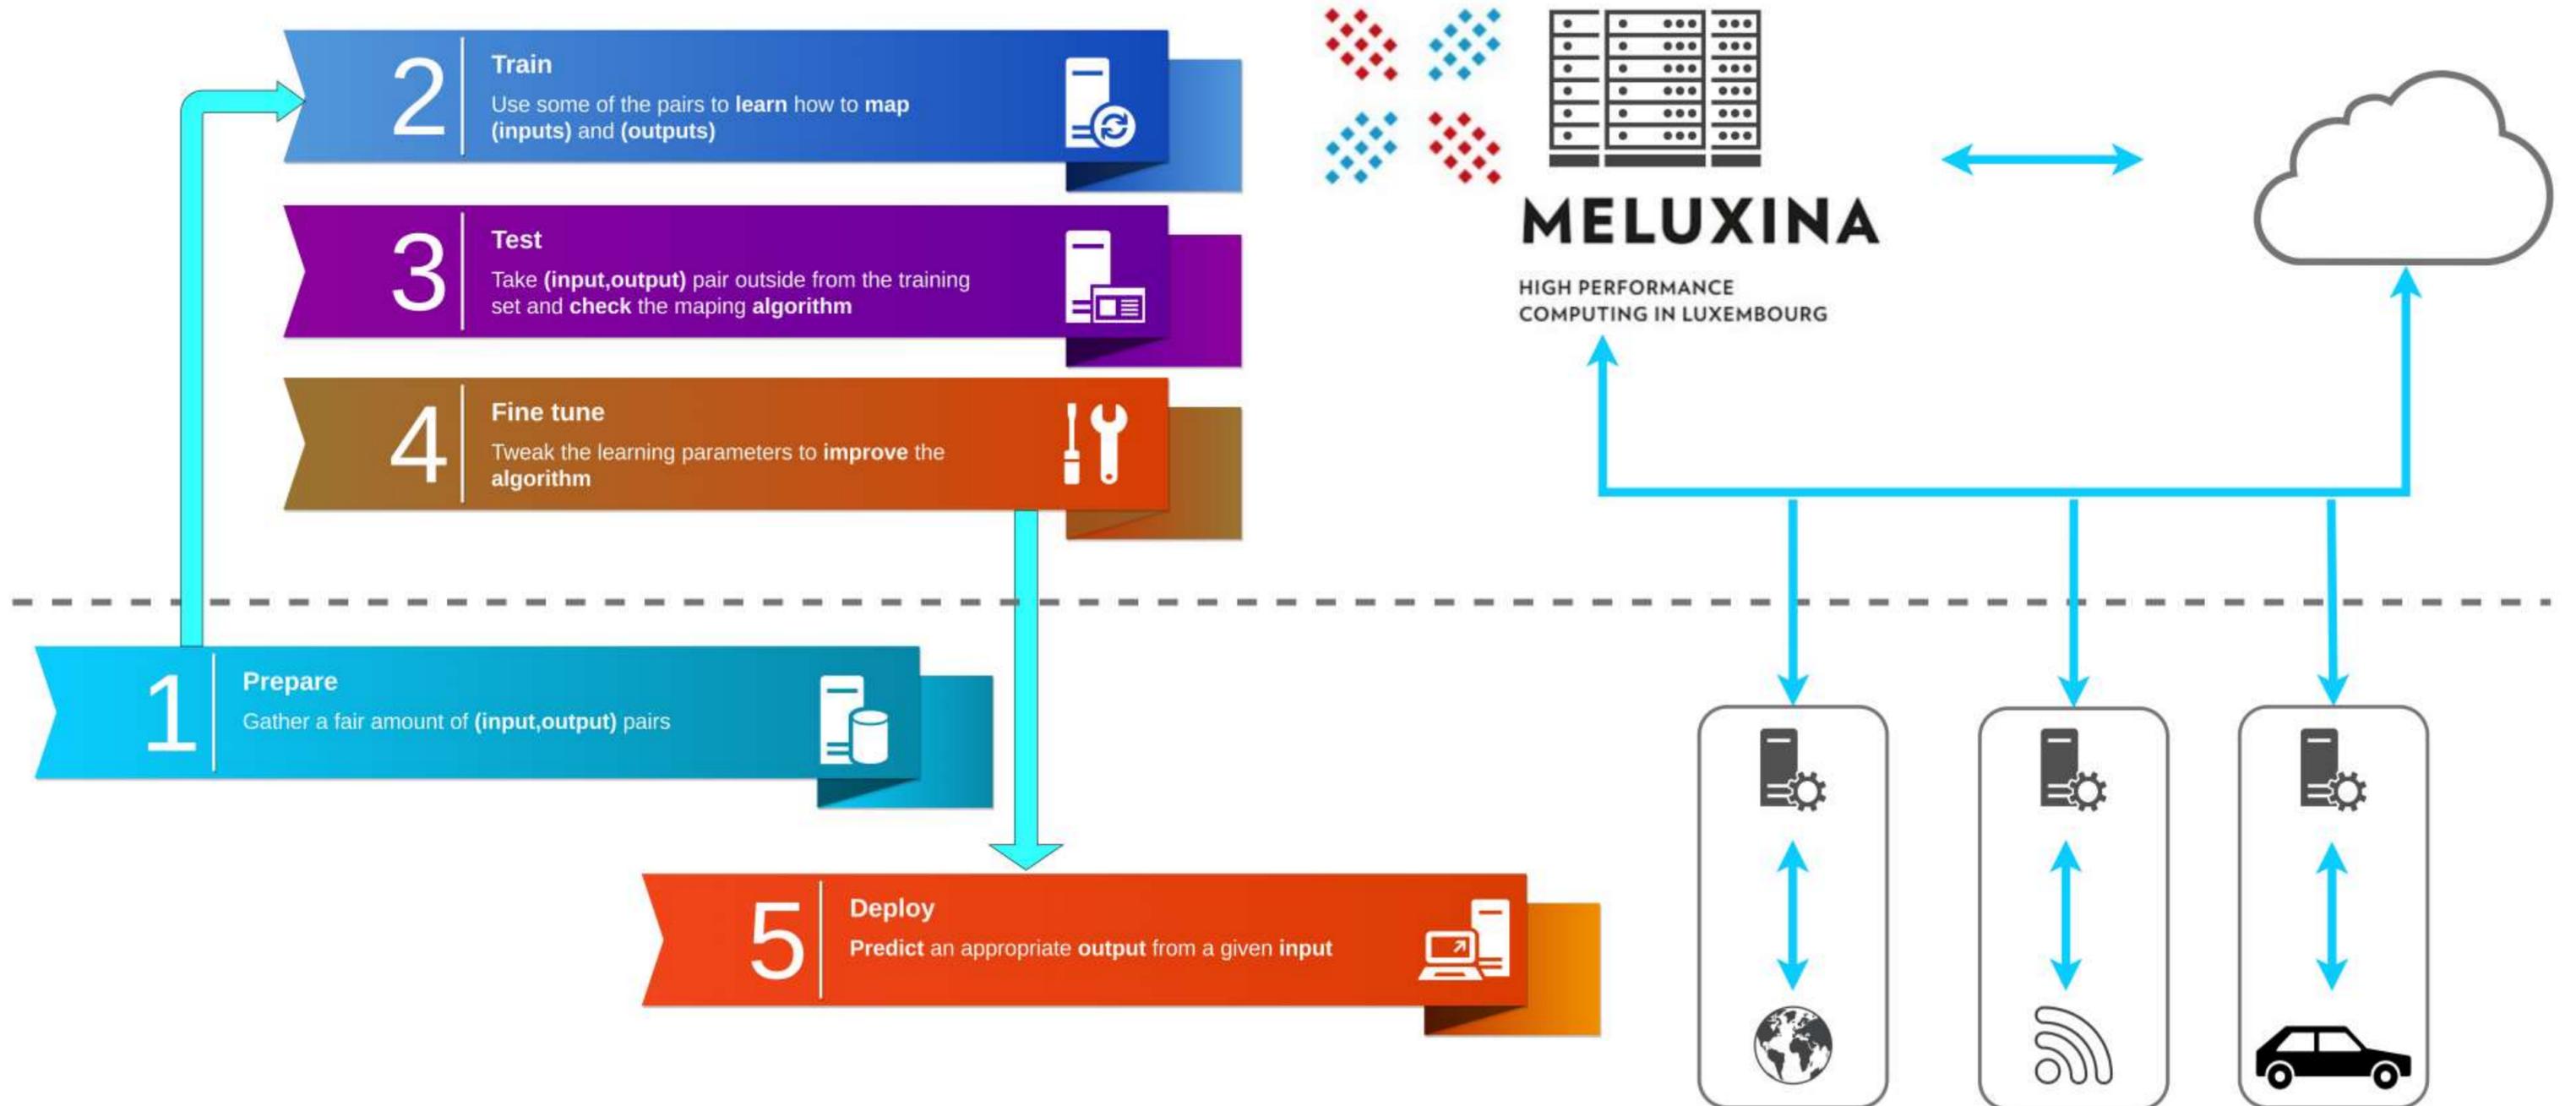
Similarities with "regular" cloud

- On-demand compute and storage
- Pay-per-use model
- Operated by a third party
- Huge but limited resources

Differences with "regular" cloud

- Mostly "batch" based (i.e., not real-time)
- Unified machine
- Located in a well identified place

CLOUD, EDGE, IoT & HPC

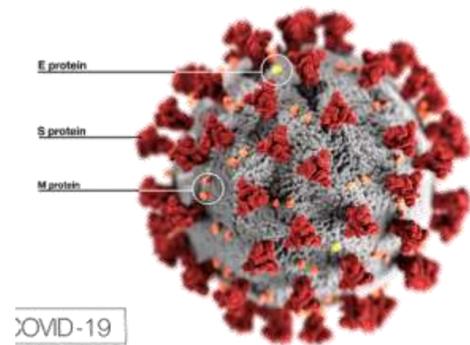


