


LUMI STATUS

DeiC 2023 Conference.



LUMI-C and LUMI-G

- LUMI-C
 - Now 2048 nodes.
 - 2x AMD EPYC 7763 – 64 cores each
 - One slingshot interconnect pr node.
- LUMI-G
 - 2928 GPU nodes
 - 4 AMD MI250x GPUs pr node (aka 8 GPUs)
 - Each with a slingshot interconnect
- Use LUMI-G for scale and LUMI-C for supplement!
 - LUMI-C queue is hugely oversubscribed.

Other LUMI news

- LUMI-O
 - Object store – less TBH cost. Use it for large data set in an S3 like way
- LUMI – on demand
 - Will be online from 13.11.
 - Type 1 like interactive HPC
 - Still early and needs some improvement.
- AI improvement
 - Effort on making it easier to do AI/ML workload.
- Soon – federated SSH
 - No more public SSH keys.
 - DeiC WAYF part of the solution use your university/wayf login (even two factor)
 - This could also be used by national HPC centers and EuroHPC centers.(if they want)

More LUMI news

- Containers
 - DeiC take leadership in developing it.
 - Focus: must be easy to use.
- DeiC taking ownership of the Danish part of LUMI (from DK uni)
 - More hands on by DeiC to make sure information goes to universities front offices.
- For up-to-date description on system use: <https://docs.lumi-supercomputer.eu/>
- LUMI will continue to run until end of 2027.

COTAINR

- New release
 - Better warning and error messages
 - Small bugfixes to improve the user experience.
- Use it for your python workload on LUMI
 - We want to extent it to R, spark and other framework used – but time is limit – so help is needed.
- Soon support for more EuroHPC facilities. But also work in progress for local AAU HPC facility.
 - But also want it to work on all the other Danish national sites.
- Working closely with AMD and HPE on creating highly optimized containers for LUMI.
- Open source - <https://github.com/DeiC-HPC/cotainr>
- Documentation on <https://cotainr.readthedocs.io/en/stable/>