VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

7TH USERS' CONFERENCE OF IT4INNOVATIONS

30–31 OCTOBER 2023
IT4INNOVATIONS, OSTRAVA

MINISTRY OF EDUCATION, YOUTH AND SPORTS

This conference is supported by the Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations project “e-INFRA CZ – LM2023054”.

2023 – UPDATES AND OPPORTUNITIES

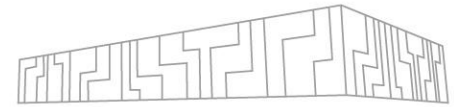

- NEW SERVICES
- NEW SYSTEMS
- LUMI-Q QUANTUM

SLURM SCHEDULER

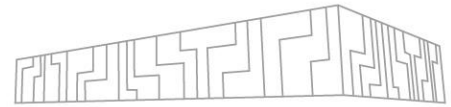

Complete switchover to Slurm

- Slurm on Barbora, Karolina and Complementary systems
- Complete integration to accounting and SCS
- Removed legacy queues, check: $ sinfo -s
- Still optimizing the queue policies: fairsharing and backfilling
- https://docs.it4i.cz/general/job-submission-and-execution/

Scheduler status view

- $ squeue [--time] –all ; squeue –u uname
- https://extranet.it4i.cz/rsweb/karolina/queues?user=uname
  (Includes electricity stats, see O. Vysocky 17:00)

SLURM SCHEDULER

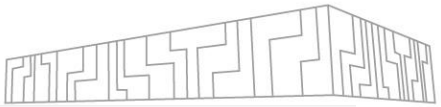

https://docs.it4i.cz/general/job-submission-and-execution/ 120%

General | Search | sccs/docs.it4i.cz

General Storage Clusters Software Clouds

Job States

Back to top

The most common job states are (in alphabetical order):

Code	Job State	Explanation
CA	CANCELLED	Job was explicitly cancelled by the user or system administrator. The job may or may not have been initiated.
CD	COMPLETED	Job has terminated all processes on all nodes with an exit code of zero.
CG	COMPLETING	Job is in the process of completing. Some processes on some nodes may still be active.
F	FAILED	Job terminated with non-zero exit code or other failure condition.
NF	NODE_FAIL	Job terminated due to failure of one or more allocated nodes.
OOM	OUT_OF_MEMORY	Job experienced out of memory error.
PD	PENDING	Job is awaiting resource allocation.
PR	PREEMPTED	Job terminated due to preemption.
R	RUNNING	Job currently has an allocation.

MPUTING

HYPERQUEUE

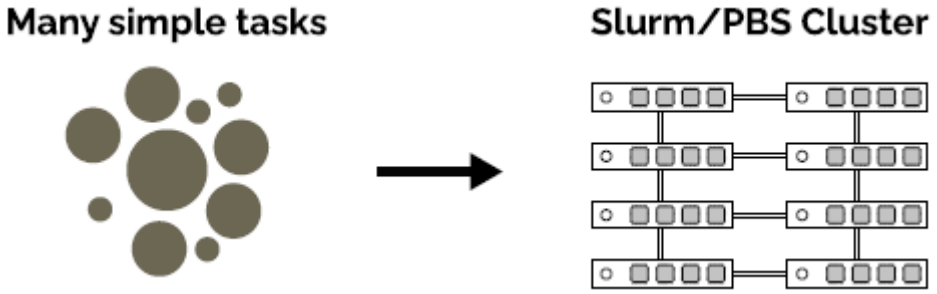

Problem:

- Large amount of jobs?
- Very large amount of small jobs?
- Only 1 or few cores per job needed?
- Complicated job dependencies?

Answer: Use hq!

https://it4innovations.github.io/hyperqueue/stable/
https://docs.it4i.cz/general/capacity-computing/#hyperqueue

E-INFRA CLOUD AT KAROLINA

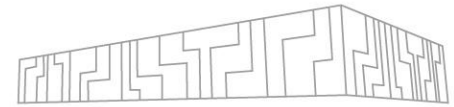

Announcing e-INFRA CLOUD service at Karolina

- Available to all active IT4I users (users attached to project)
- 22 Cloud nodes (configuration like cn[001-720])
- https://horizon.ostrava.openstack.cloud.e-infra.cz/
- Small resources available by default, including IP address
- Connectivity to Karolina login
- https://docs.it4i.cz/cloud/einfracz-cloud/
- IF more resources needed, **apply!**