CREATING COMPETITIVE ADVANTAGES

Through High-Performance Computing

Improve time to market

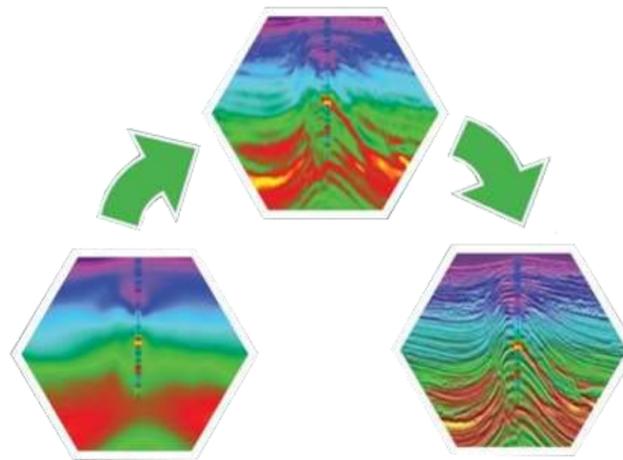
Beat competition with new and better products
Race to find COVID-19 medication



Source: Elsevier

Tackle larger problems

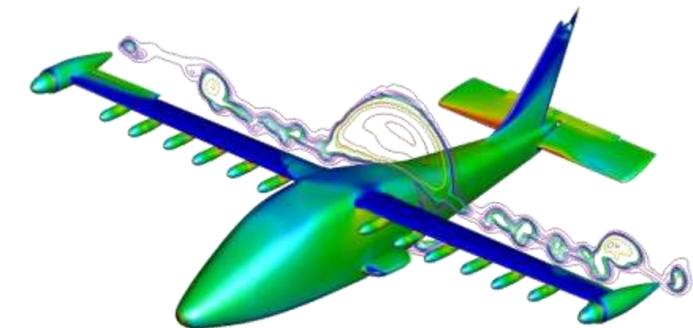
Assimilate more data
Improve physics & algorithms
Increase resolution



