



# HIGH PERFORMANCE

Computing & Data Analytics

**Cloud is in the Air**

**From data storage to data analytics:  
bridging Cloud and HPC, MeluXina use cases**

# MeluXina: part of the European Digital Decade and part of the new Luxembourg collaborative digital ecosystem

“... significant milestone in EuroHPC's operations, bringing us one step closer to our ambition of making Europe a global leader in high performance computing. By early next year, the **MeluXina supercomputer will boost European innovation and competitiveness, and support European researchers and industry, wherever they are in Europe.**”

MeluXina: a new EuroHPC world-class supercomputer in Luxembourg  
[eurohpc-ju.europa.eu/news/meluxina-new-eurohpc-world-class-supercomputer-luxembourg](https://eurohpc-ju.europa.eu/news/meluxina-new-eurohpc-world-class-supercomputer-luxembourg)



# LuxProvide's Mission

Leading-edge national supercomputing and data infrastructure

Provide HPC, HPDA, Big Data & AI commercial services

Empower Luxembourg digital ecosystem



Co-funded by the Government of Luxembourg



**EuroHPC**  
Joint Undertaking

Hosting entity for EuroHPC Joint Undertaking



Part of the European Competence Centers initiative



# EuroHPC JU - 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**: LuxProvide, LuxInnovation and the University of Luxembourg
- Currently working on the program for the next 1.5 years

## ... 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 for all audiences

## Coordination

- CASTIEL project will promote interaction & exchanges across NCCs



# Luxembourg Solution Specificities

## Part of the European Petascale Supercomputing Solutions

- 35% of MeluXina will serve EuroHPC projects
- **Increasing visibility** of Luxembourg's capabilities and competences

## The remaining 65% will be dedicated to

- Industry
- Research
- Public administration

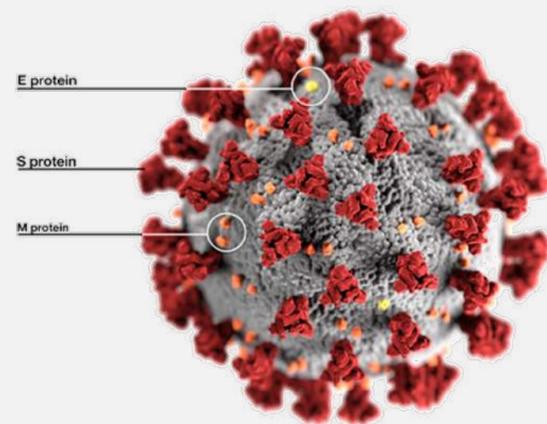
## Our customer-oriented approach implies to

- **be aware of industry-related constraints**
- bring the highest security and **certification standards**
- **serve customers** on a pay per use basis
- **be agile** and feature a fast decision process

# Create Competitive Advantage with High-Performance Computing

## Improve time to market

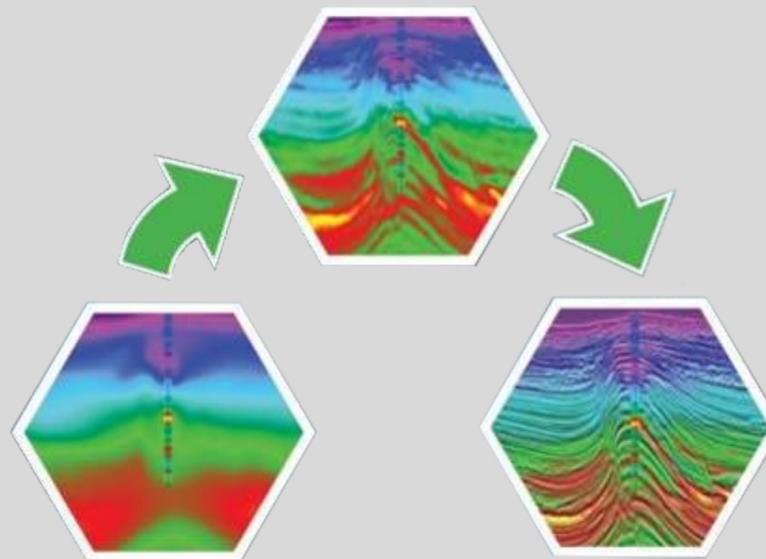
Race to find COVID-19  
medication



Source: Elsevier

## Tackle larger problems

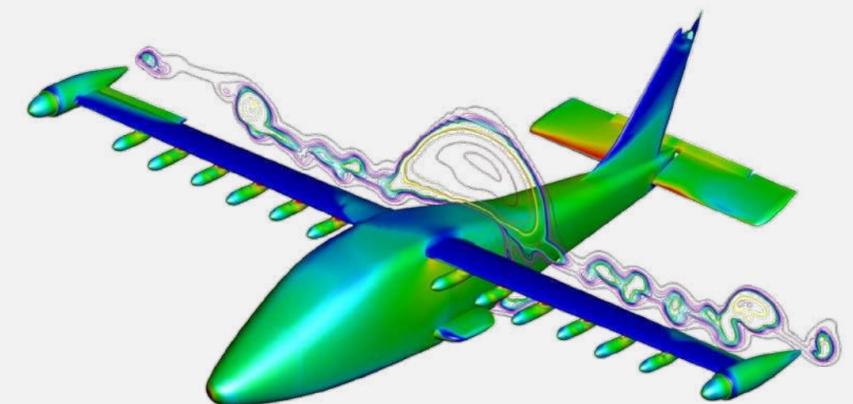
Assimilate more data  
Improve physics & algorithms  
Increase resolution



Source: The Leading Edge

## Save costs

Reduce needs for costly tests  
Test more configurations

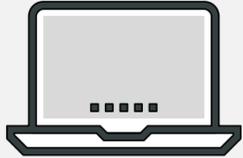


Source: NASA

# MeluXina, in Numbers

Compute  
**18 PFlops**

18,000,000,000,000,000  
Floating Point Operations / Second  
Peak aggregated performance

faster than  
**17,000 x** 

Storage  
**21 PB**

21,000,000,000,000,000 Bytes

more than  
**21,000 x** 

Interconnect  
**HDR 200G**

200-400 InfiniBand Gigabit/s links  
130 ns Latency

better than  
**768,000 x** 

# Cost distribution

## Classical IT

- 1 Engineer costs **100.000 EUR per year**
- 1 PC costs **2.000 EUR over a 4 years** lifetime

## HPC

- 1 HPC specialist costs **150.000 EUR per year**
- 1 HPC system costs **30.000.000 EUR over a 5 years** lifetime

The problematic of **minimizing the engineering cost** is turned into the problematic of **maximizing the machine usage**

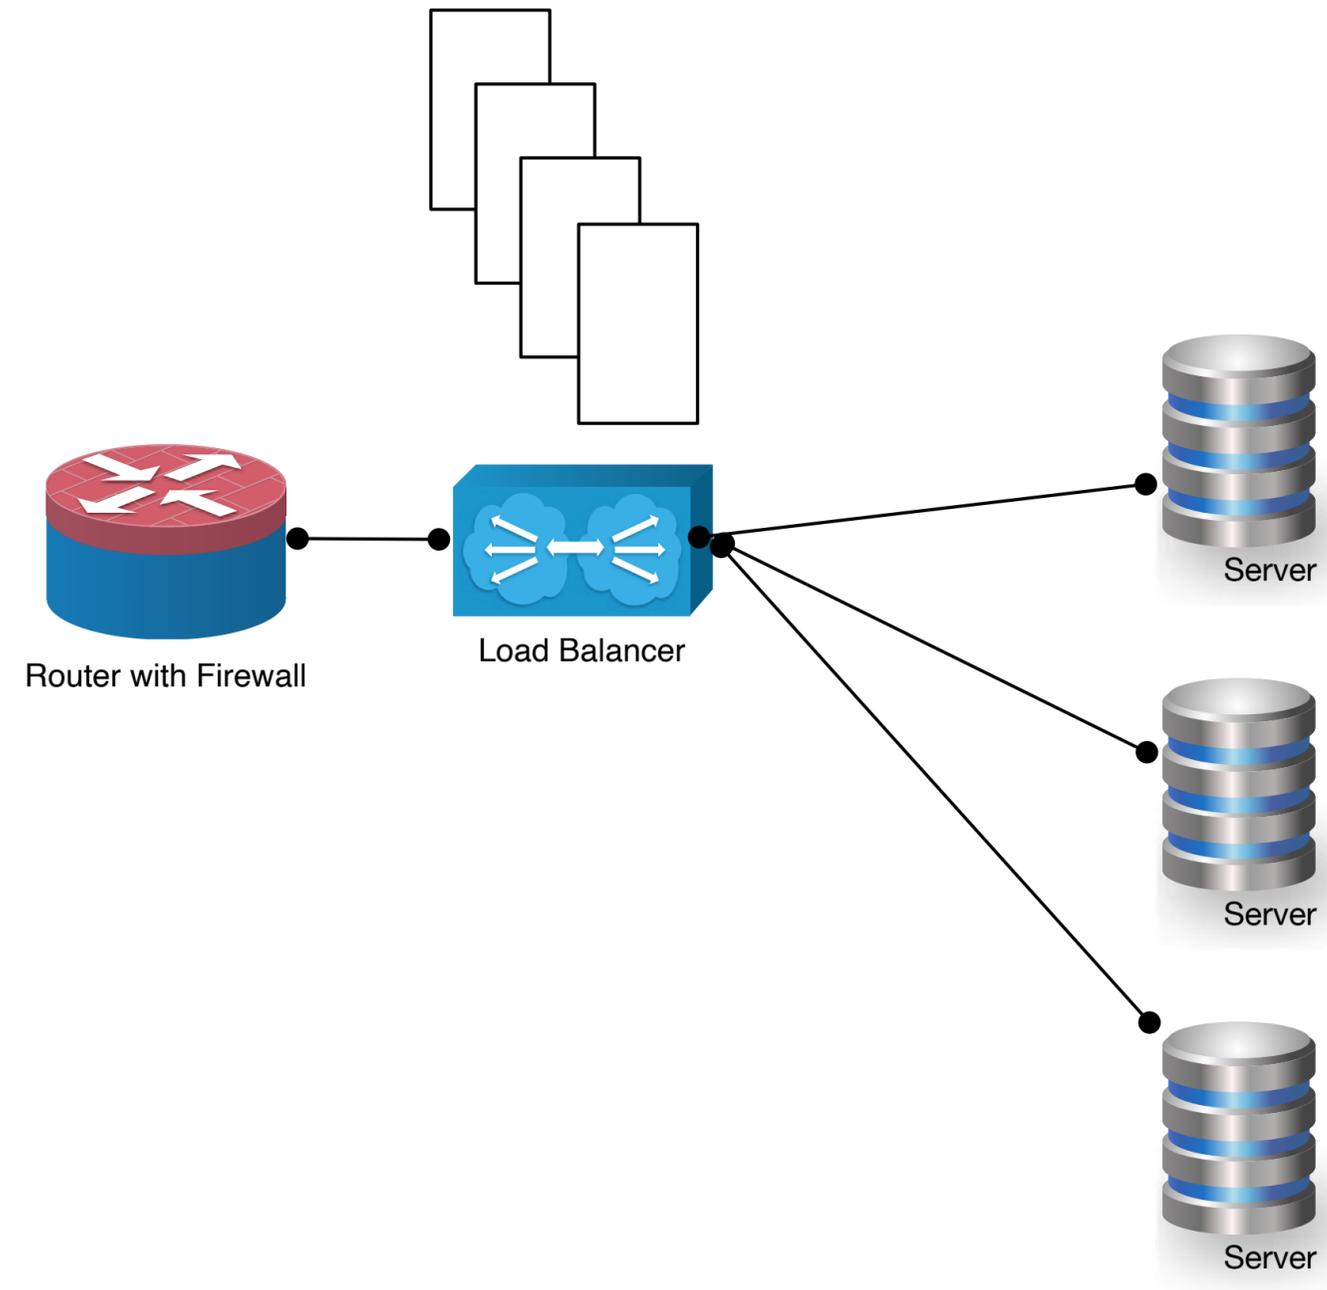
## Similar mindset as sending a spacecraft to Mars

