LUMI

The EuroHPC pre-exascale system of the North

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

Adjunct Professor, University of Helsinki

16.9.2021

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

Countries which have signed the EuroHPC Declaration

LUMI Consortium countries

CSC Datacenter in Kajaani

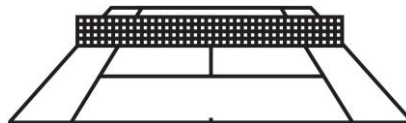# LUMI: one of the fastest supercomputers in the world

- 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)

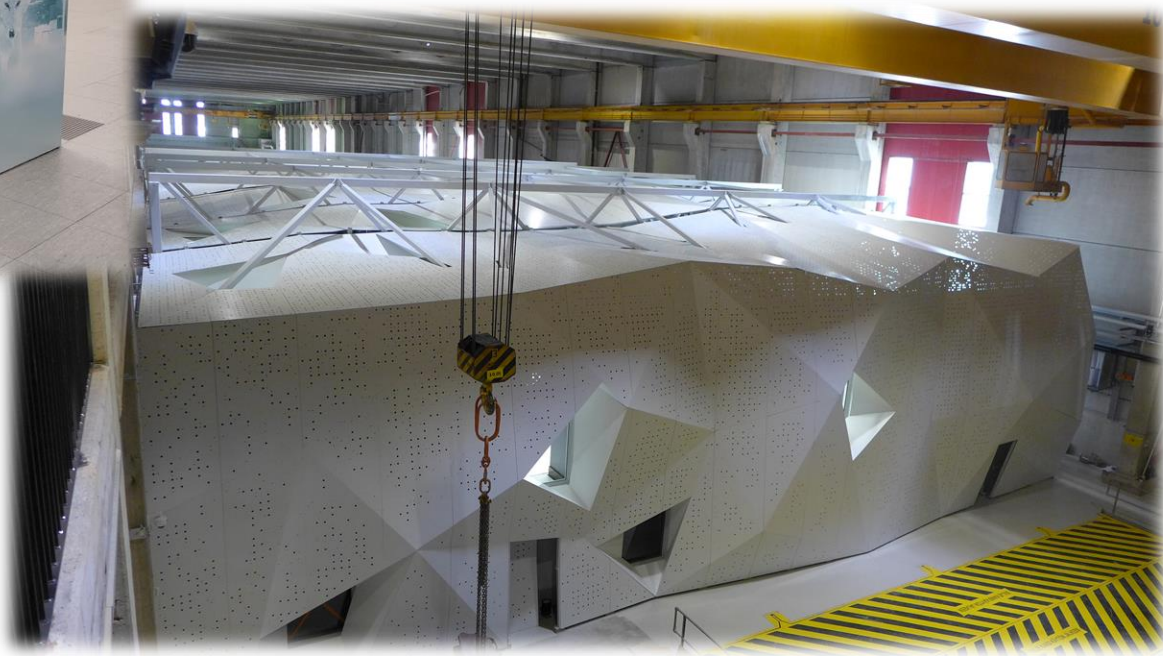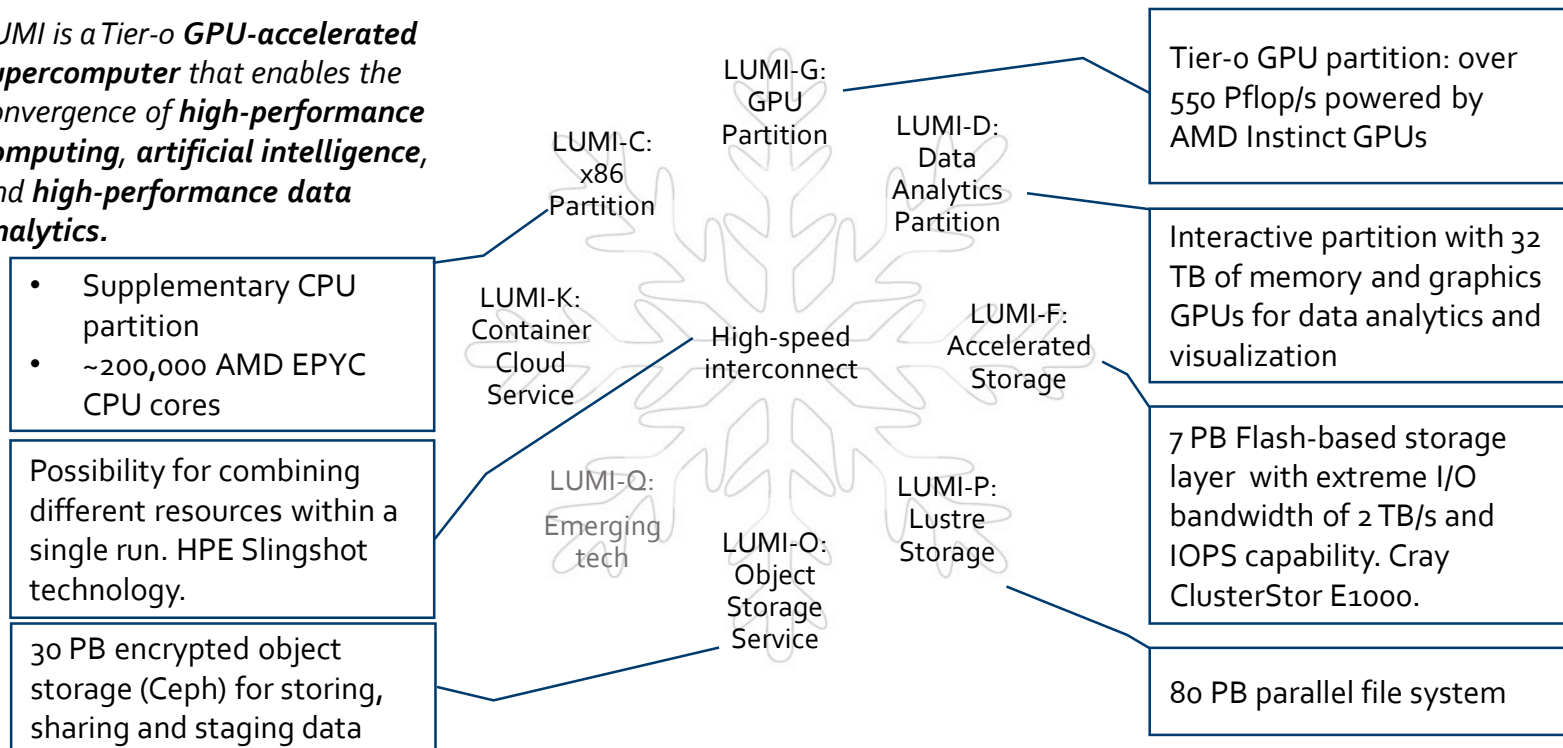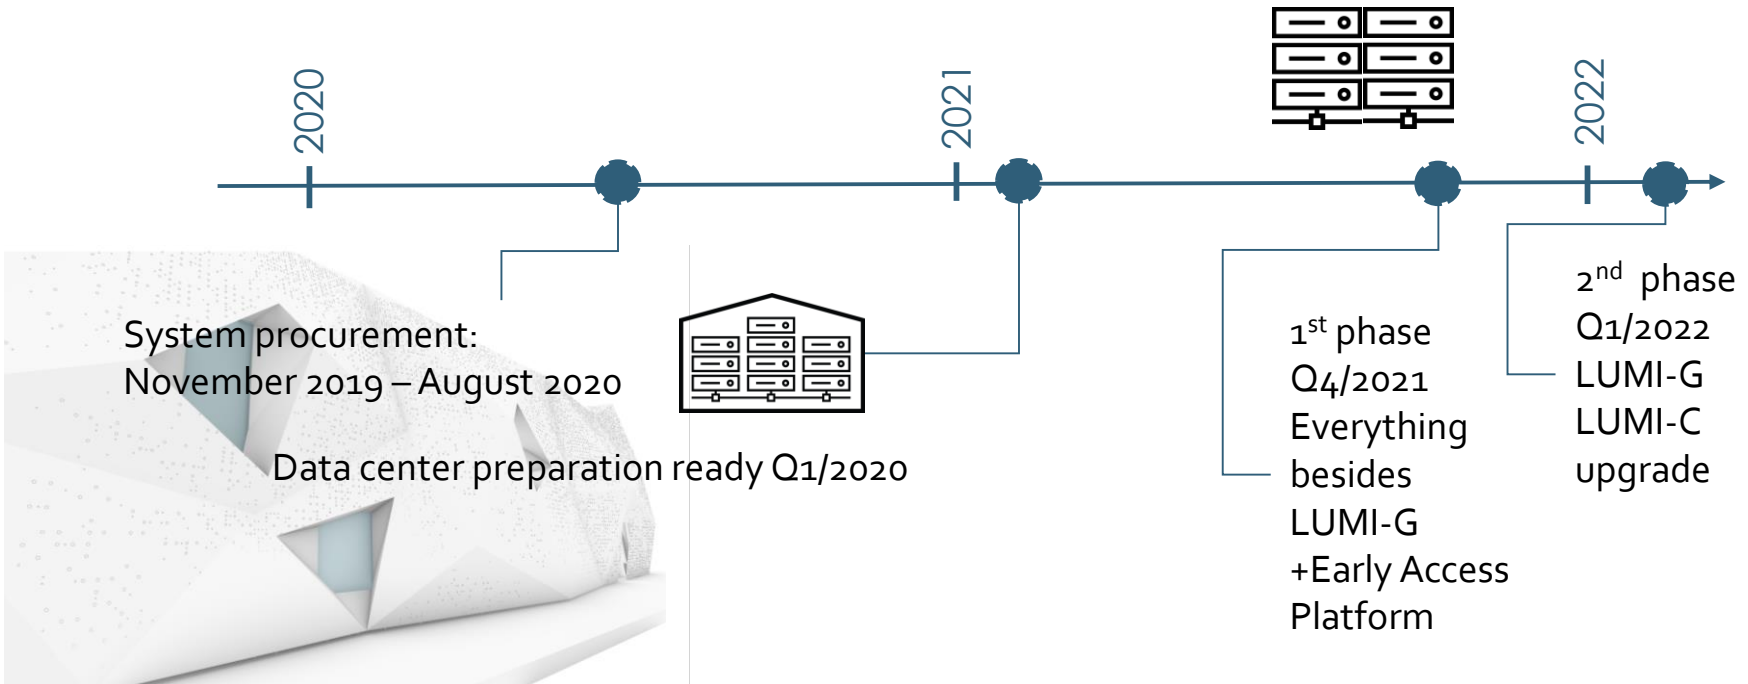|			
1 system **550** **Pflop/s** Peak Performance	**Computing power equivalent to** **1 500 000** **Modern laptop computers**	 **Size of a tennis court**	**Modern platform for** **High-performance computing, Artificial intelligence, Data analytics** **Based on GPU technology**

