

# LUMI

A white 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 sense of a high-tech, data-driven world. The wolf is looking slightly to the right, and its fur is rendered with fine detail. The overall color palette is dominated by dark blues and greys, with the white of the wolf providing a strong contrast.

**The European flagship supercomputer  
of the North**

8.4.2021

**Dr. Pekka Manninen**  
Director, LUMI Leadership Computing Facility  
CSC – IT Center for Science, Finland  
Adjunct Professor, University of Helsinki

# The EuroHPC Initiative

- The **EuroHPC Joint Undertaking** will pool 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 CO<sub>2</sub> footprint: 13500 tons reduced every year

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

Elevated security standards guaranteed by ISO27001 compliancy




# 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

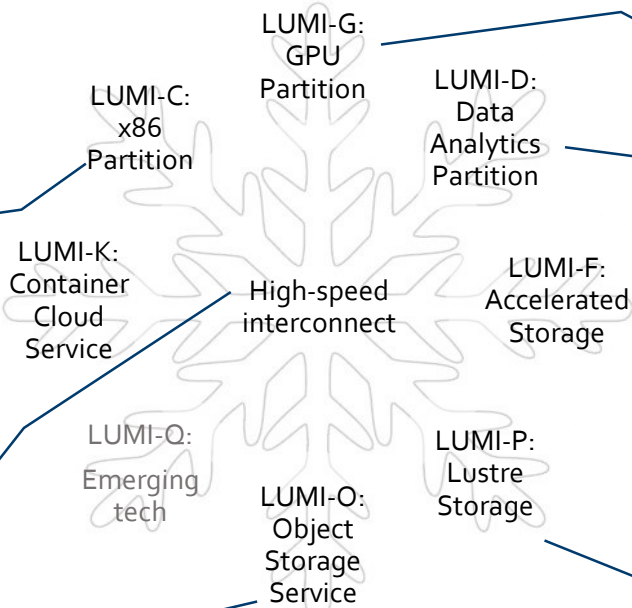
# LUMI, the Queen of the North

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 "Milan" 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