E-INFRA CLOUD AT KAROLINA

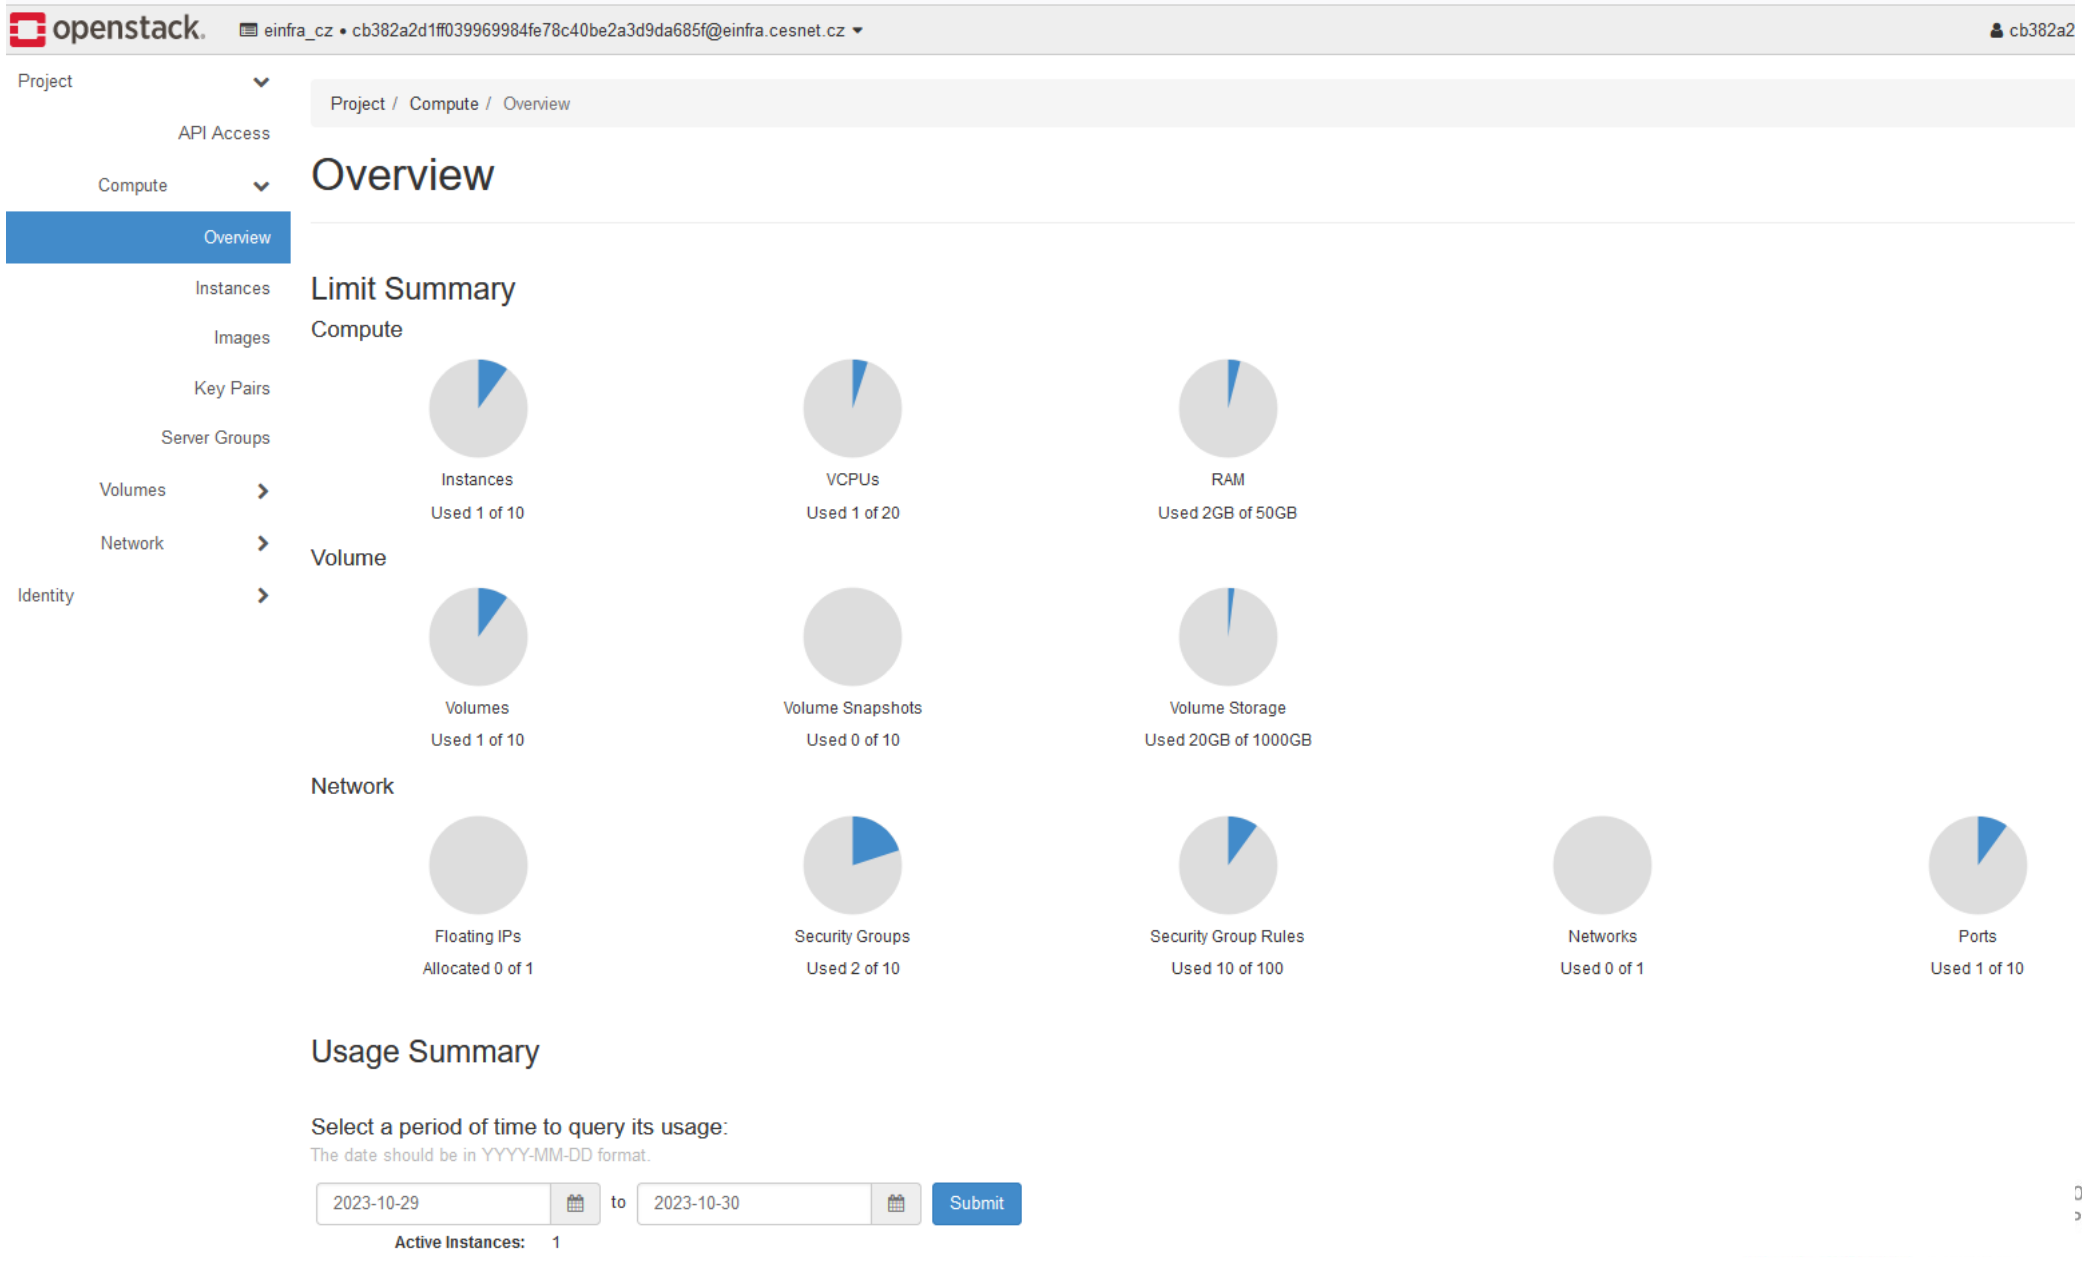

ONS
PERCOMPUTING

IT4I DATA TRANSFERS AND STORAGE

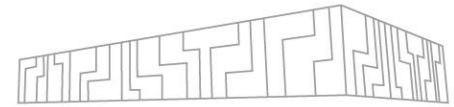

Improved connectivity to LUMI supercomputer

- Updated network parameters on Karolina and LUMI system
- New routing between IT4I and CSC proposed
- 3x speedup
- Use parallel transfers:
  https://docs.it4i.cz/general/shell-and-data-access/?h=xargs#data-transfer

