




8TH USERS' CONFERENCE OF IT4INNOVATIONS

BRANISLAV JANSÍK

**4 NOVEMBER 2024
IT4INNOVATIONS, OSTRAVA**

2023 – UPDATES AND OPPORTUNITIES




| INTERACTIVE/INSTANCE BASED
COMPUTING

| SERVICES and NEWS

| SMALL CLUSTER III

| LUMI-Q QUANTUM


E-INFRA CLOUD AT KAROLINA



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



einfra_cz • cb382a2d1ff039969984fe78c40be2a3d9da685f@einfra.cesnet.cz

cb382a2

Project

Project / Compute / Overview

API Access

Compute

Overview

Overview

Instances

Images

Key Pairs

Server Groups

Volumes

Network

Identity

Limit Summary

Compute



Instances
Used 1 of 10



VCPUs
Used 1 of 20



RAM
Used 2GB of 50GB



Volumes
Used 1 of 10



Volume Snapshots
Used 0 of 10



Volume Storage
Used 20GB of 1000GB

Network



Floating IPs
Allocated 0 of 1



Security Groups
Used 2 of 10



Security Group Rules
Used 10 of 100



Networks
Used 0 of 1



Ports
Used 1 of 10

Usage Summary


Select a period of time to query its usage:

The date should be in YYYY-MM-DD format.

2023-10-29 to 2023-10-30

Active Instances: 1

KAROLINA GUI - OPENONDEMAND



https://ood-karolina.it4i.cz/pun/sys/dashboard/batch_connect/sessions

Jobs Clusters Interactive Apps Simulation My Interactive Sessions Help Logged in as jan

Session was successfully

Home / My Interactive

Interactive Apps

- Desktops
 - Karolina Login Mate
 - Karolina Login XFCE
 - Gnome Desktop
- GUIs
 - ANSYS
 - Blender
 - ParaView
 - TorchStudio
- Servers
 - Code Server
 - Jupyter (+Julia)
 - MATLAB
 - TensorBoard

(Julia) (1858092) 1 node | 128 cores | Running

2.karolina.it4i.cz

24-11-04 08:50:12 CET


ng: 55 minutes

33c9df-e96a-4736-bf68-96522204ee29

to Jupyter

Delete

KAROLINA GUI - OPENONDEMAND



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

Browser address bar: https://ood-karolina.it4i.cz/pun/sys/dashboard/batch_connect/sessions



KAROLINA Files Jobs Clusters Interactive Apps My Interactive Sessions Help Logged in as jansik Log Out

Session was successfully created. ✕

Home / My Interactive Sessions

Interactive Apps

Desktops

-  Karolina Login Mate
-  Karolina Login XFCE

Karolina Login Mate (64897.host) 1 node | 53 cores | Running

Host: [>_login1.karolina.it4i.cz](#) Delete

Created at: 2022-11-03 08:55:16 CET


Time Remaining: 1 hour and 35 minutes

Session ID: [dfd5e0d0-d806-49b2-90e1-0e7d0a7f4786](#)

Compression 0 (low) to 9 (high) | Image Quality 0 (low) to 9 (high)

Launch Karolina Login Mate View Only (Share-able Link)

KAROLINA GUI - OPENONDEMAND



The screenshot displays the MATLAB R2021a academic use interface. The main window shows the Command Window with the following code:

```
>> na = linspace(-10,10,1000);
xt = exp(-t./10).*sin(5*t);
yt = exp(-t./10).*cos(5*t);
p = plot3(xt,yt,t);
na = linspace(-10,10,1000);
t = linspace(-10,10,1000);
xt = exp(-t./10).*sin(5*t);
```

