

LUMI

A white Arctic wolf is the central focus, standing in a futuristic, blue-toned digital environment. The background is filled with vertical lines, data streams, and server racks, creating a high-tech, cybernetic atmosphere. The wolf is looking slightly to the right with a calm expression.

**The pan-European supercomputer
of the North**

Dr. Pekka Manninen
Director, LUMI
CSC – IT Center for Science, Finland

Outline

- The background story
- The opportunities and benefits for research, development and innovation offered by LUMI
- How to access LUMI?
- Q&A

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 2023-2024 and 2026-2027

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: Only one 2 min outage in 38 years

100% free cooling available, PUE 1.03

Waste heat reuse: effective energy price 35 €/MWh,
negative CO₂ footprint: 13500 tons 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



Benefits and opportunities for R&I by 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 **375 petaflop/s** makes the system one of the world's fastest
 - Partial system listed 05/22 with 152 Pflop/s, #3 Top500
 - #3 also in Green500 and HPCG



1 system


375
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

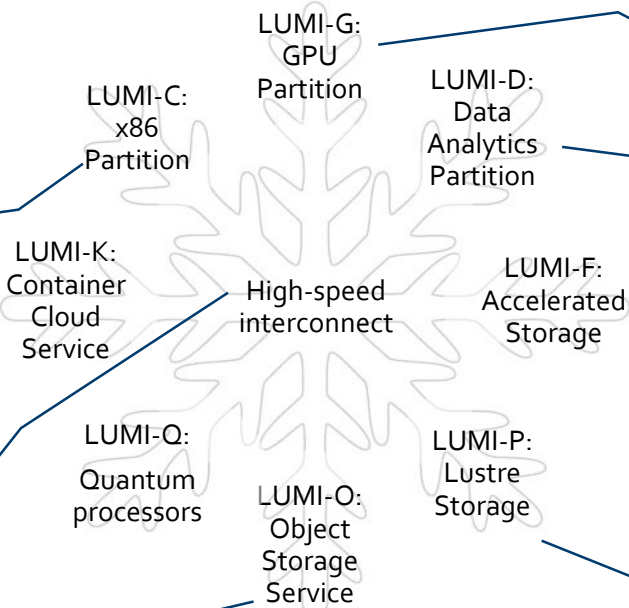
LUMI, the Queen of the North

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



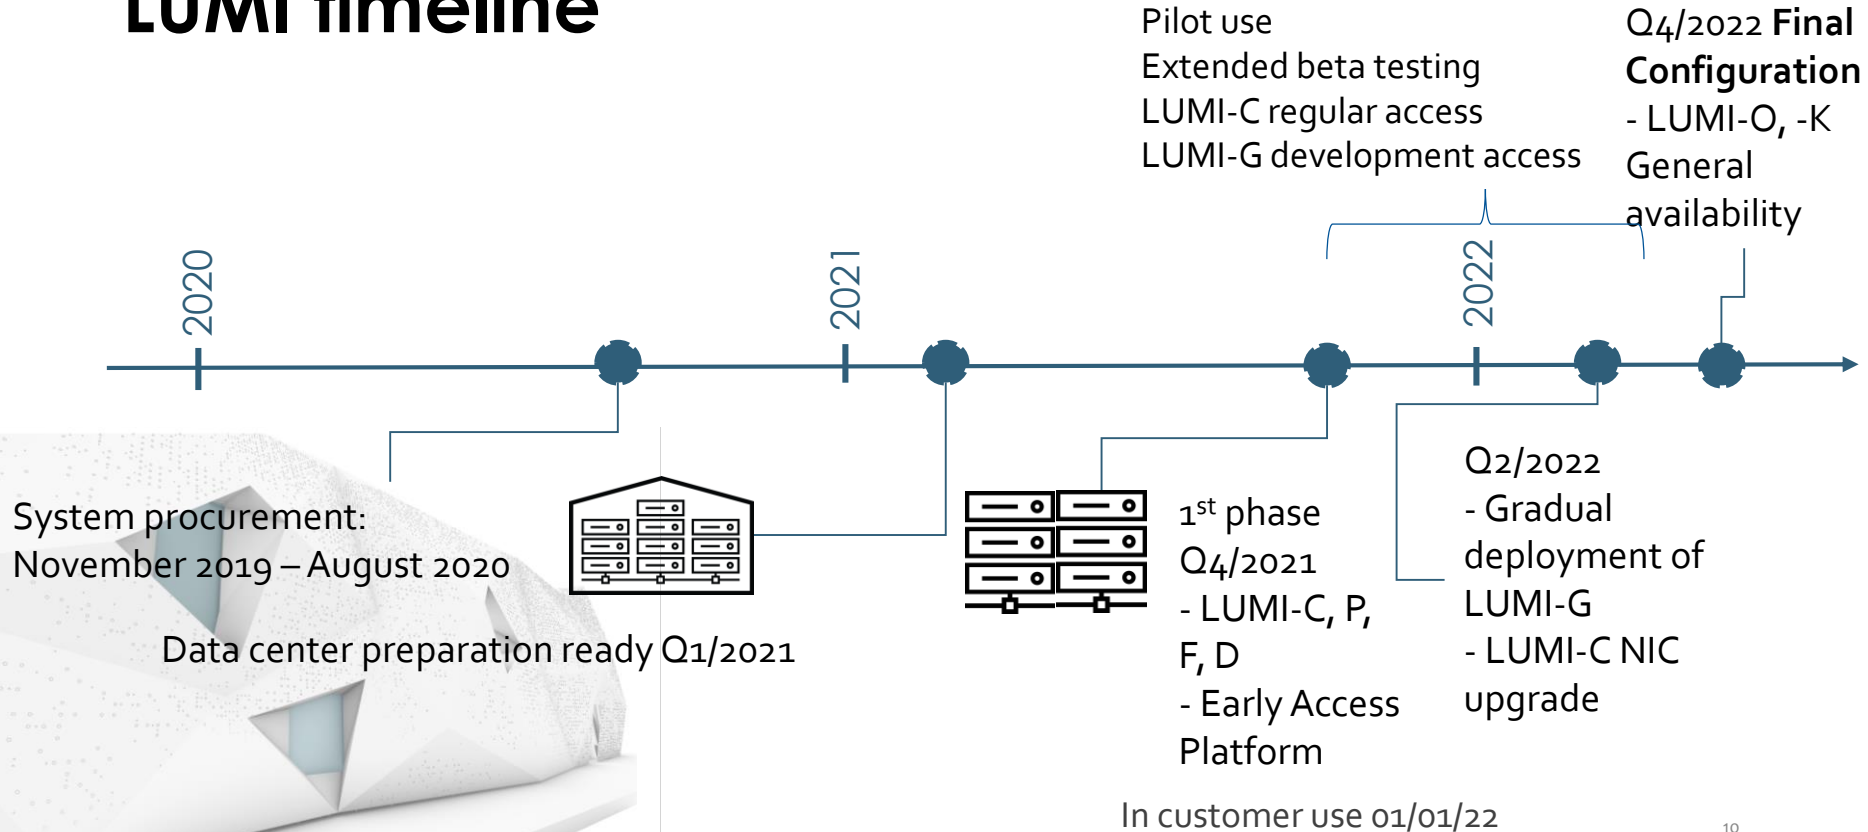
2560 nodes with 4 x MI250X + 1 x AMD Trento processor, 512 GB host memory and 512 GB device memory (HBM2)
4 x 200 Gbit/s NIC
Infinity Fabric

