The EuroHPC Flagship Supercomputer of the North

Dr. Pekka Manninen
Director, Science and Technology
Advanced Computing Facility
CSC – IT Center for Science, Finland

Outline

- EuroHPC initiative & LUMI
- LUMI technical overview
- LUMI science cases

LUMI

The EuroHPC Initiative

- The **EuroHPC Joint Undertaking** pools EU and national resources in high-performance computing (HPC)
  - **acquiring and providing a world-class supercomputing and data infrastructure** for Europe's scientific, industrial and public users
  - supporting an ambitious **research and innovation agenda**
- The EuroHPC declaration has been signed by **32 European countries**
- The first generation of EuroHPC systems announced in June 2019
  - 3 pre-exascale systems to Finland, Italy and Spain
  - 5 petascale systems to Czech Republic, Bulgaria, Luxembourg, Portugal and Slovenia
- Next generations of systems planned for 2024-2025 and 2027-2029

LUMI Consortium

- Unique consortium of 10 countries with strong national HPC centers
- The resources of LUMI will be allocated per the investments
- The share of the EuroHPC JU (50%) will be allocated by a peer-review process (cf. PRACE Tier-0 access) and available for all European researchers
- The shares of the LUMI partner countries will be allocated by local considerations and policies – seen and handled as extensions to national resources

