The European flagship supercomputer of the North – a status update

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

Adjunct Professor, University of Helsinki

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: one of the fastest supercomputers in the world

LUMI

- LUMI will be an **HPE Cray EX** supercomputer manufactured by **Hewlett Packard Enterprise**
- Peak performance over **550 petaflop/s** makes the system one of the world's fastest
  - Fastest today is Fugaku supercomputer in Japan with 513 petaflop/s, second fastest Summit in USA with 200 petaflop/s)

1 system

550 Pflop/s

Peak Performance

Computing power equivalent to

1 500 000

Modern laptop computers

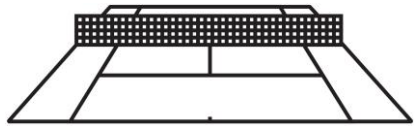

Size of a tennis court

Modern platform for

High-performance computing, Artificial intelligence, Data analytics

Based on GPU technology


RENFORSIN RANTA
RATA
LUMI
KONE

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

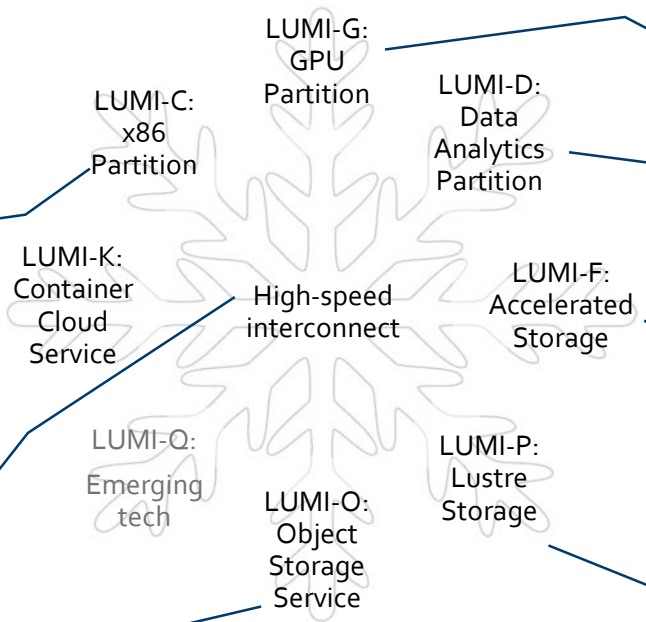

Tier-0 GPU partition: over 550 Pflop/s powered by AMD Instinct GPUs

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

7 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 timeline

LUMI

Pilot use
Extended beta testing
LUMI-C regular access
LUMI-G development access

Q2/2022 **Final Configuration**
- LUMI-O, (K)
General availability

2020

2021

2022

System procurement:
November 2019 – August 2020

Data center preparation ready Q1/2020

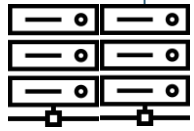

1st phase
Q4/2021
- LUMI-C, P, F, D
- Early Access Platform

Phase 1.5
Q1/2022
- 6 cabinets of LUMI-G
- LUMI-C NIC upgrade

LUMI

Enhanced user experience

- In addition to traditional CLI, we wish to support high-level interfaces on LUMI, i.e. seamlessly integrate Jupyter Notebooks, Rstudio and such to back-end to LUMI
  - Open OnDemand piloted, deployment on Lumi Q2/22
- Large software budget will enable a rich stack of pre-installed software
  - ISV applications is being procured and software being installed at the moment
- Datasets as a Service: curated large reference datasets available and maintained
- Support for handling sensitive (GDPR subjected, IP-closed, etc) data
  - Progressing in collaboration with the Elixir community

LUMI user support

- LUMI user support and a centralized helpdesk 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
- Helpdesk open, see **https://www.lumi-supercomputer.eu/user-support/**
- ”Level 3” support (e.g. application enabling, methodology support) via local centers as well as the EuroHPC Competence Centers

LUMI Phase 1 pilot projects

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

- Belgium: **M. Rasquin, T. Toulorge** and **K. Hillewaert ,** Direct Numerical Simulations of a Smoothed Backward Facing Step Featuring Incipient Separation
- Belgium: **G. Lapenta,** LIFTHRASIR: LumI First Tests of the HaRdware Architecture by Simulations of Interplanetary Regions
- Czech Republic: **D. Legut ,** Understanding the Physics of Phonons utilizing electronic structure and atomistic calculations employing state-of-the-art methods
- Czech Republic: **D. Biriukov,** Interaction network in extracellular space: an all-atom simulation model of the glycocalyx and cell membrane
- Denmark: **O. Christiansen** and **J. Elm:** High Performance Computing Quantum Chemistry on LUMI
- Denmark: **G. Frølund Pedersen** and **O. Franek** : Perfect Antennas for Reconfigurable Intelligent Surfaces
- Estonia: **T. Laisk** and **R. Mägi:** Genetic Susceptibility Factors of Different Traits
- Estonia: **G. Väli,** High-Resolution Ocean Model for Main Basins of The Baltic Sea
- Finland: **M. Palmroth,** LUmi – Carrington Kinetic simulations
- Finland: **S. Kaptan,** Molecular Basis of Intelligence
- Iceland: **H. Jónsson,** Computational Chemistry
- Iceland: **H. Úlfarsson,** Algorithmic Mathematics
- Norway: **M. Carlsson,** Solar Atmospheric Modelling
- Norway: **M. Bentsen,** Computing for nationally coordinated NorESM experiments
- Poland: **G. Wlazłowski,** Supersolidity in ultracold Fermi gas
- Poland: **J. M. Bujnicki,** Drugging the genomic RNA of SARS-CoV-2 with small molecules
- Sweden: **O. Agertz** and **F. Renaud,** The baryon cycle in colliding galaxies
- Sweden: **P. Schlatter** and **J. Vincent** and **J. Gong,** Turbulence data generation on Boeing hump
- Switzerland: **N. Marzari,** High-throughput computing

Getting LUMI resources

- 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 countries receive shares of these pools per their share of the TCO
- Resources brokered in terms of
  - Preparatory access projects (XS) – single-PI
  - Development access projects (S) – single-PI
  - General access (Tier-1) projects (M) – single-PI
  - Extreme scale (Tier-0) projects (L) – single-PI, should be mostly GPU hours
- Researchers affiliated to Danish institutions can apply from the EuroHPC allocation or from Denmark's own allocation

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, Reveal, debuggers,…
  - CPE Deep Learning Plugin
- More information:
  https://www.lumi-supercomputer.eu/may-we-introduce-lumi/

Preparing applications and workflows for LUMI

LUMI

- Remember the possibility of combining CPU and GPU nodes within one job – perhaps only part of the application needs to be GPU-enabled
- Consider writing your application on top of modern frameworks and libraries
  - Kokkos, Alpaka etc, or domain-specific frameworks
- Convert CUDA codes to HIP, OpenACC codes to OpenMP5
  - HIPify tools can automatize the effort
  - We have observed this to be very straightforward, more so than we expected
- LUMI 1st phase features a code porting platform (MI100 GPUs)

Concluding remarks

- **EuroHPC era: Unprecendent 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 will be a GPU system,** which needs some preparatory work – but it will be 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
  - **We are getting there**! LUMI-C in production soon, LUMI-G in Q2/22