

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 overall aesthetic is clean and modern, with a strong emphasis on the color blue.

The pan-European supercomputer of the North

7.10.2020

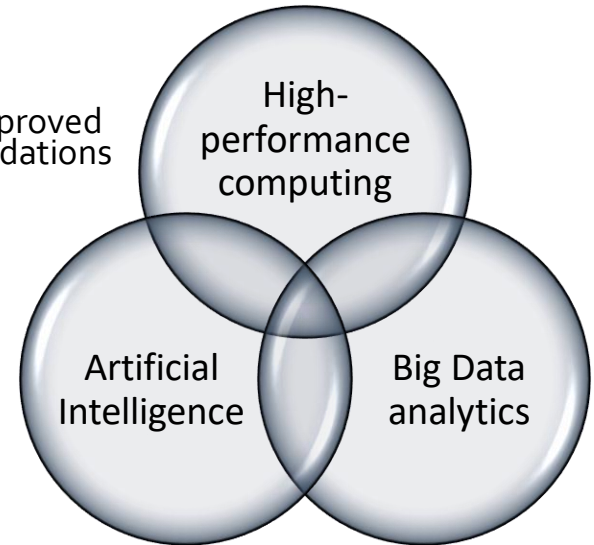
Dr. Pekka Manninen
Director, LUMI Leadership Computing Facility
CSC – IT Center for Science, Finland
Adjunct Professor, University of Helsinki

In the era of data-intensive computing

- The converge of high-performance computing, artificial intelligence and big data will be key for solving the great scientific and societal challenges
 - Understanding climate change, treatments of diseases, finding improved materials, finding clean sources of energy, understanding the foundations of reality,...
- For enhancing this kind of cross-fertilisation,
 - We need interconnect efficiency at all levels.
 - We need extreme I/O bandwidth.
 - We need extreme computing capacity.

We need supercomputers more than ever.

Today, we will discuss LUMI, one of the most powerful supercomputers on Earth – designed for the converge of HPC, AI and Big Data.



Outline

- The long story of LUMI and EuroHPC (made short)
- The opportunities and benefits for research, development and innovation offered by LUMI
- How to prepare for LUMI?
- Q&A

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

Why did Finland choose to apply for LUMI?

- The Finnish share of LUMI resources strengthens the national infrastructure for data management and computing, and thereby the capacities of the Finnish research community
- Finland has not got an international research infrastructure of this scale and visibility previously
 - Also a great asset in terms of employment and economy: LUMI is an investment of more than 200 M€ to Kajaani
- EuroHPC was seen as continuation of the long tradition of collaboration in scientific computing, which has been greatly beneficial for research in Finland

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

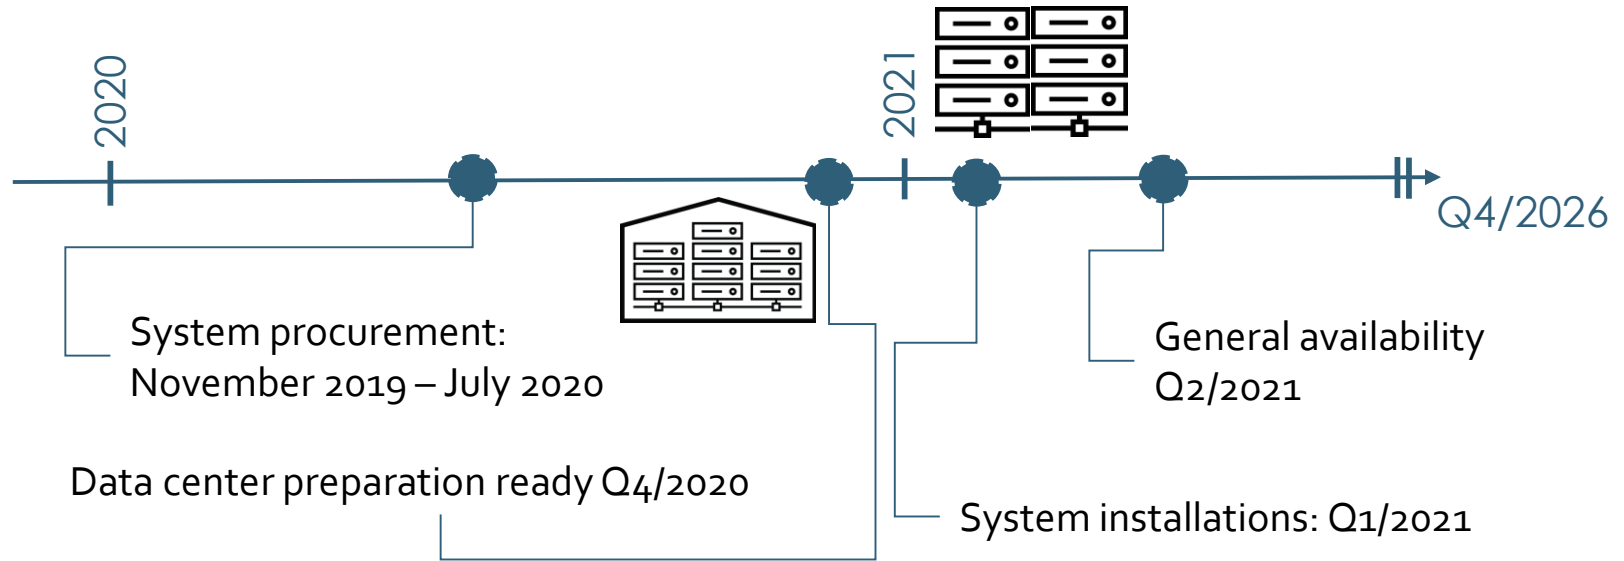
Extreme connectivity: Kajaani DC is a direct part of the Nordic backbone.
4x100 Gbit/s to GÉANT in place, can be easily scaled up to multi-terabit level