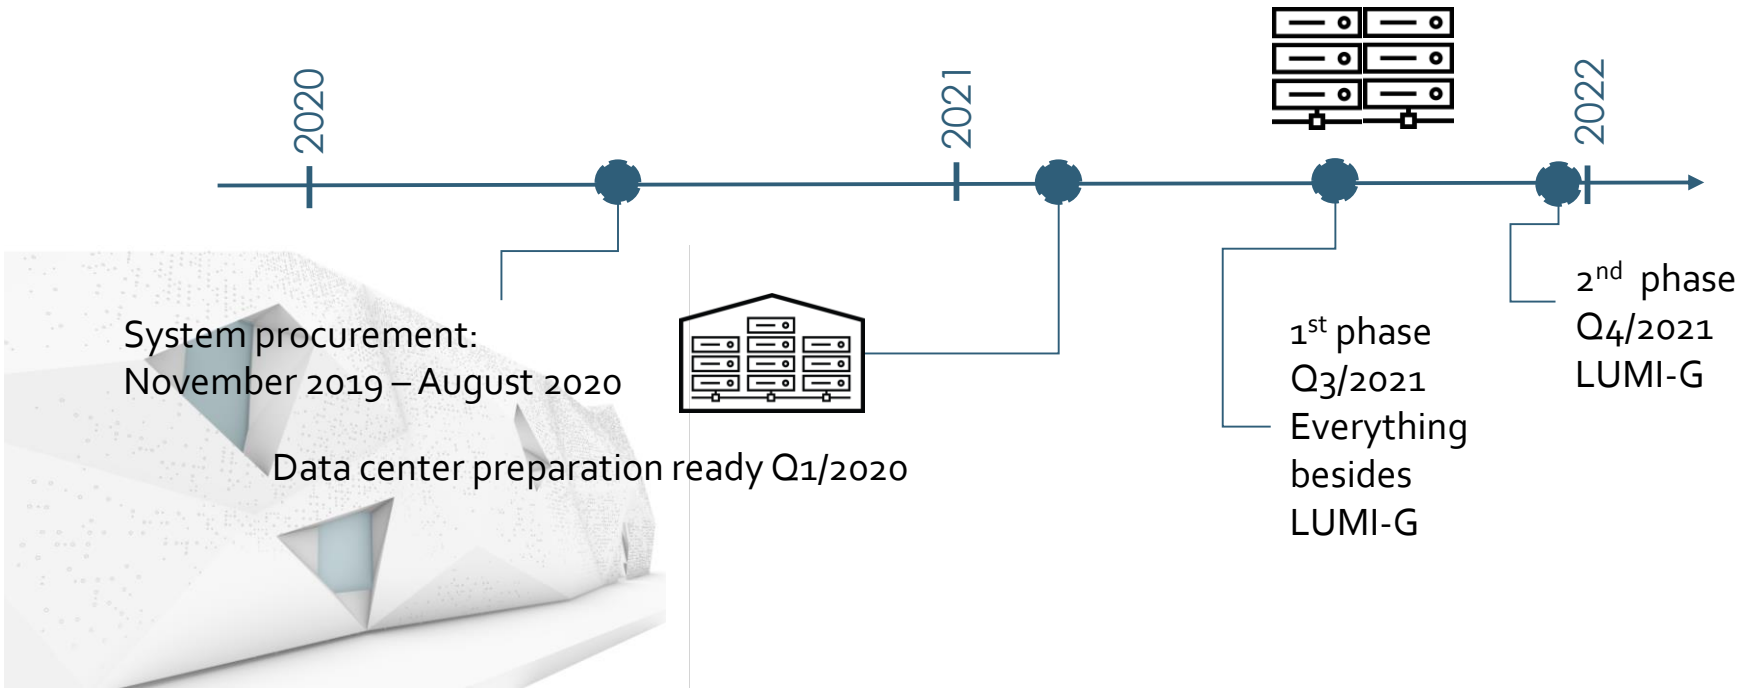
Tier-0 GPU partition: over 550 Pflop/s powered by AMD Instinct GPUs

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



# 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**.

- **More precise climate models** and the interconnection of different climate models: How will living conditions change when the climate is warming?
- Shedding light on the causes of illness and improving **personalized medicine**: sequencing and analyzing of full genomes combined with data analysis and correlations to clinical data
- **Artificial intelligence** (deep learning): analyzing large data sets (simulated and measured) and reanalyzing e.g. in atmospheric science, environmental science, climate modelling, material science and linguistics
- **Self-driving cars and vessels**: the study of related algorithms with previously unprecedented computing power
- **Social sciences**: large-scale data set 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.

# How to prepare for LUMI?

- Thinking projects and use cases for Lumi
  - Cases for Tier-0 grand challenges
  - Combining simulation and AI methods within the same workflow
- There is a vast pool of GPU-enabled community codes
  - See if your favorite software suite already has been enabled- if not, consider moving to a competing package that is
  - Perhaps only part of the application needs to be GPU-enabled, rest running on the CPU nodes
- Modernizing applications and GPU-enabling them
  - "even if it works, fix it"

# 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 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



# Pilot use

- The four-week pilot use periods (which coincide with the 28-day availability tests that are a part of the acceptance of the installations) should start
  - Early September 2021 for the phase 1
  - in late December 2021 for the phase 2
- 2 pilot projects from each LUMI consortium country for both phases
- The selected pilot projects can (and should) compute as much as they can during the pilot phases, no accounting will take place. No limitations requested resources.
- The pilots will have the following profiles
  - data-intensive computing, high-throughput computing and high-performance data analytics (bioinformatics, digital humanities, materials science, astronomy etc) for the phase 1
  - highly scalable GPU applications (“hero runs”) for the phase 2.