Source: The Leading Edge

Save costs

Reduce needs for costly tests
Safe test environments
Test more configurations



Source: NASA

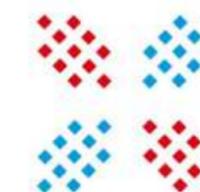
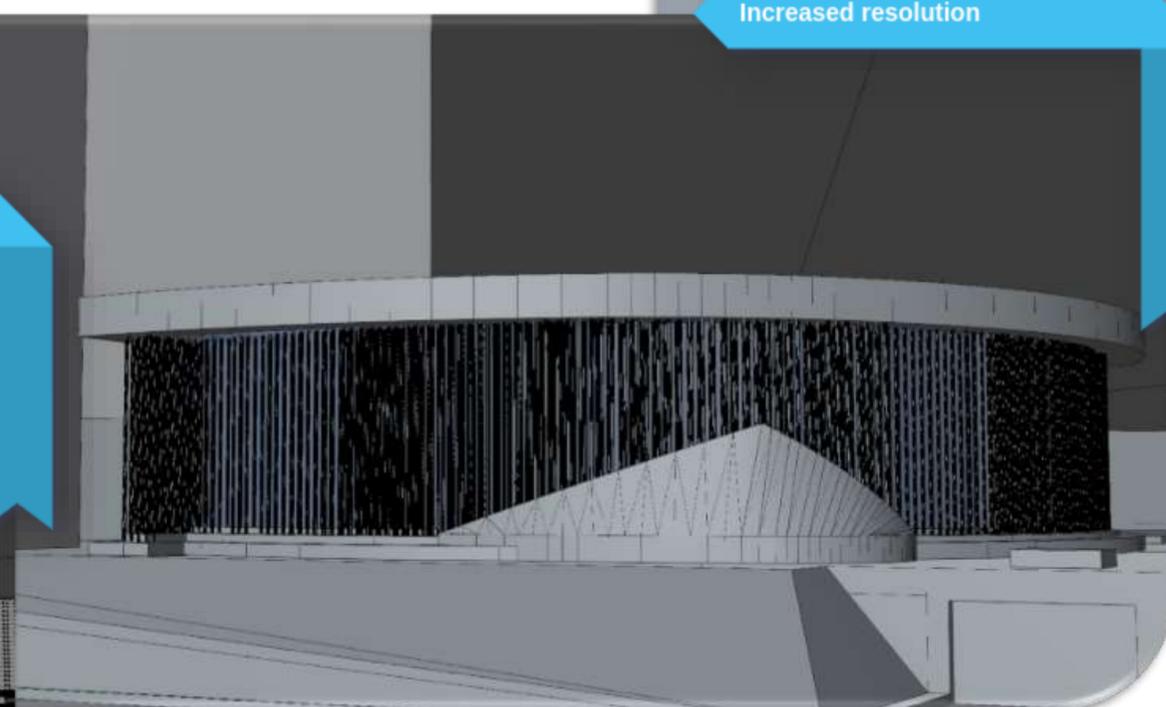
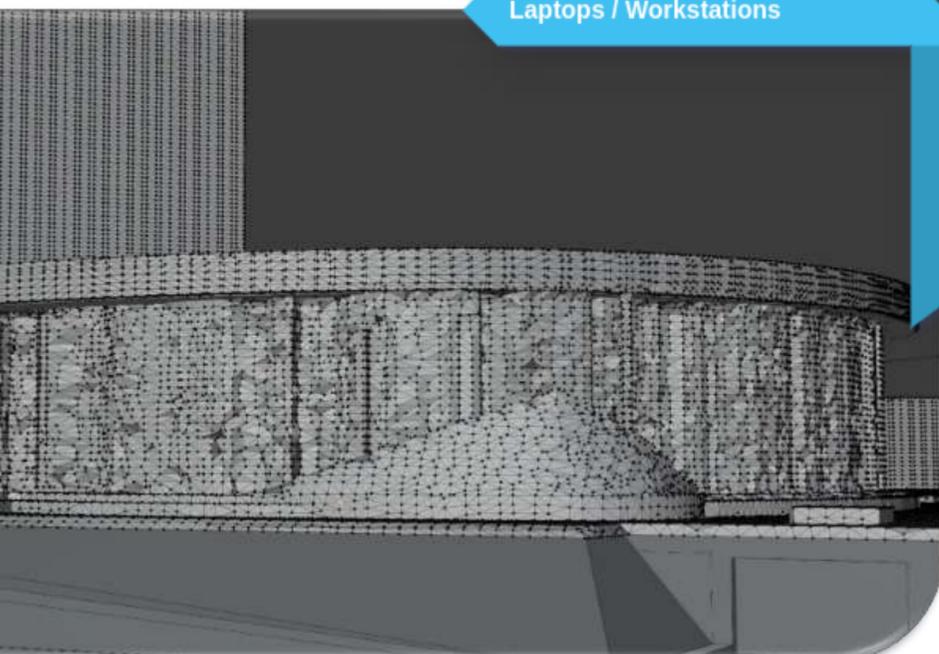
BEYOND LIMITATIONS WITH MELUXINA

- Increase resolution, granularity, problem's size
- Explore multiple meteorological conditions
- Parallel execution to reduce time to solution
- Post-process huge amount of data
- Train Artificial intelligence (AI) at light speed