- no space for low quality software: performance for HPC and bug free for Space

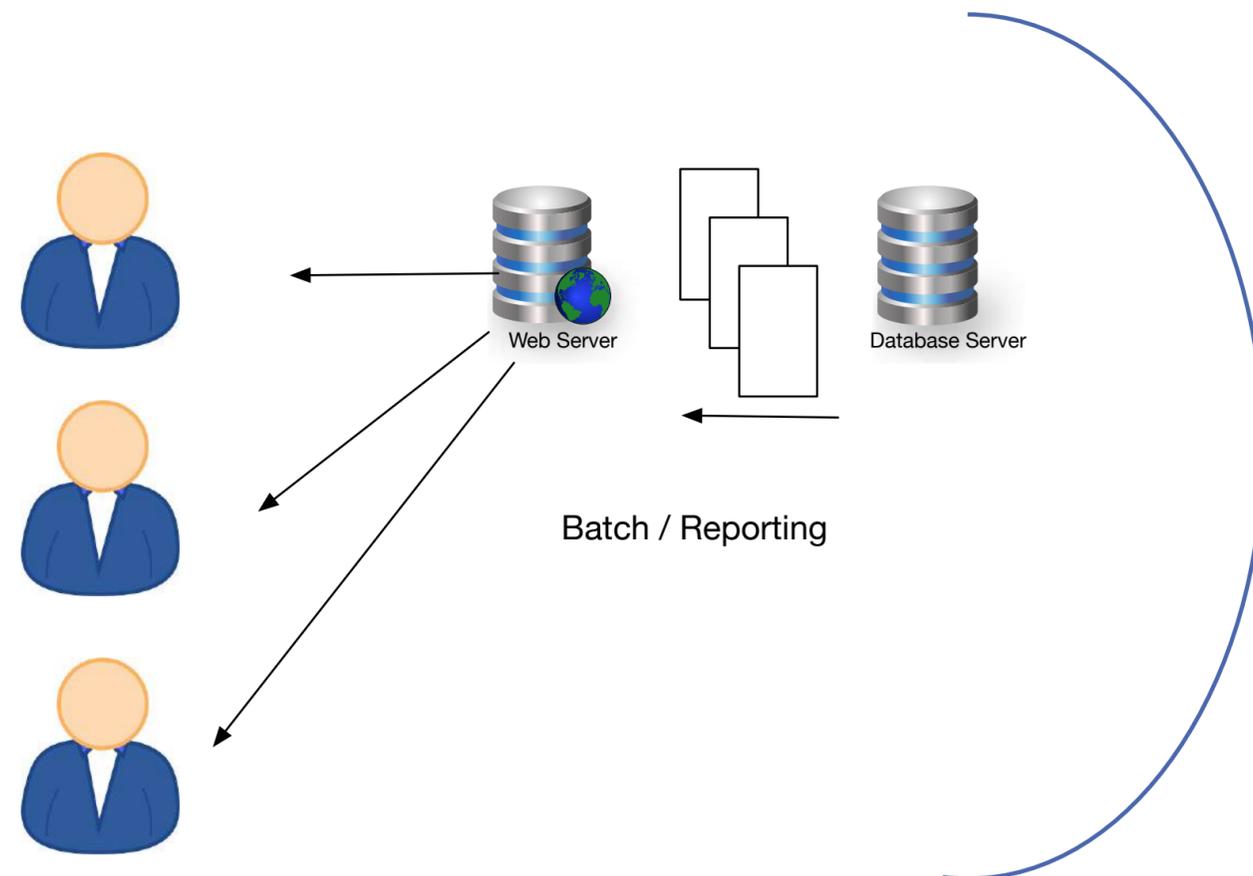
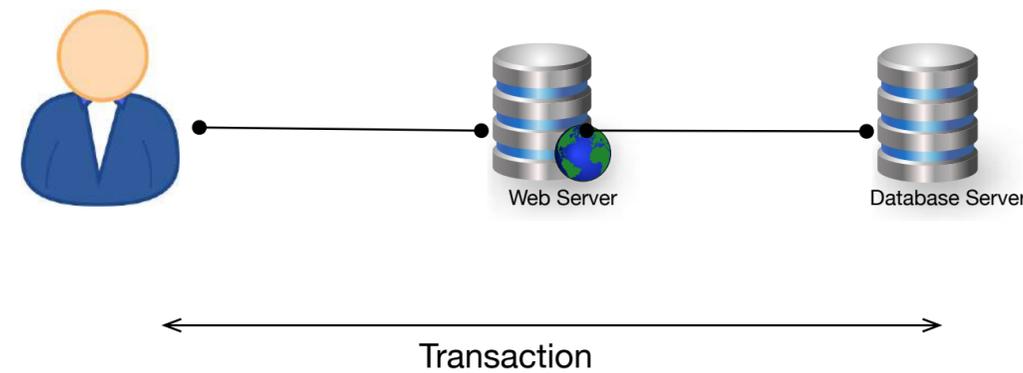
**Huge investments have been done for optimising HPC software** for over the last 40 years

