

LUMI supercomputer: Benefits and opportunities for research, development & innovation

Dr. Pekka Manninen

Director, LUMI

CSC – IT Center for Science, Finland

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

Enabler of world-class scientific breakthroughs

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: How will living conditions change when the climate is warming?
- **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.
- **Life sciences:** enabling calculation of protein function, structural protein-protein interactions.
- **Materials science:** quantum-mechanical simulations with global impact are development of better energy storage materials, more efficient solar cells, and better catalyst materials.
- **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. pandemic.

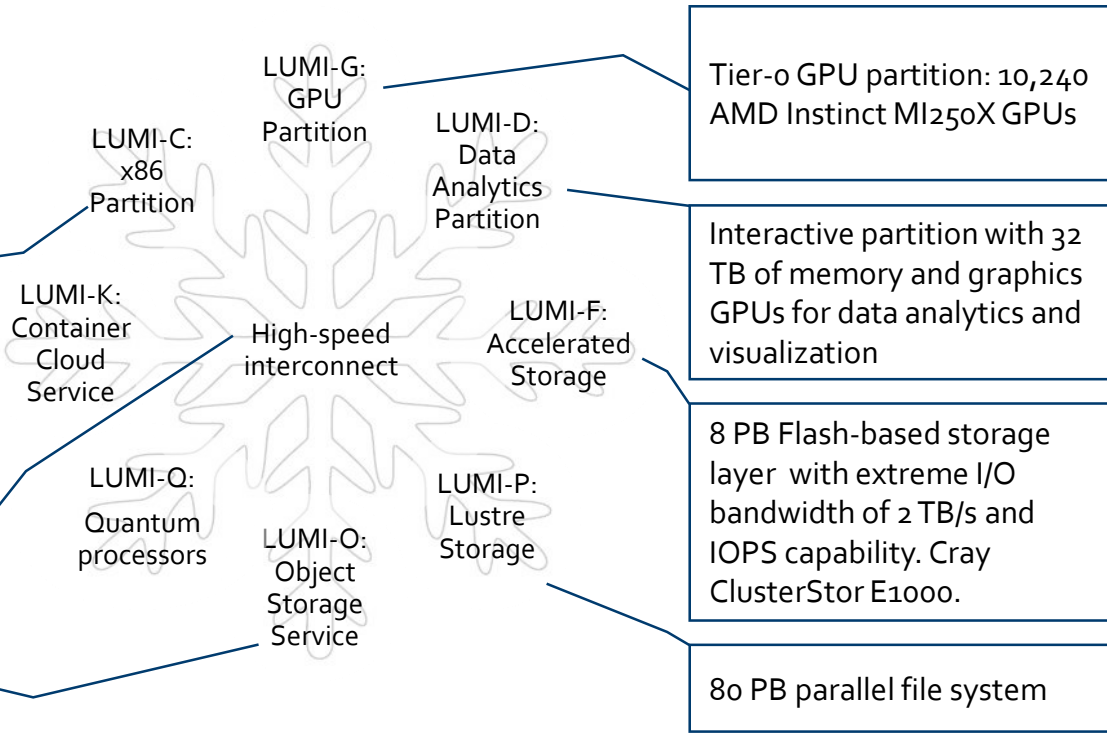
LUMI, the Queen of the North

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

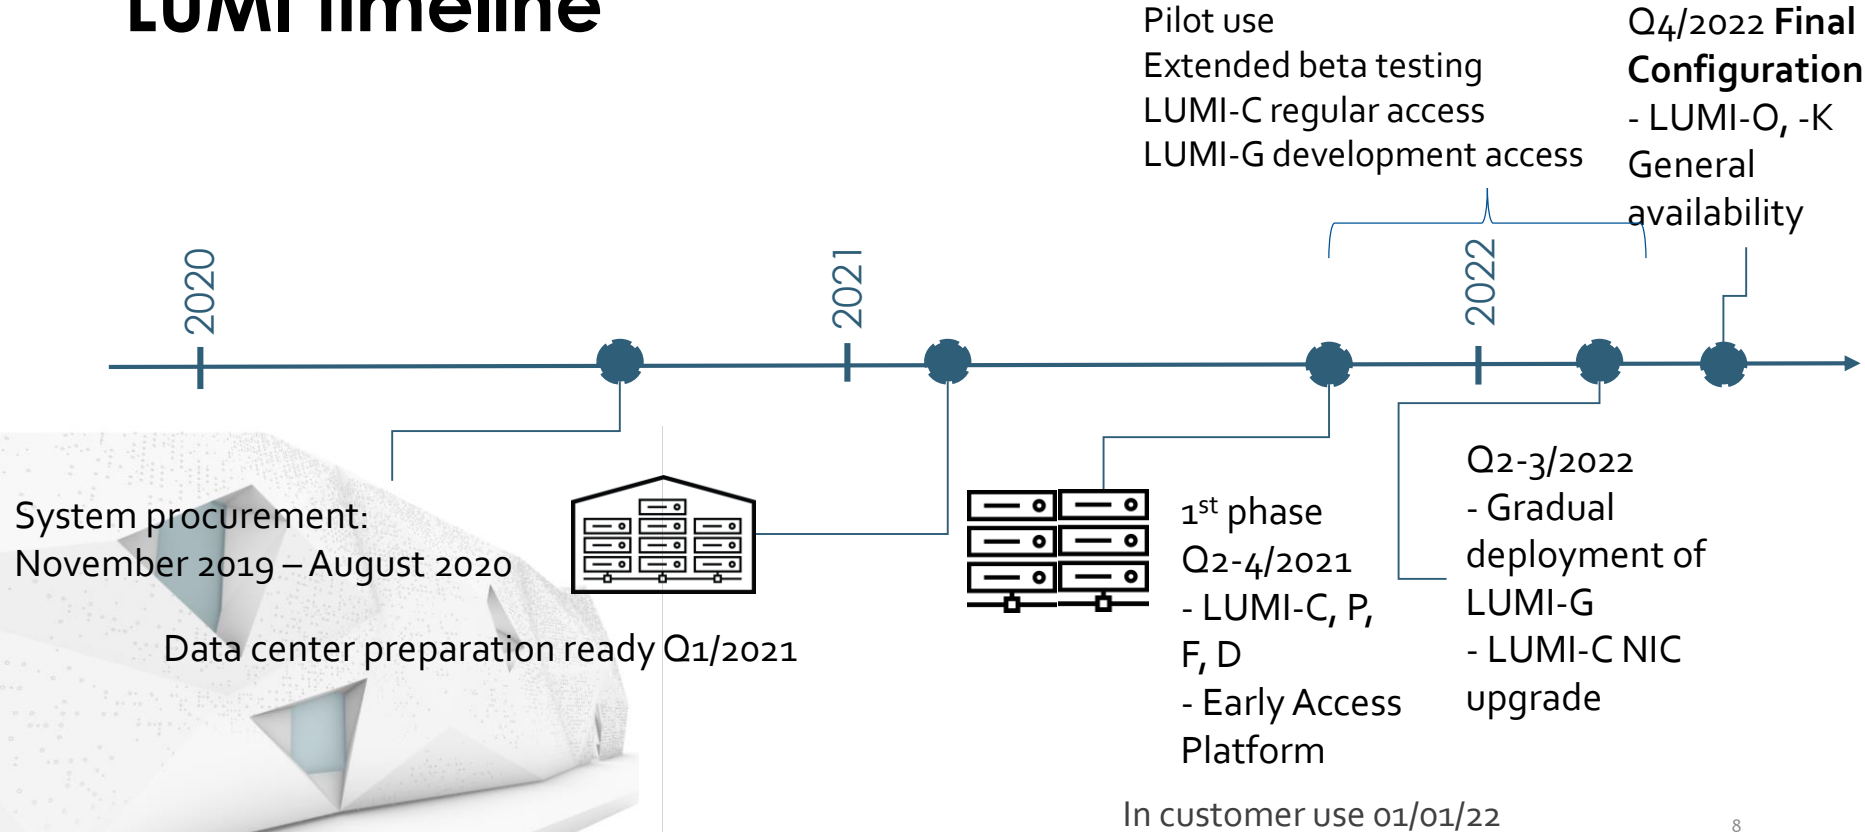
- Supplementary CPU partition
- ~200,000 AMD EPYC CPU cores

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

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



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 (Q4/22)
- A rich stack of pre-installed software (Q2/22)
- Datasets as a Service: curated large reference datasets available and maintained (Q1/23)
- Support for handling data needing elevated security (GDPR subjected, IP-closed, etc) (Q2/23)

LUMI features for AI workflows

- PyTorch and Tensorflow available as optimized multi-GPU versions
- Possibility for running Notebooks and Julia on compute partitions [wip]
- LUMI-K – a Kubernetes cluster having access to the parallel filesystems [wip]
 - Interactive + batch job orchestration
- Model and dataset repositories [wip]
- CI/CD pipelines for model and dataset versioning [wip]
- Interactive LUMI-D partition for visualisation of the progress and results of batch job runs

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)

Enterprise use of LUMI resources

- New R&D opportunities are available utilizing world class HPC resources combined with data-analytics and AI
- Novel co-operation possibilities for enterprises with HEIs and research centers
- Up to 20% of LUMI's capacity is reserved for enterprises
 - Pay-per-use access via EuroHPC JU or consortium countries
 - Some consortium countries have also other allocation mechanisms for industry
- Computing resources can also be used in co-innovation projects of enterprises and academia
 - Resources applied by academic PIs in these cases, free of charge

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 will be a robust production system, and not experimental or esoteric in any manner
- Harnessing the largest supercomputer systems is not trivial, and needs a lot of focused effort – but it will pay off