

SUBMISSION SYSTEM FOR ATLAS JOBS USED ON LUMI HPC



M. Svatoš, J. Chudoba, P. Vokáč

8th Users Conference of IT4Innovations

4.-5.11.2024

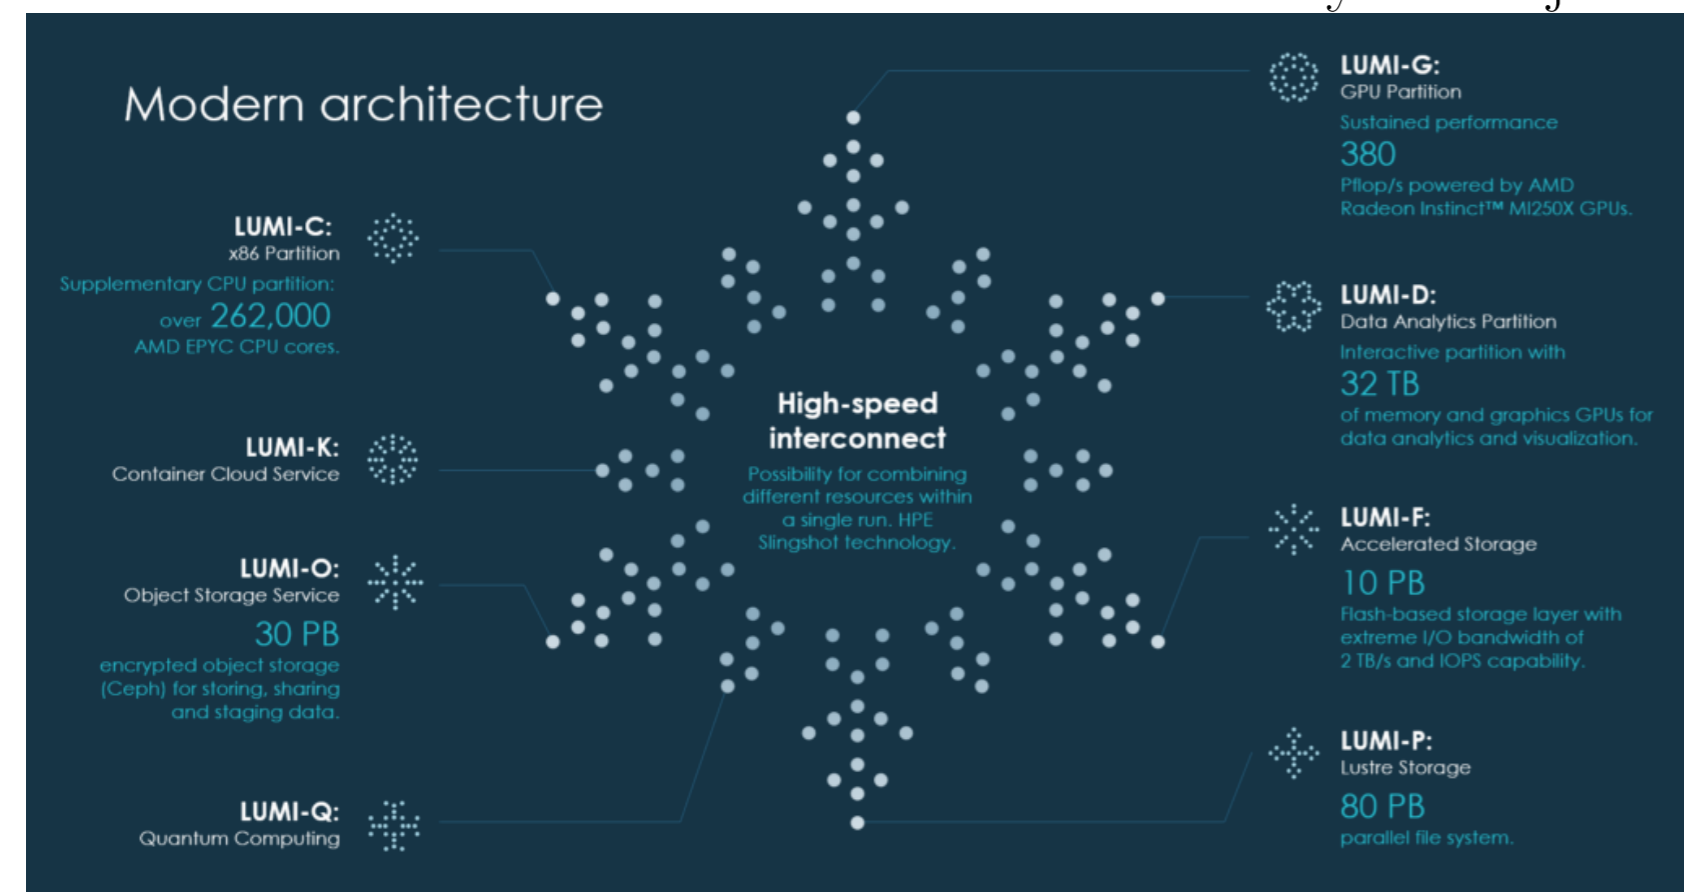


LUMI



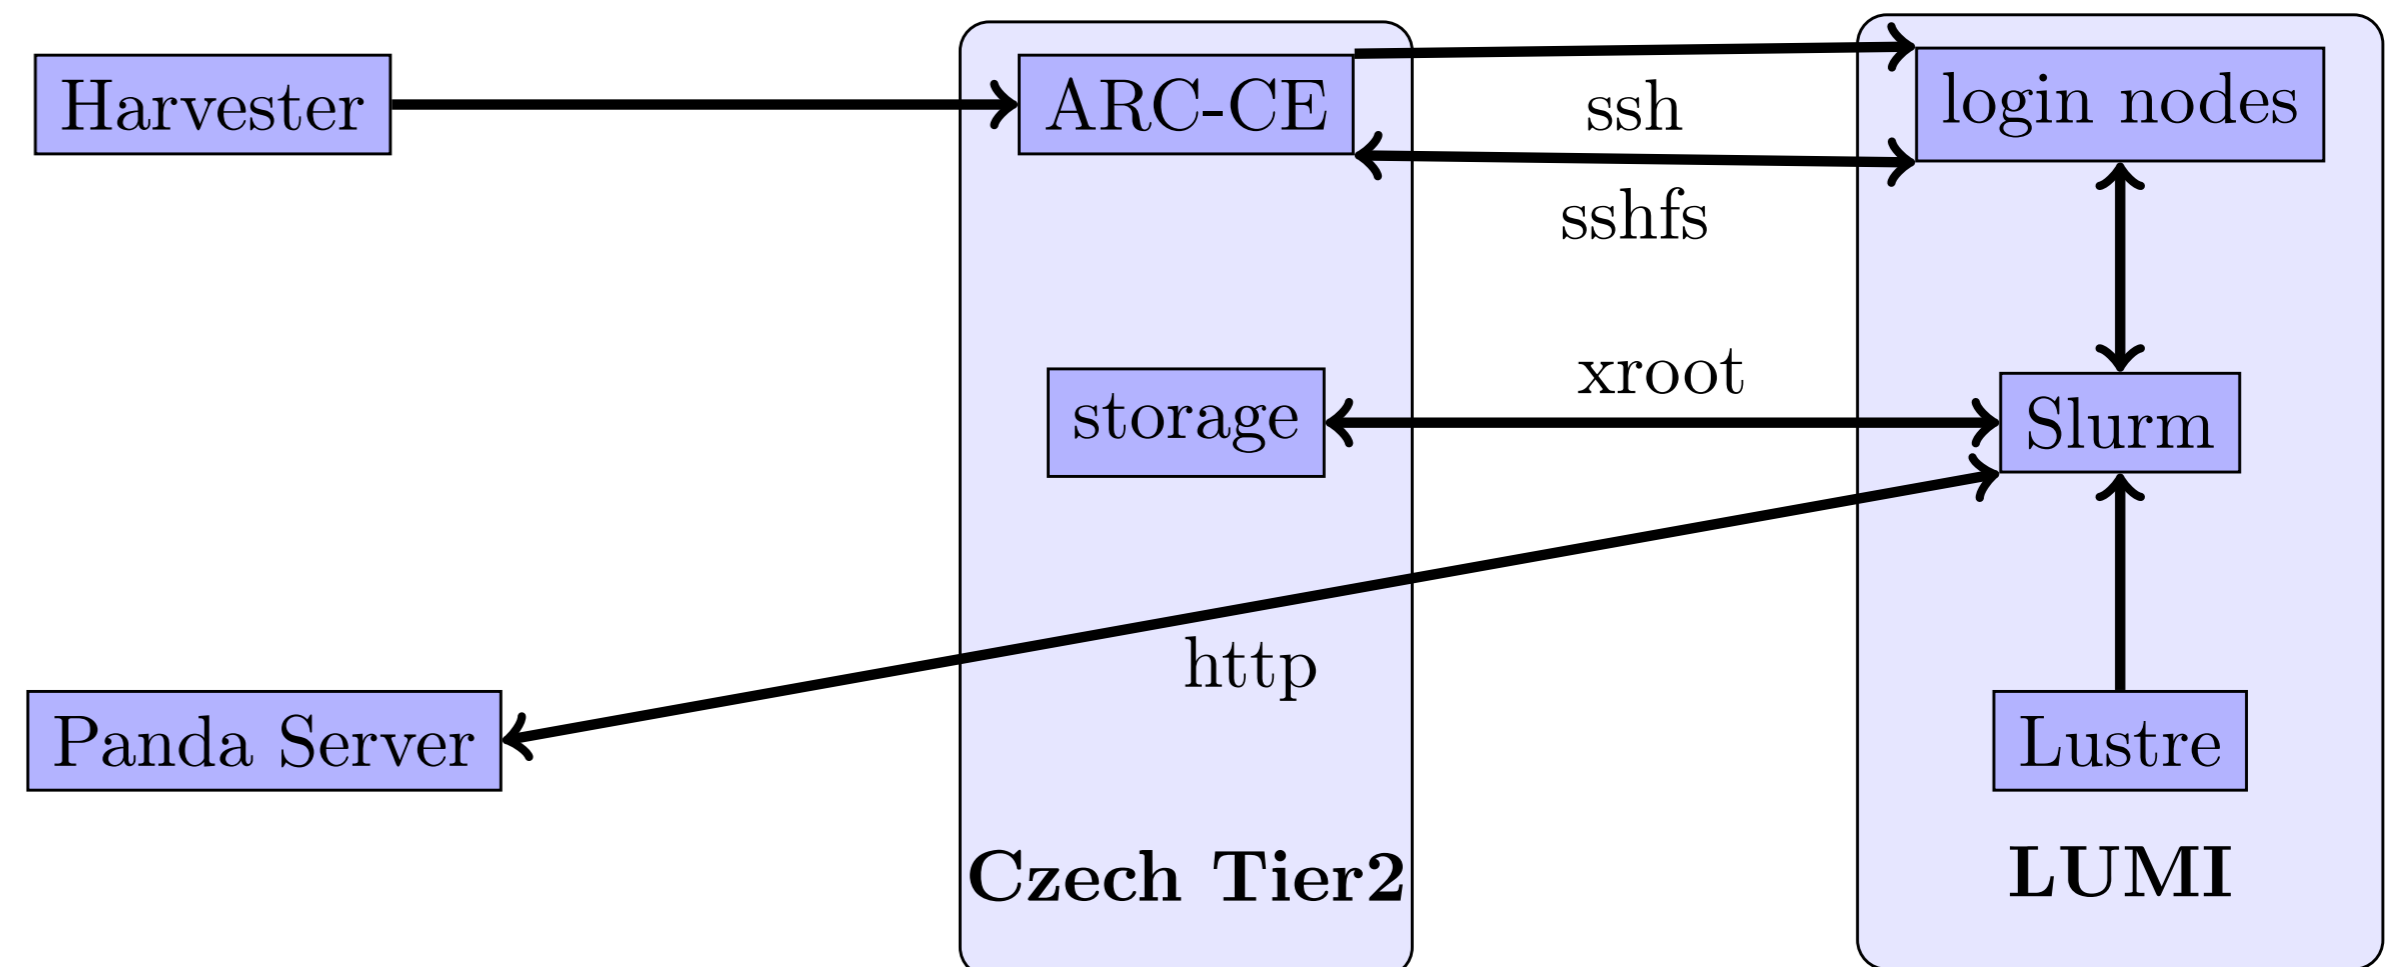
LUMI (Large Unified Modern Infrastructure) is a pan-European pre-exascale supercomputer located in data center in Kajaani, Finland. It was number five on the Top500 list (published in May 2024) and number twelve on the Green500 list (published in May 2024).