- but: 90% of the software running on HPC is written in Fortran

# Classical Computer Cluster

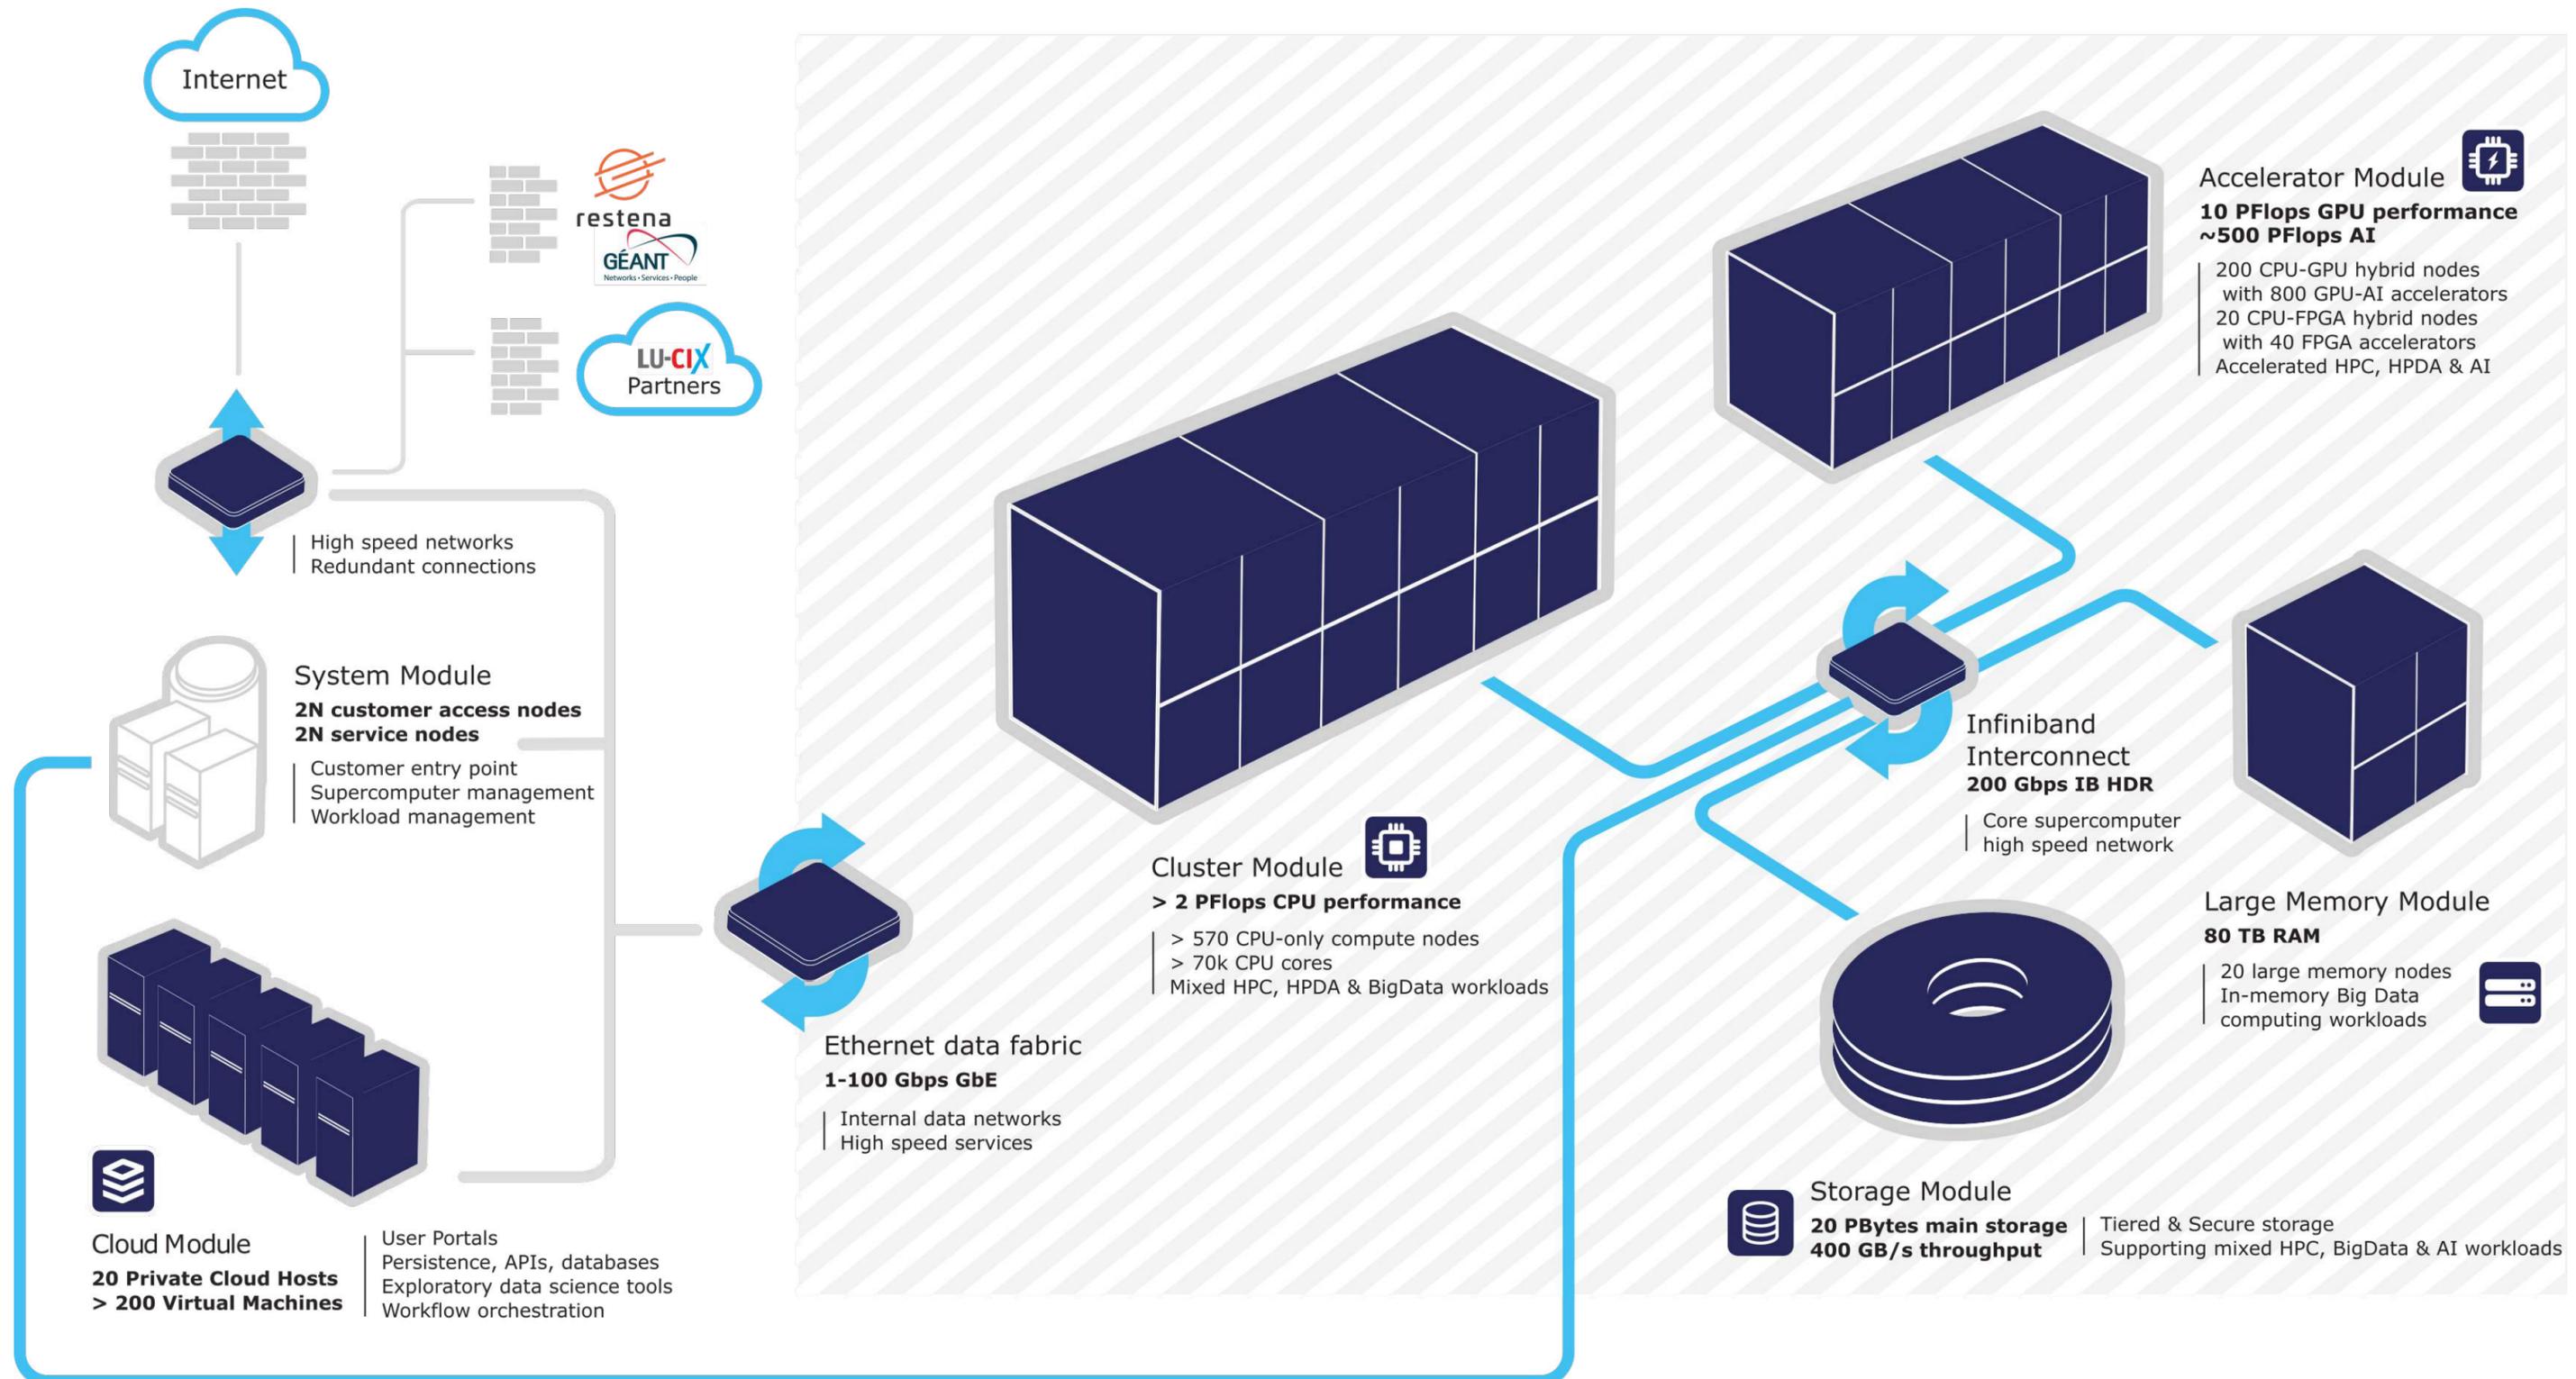


# Interactive (Synchronous) or Batch mode (Asynchronous)



**HPC has more of a batch-oriented approach**

# MeluXina & its Low Latency Interconnect



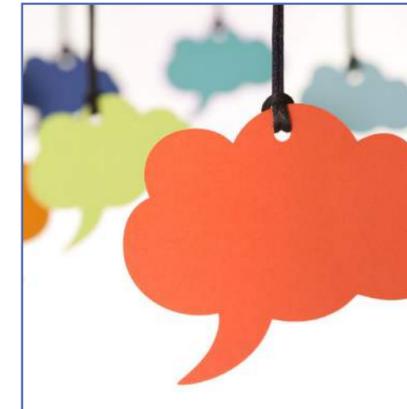
# Workloads



**MeluXina:**  
Making new things possible

## Great match

- Batch processing of large-scale workloads  
Run applications that wouldn't fit in lesser systems
- Efficient execution of intensive and highly parallel applications  
Reduce computing time from months to weeks
- Efficient analysis of huge quantities of data  
Compute on petabytes to *discover* or *predict*
- Development of HPC, HPDA & AI applications  
Develop using a specialized software environment
- Simulation, modeling & learning tasks  
Tackle grand societal challenges