Elevated security standards guaranteed by ISO27001 compliancy



Benefits and opportunities for R&I by LUMI

LUMI Timeline



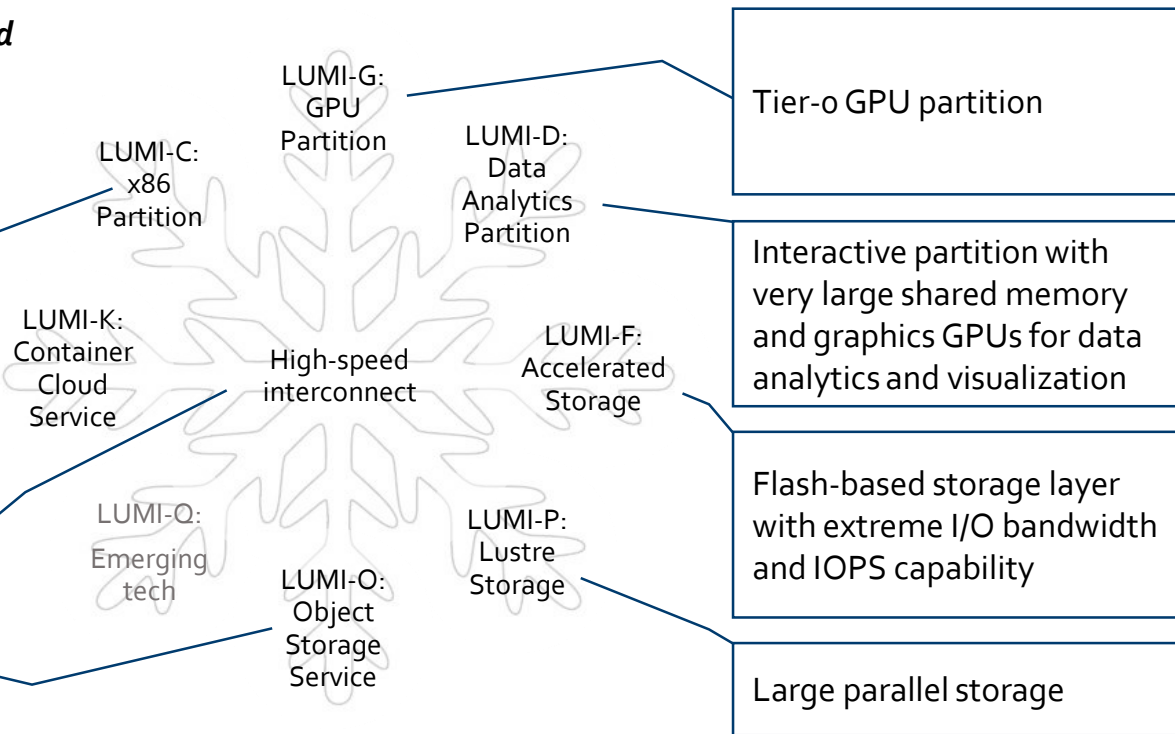
LUMI System Architecture

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

- Supplementary "Tier-1" CPU partition
- M, L and XL memory nodes

Possibility for combining different resources within a single run

Encrypted object storage (Ceph) for storing, sharing and staging data

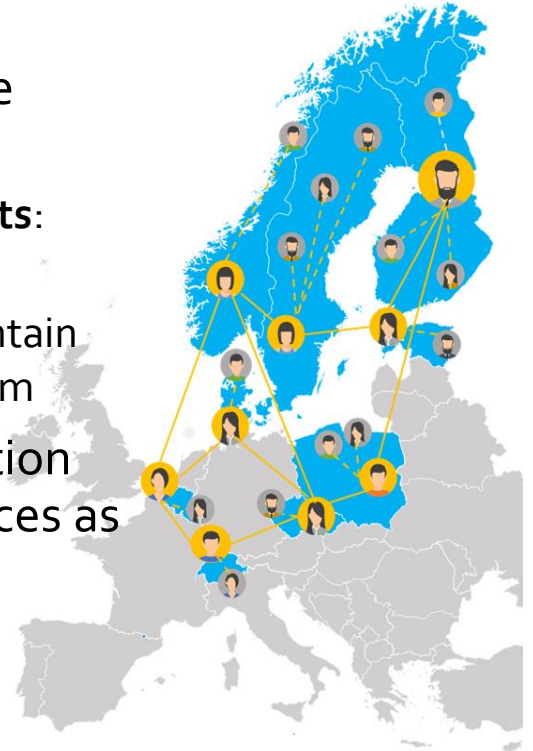


Enhanced user experience

- 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
 - Ultimately the LUMI resources being an extension to your laptop
- Large software budget will enable a rich stack of pre-installed software
- Datasets as a Service: curated large reference datasets available and maintained
- Interactive resources for steering batch job simulations

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



Enterprise use of LUMI resources

- Up to 20% of LUMI's capacity is reserved for industry and SMEs
- New RDI opportunities are available utilizing world class HPC resources combined with data-analytics and AI
- Novel co-operation possibilities for enterprises with universities and research centers
- Business Finland's Business AI program provides free-of-charge computing resources for startups and SMEs for significant research projects
- Computing resources can also be used in co-innovation projects of companies and academia
 - Resources arranged through academic PIs in these cases

LUMI catalyzing international collaboration

- LUMI offers world-class computing and data management resources for Finnish HEIs, research institutions and enterprises
- LUMI can be used as an asset when recruiting researchers
- The LUMI resources can be used to secure participation and major roles in international projects
- LUMI consortium will actively promote collaboration opportunities
- The distributed user support model will bring new competencies available for Finnish users

Finnish scientific computing infrastructure

LUMI

National (DL2021)



