

# LUMI



**Kurt Lust**  
LUMI User Support Team (LUST)  
University of Antwerp

**Jørn Dietze**  
LUMI User Support Team (LUST)  
UiT The Arctic University of Norway

## Welcome and Introduction

May 2024

# Aim of the course

- An introduction to working on LUMI for people who already have some HPC knowledge
  - You should have had some introductory HPC training in your local organisation. E.g., some familiarity with batch processing is expected
  - Neither do we have the team on-site and on-line to answer very technical questions about, e.g., GPU programming on AMD either
    - 4-day trainings with 4-6 experts from HPE and AMD. On-site presence recommended (next one likely in the week of 28/10 in Amsterdam)
    - Monthly virtual user coffee break also offers some opportunities to get in touch with experts from HPE and AMD
- Know enough to know where to (not) look for more information
- More relaxed version of our 1-day course, with some new topics added (e.g., access via Open OnDemand, object storage and containers for AI)

# Practicals

- As announced, the course is on-site first
  - So a good experience for the on-site participants is the priority
  - Networking with colleagues, with the LUMI team and (for locals) with the local organisation is also important
- On-site presence from LUST:
  - Kurt Lust, Belgium
  - Jørn Dietze, Norway
  - Henk Dreuning, SURF
- On-line team to answer questions on HedgeDoc

# Practicals (2)

- On-site: Feel free to ask questions in the sessions and speak to us during breaks.
- HedgeDoc questions:
  - Questions are anonymous. But there is a limit to how much we can answer to such questions.
  - Please stay to the topic of the talk with your questions.  
The course is not meant to quickly give answers to all questions in the first hour after which you can leave.

# HedgeDoc demo

# Practicals (3)

- Course materials will be made available in the [LUMI training materials](https://lumi-supercomputer.github.io/LUMI-training-materials) archive site at [lumi-supercomputer.github.io/LUMI-training-materials](https://lumi-supercomputer.github.io/LUMI-training-materials).
  - Exercises during the course
  - PDF of the slides
  - Notes for most talks
  - Video recordings some time after the course (if they succeed)

## Welcome

Welcome to the LUMI supercomputer user guide. To navigate this guide, select a category from the navigation bar at the top of the page or use the search function.

You have not connected to LUMI yet? Please visit the first steps section to get started.

[→ First steps](#)

[LUMI helpdesk](#) [LUMI status](#) [LUMI events](#) [LUMI training materials](#)

**Discover the LUMI Hardware**

**Submitting a Job**

**Storage**



# Content

- Day 1: Building blocks before we can run
  - LUMI architecture
  - LUMI system software and programming environment
  - How do we offer and access application software?
  - How can we log on to the system and transfer data?
- Day 2
  - Morning: Running jobs on LUMI
  - Afternoon: Working around problems on LUMI:
    - Dealing with the Lustre file system
    - Containers (which is also somewhat dealing with the file system)
    - Working with the support team



# Acknowledgements

- Course based on materials of our 1-day training and materials under development for a training in Belgium
- Development partly done in the framework of the VSC Tier-0 support project, funded by the Research Foundation – Flanders (FWO) as part of the VSC project
- Thanks also to my colleagues in LUST as I borrow material from presentations by Maciej Szpindler, Jørn Dietze and René Løwe Jacobsen
- Other VSC and CÉCI training materials also offered inspiration

LUMI

**Enjoy the course!**





# LUMI



**Kurt Lust**  
LUMI User Support Team (LUST)  
University of Antwerp

**Jørn Dietze**  
LUMI User Support Team (LUST)  
UiT The Arctic University of Norway

## Day 2 Introduction

May 2024

# Lessons from day 1

- LUMI architecture
  - Both CPU and GPU nodes are very hierarchical
  - Proximity of resources is important
    - Proper process and thread distribution needed
  - Powerful interconnect, but different from “the standard” and new
    - Some libraries are different from standard distributions
- HPE Cray Programming Environment
  - Essential part of the system software on an HPC cluster
  - Usually no ABI standard for packages: OpenMP runtime, MPI, ...
  - Therefore bringing in binaries not always straightforward
  - Restricted Linux on the compute nodes
    - Improved scalability for large applications
    - But it also implies some software may not work

# Lessons from day 1 (2)

- Accessing LUMI
  - Not different from most HPC clusters
  - Organisation in users and projects if fairly typical for tier-1/tier-0 systems
  - But the distributed nature of allocation management is very atypical and may be confusing
- Data transfer to/from LUMI
  - Users experience it as a pain point
  - Limited protocols
  - Designers of LUMI focused strongly on the object storage as an intermediate station

# Lessons from day 1 (3)

- Modules: LMOD
  - Installed and available modules
  - Search commands for modules
- Application software on LUMI: Unconventional for many
  - Taking diversity into account
  - Need to react quickly to system updates
  - [Supportability by a small central support staff](#)
  - Evolution towards custom environments
  - EasyBuild (main tool) and Spack (secondary tool) + [containers \(this afternoon\)](#)

# This morning: Running jobs

- Slurm is not a very good resource manager/scheduler but it is the best production quality one we have...
- Slurm introduction
  - No two Slurm systems are configured the same so useful even if you are a Slurm expert!
- Process and thread distribution and binding
  - Not only the work of Slurm
  - Can have a large influence on performance of an application
  - Unfortunately there is no uniformly best way to do it so we cannot preconfigure it for you...
    - And we'll run badly into Slurm limitations

# Afternoon: Working around problems

L U M I

- Large parallel file systems can be your friend and your enemy
  - Your friend if you have a good HPC-friendly code and know how to use it...
  - But your enemy in many cases.
  - Basically, some file system parameters don't scale well with the size of the machine so we need to use the file system properly
- An important use of containers on LUMI is actually helping the file system deal with HPC-unfriendly software packaging
  - Pitfalls of containers on LUMI
  - What we do to make life a bit easier, especially for AI
- How to get the most out of LUMI support?



LUMI

**Have another interesting day!**