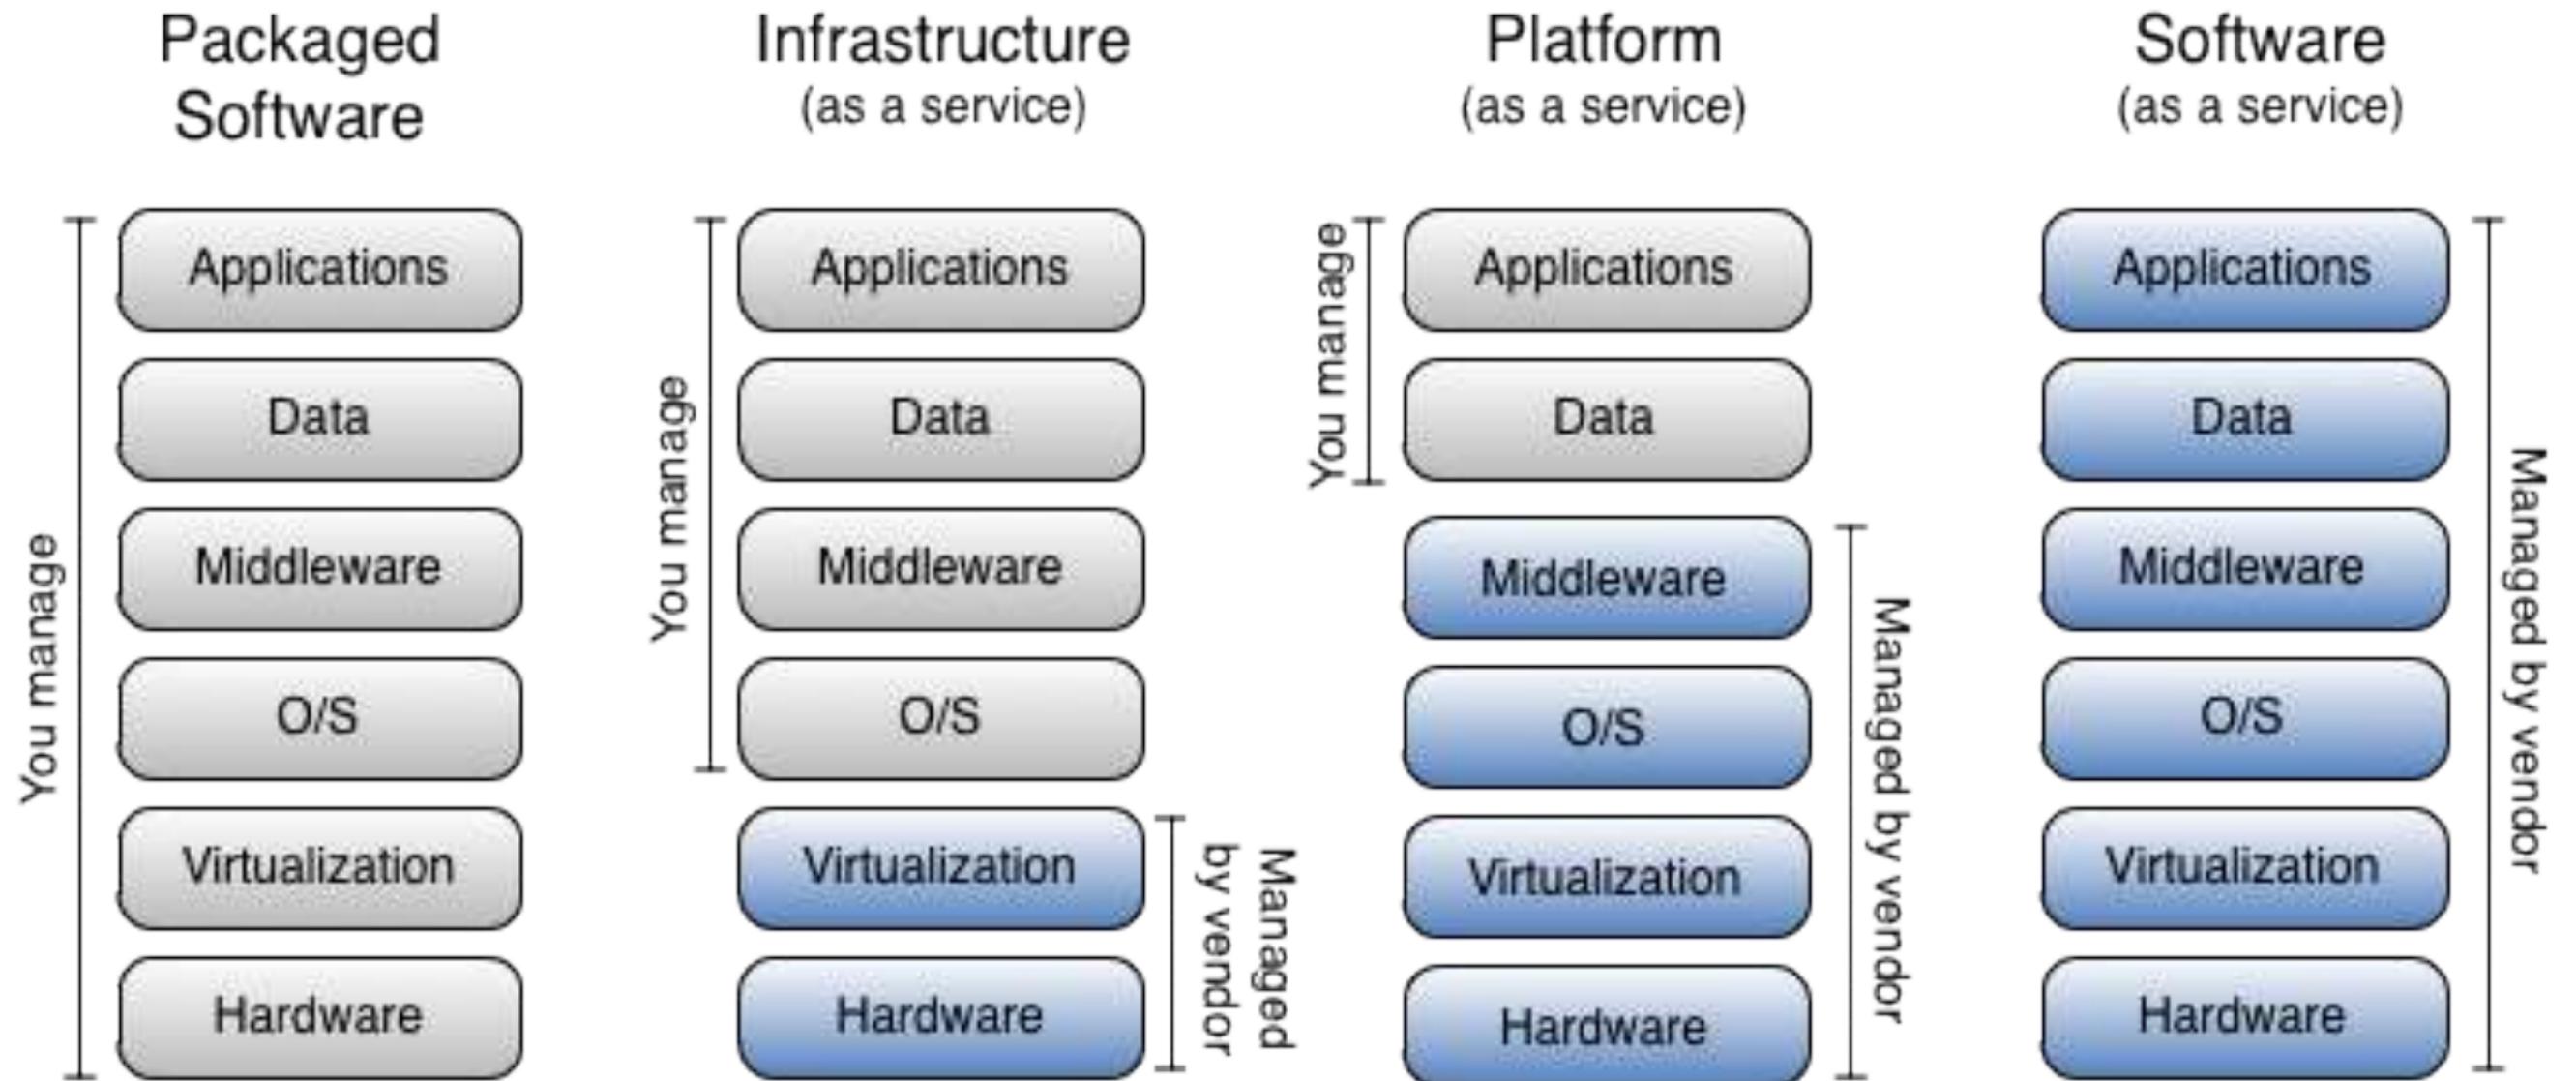


**Other platforms:**  
Complementing our HPC, HPDA  
& AI capabilities

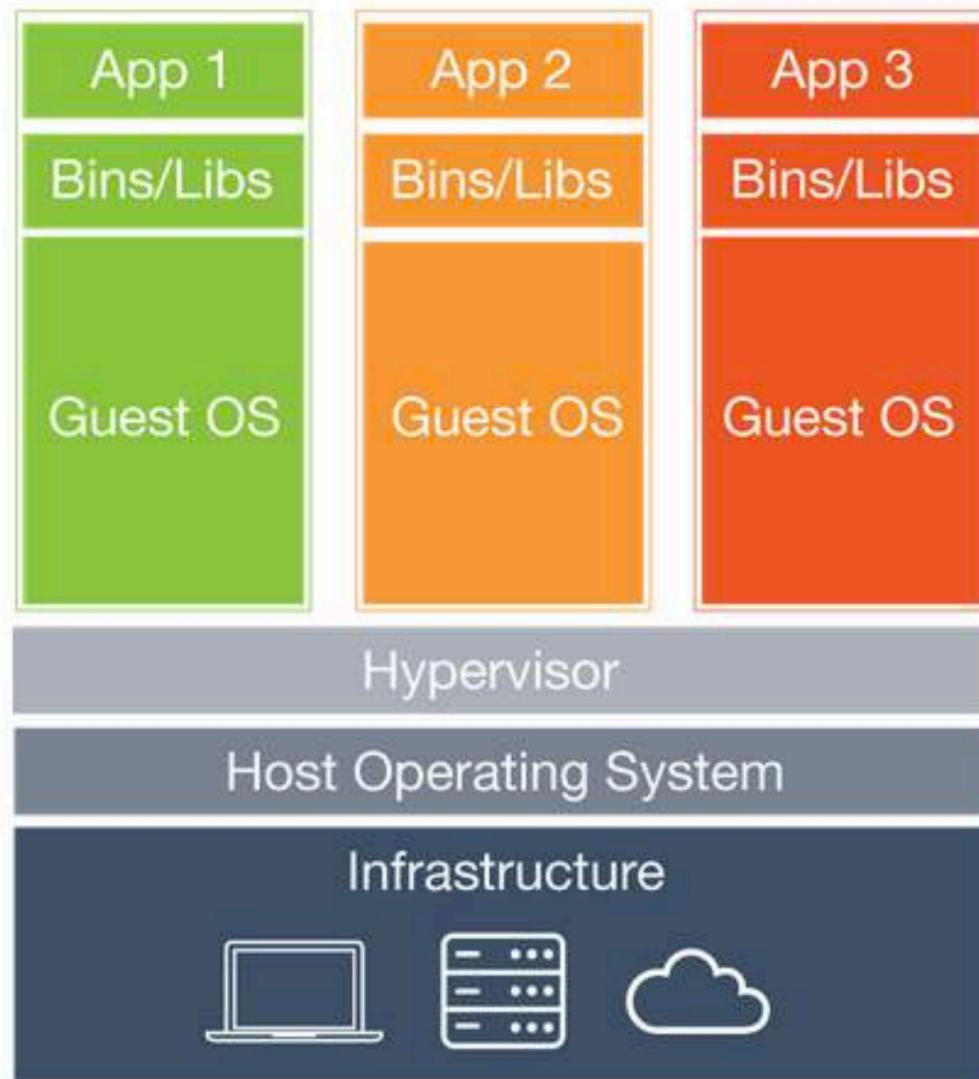
## Other platforms

- Real-time mission critical applications  
Processes requiring dedicated resources
- Data storage & transfer without computation  
Solutions for long term storage, video streaming
- Web applications  
Workflows performing minimal computation
- Remote desktop & application development  
General-purpose interactive development
- Distributed but not parallel workloads  
Execution of uncoupled tasks