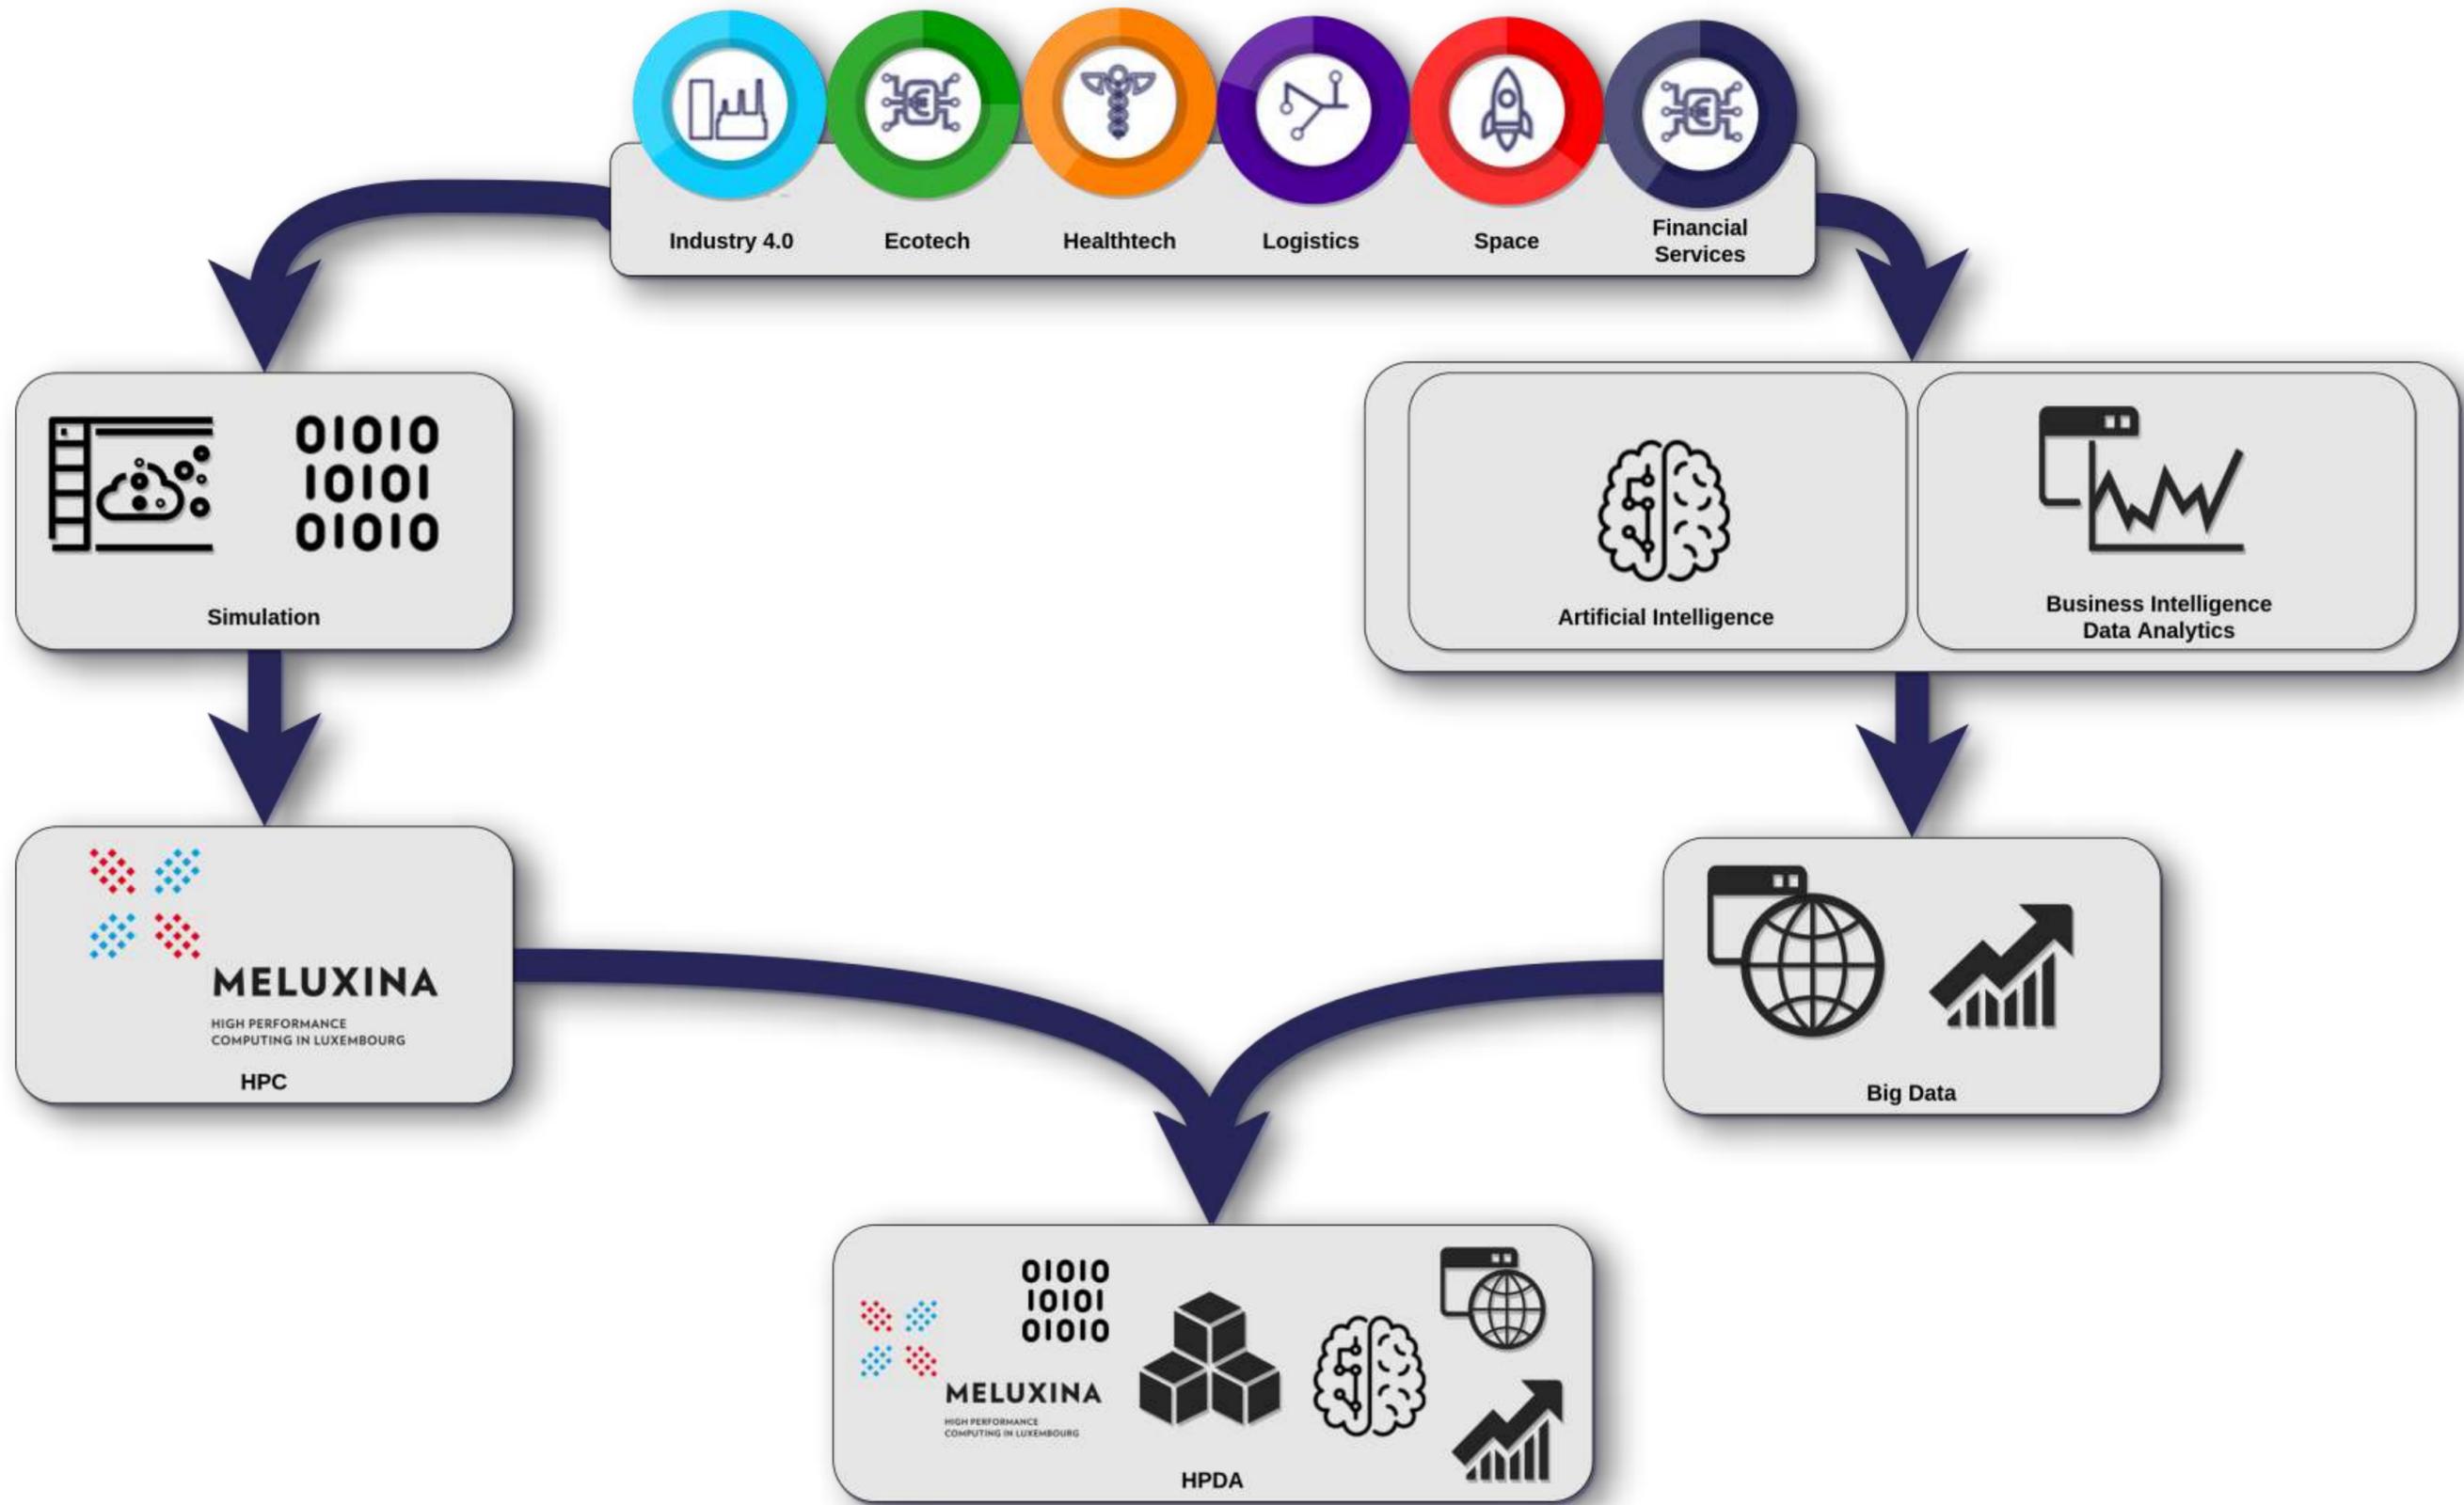
Increased resolution

Laptops / Workstations



MELUXINA

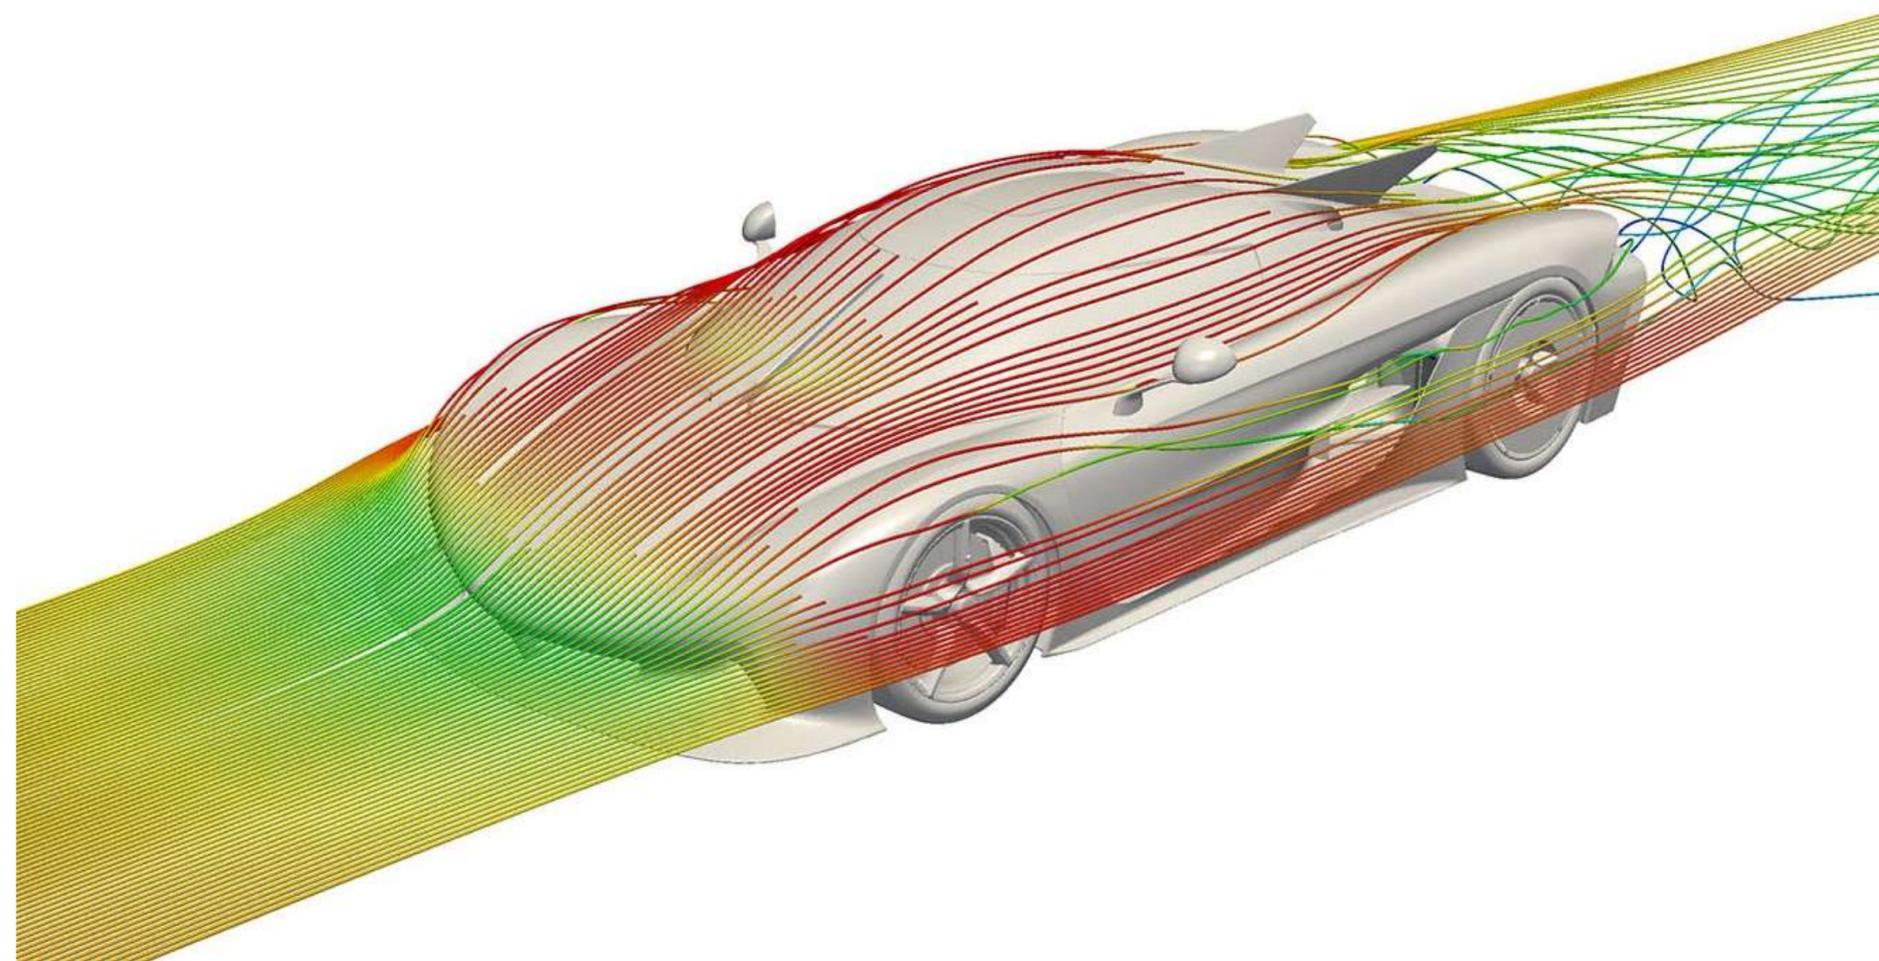
HIGH PERFORMANCE
COMPUTING IN LUXEMBOURG



KOENIGSEGG

Sports-car aerodynamics

HPC-based computational fluid dynamics (CFD) simulations help to build First “megacar” with one-megawatt power and power-to-weight ratio of 1:1



Benefits

30%
in design costs

50%
less wind tunnel testing

60%
in prototyping costs

30%
shorter time to market

Return on investment in
3 month



FORTISSIMO

COPERNICUS & PROVIDERS

Europe's eyes on Earth