LUMI Datacenter in Kajaani

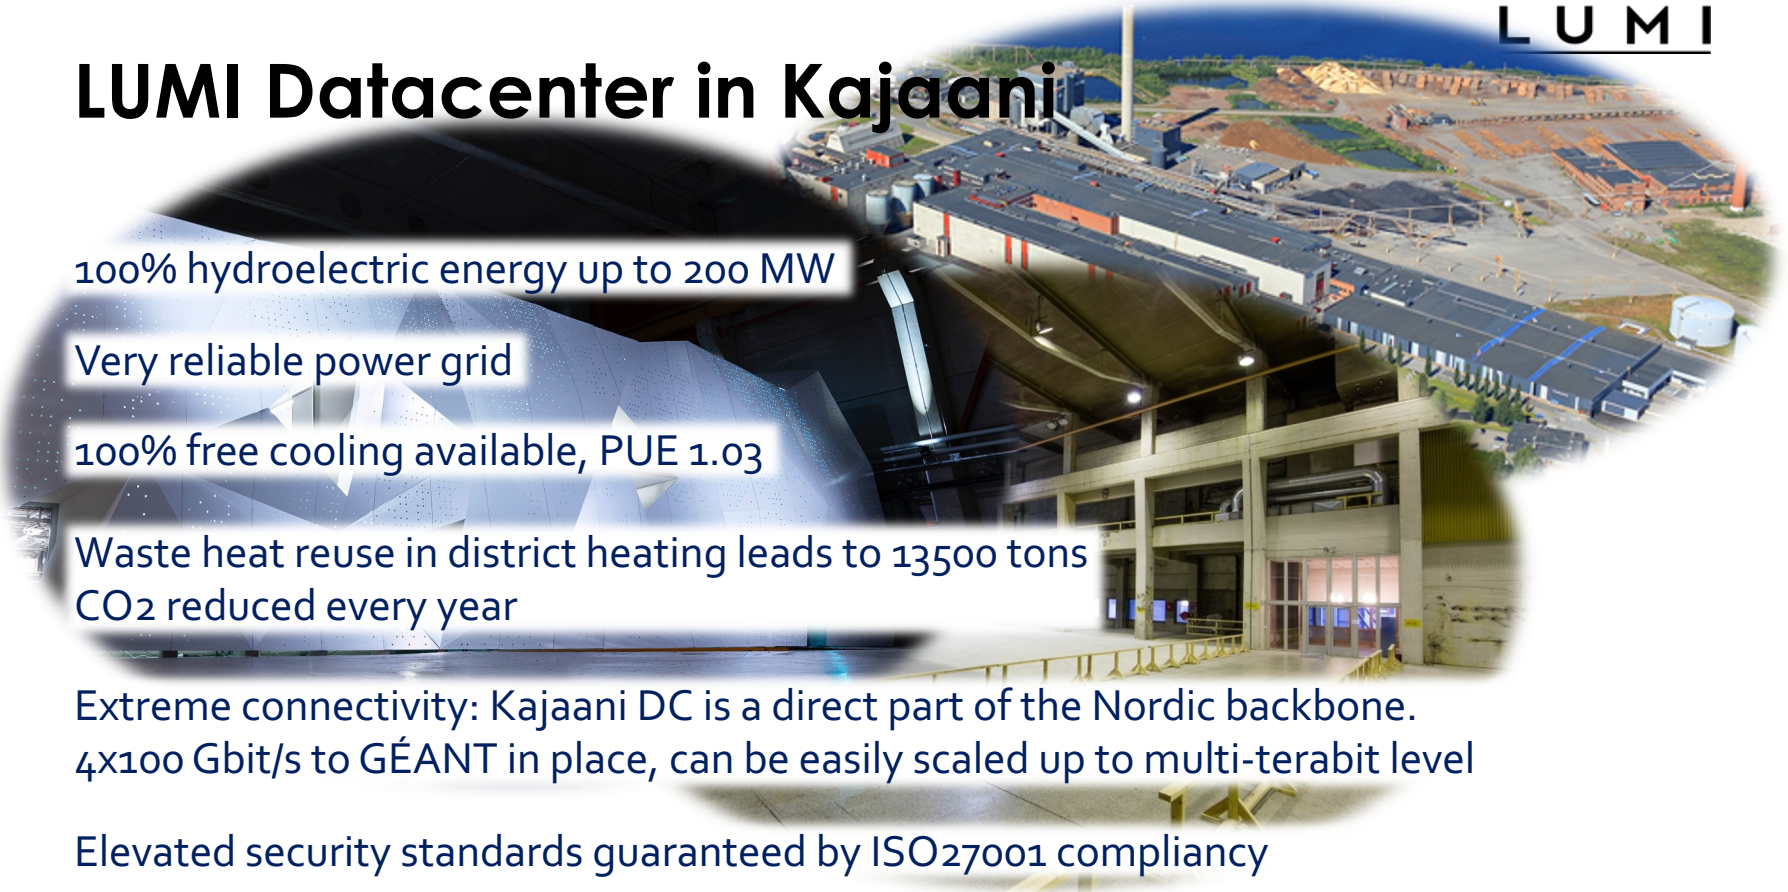

100% hydroelectric energy up to 200 MW

Very reliable power grid

100% free cooling available, PUE 1.03

Waste heat reuse in district heating leads to 13500 tons CO_2 reduced every year

Extreme connectivity: Kajaani DC is a direct part of the Nordic backbone. 4x100 Gbit/s to GÉANT in place, can be easily scaled up to multi-terabit level

Elevated security standards guaranteed by ISO27001 compliancy

An overview on LUMI technical architecture

LUMI

LUMI: one of the fastest supercomputers in the world

- LUMI is an **HPE Cray EX** supercomputer manufactured by **Hewlett Packard Enterprise**
- HPL performance over **309 petaflop/s** makes the system one of the world's fastest
  - #3 in Top500, #2 HPL-MxP, #3 HPCG

1 system

309 Pflop/s

Sustained performance

Computing power equivalent to

1 500 000

Modern laptop computers

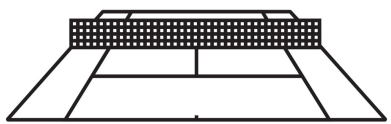

Size of two tennis courts

Modern platform for

High-performance computing, Artificial intelligence, Data analytics

Based on GPU technology

LUMI, the Queen of the North

LUMI is a Tier-0 ***GPU-accelerated supercomputer*** *that enables the convergence of* ***high-performance computing***, ***artificial intelligence***, *and* ***high-performance data analytics.***