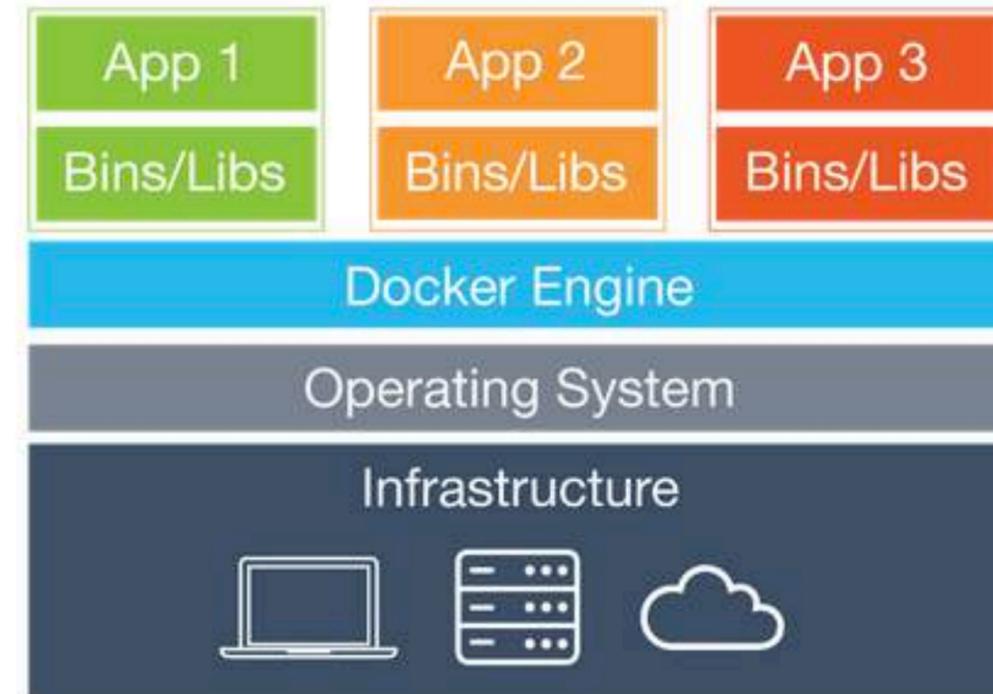
# Cloud Stacks



# VMs vs Containers



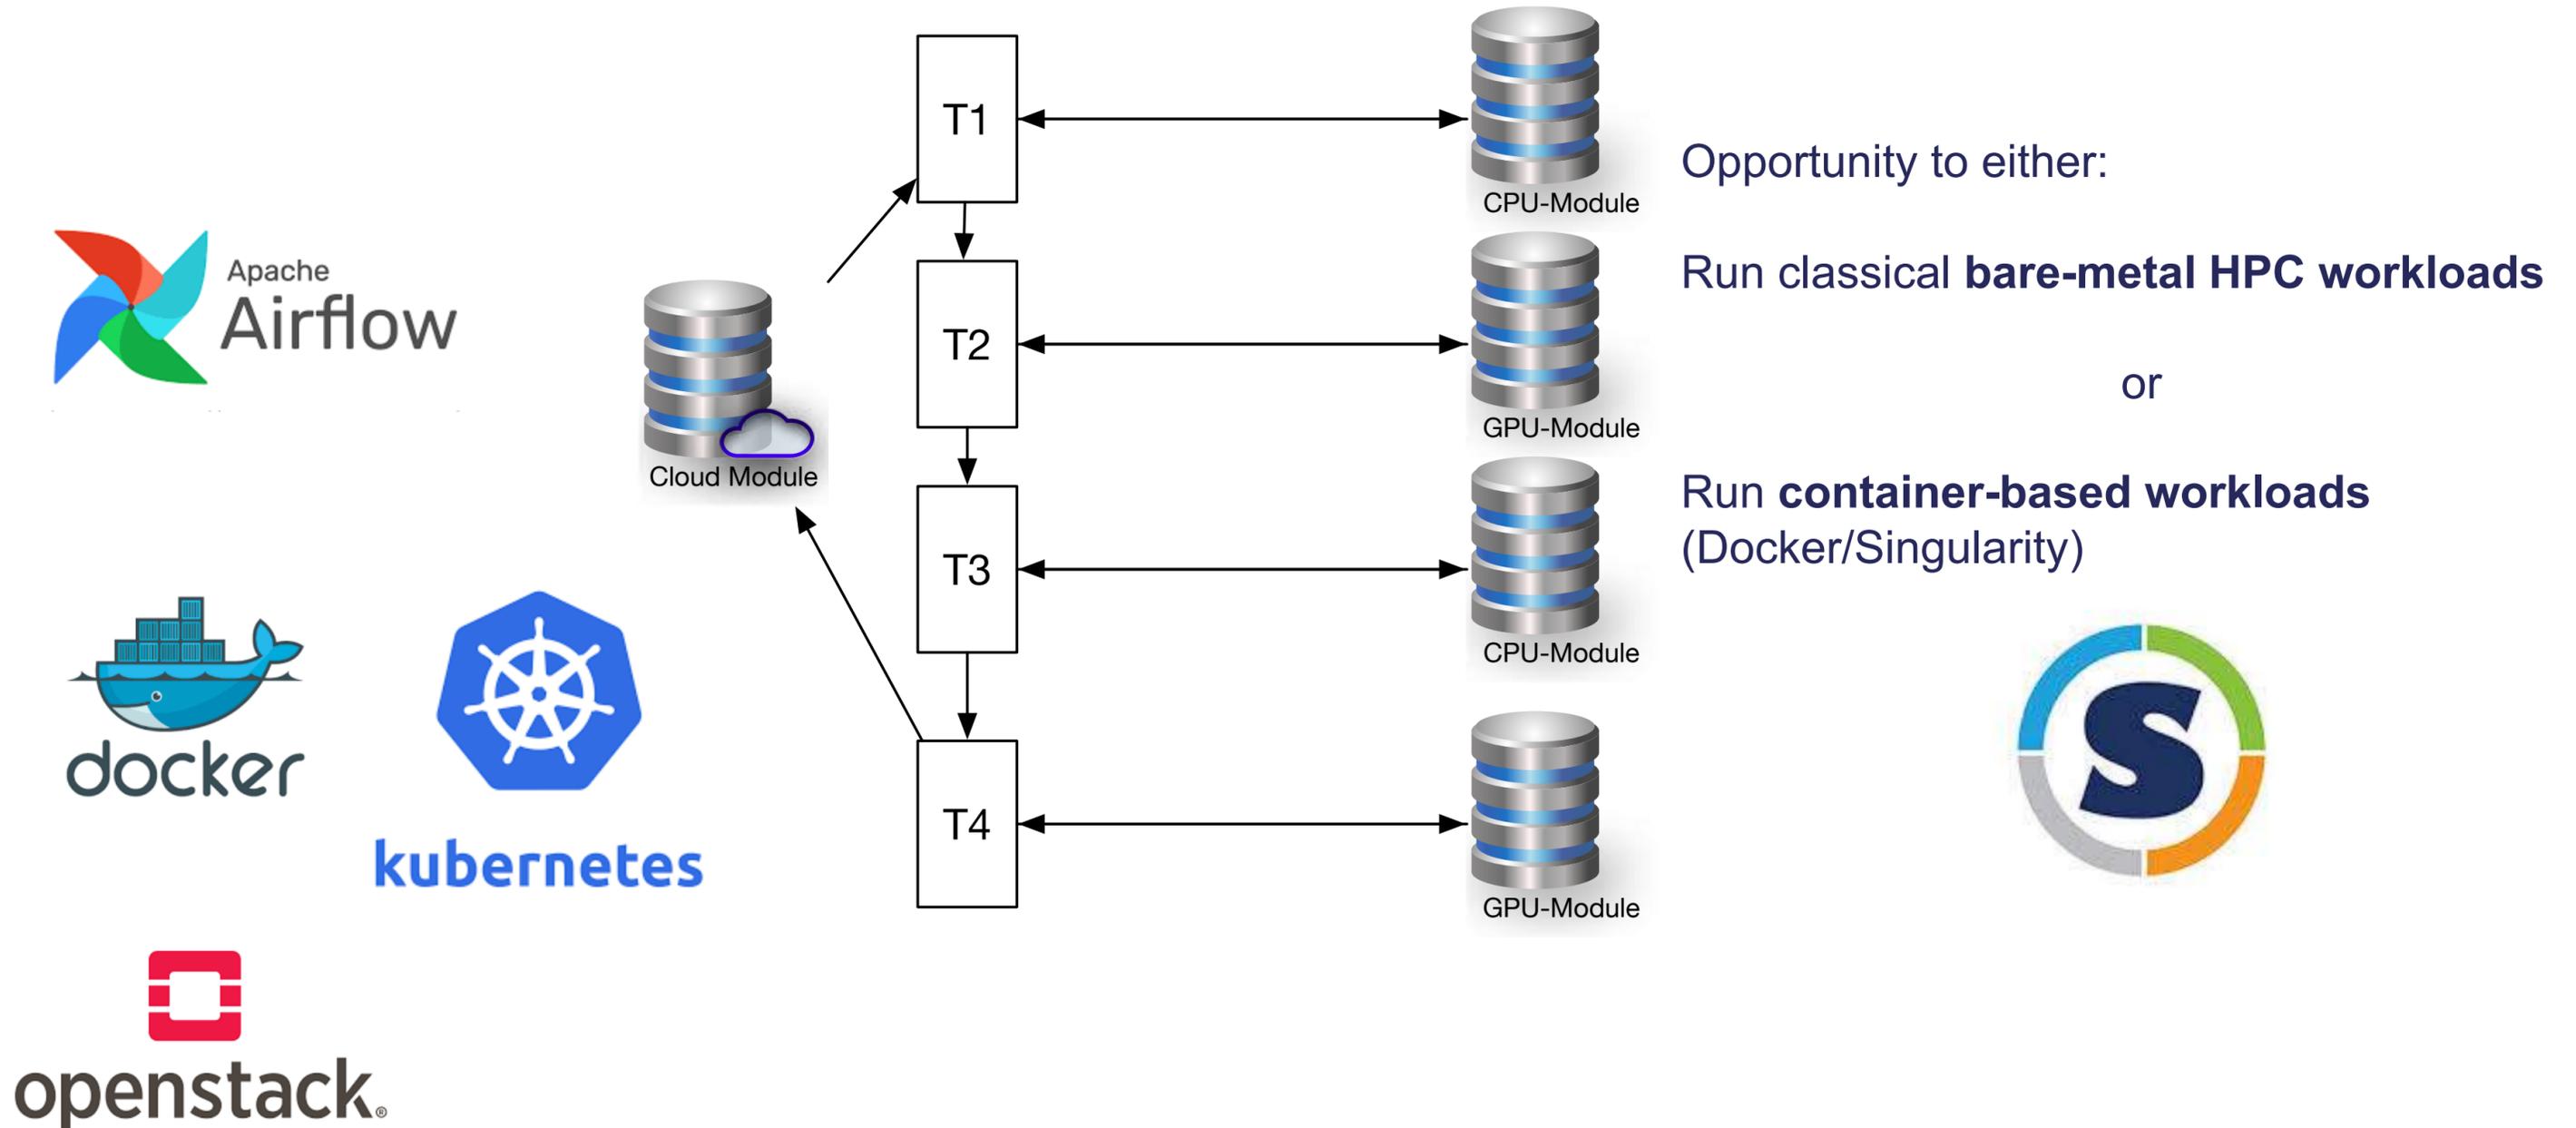
Virtual Machines



Containers

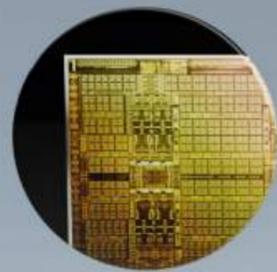
© Docker

# Cloud - HPC orchestration - example of a (scientific) workflow -

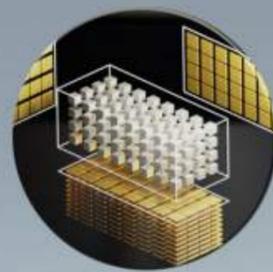


# 800 A100 Ampere GPUs for boosting AI

GREATEST GENERATIONAL LEAP – 20X VOLTA



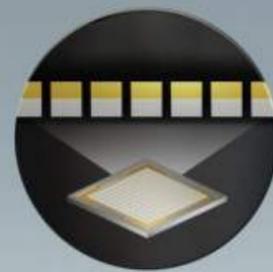
54 BILLION XTORS



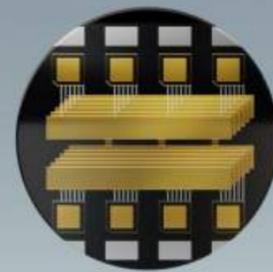
3<sup>RD</sup> GEN TENSOR CORES



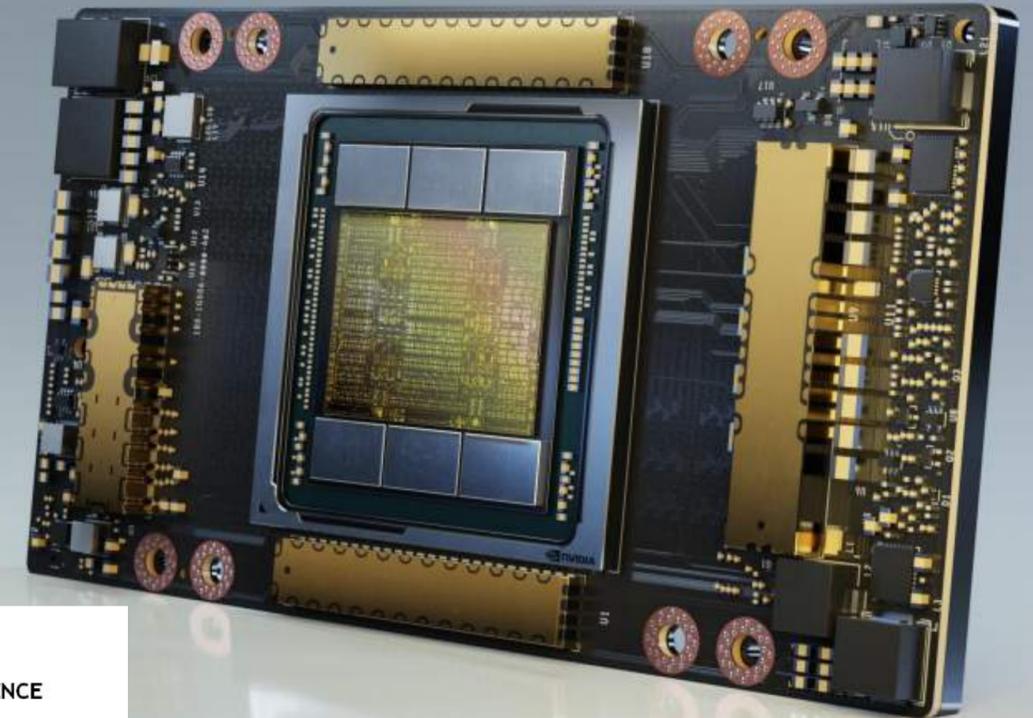
SPARSITY ACCELERATION



MULTI INSTANCE GPU

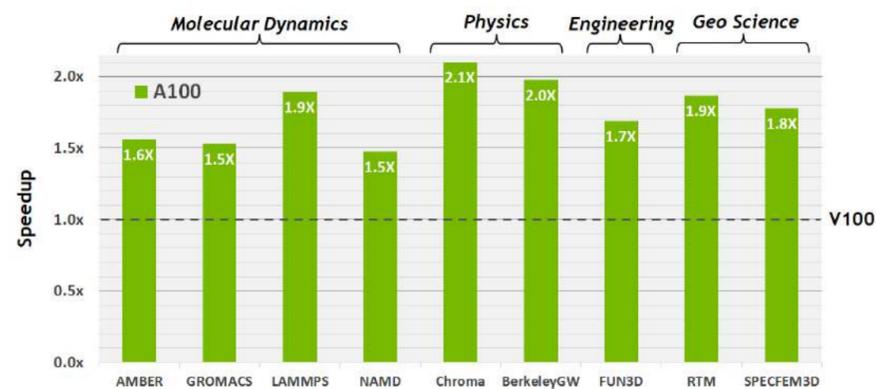


3<sup>RD</sup> GEN NVLINK & NVSWITCH



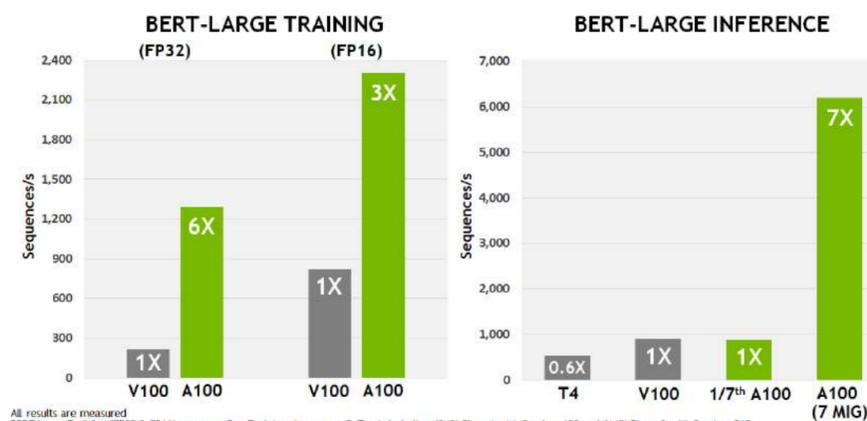
| TSMC 7N | 40GB Samsung HBM2 | 3<sup>rd</sup> gen Tensor Core GPU | 600 GB/s NVLink

## ACCELERATING HPC



All results are measured  
 Except BerkeleyGW, V100 used is single V100 SXM2, A100 used is single A100 SXM4  
 More apps detail: AMBER based on PME-Celulose, GROMACS with STMY (H-bond), LAMMPS with Atomic Fluid LJ-2.5, NAMD with v3.0a1 STMY\_NVE  
 Chroma with szsc21\_24\_128, FUN3D with dpw, RTM with Isotropic Radius 4 1024^3, SPECFEM3D with Cartesian four material model  
 BerkeleyGW based on Chi Sun and uses 8xV100 in DGX-1, vs 8xA100 in DGX A100

## UNIFIED AI ACCELERATION



All results are measured  
 BERT Large Training (FP32 & FP16) measures Pre-Training phase, uses PyTorch including (2/3) Phase1 with Seq Len 128 and (1/3) Phase 2 with Seq Len 512,  
 V100 is DGX1 Server with 8xV100, A100 is DGX A100 Server with 8xA100, A100 uses TF32 Tensor Core for FP32 training  
 BERT Large Inference uses TRT 7.1 for T4/V100, with INT8/FP16 at batch size 256. Pre-production TRT for A100, uses batch size 94 and INT8 with sparsity