16.9.2021

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-C: x86 Partition

LUMI-K: Container Cloud Service

LUMI-Q: Emerging tech

LUMI-G: GPU Partition

LUMI-D: Data Analytics Partition

High-speed interconnect

LUMI-F: Accelerated Storage

LUMI-O: Object Storage Service

LUMI-P: Lustre Storage

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

2020

2021

2022

System procurement:
November 2019 – August 2020

Data center preparation ready Q1/2020

1st phase
Q4/2021
Everything
besides
LUMI-G
+Early Access
Platform

2nd phase
Q1/2022
LUMI-G
LUMI-C
upgrade

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

- Large software budget will enable a rich stack of pre-installed software

- Datasets as a Service: curated large reference datasets available and maintained

- Support for handling sensitive (GDPR subjected, IP-closed, etc) data

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
- "Level 3" support (e.g. application enabling, methodology support) via local centers as well as the EuroHPC Competence Centers

L U M I

LUMI programming environment

- ROCm (Radeon Open Compute)
  - Usual set of accelerated scientific libraries (BLAS, FFT etc)
  - Usual machine leaning 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

- LUMI 1st phase will come with a code porting platform (MI100 GPUs)
  - HIP porting can be done already now on Nvidia GPU platforms

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
  - It is time already to start preparing for the LUMI era

L U M I

Dr Pekka Manninen

Director

LUMI Leadership Computing Facility

CSC – IT Center for Science Ltd

pekka.manninen@csc.fi

tel. +358 50 3812 831

Follow us

Twitter: @LUMIhpc

LinkedIn: LUMI supercomputer

YouTube: LUMI supercomputer

www.lumi-supercomputer.eu
contact@lumi-supercomputer.eu

EuroHPC Joint Undertaking

Leverage from the EU 2014–2020

European Union European Regional Development Fund

Kainuun liitto