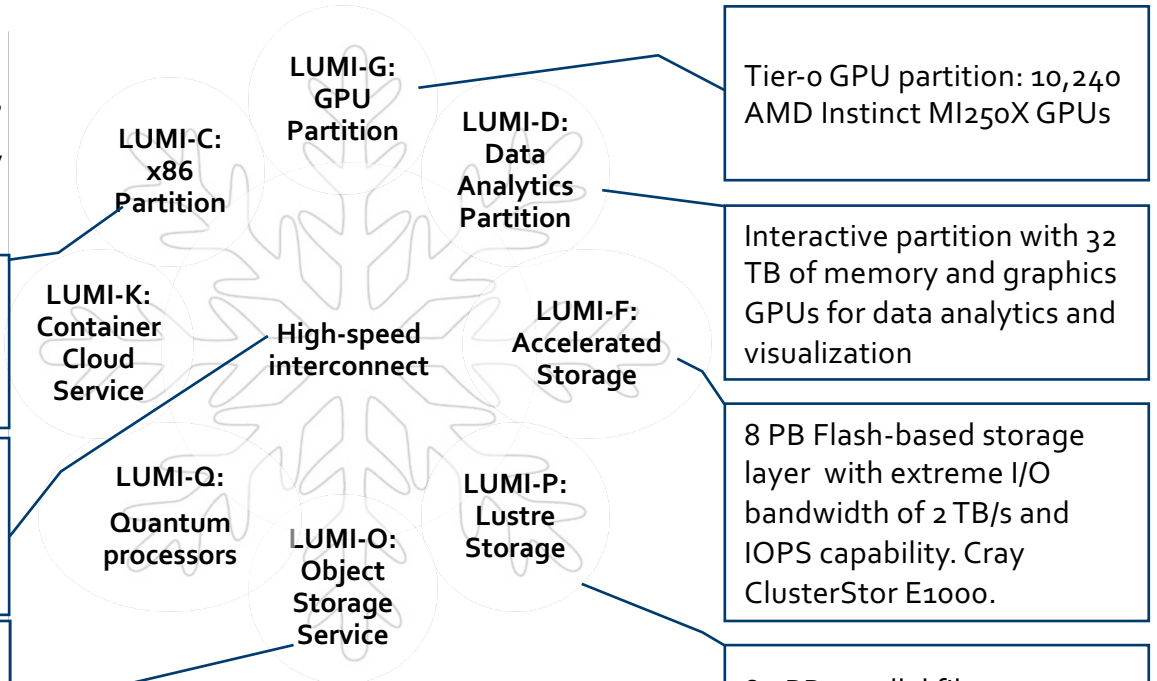

Supplementary CPU partition, ~200,000 AMD EPYC CPU cores

Possibility for combining different resources within a single run. HPE Slingshot technology.

30 PB encrypted object storage (Ceph) for storing, sharing and staging data

Tier-0 GPU partition: 10,240 AMD Instinct MI250X GPUs

Interactive partition with 32 TB of memory and graphics GPUs for data analytics and visualization

8 PB Flash-based storage layer with extreme I/O bandwidth of 2 TB/s and IOPS capability. Cray ClusterStor E1000.

80 PB parallel file system

LUMI compute node configurations

LUMI-G

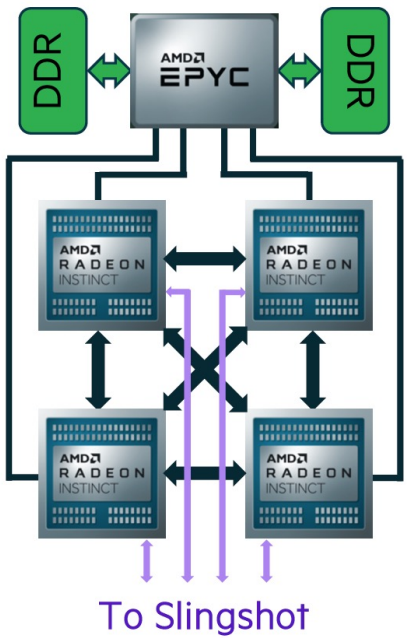

2560 nodes with 4 x MI250X + 1 x AMD Trento processor, 512 GB host memory and 512 GB device memory (HBM2)
4 x 200 Gbit/s NIC
Infinity Fabric

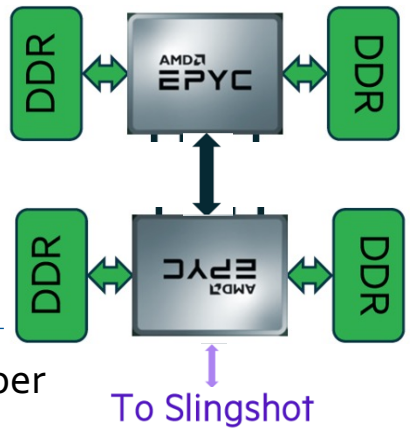

2x 64-core AMD Milan processors per node
1376 nodes with 256 GB, 128 with 512 GB and 32 with 1 TB
1 x 200 Gbit/s NIC