## MELUXINA: Luxembourg's 500 PetaFlops AI supercomputer

HIGH PERFORMANCE COMPUTING IN LUXEMBOURG

# Basic Schema for Machine Learning



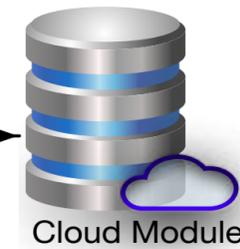
Google Cloud



Data ingestion  
Inferencing



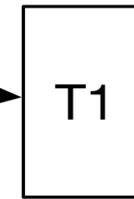
IoT



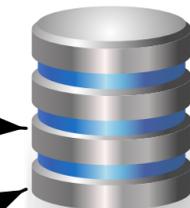
Cloud Module



Storage



T1



GPU-Module

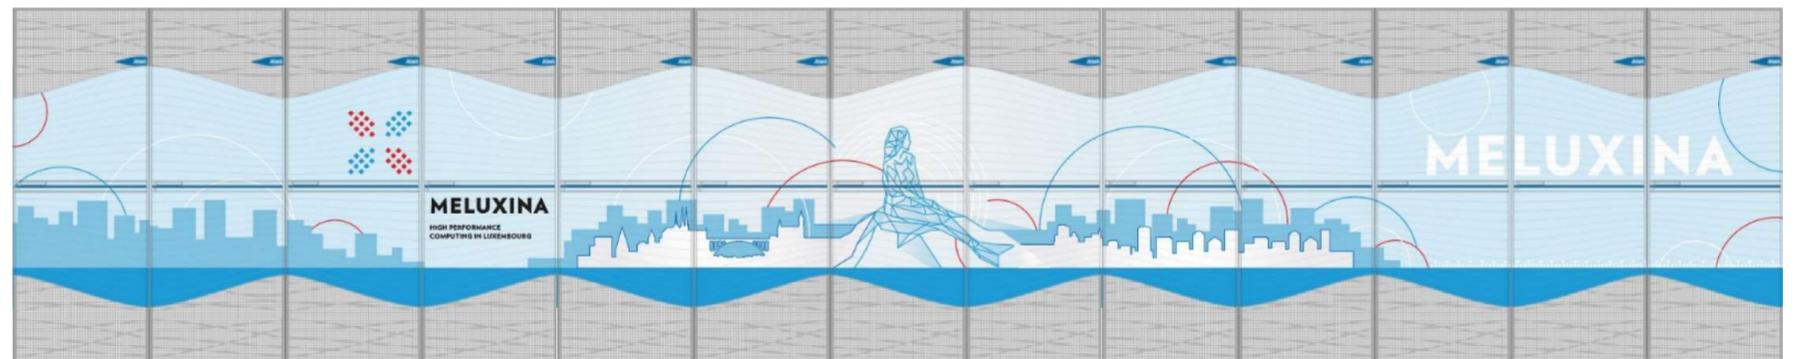


Infiniband Fabric  
Interconnect

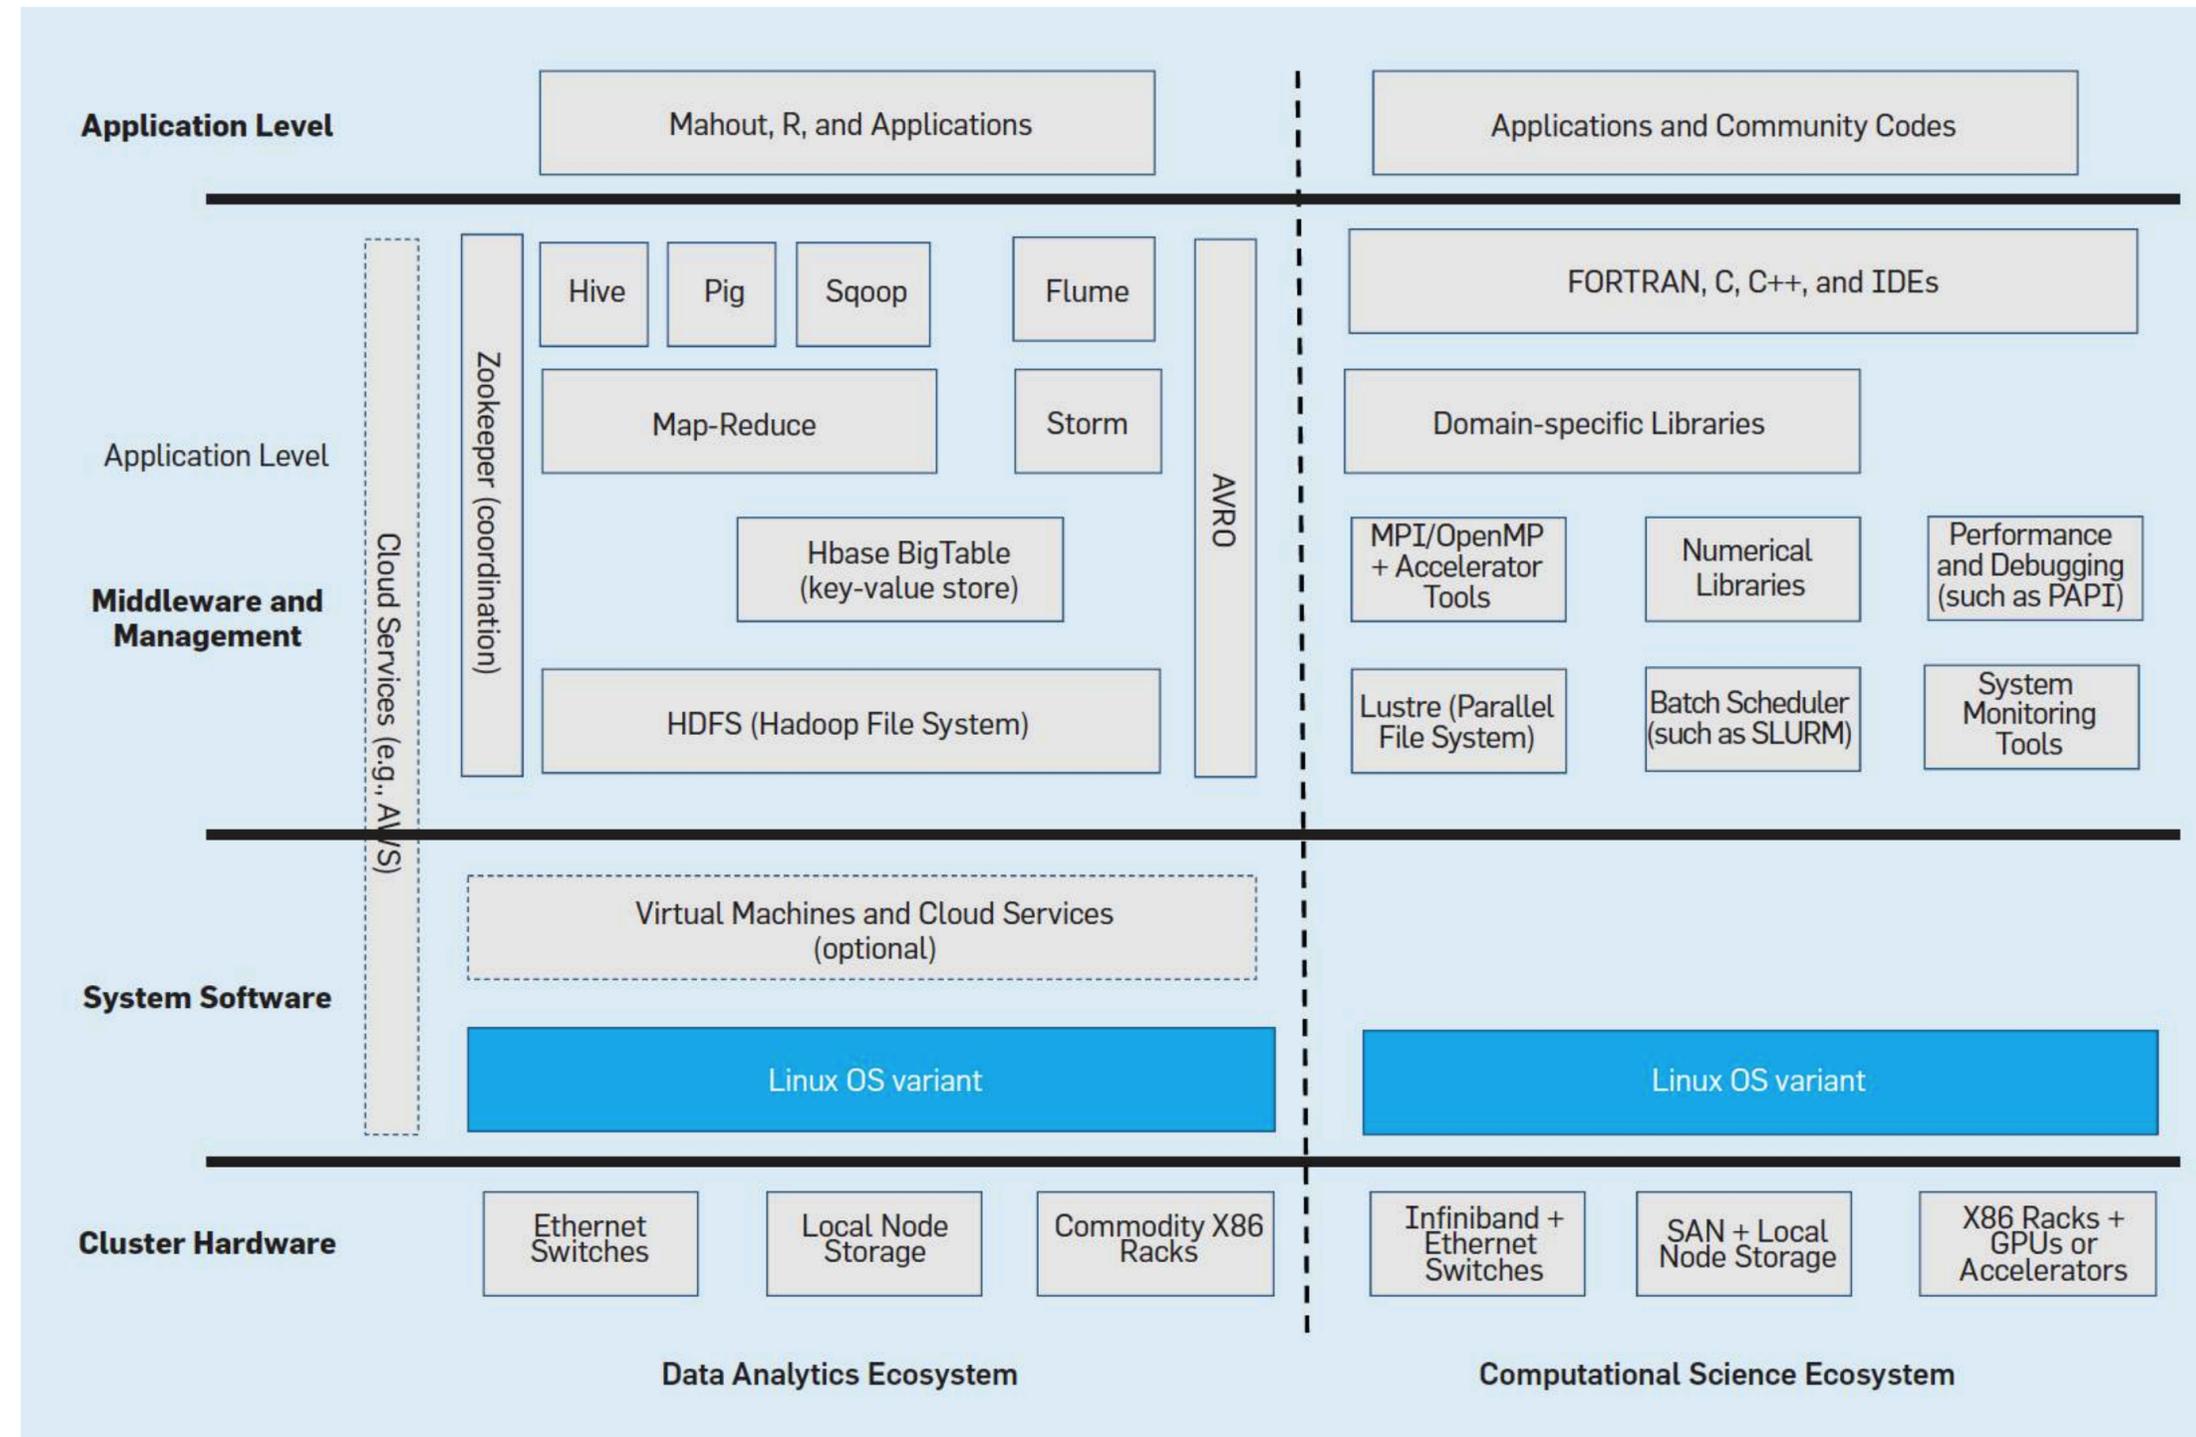
Learning & model building  
Inferencing

1. Download your data
2. Process them using the GPU Module
3. Upload the results

The InfiniBand interconnect allows **low latency** and **high throughput** between all parts of MeluXina.



# Big Data vs HPC Toolset/Software Stack



*Data analytics and computing ecosystem compared.* Daniel A. Reed and Jack Dongarra. **Scientific discovery and engineering innovation requires unifying traditionally separated high-performance computing and big data analytics.** Communications of the ACM July 2015.

# The *MUSE*: MeluXina User Software Environment

## Software stack(s)

- Rich software environments with MeluXina architecture optimizations
- Enable wide variety of workloads, provide both Open-Source & Commercial tools
- Specialized middleware on HPC and Cloud Modules

## Development & Performance tools

- AOCC, GCC, Intel, NVIDIA HPC SDK (incl. PGI), ARM Forge

## MPI suites enabling scalable parallel executions

- OpenMPI, Intel MPI, ParaStationMPI

## Numerical, data & parallel libraries

- BLIS, Intel MKL, TBB, FFTW, OpenBLAS, ScalaPACK
- cuBLAS, cuFFT, cuDNN, TensorRT, NCCL, Kokkos, HDF5, netCDF

## Frameworks for Machine Learning, Deep Learning & Big Data

- PyTorch, TensorFlow, Horovod, Keras, Apache Spark

## Applications

- Visualisation: VMD, VisIT, ParaView
- **Domain-specific and cross-domain applications**

AMD intel

PGI

NVIDIA  
CUDA



ParaStation  
MPI



Open  
BLAS

TensorFlow

PyTorch

Keras Apache Spark



# Conclusion

- **A bridge between Cloud and HPC** has been established
  - by having MeluXina incorporate a **Cloud Module** with access to high-speed fabrics
- **Edge computing** shall nicely complement the current solution
- The Cloud Module allows to **orchestrate HPC workflows**
- **Singularity** allows to run Docker **Containers** on the HPC Modules
- Next evolutions should involve
  - Convergence of Kubernetes and HPC for **higher workload flexibility**
  - The rewriting of Big Data tools/libraries for HPC for **added workflow acceleration**

# Improve Your Business with Proven Combinations of Services

## Compute Resources

- On-demand reservation
- Customer-oriented packages
- Highest security
- GDPR compliance