S3 Services

- S3 access to PROJECT and CL3 (total capacity 33PB)
  https://docs.it4i.cz/storage/proj4-storage/

KAROLINA GUI - OPENONDEMAND

FortiClient VPN + **https://ood-karolina.it4i.cz/**

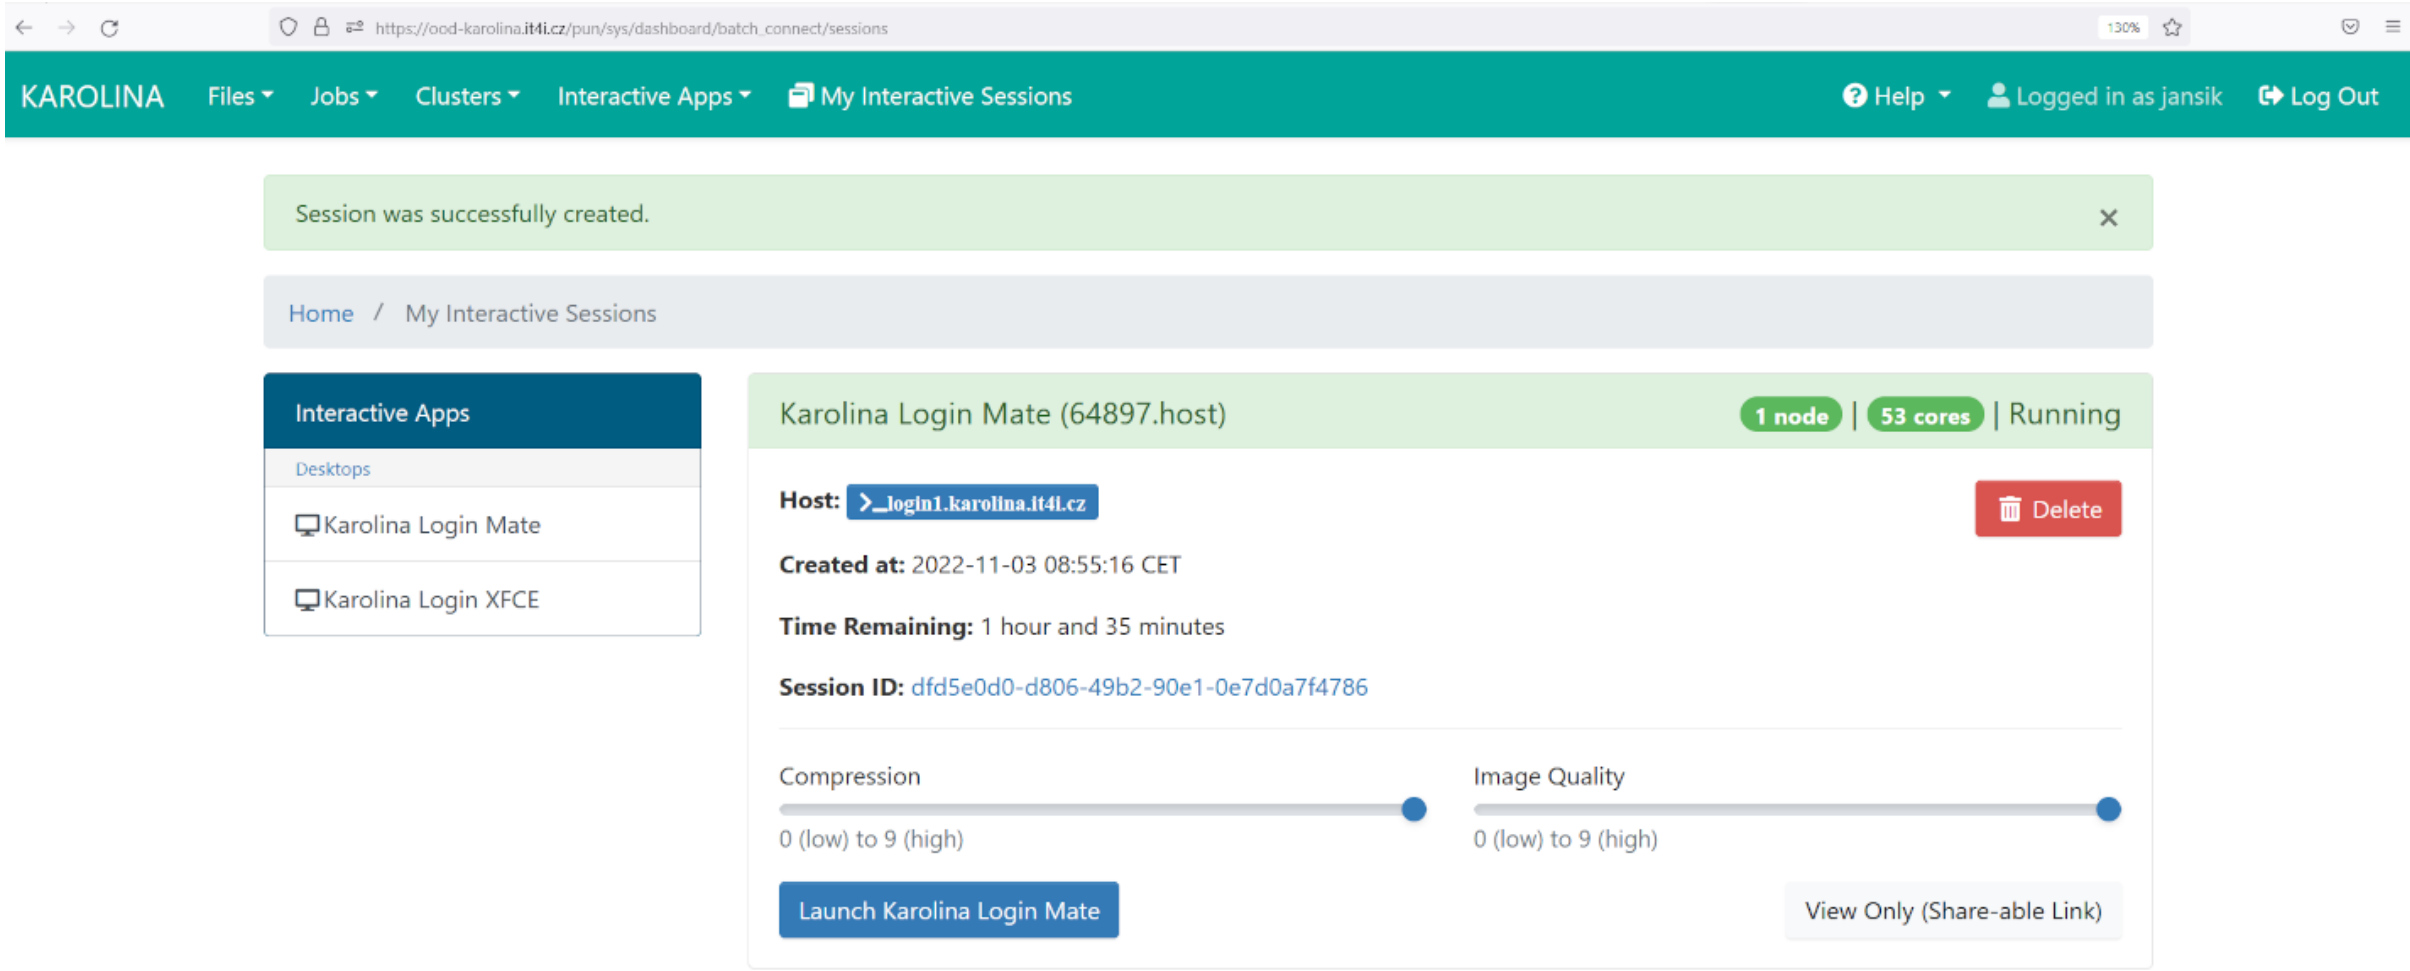

KAROLINA GUI - OPENONDEMAND

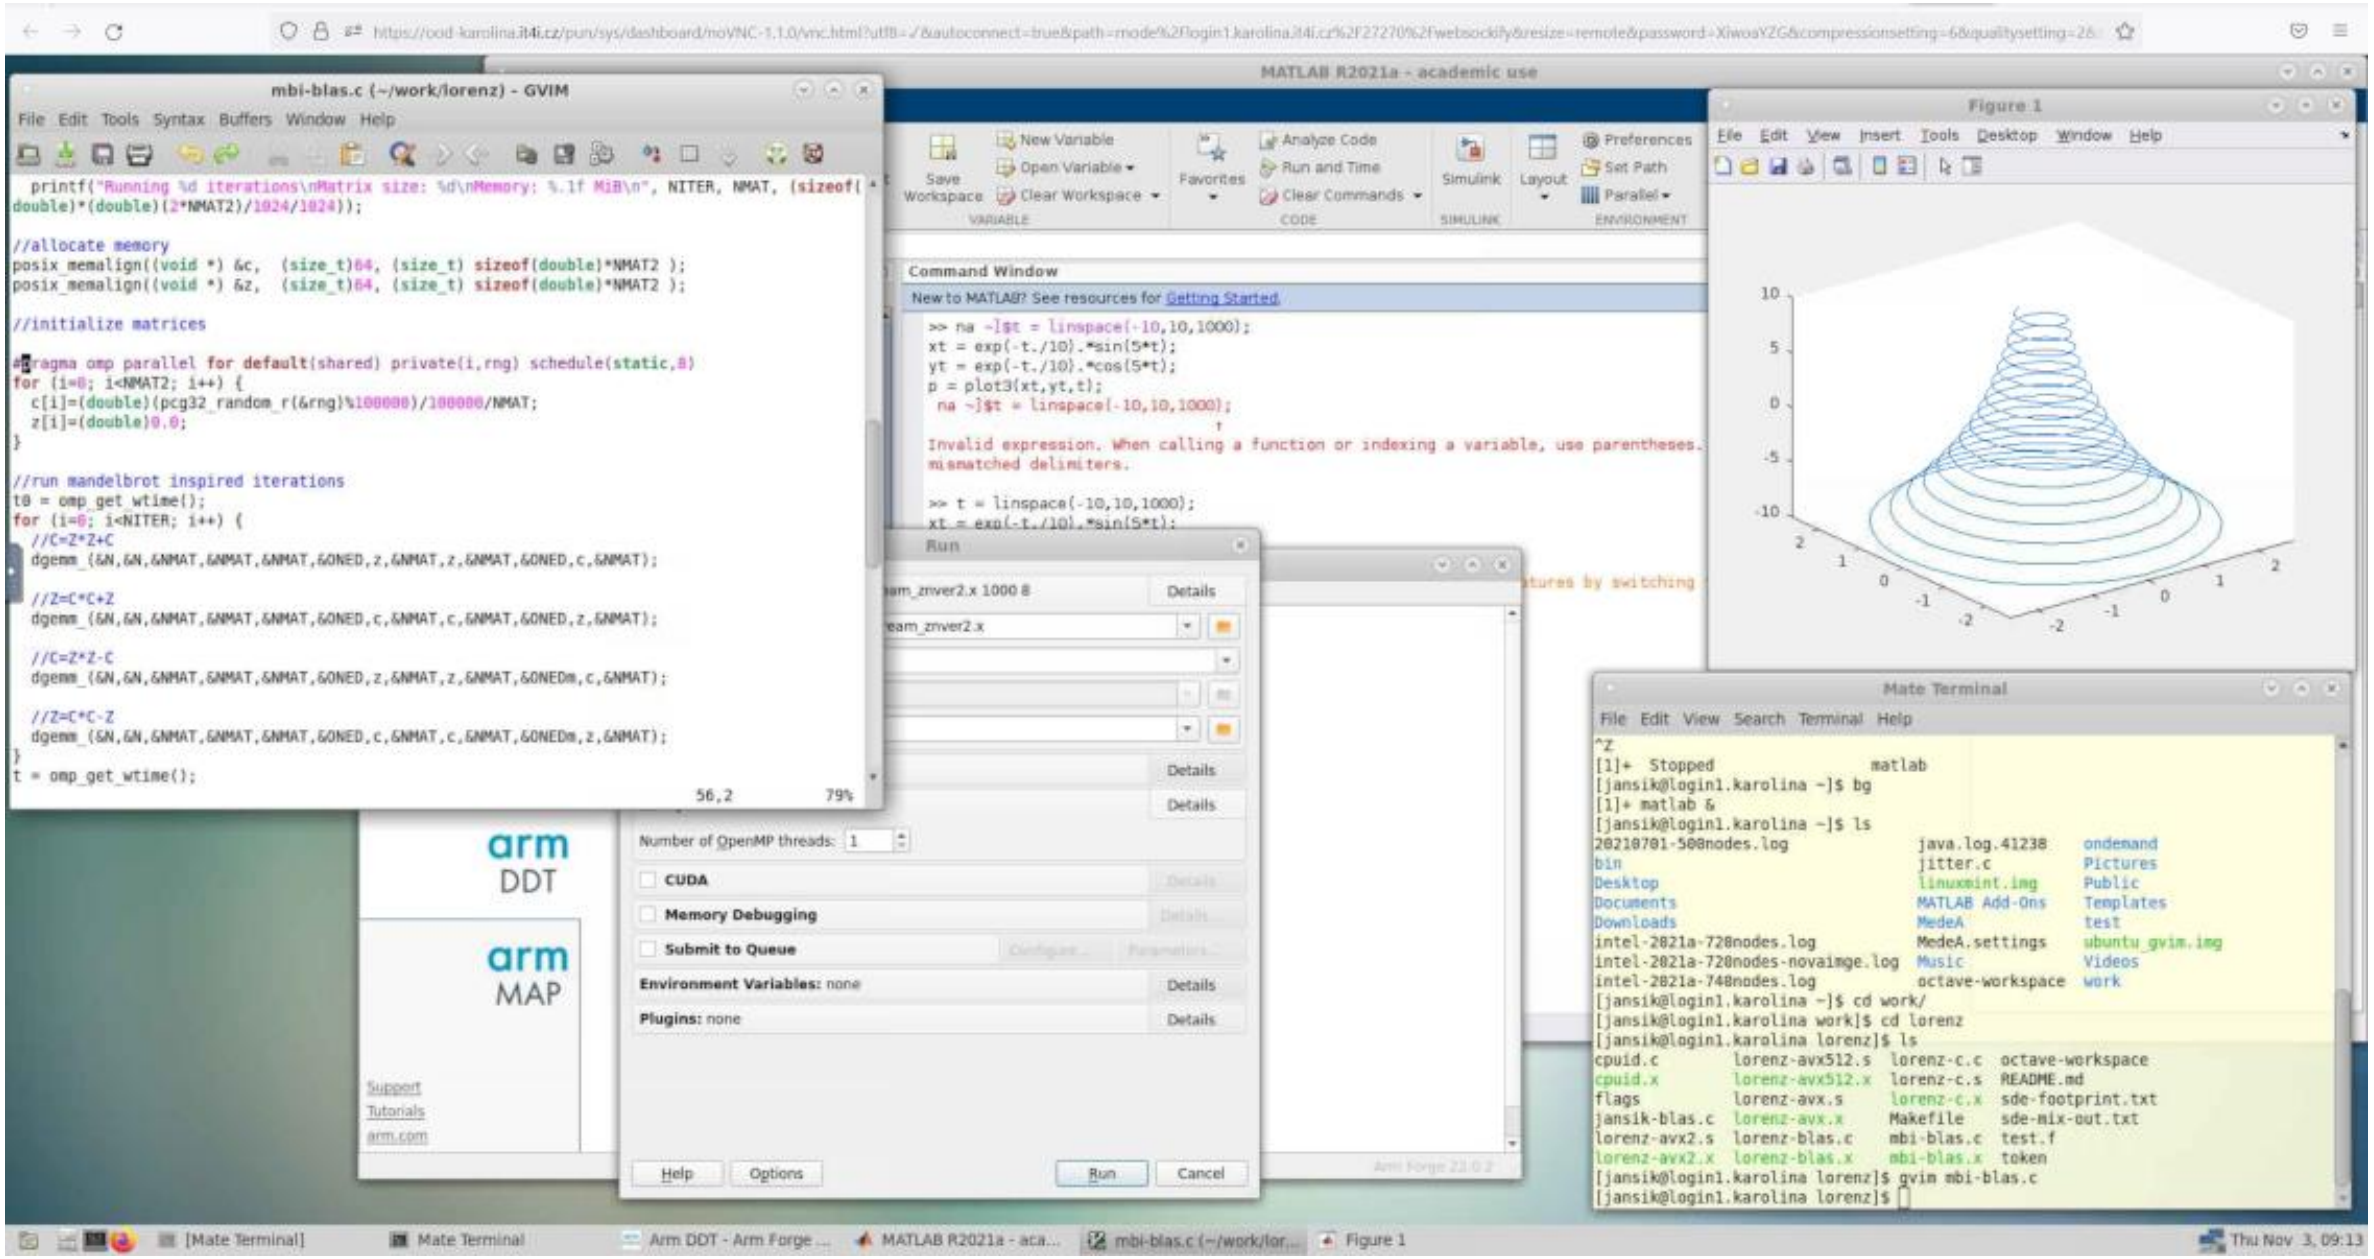

COMPLEMENTARY SYSTEMS

Targeted at perspective technology
to provide „innovator/early adopter“ advantage in the Rogers curve sense.