EuroHPC/LUMI

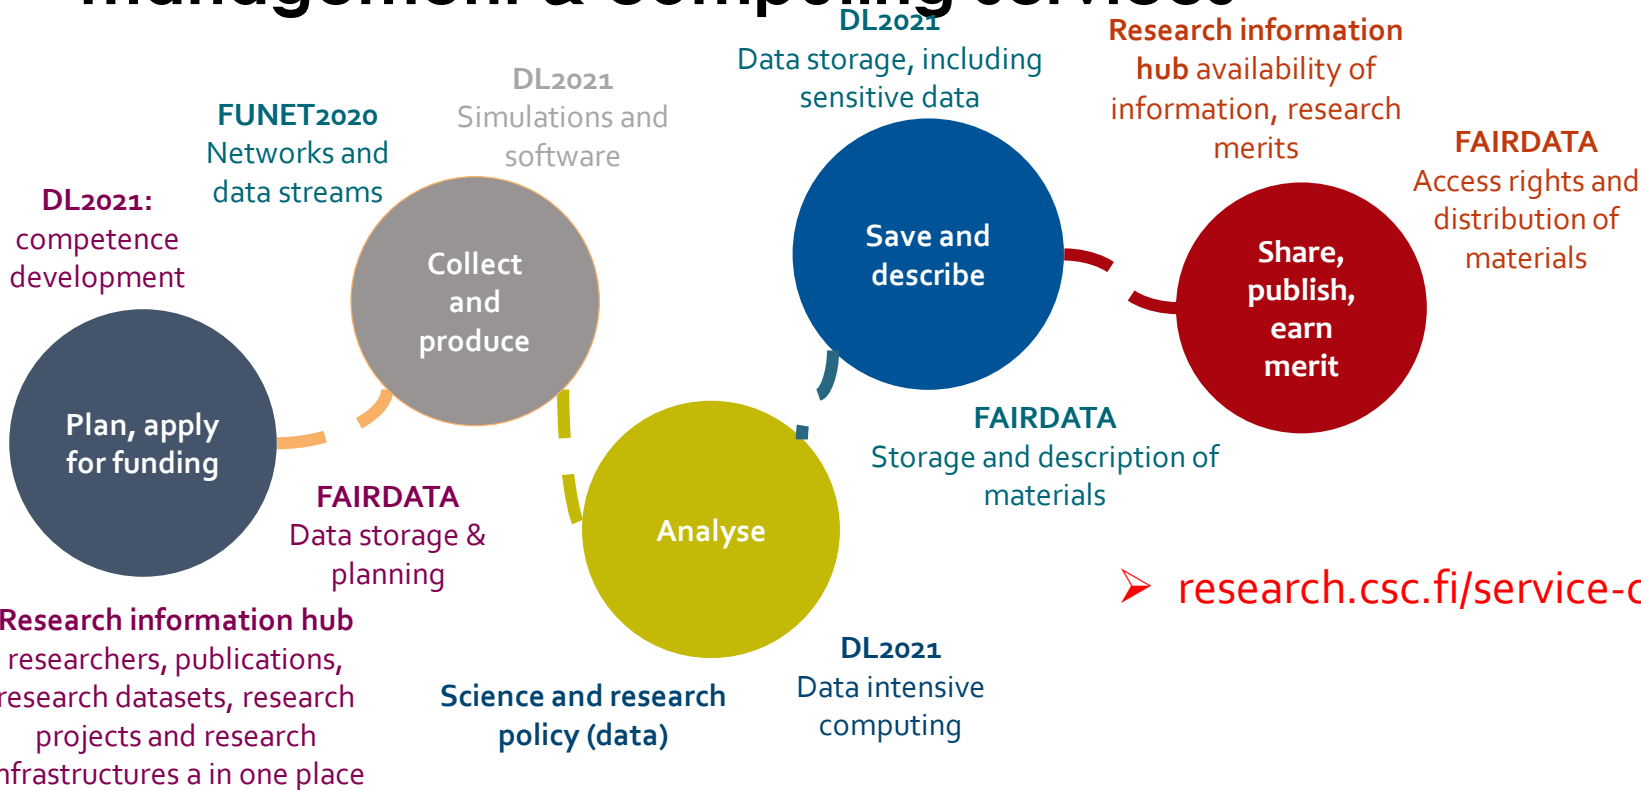


For all use cases in scientific computing in Finland

- Large (tier-1) scale simulations
- High throughput computing
- High-performance data analytics
- Data streams

- Most resource-intensive projects (Tier-0)
- Science and innovation policy priorities (flagships, CoEs)
- Artificial intelligence
- Support for large research infras
- Collaboration with industry
- International collaboration

CSC offers wide selection of data management & computing services



➤ research.csc.fi/service-catalogue

How to prepare for LUMI?

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
 - Community access projects (XL) – multi-PI, multi-year

LUMI resource channels in Finland

Access type Brokerer	GPU node hours	CPU node hours	Storage hours
Extreme Scale (Tier-0) CSC / Scientific Customer Panel	40%	15%	20%
General Access (Tier-1) CSC / Resource Allocation Group	10%	30%	25%
Community access CSC / Resource Allocation Group	25%	30%	30%
Innovation funding for SMEs Business Finland	20%	20%	20%
Director's share LUMI	5%	5%	5%

Exact percentages still tentative

LUMI capacities, a brief summary

- Extreme computing capacity based on LUMI-G and LUMI-C partitions
 - LUMI queue policies will support jobs from single CPU core or a GPU to 50% of the nodes, even 100% with special arrangements
 - Jobs can combine resources from both sides within a workflow, even within the same executable
- Interactive use (visualization, data analysis, pre/post processing,..) on LUMI-D
- Broad stack of pre-installed scientific software, databases 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

How LUMI-ready is your computation?

- Thinking projects and use cases for Lumi
 - Cases for Tier-0 grand challenges