Huge EU program to collect large satellite/in-situ database and provide free and open information services.

Several Start-ups and SMEs exploit Copernicus data using HPC, HPDA or IA to provide real business: Agriculture, Ocean Monitoring, Air Quality, Oil&Gas, Urban Monitoring, Renewable Energies, ...etc



Benefits

Beyond 2020 Copernicus related project are set to generate €1 BILLION of revenue and create 4,000 jobs/year

20 Start-ups incubations / year ~ €1 million / year
(Copernicus Incubation – A boost to your business journey (copernicus-incubation.eu))

TEXA

Diagnostic Instruments for Motor Vehicles

On-board diagnostics system of sensors

Data gathered and processed by **HPC provider**

Predictive maintenance from how the vehicle is driven, failure patterns, and overall health of the vehicle.



Benefits

Better prediction of failures

**New services to produce
€1.2M revenue
over the next 3 years**



FORTISSIMO

KALEIDO IDEAS & LOGISTICS

Worldwide Logistics and Intermodal Transport

Logistics represent up to a 20% of manufacturing costs and **marine freight** up to 90% of transport costs

Estimation of marine freight costs using **deep learning**



Benefits

Ability to predict freight costs

Faster quotations, improve performance

by 5%

Quick access to prediction produces a return of €75K during a two-year amortization period



FORTISSIMO

HPC & HPDA: GO / NO GO CRITERIA

Go



- Dataset size larger than several GB
- Long running, compute intensive workloads
- Necessity to scale to stay ahead of competitors

Later



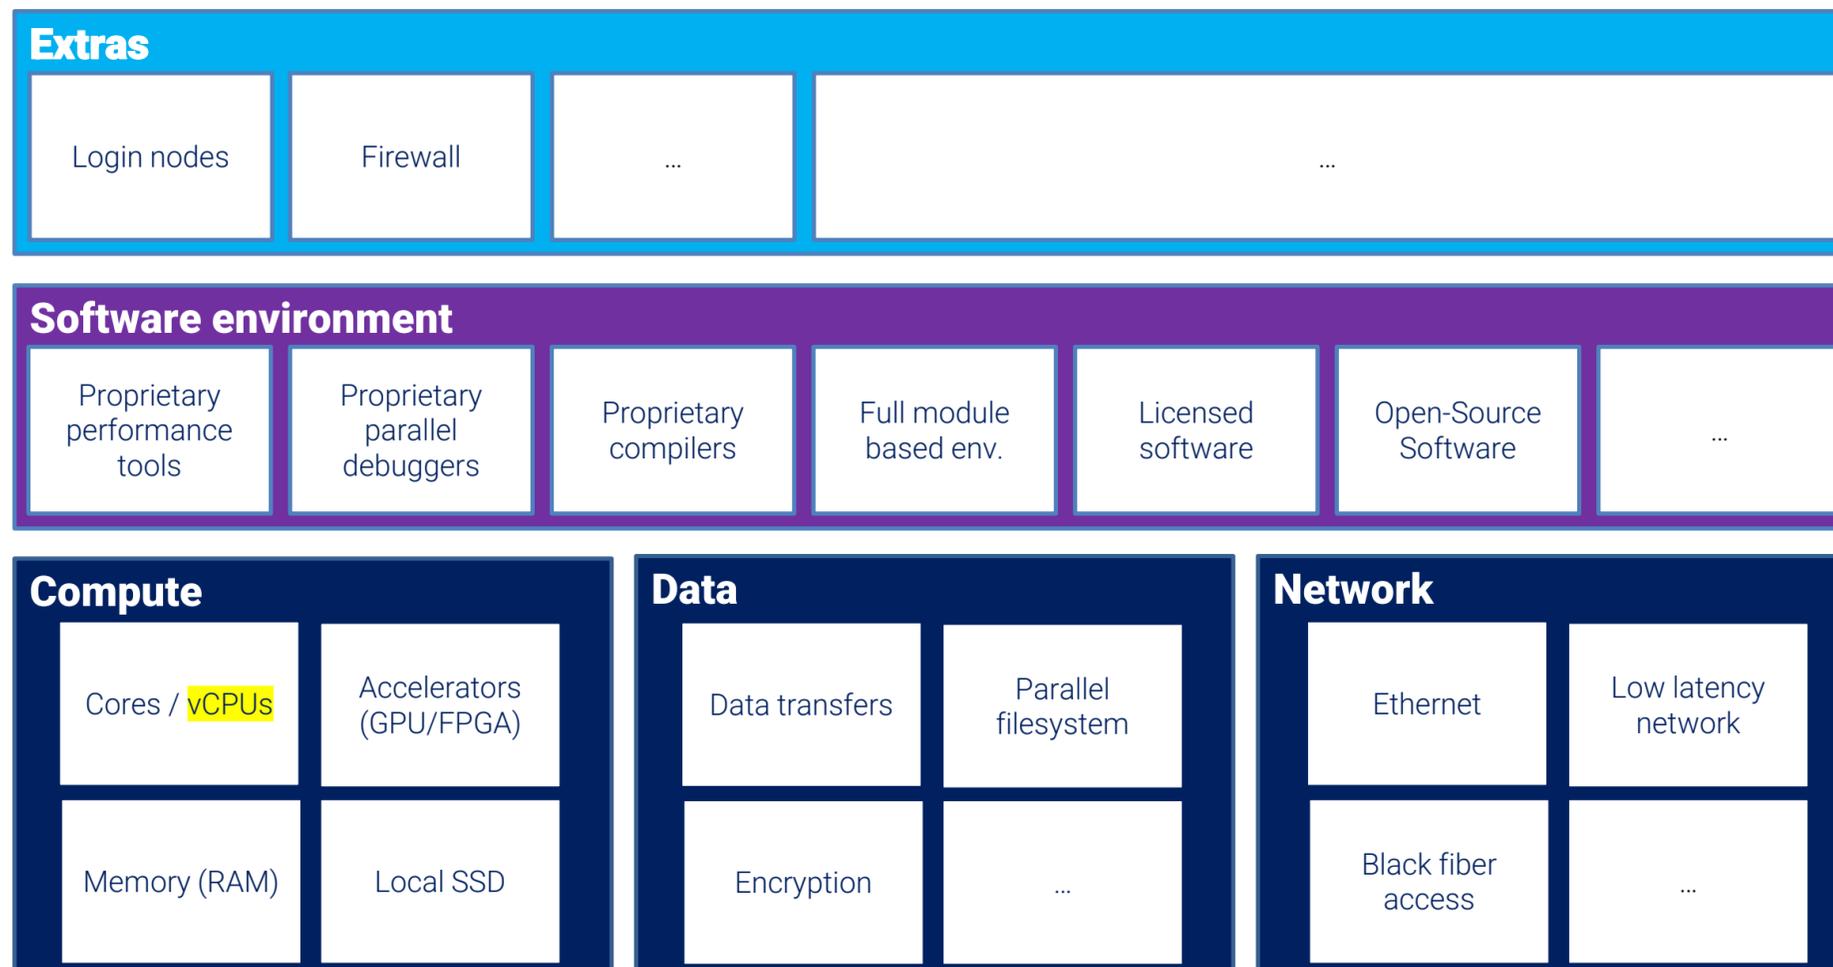
- Limited but growing dataset
- Algorithms with low processing intensity (time and resources)
- Design a growth strategy with HPC in mind