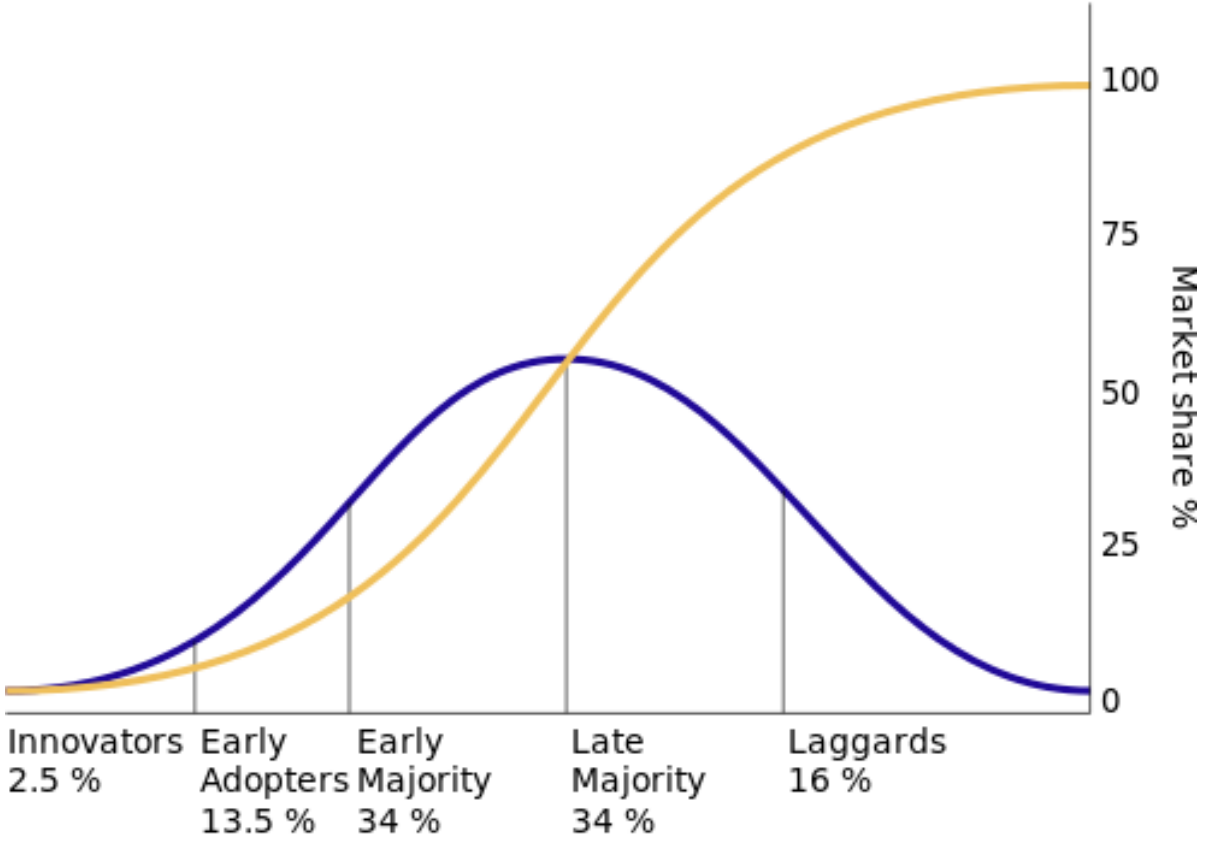

COMPLEMENTARY SYSTEM 2

Complementary Systems 2

Second stage of complementary systems implementation comprises of these partitions:

- compute partition 6 - based on ARM technology + CUDA programmable GPGPU accelerators on ampere architecture + DPU network processing units
- compute partition 7 - based on IBM Power10 architecture
- compute partition 8 - modern CPU with a very high L3 cache capacity (over 750MB)
- compute partition 9 - virtual GPU accelerated workstations

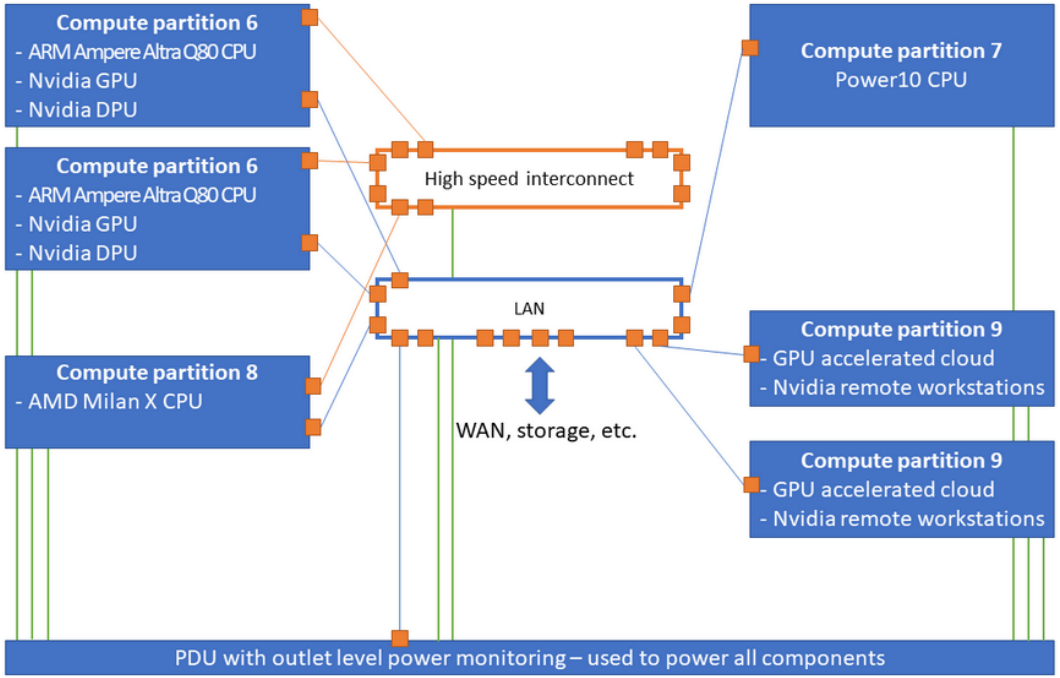

NEW SYSTEMS, MCIII

Compute		180	nodes
Home		25	TB
Scratch		300	TB
Login		2	nodes
Mgmt		2	nodes
Infiniband	200GBps	194	ports
LAN		196	ports

HBM Memory
x86_64 architecture

HBM2 mem 64 GB, 1,2 TB/s, 1,2 PFlop/s

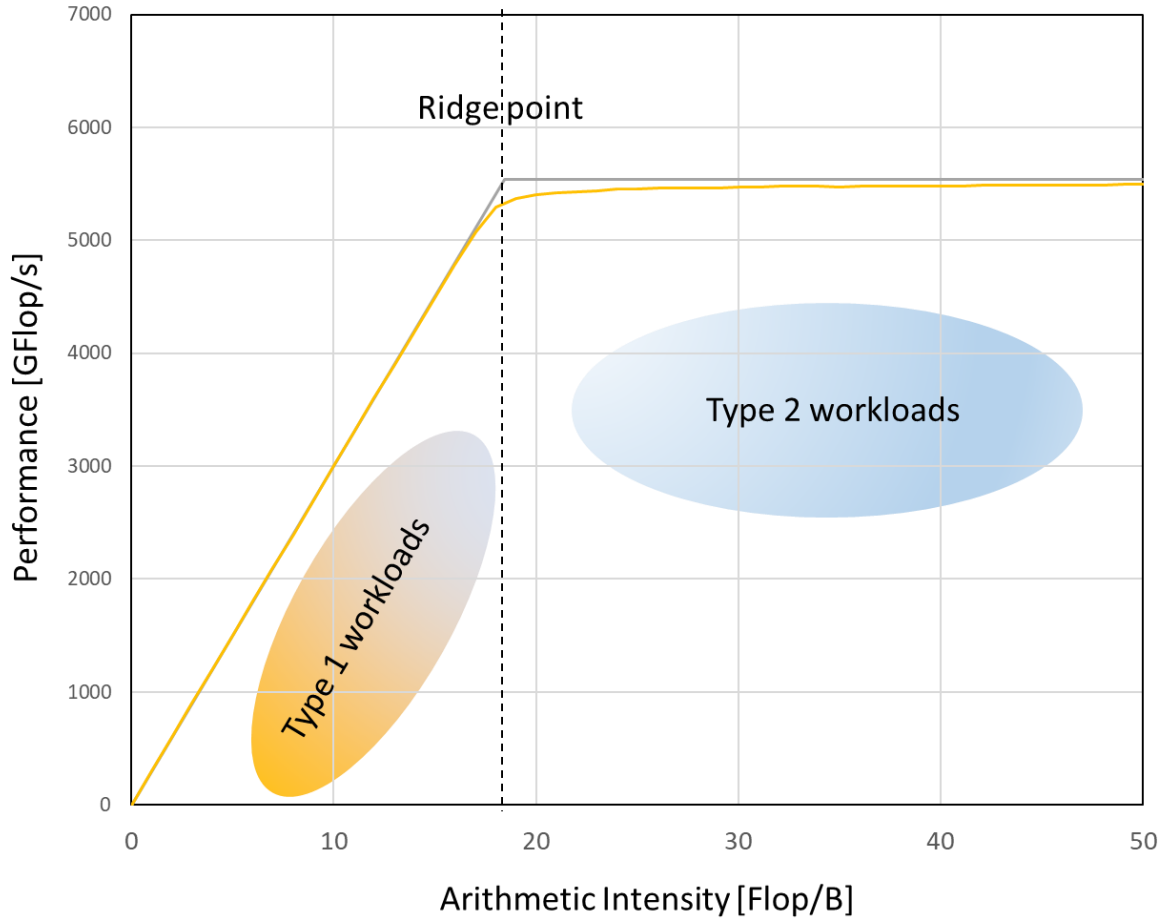

NEW SYSTEMS, VCIII

CPU	300	nodes
Accel	74	nodes
Home	25	TB
Scratch	2000	TB
Login	4	nodes
Mgmt	4	nodes
Infiniband	530	ports
LAN	460	ports

