

LUMI



LUMI User Support Team

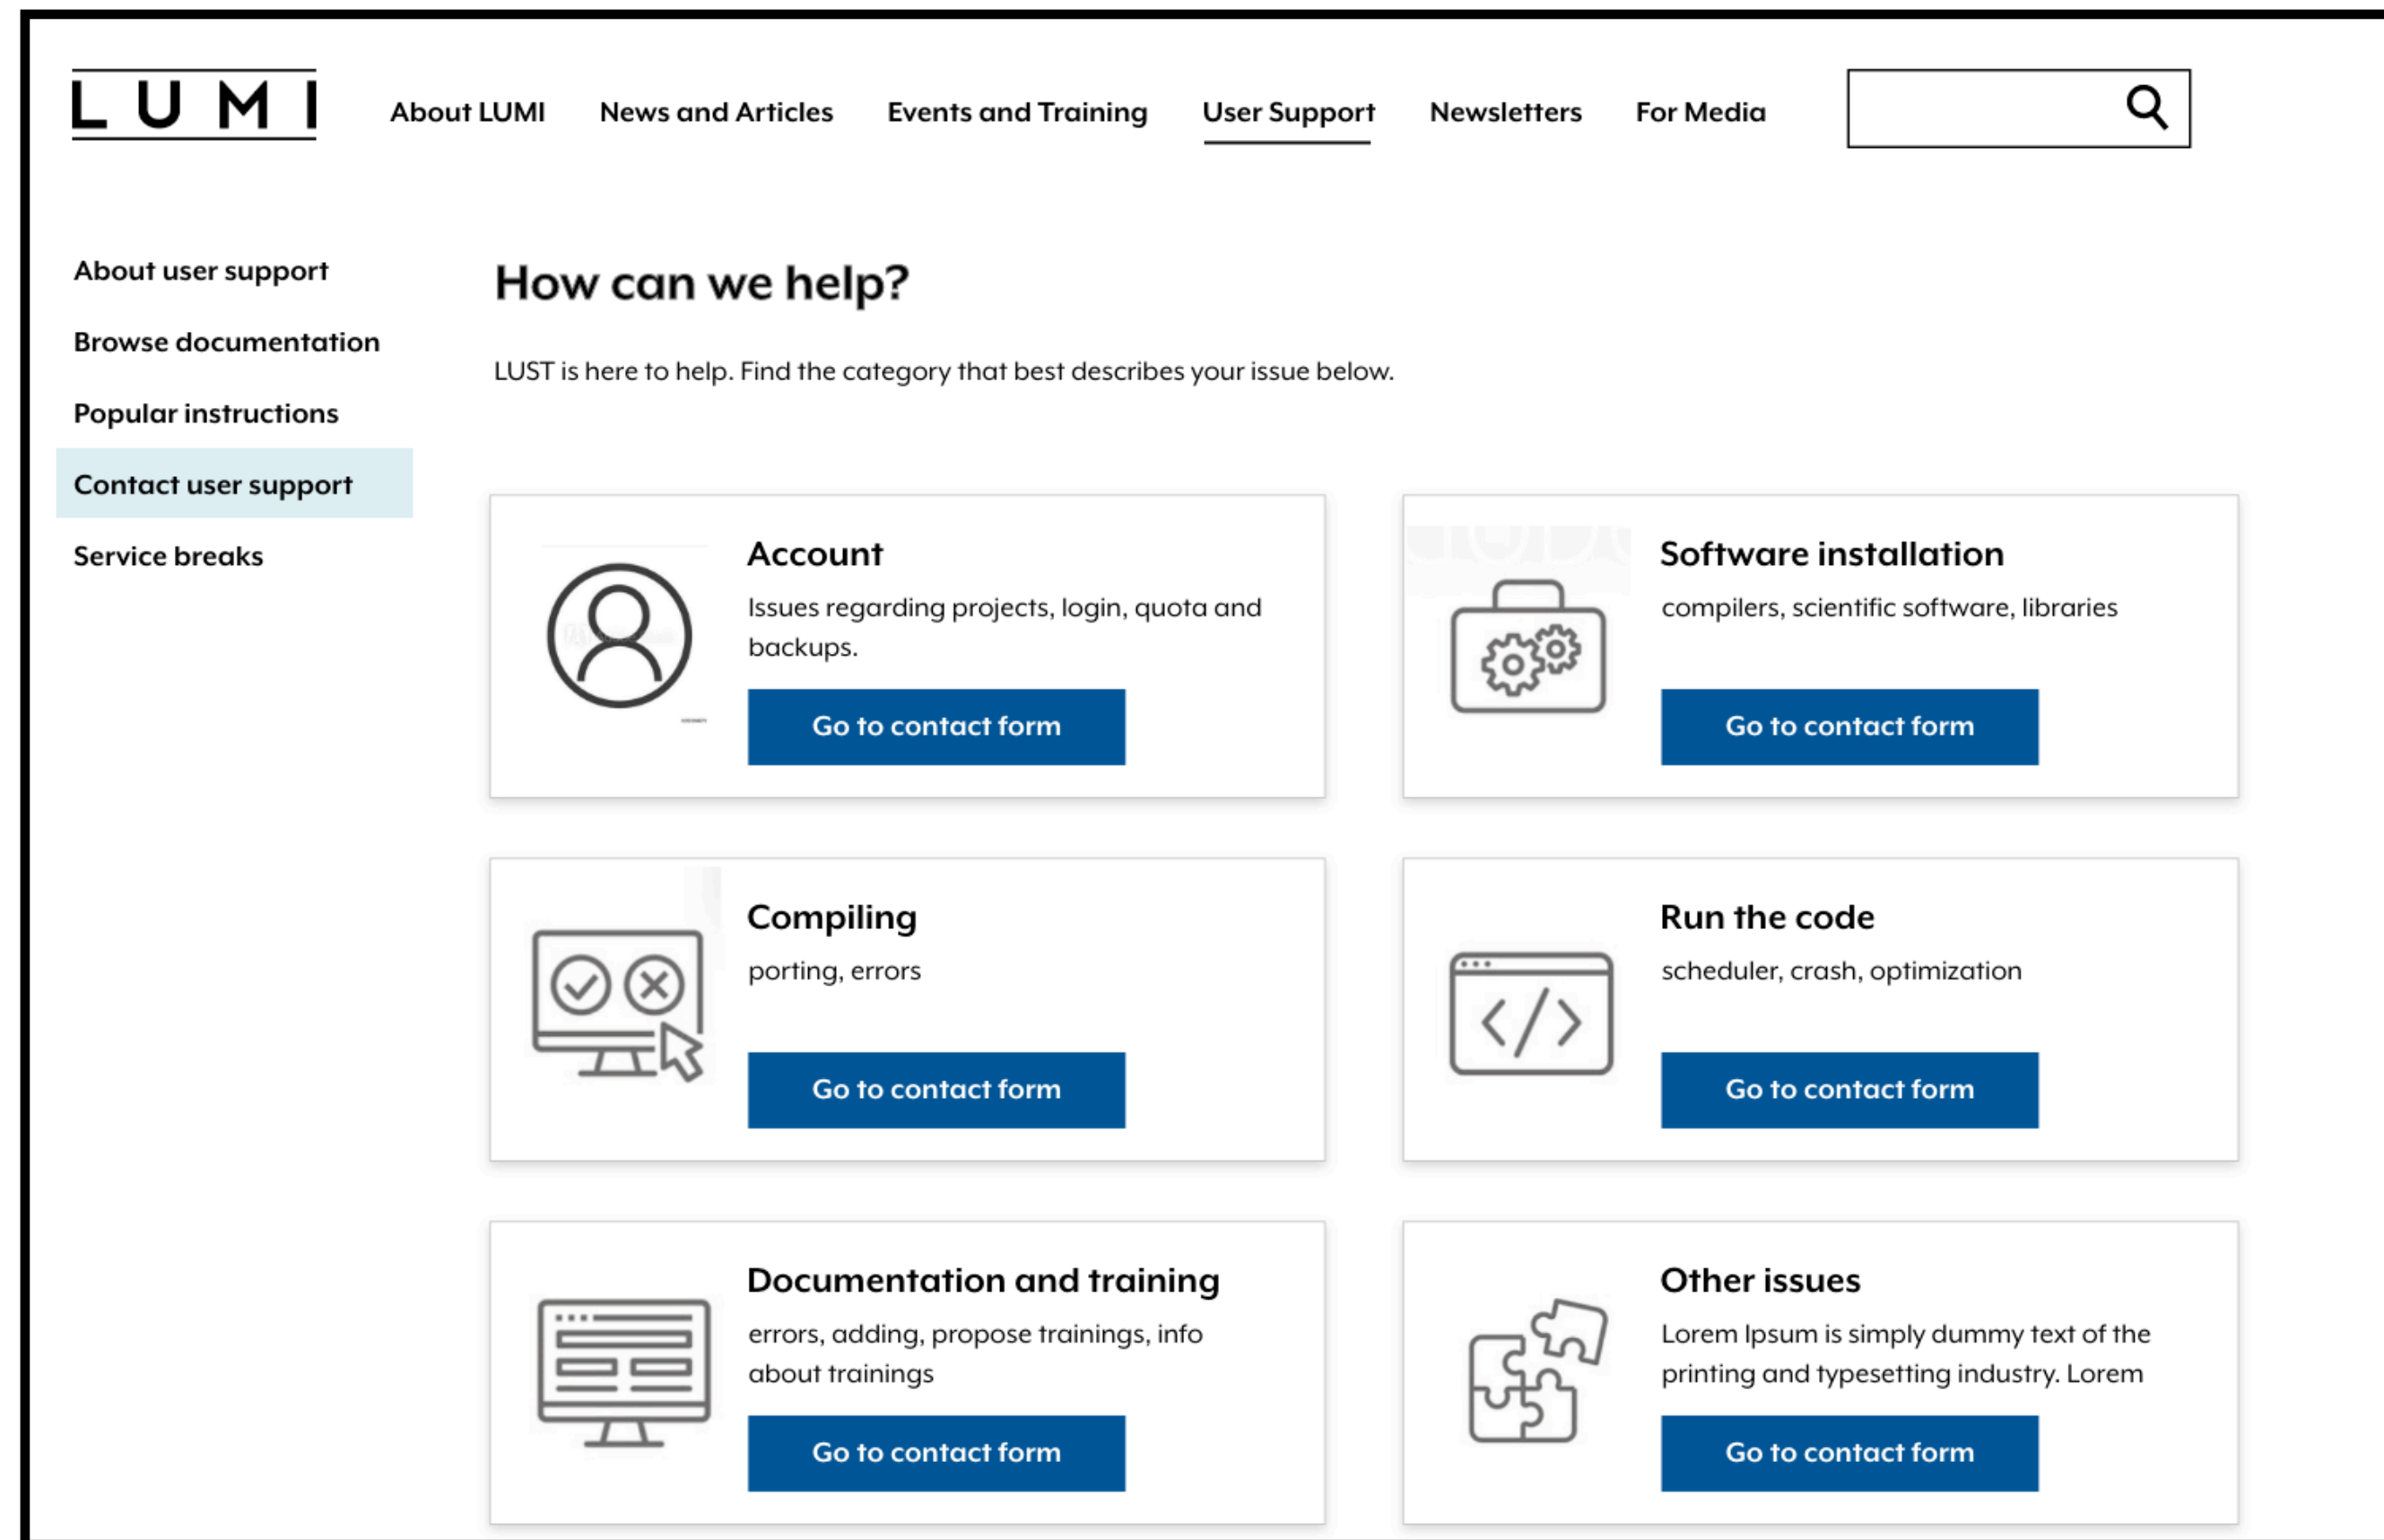
LUMI will have a support team

Current **8 FTE** with “application expert” background
Each country nominates one person.