No Go



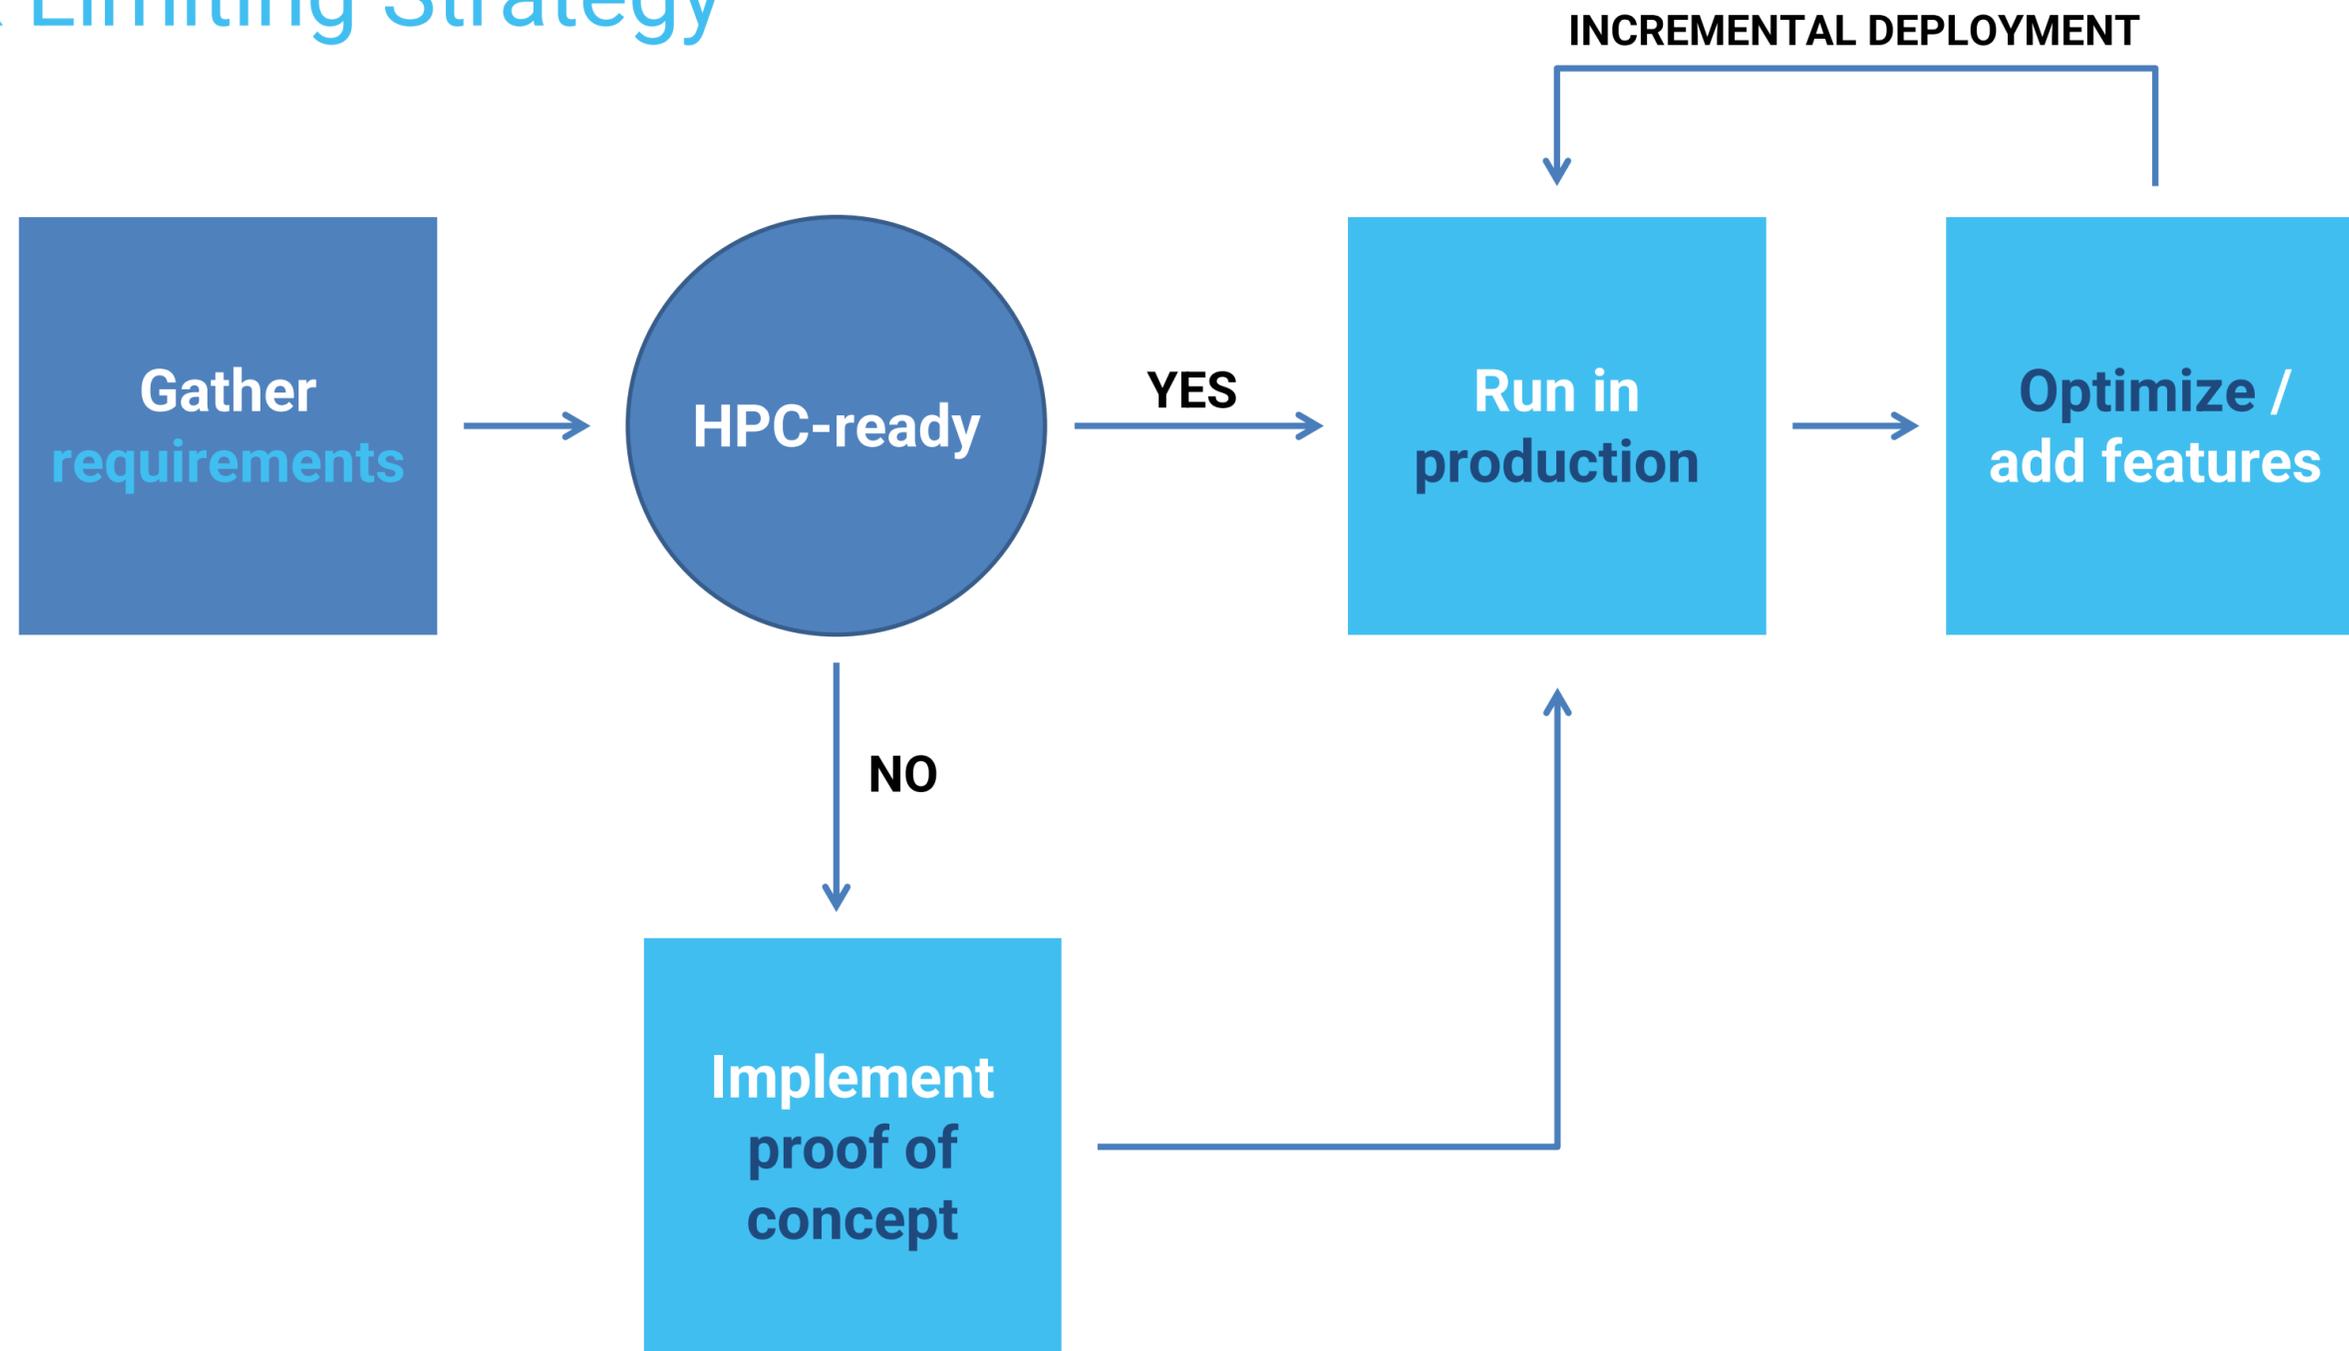
- Interactive applications (web apps, gaming, ...)
- Purely database-driven applications
- In general, deploying real-time applications

ESTIMATING COSTS



AGILE IMPLEMENTATION AND DEPLOYMENT

A Risk Limiting Strategy





Thank you!