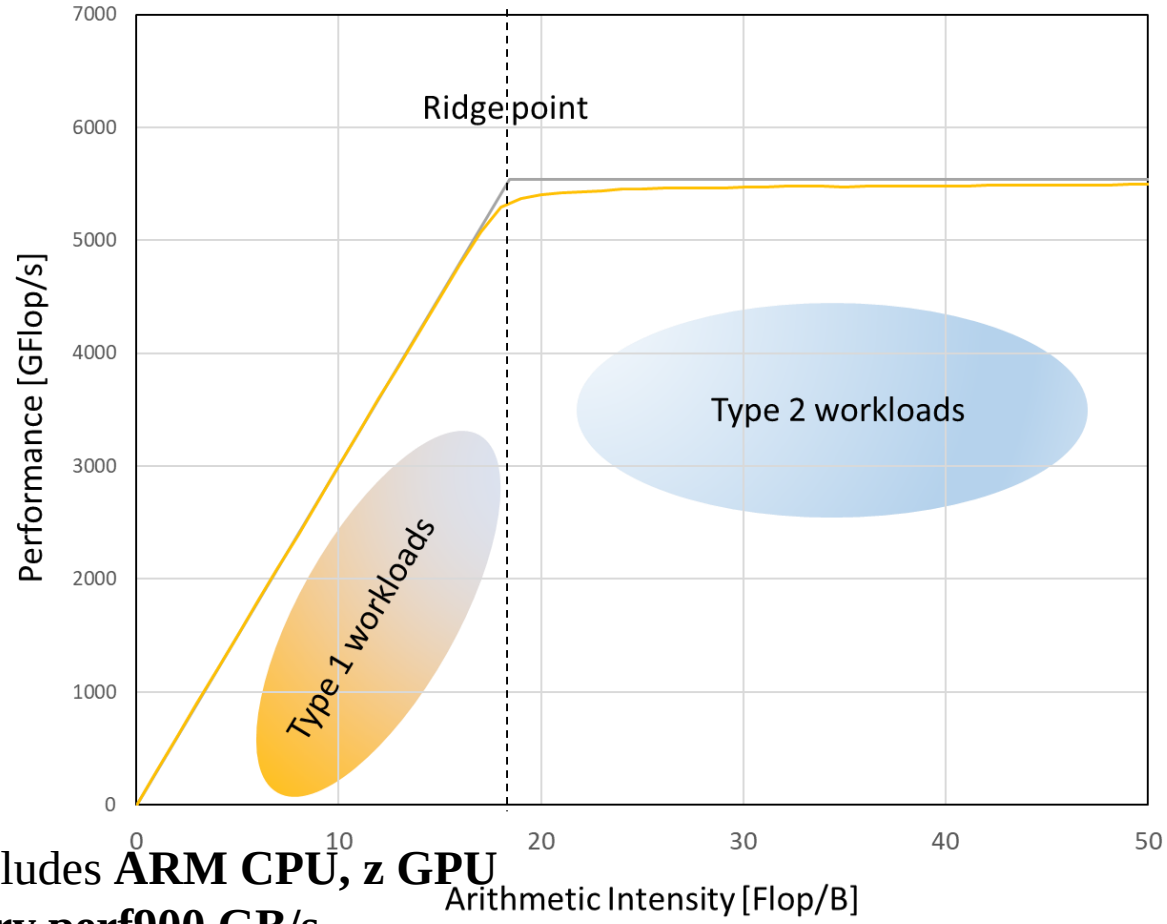

x86_64 architecture
ARM architecture
NVIDIA Architecture
GRACE+HOPPER

4x Grace Hopper NVIDIA, includes **ARM CPU, z GPU memory perf 4 TB/s** a **z memory perf900 GB/s**.
Total **18 PFlop/s**.

LUMI-Q PROGRESS

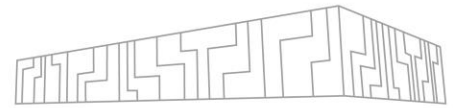

- LUMI-Q Hosting agreement (procurement and hosting) signed June 2023
- LUMI-Q Operational costs grant agreement in final stages before signature
- LUMI-Q Consortium agreement provided to partners, collecting feedback to v2, up to 3.11.2023
- LUMI-Q Integration costs grant application submitted 28.9.2023, project name EuroQHPC-Integration. The project includes all 6 hosting entities, led by GENCI
- LUMI-Q RFP prepared and submitted to EuroHPC JU on 16.10.2023.

LUMI-Q TECHNOLOGY

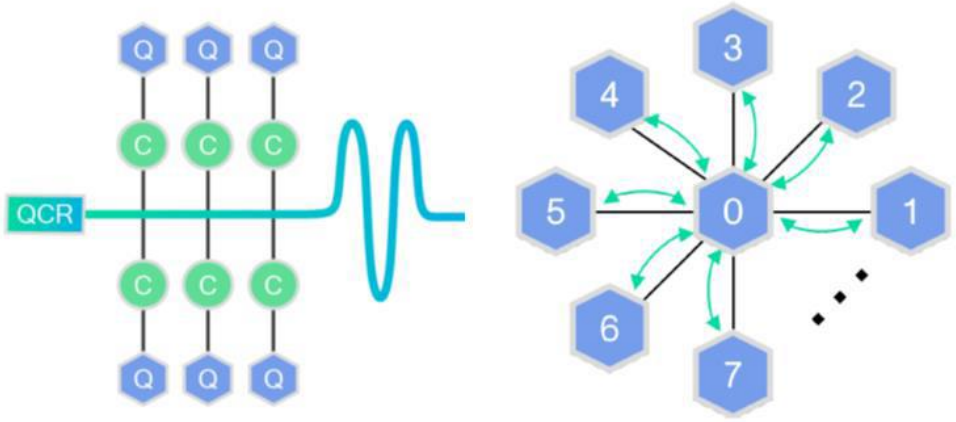

- Star-shape qubit topology, one-to-all qubit connectivity,
- Strong reduction of SWAP operations count
- High complexity quantum algorithms possingle
- Massive advancement compared to anything presently availabe
- Increased performance

Operation	SWAP ops 12 qubits in 3x4 grid	SWAP ops 12 qubits LUMI-Q
Two qubit gates	11	2
Two qubit gates, all pairs	45	26
3-qubit GHZ state	11	2
One simulation step of all-to-all spin system	45	26