The Figure window shows a 3D plot of a Lorenz attractor. The terminal window shows the following commands and output:

```
^Z
[1]+ Stopped matlab
[jansik@login1.karolina ~]$ bg
[1]+ matlab &
[jansik@login1.karolina ~]$ ls
20210701-500nodes.log      java.log.41238      ondemand
bin                        jitter.c            Pictures
Desktop                   linuxmint.img      Public
Documents                  MATLAB Add-Ons     Templates
Downloads                  Medea              test
intel-2021a-720nodes.log  Medea.settings    ubuntu_gvim.img
intel-2021a-720nodes-novainge.log Music               Videos
intel-2021a-740nodes.log  octave-workspace  work
[jansik@login1.karolina ~]$ cd work/
[jansik@login1.karolina work]$ cd lorenz
[jansik@login1.karolina lorenz]$ ls
cpuid.c      lorenz-avx512.s  lorenz-c.c  octave-workspace
cpuid.x      lorenz-avx512.x  lorenz-c.s  README.md
flags        lorenz-avx.s     lorenz-c.x  sde-footprint.txt
jansik-blas.c lorenz-avx.x     Makefile    sde-mix-out.txt
lorenz-avx2.s lorenz-blas.c    mbi-blas.c  test.f
lorenz-avx2.x lorenz-blas.x    mbi-blas.x  token
[jansik@login1.karolina lorenz]$ gvim mbi-blas.c
[jansik@login1.karolina lorenz]$
```

The code editor window shows the following code:

```
printf("Running %d iterations\nMatrix size: %d\nMemory: %.1f MiB\n", NITER, NMAT, (sizeof(double)*NMAT2)*(double)(2*NMAT2)/1024/1024);

//allocate memory
posix_memalign((void *) &c, (size_t)64, (size_t) sizeof(double)*NMAT2);
posix_memalign((void *) &z, (size_t)64, (size_t) sizeof(double)*NMAT2);

//initialize matrices
#pragma omp parallel for default(shared) private(i, rng) schedule(static, 8)
for (i=0; i<NMAT2; i++) {
    c[i]=(double)(pcg32_random_r(&rng)%100000)/100000/NMAT;
    z[i]=(double)0.0;
}

//run mandelbrot inspired iterations
t0 = omp_get_wtime();
for (i=0; i<NITER; i++) {
    //Z=Z+Z
    dgemv_(&N,&N,&NMAT,&NMAT,&NMAT,&ONED,z,&NMAT,z,&NMAT,&ONED,c,&NMAT);


    //Z=C+C+Z
    dgemv_(&N,&N,&NMAT,&NMAT,&NMAT,&ONED,c,&NMAT,c,&NMAT,&ONED,z,&NMAT);

    //Z=C+Z-C
    dgemv_(&N,&N,&NMAT,&NMAT,&NMAT,&ONED,z,&NMAT,z,&NMAT,&ONED,c,&NMAT);

    //Z=C+C-Z
    dgemv_(&N,&N,&NMAT,&NMAT,&NMAT,&ONED,c,&NMAT,c,&NMAT,&ONED,z,&NMAT);
}
t = omp_get_wtime();
```

The GUI also shows the Arm DDT and Arm MAP logos, and the Arm Forge 23.0.2 version information.

KAROLINA GUI - OPENONDEMAND



https://ood-karolina.it4i.cz/node/cn002.karolina.it4i.cz/22132/notebooks/work/Jupyter/Untitled.ipynb 130% ☆



File Edit View Run Kernel Settings Help

Trusted

Code

JupyterLab Python 3 (ipykernel)

```
[1]: print("Hello world")
```

```
Hello world
```

```
[2]: from mpi4py import MPI
```

```
[3]: import sys
```

```
[4]: size = MPI.COMM_WORLD.Get_size()
```

```
[5]: pip install mpi4py
```

```
[1730706673.355375] [cn002:593801:0]          ib_md.c:1232 UCX  WARN  IB: ibv_fork_init() was disabled or failed, yet a fork() has been issued.  
[1730706673.355385] [cn002:593801:0]          ib_md.c:1233 UCX  WARN  IB: data corruption might occur when using registered memory.  
Defaulting to user installation because normal site-packages is not writeable  
Requirement already satisfied: mpi4py in /home/jansik/.local/lib/python3.11/site-packages (3.1.6)
```

```
[notice] A new release of pip is available: 24.1.2 -> 24.3.1
```

```
[notice] To update, run: pip install --upgrade pip
```

```
Note: you may need to restart the kernel to use updated packages.
```


```
[6]: from mpi4py import MPI
```

```
[7]: size = MPI.COMM_WORLD.Get_size()
```


```
[8]: size
```

```
[8]: 1
```

```
[ ]:
```



HYPERQUEUE




Problem:

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

Many simple tasks



Slurm/PBS Cluster





Answer: Use hq!