---

## Data Resources

- Ultrafast data tiers
- Parallel data systems
- Long-term archival capabilities

## Training

- Awareness
- Beginner & advanced course
- Domain-specific workshops

---

## Support

- Specialized guidance
- Account management
- Data retrieval
- Task scheduling

## Solution Engineering

- HPC / Data expertise
- Software engineering
- Business intelligence
- Consulting

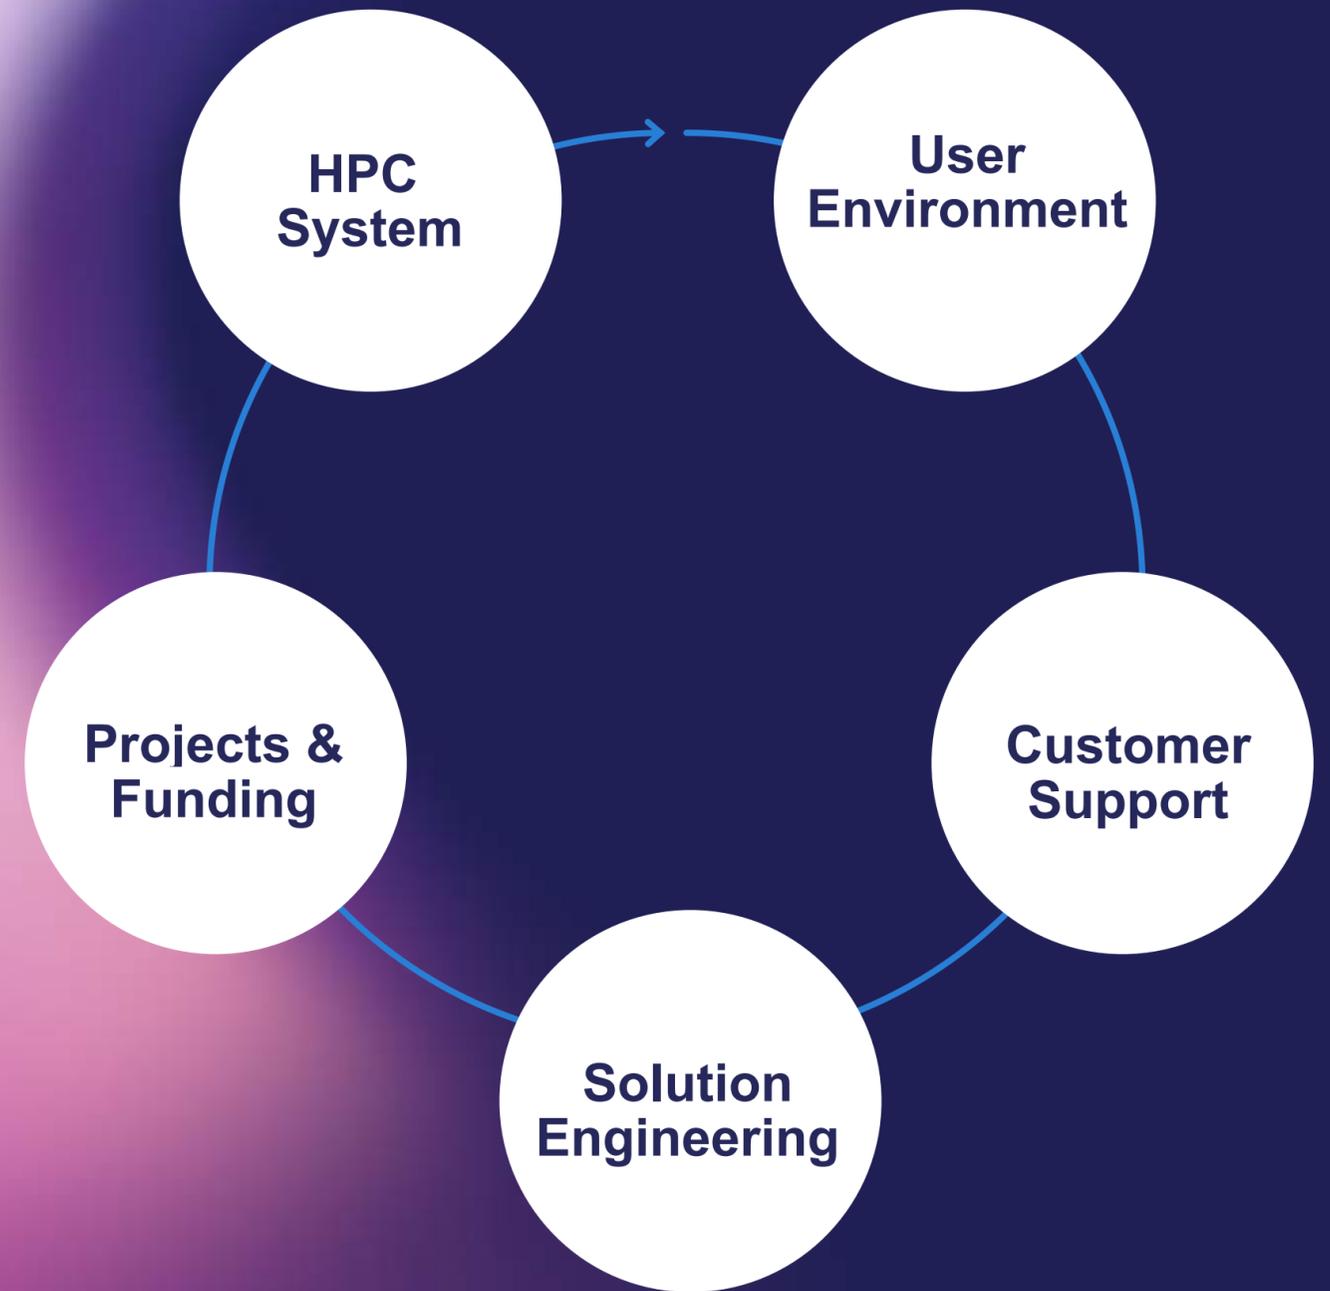
---

## Cloud Access

- PaaS and SaaS
- Remote visualization
- Specialized software
- Dedicated VMs



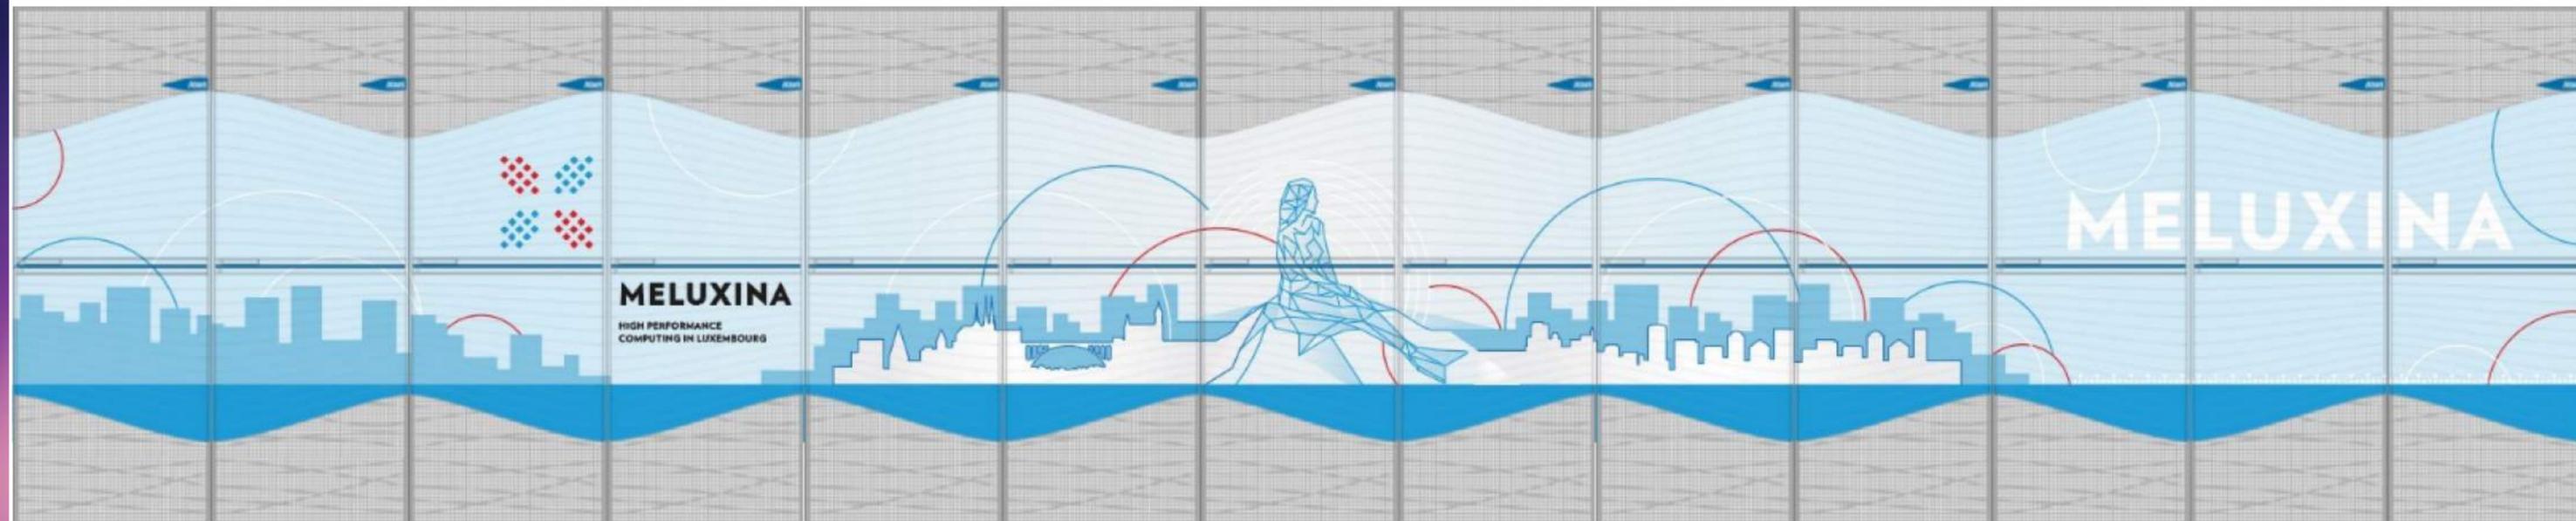
Luxembourg's one-stop shop for  
**high performance**  
computing and data analytics



# Thank you for your kind attention!

Sparked your interest?

For more information, reach us at [info@lxp.lu](mailto:info@lxp.lu)

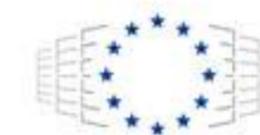


## Acknowledgements & Disclaimer

The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation program, as well as the Grand Duché du Luxembourg.

This publication only reflects the authors' view and the EuroHPC Joint Undertaking is not responsible for any use that may be made of the information it contains.

Names and logos used in this presentation are property of their respective owners.



**EuroHPC**  
Joint Undertaking