LUMI-Q EUROHPCQS-INTEGRATION

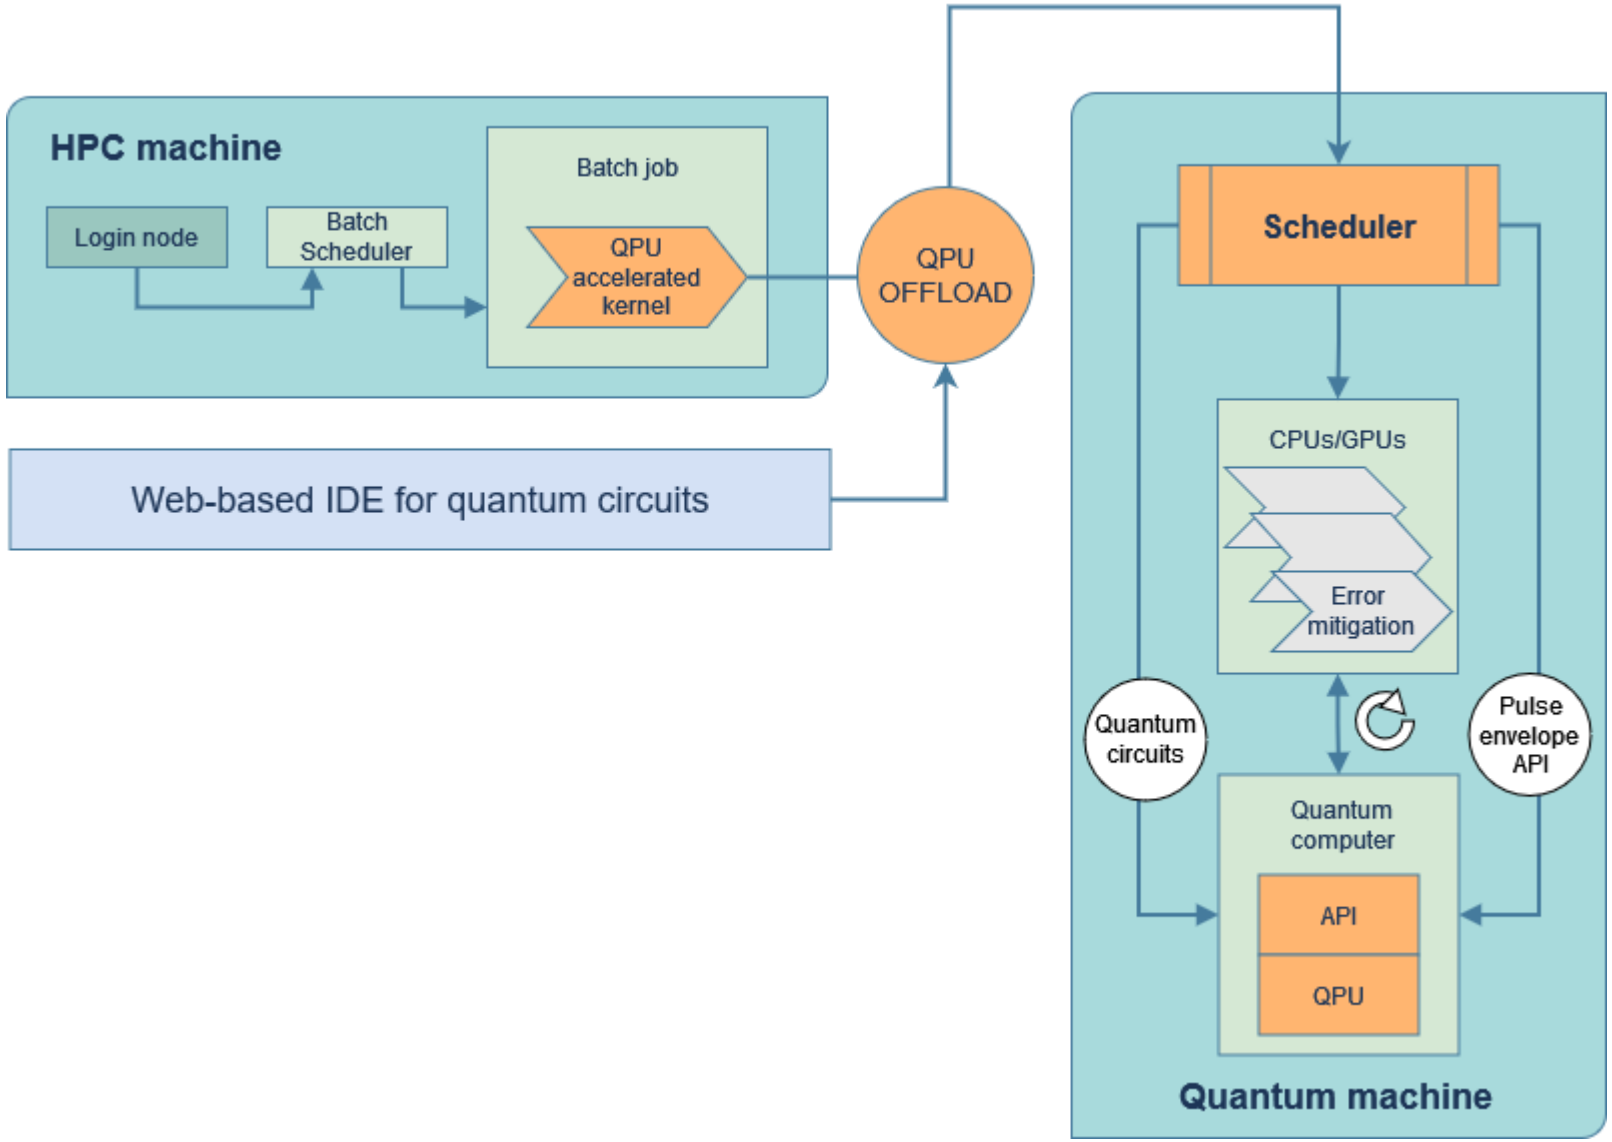

TRAINING

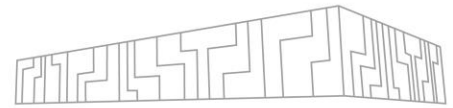

https://events.it4i.cz/

Nov 21

[ONLINE] IO-SEA: Automated Workflows and Benchmarks with JUBE in IO-SEA and Beyond

Nov 08

[HYBRID] Parallel Computing with MATLAB and Scaling MATLAB Code to the HPC Cluster (EuroCC)

Branislav Jansík
branislav.jansik@vsb.cz

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

IT4Innovations National Supercomputing Center
VSB – Technical University of Ostrava
17. listopadu 2172/15
708 00 Ostrava-Poruba, Czech Republic
www.it4i.cz