LUMI (an HPE Cray EX system) has a sustained computing power of 380 petaflops (HPL, High-Performance Linpack). It is using 100% hydro-powered energy and its waste heat is used to heat hundreds of households in the city of Kajaani.




LUMI provides a CPU partition, a GPU partition, a cloud services and many other computing resources.

Submission System



- The ARC-CE (located at Czech Tier2) receives a pilot job, translates the job description into script that can be run in the Slurm batch system, puts necessary files into a folder shared with the HPC via sshfs, and submits the job via ssh connection to a login node.
- When the batch job starts, pilot contacts panda server to receive a payload job.
- If it receives a payload job, it gets input file(s) from Czech Tier2 storage, starts the calculation (in software container), and sends outputs to Czech Tier2 storage when the payload finishes.
- If the pilot can expect that another payload would finish, it requests one.


Allocation



- The supercomputer is hosted by the **LUMI consortium** including eleven European countries (Finland, Belgium, the Czech Republic, Denmark, Estonia, Iceland, the Netherlands, Norway, Poland, Sweden, and Switzerland).
- Half of the LUMI resources belong to the EuroHPC Joint Undertaking, and the other half of the resources belong to the LUMI consortium countries. The shares for each of the countries are allocated according to local considerations and policies, i.e. LUMI is seen and handled as an extension to national resources.
- The allocation was granted from the Czech Republic share.

Performance (1.6.-15.8.2024)

The plots are coming from ATLAS job accounting dashboard.



Summary

- The environment of LUMI is different from the usual WLCG site (the OS is SUSE, not RHEL; user namespaces are disabled; etc.), making the job running difficult.
- We were able to overcome the differences and successfully run ATLAS jobs at LUMI HPC.
- The annual allocation can be consumed within a few months as it is relatively small and meant for testing purposes.

Acknowledgement

We acknowledge VSB – Technical University of Ostrava, IT4Innovations National Supercomputing Center, Czech Republic, for awarding this project access to the LUMI supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CSC (Finland) and the LUMI consortium through the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (grant ID: 90254). The work was also supported by the project CERN-CZ LM2023040.