LUMI-C

LUMI timeline

System procurement:
November 2019 – August 2020

Data center preparation ready Q1/2021

2020

2021

2022

Q2-3/2022
- Gradual deployment of LUMI-G
- LUMI-C NIC upgrade

2023

LUMI-O
In customer use 12/04/23

2024

1st phase
Q4/2021
- LUMI-C, P, F, D
- Early Access Platform

In customer use 01/01/22

2nd phase

In customer use 19/12/22

3rd phase
07/23

LUMI

Enhanced user experience

- High-level interfaces on LUMI: Jupyter Notebooks, Rstudio and such to back-end to LUMI compute nodes (04/23)
- A rich stack of pre-installed software
- Datasets as a Service: curated large reference datasets available and maintained
- Support for handling data needing elevated security (GDPR subjected, IP-closed, etc) (12/23)

LUMI capacities, a brief summary

- Extreme computing capacity based on LUMI-G and LUMI-C partitions
  - LUMI queue policies will support jobs from single node to 50% of the nodes, even 100% with special arrangements
  - Jobs can combine resources from both within the same workflow, even within the same executable
- Interactive use (visualization, data analysis, pre/post processing,..) on LUMI-D
- Broad stack of pre-installed scientific software and datasets, both commercial and community
- Sharing datasets over LUMI-O service
- Running microservices on LUMI-K
- Exploring the quantum computing world with LUMI-Q

LUMI science cases

Enabler of world-class scientific breakthroughs

LUMI

LUMI is designed as a 'Swiss army knife' targeted for **a wide spectrum of use cases and user communities**.

- **Climate research:** More precise climate models and the interconnection of different climate models – digital twins of Earth ICON
- **Data science**: analyzing and re-analyzing large data sets (simulated and measured) e.g. in atmospheric science, environmental science, climate modelling, material science and linguistics.
- **Plasma physics**: Predicting and preparing for the societal effects of extreme space weather events. Multi-scale modeling of fusion reactors. Vlasiator, GENE
- **Life sciences:** enabling calculation of protein function, structural protein-protein interactions. Gromacs
- **Materials science:** quantum-mechanical simulations with global impact are development of better energy storage materials, more efficient solar cells, and better catalyst materials. CP2K, GPAW
- **Humanities and social sciences:** Natural language processing. Large-scale data analytics from social networks and the modelling of complex societal phenomena.
- Fast-track for **urgent computing** needs in time- and mission-critical simulations, e.g. related to national or EU security, or other major crisis e.g. pandemia.

Early showcases: Large Language Models

LUMI

- Several ongoing Lumi projects to train GPT-3 level language models of various European languages
  - Finnish, Swedish, Norwegian, Estonian, English
- For instance, the pilot use of Lumi pre-trained a 13B parameter GPT-3 and a 176B parameter Bloom model – by far the largest language model of Finnish to date
  - Exhausting all digital material in Finnish
- We are working on to provide an API for instructional LLM and open-source foundational LLMs
- Democratization of generative AI

https://www.lumi-supercomputer.eu/research-group-created-the-largest-finnish-language-model-ever-with-the-lumi-supercomputer/

LUMI Phase 2 pilot projects

See https://www.lumi-supercomputer.eu/second-pilot-phase-projects-selected/

The projects for the pilot phase of LUMI-G were chosen from the LUMI consortium countries (max 3 per consortium country), presenting various disciplines

- Astrophysics
- Natural language processing
- Materials science
- Biophysics
- Deep learning
- Climate modeling
- Computational fluid dynamics

LUMI Applications Readiness Program

- Porting applications as a joint project between the developer group, HPE, AMD and LUST
- First round applications
  - Quantum Espresso (materials science)
  - Megatron-LM (linguistics)
  - GPAW (materials science)
  - tmLQCD (particle physics)
  - Kaldi (speech recognition)
- Next ARP round coming soon

Concluding remarks

- Europe back in HPC landscape with **EuroHPC Joint Undertaking**
- **LUMI, the Queen of the North**: leadership-class resource designed for a broad range of user communities and workloads, with an enhanced user experience
  - In full customer use since December 2022
- LUMI's capabilities already put into use in societally important science initiatives