 - Combining simulation and AI methods within the same workflow
- A quick poll: which user type are you (mostly)?
 - Load and go: I'm using prebuilt or proprietary software
 - Build and go: I'm building a third-party software myself, perhaps with small modifications
 - Code and go: I'm developing my own software
- Does the software you are using have a multi-GPU version?
 - 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 an alternative package that has

Some considerations when aiming at large-scale GPU applications

- Remember the possibility of combining CPU and GPU nodes within one job – perhaps only part of the application needs to be GPU-enabled
- Refactorize and modernize your code
 - *“Even if it works, you can still fix it”*
- Employ modern frameworks and libraries
 - General: Kokkos (Sandia NL), RAJA (LLNL), Alpaka (TU Dresden),...
 - Domain-specific: GridTools (ETH Zurich), Runko (Nordita/Flatiron), DISPATCH (U Copenhagen), DEDALUS (Flatiron),...
- Write a well-scaling MPI code first and accelerate it with GPUs, not vice versa
 - OpenMP5 is a directive-based open standard with portable performance on various architectures

Concluding remarks

- **EuroHPC era: Unprecedented amount of computational resources and capabilities** available for European research & innovation
 - Complemented by competence building and user support activities
 - Synergistic with the up-to-date national data infrastructure
- **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

LUMI



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



The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation programme, as well as the of Participating States FI, BE, CH, CZ, DK, EE, IS, NO, PL, SE.

Leverage from
the EU
2014–2020



European Union
European Regional
Development Fund



Kainuun liitto