+ support from HPE Center of Excellence, **4 FTE**

- There will be a support portal on the web page.
- Mondays-Fridays 08.00-18.00 CET
- You can ask for help just like you are used to when running at a SNIC supercomputing center.
- We will also write documentation for the system, and work **proactively** with support.

<https://lumi-supercomputer.eu>



What we can help with

- **Installing software:**
 - Compilers, choosing flags, fixing errors
 - Finding the right library to speed up the computation
 - Smaller code fixes.
- **Troubleshooting:**
 - Why did my program crash?
 - Investigating numerical issues like instability and bad convergence
- **Best practice advice:**
 - How to use scientific software, which algorithms, best convergence parameters etc
- **HPC workflow / logistics:**
 - How to use the batch system
 - Moving and staging data and calculations.

Please ask for help!

- The LUMI Support might have seen the same problem that you have 10 times before and we can solve it in 5 minutes.
- We are also interested in feedback.
Annoyed about something?
Tell us.

AMD GPU support

Code	Domain	Comment
CP2K	Materials science	Part of LUMI acceptance test
Gromacs	Molecular dynamics	Part of LUMI acceptance test
LAMMPS	Molecular dynamics	version Oct. 2020+
ICON	Weather and climate	Part of LUMI acceptance test
GridTools (COSMO)	Weather and climate	Part of LUMI acceptance test
NEK5000	CFD solver	W.I.P. Supported by ENCCS
VeloxChem	Quantum chemistry	W.I.P. Supported by ENCCS
Tensorflow	Deep learning	AMD has a fork
PyTorch	Deep learning	v1.8.0+



Search docs

RELEASE DOCUMENTATION

AMD ROCm™ Release Notes v4.1

What's New in This Release and Other Features

Driver Compability Issue in ROCm v4.1

Deploying ROCm

Deprecations

AMD ROCm Version History

DISCLAIMER

INSTALL ROCM


ROCm Installation

Multi Version Installation

HIP Installation

Using CMake with AMD ROCm

More Multimedia Installation



Welcome to AMD ROCm Platform

AMD ROCm is the first open-source software development platform for HPC/Hyperscale-class GPU computing. AMD ROCm brings the UNIX philosophy of choice, minimalism and modular software development to GPU computing.

Since the ROCm ecosystem is comprised of open technologies: frameworks (Tensorflow / PyTorch), libraries (MIOpen / Blas / RCCL), programming model (HIP), inter-connect (OCD) and up streamed Linux® Kernel support – the platform is continually optimized for performance and extensibility. Tools, guidance and insights are shared freely across the ROCm GitHub community and forums.

Note: The AMD ROCm™ open software platform is a compute stack for headless system deployments. GUI-based software applications are currently not supported.

2020: AMD ROCm™ 4.0
Complete Exascale Solution for ML/HPC

How to start preparing

Migration Path to LUMI

Today	LUMI	Comments
Intel Compilers	Cray PE, GCC, AOCC, clang	LUMI has AMD CPUs. Intel compilers may work, but not optimally. This includes Intel MKL library.
CUDA C/C++	HIP C/C++	≈ 75% of CUDA code can be converted to HIP automatically with "hipify" tools
OpenACC	OpenMP (offload)	Unclear support, maybe only in Cray compilers?
OpenMP (offload)	OpenMP (offload)	Preferred for LUMI. Note that OpenMP 5.0 is not fully implemented yet.
OpenCL	OpenCL	AMD's ROCm has OpenCL runtime.
SYCL	HipSYCL	Research project at Heidelberg University

If you have not run on a Cray supercomputer before, try to get access to one before. In Sweden, we have the *Beskow* system at PDC, and the upcoming *Dardel* system, which is very similar to LUMI.

Please ask me directly*

*until the support pages are up

e-mail: ypetla@kth.se