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

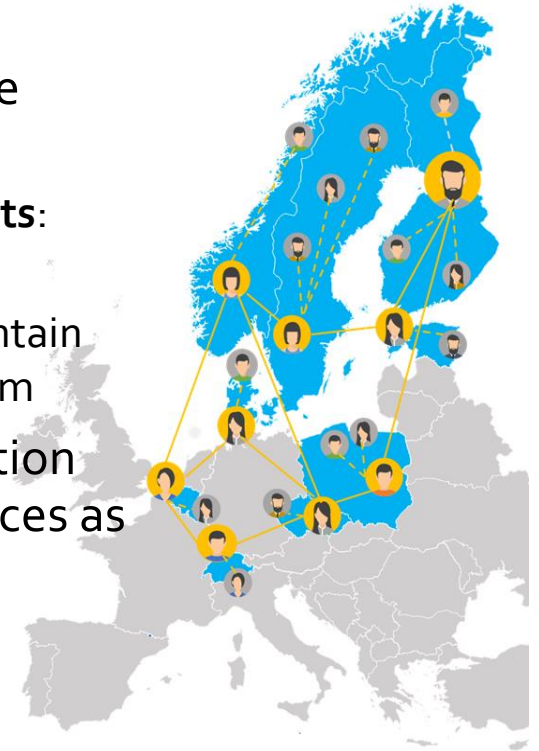


Enhanced user experience

- In addition to traditional CLI, high-level interfaces on LUMI, i.e. seamlessly integrate Jupyter Notebooks, Rstudio and such to back-end to LUMI compute nodes (Q3/22)
- A rich stack of pre-installed software (Q2/22)
- Datasets as a Service: curated large reference datasets available and maintained
- Support for handling data needing elevated security (GDPR subjected, IP-closed, etc) (Q2/23)

LUMI user support

- LUMI user support and a centralized help-desk by the distributed LUMI User Support Team
 - The model is based on a network of **dedicated LUMI experts**: each partner will provide one full-time person for the task
 - User Support Team will also provide end-user training, maintain the software portfolio and user documentation of the system
- CSC will be providing “Level 3” support (e.g. application enabling, methodology support) via its existing services as well as the EuroHPC Competence Center



How to access LUMI?



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

Getting LUMI resources

- European researchers can apply for LUMI resources via EuroHPC calls
- Researchers in the LUMI consortium countries can additionally apply from local resource providers
 - See www.lumi-supercomputer.eu/get-started
- LUMI resources are allocated in terms of GPU-hours, CPU-core-hours, and storage hours
 - Each project applies and gets a combination of this
 - No dedicated hardware - all users can access the whole system within the batch job policies
 - All consortium countries receive shares of these pools per their share of the TCO
- Resources brokered in terms of
 - Preparatory access projects (XS)
 - Development access projects (S)
 - General access (Tier-1) projects (M)
 - Extreme scale (Tier-0) projects (L) (should be mostly GPU hours)

LUMI programming environment

- ROCm (Radeon Open Compute)
 - Usual set of accelerated scientific libraries (BLAS, FFT etc)
 - Usual machine learning frameworks and libraries (Tensorflow, PyTorch etc)
 - Compilers for the GPUs
- Cray Programming Environment (CPE) stack
 - Cray Compiling Environment, LibSci libraries, CrayPAT, debuggers,...
 - CPE Deep Learning Plugin

Preparing applications and workflows for LUMI

- Remember the possibility of combining CPU and GPU nodes within one job – perhaps only part of the application needs to be GPU-enabled
- Convert CUDA codes to HIP
 - HIPify tools can automatize the effort (~25% code needs manual work)
- In case of major rewrites: Consider writing your application on top of modern frameworks and libraries
 - Kokkos, Alpaka etc, or domain-specific frameworks

Concluding remarks

- **EuroHPC era: Unprecedented amount of computational resources and capabilities** available for European research & innovation
 - Complemented by competence building and user support activities
- **LUMI, the Queen of the North:** leadership-class resource designed for a broad range of user communities and workloads, with an enhanced user experience
 - **LUMI is a GPU system**, which needs some preparatory work – but it is also a robust production system, and not experimental or esoteric in any manner
- **Modernizing HPC applications** for harnessing the largest systems is not trivial, and needs a lot of focused effort – but it will pay off