<https://it4innovations.github.io/hyperqueue/stable/>

<https://docs.it4i.cz/general/capacity-computing/#hyperqueue>

HYPERQUEUE TASK DEPENDENCIES



HyperQueue

Overview

Installation

Getting Started >

Examples >

Deployment >

Task computation >

CLI >

Python API >

Getting started

Client

Submitting jobs

[Dependencies](#)

API reference

Events

FAQ

Comparison With Other Tools

Task dependencies



Table of contents

Defining dependencies

One of the most useful features of the HyperQueue Python API is that it allows you to define dependencies between individual tasks of a job.

If a task **B depends** on task **A**, then **B** will not be executed until **A** has (successfully) finished. Using dependencies, you can describe arbitrarily complex DAG (directed acyclic graph) workflows.

Notice

HyperQueue jobs are independent of each other, so dependencies can only be specified between tasks within a single job.

Defining dependencies

To define a dependency between tasks, you will first need to store the `Task` instances that you get when you create a `task`. You can then use the `deps` parameter when creating a new task and pass an existing task instance to define a dependency:

```
from hyperqueue import Job

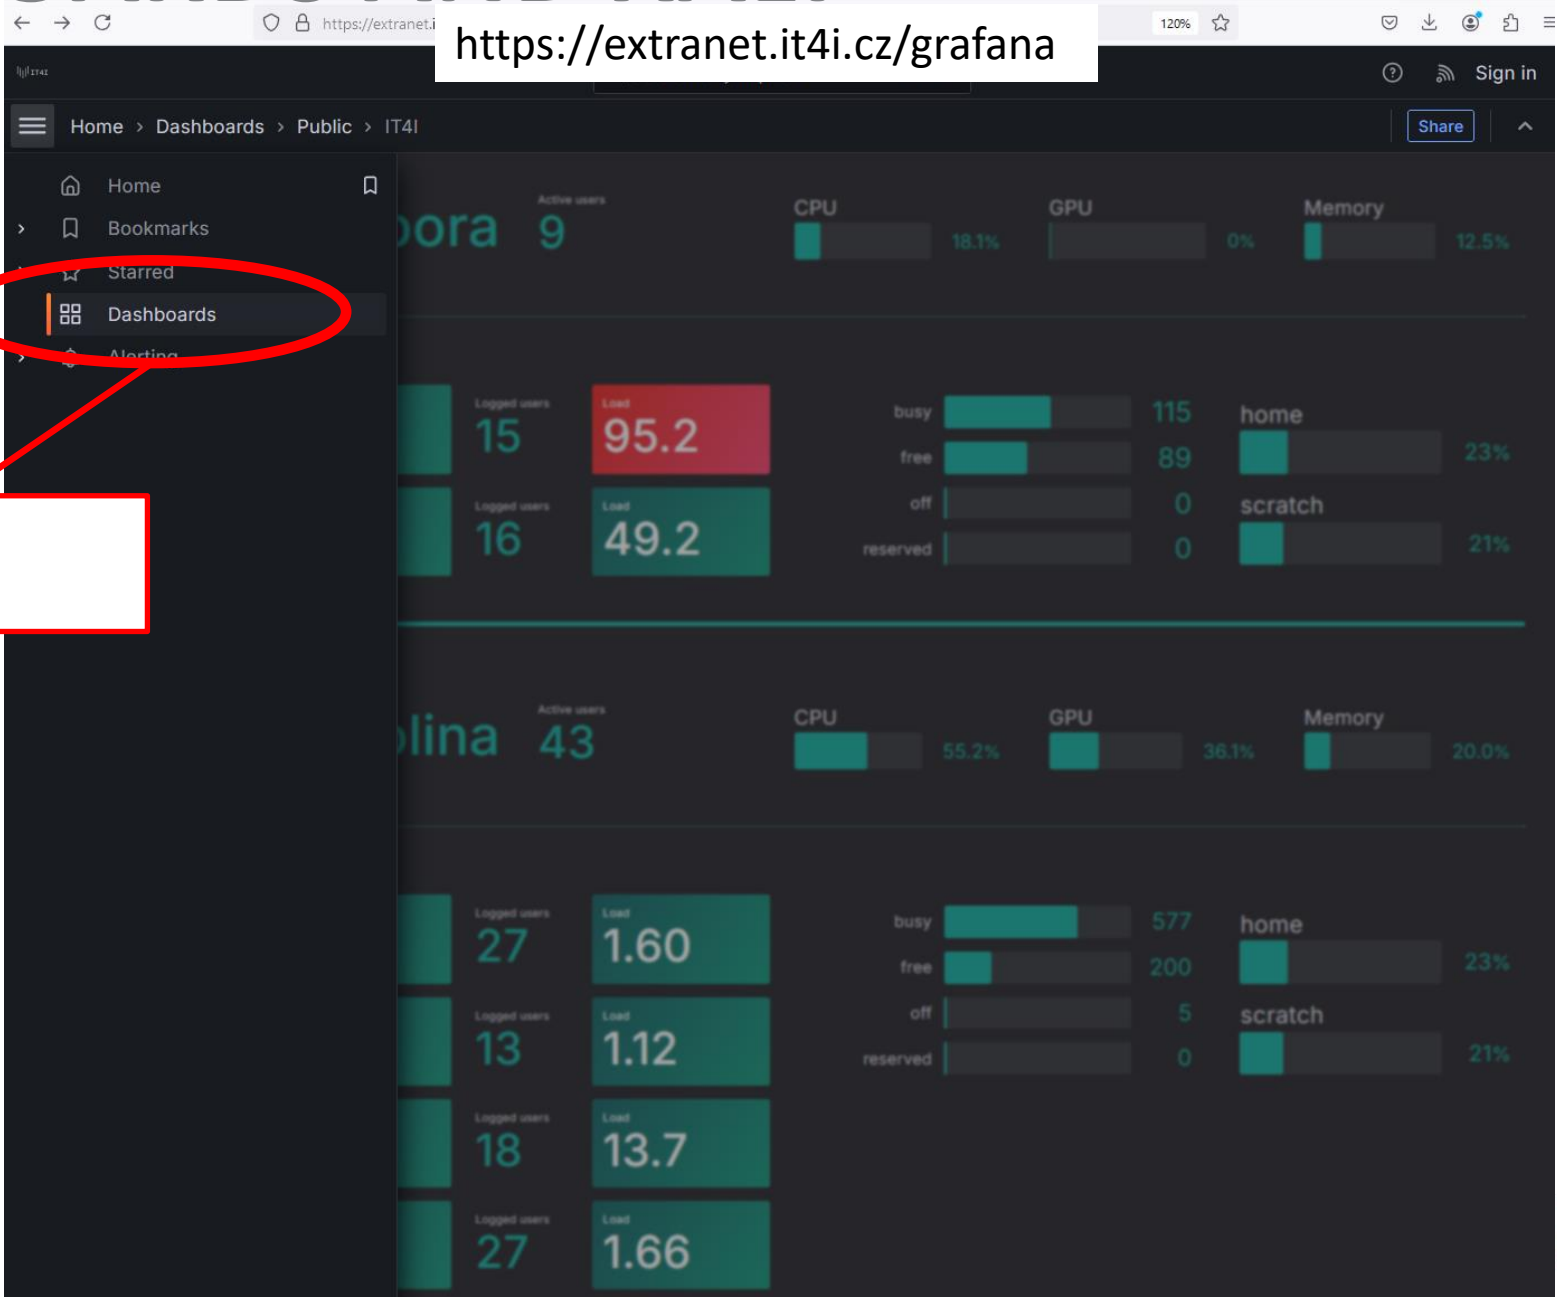

job = Job()

# Create a first task that generates data
task_a = job.program(["generate-data", "--file", "out.txt"])


# Create a dependent task that consumes the data
job.program(["consume-data", "--file", "out.txt"], deps=[task_a])
```

The second task will not be started until the first one successfully finishes.

DASHBOARDS AND XALT

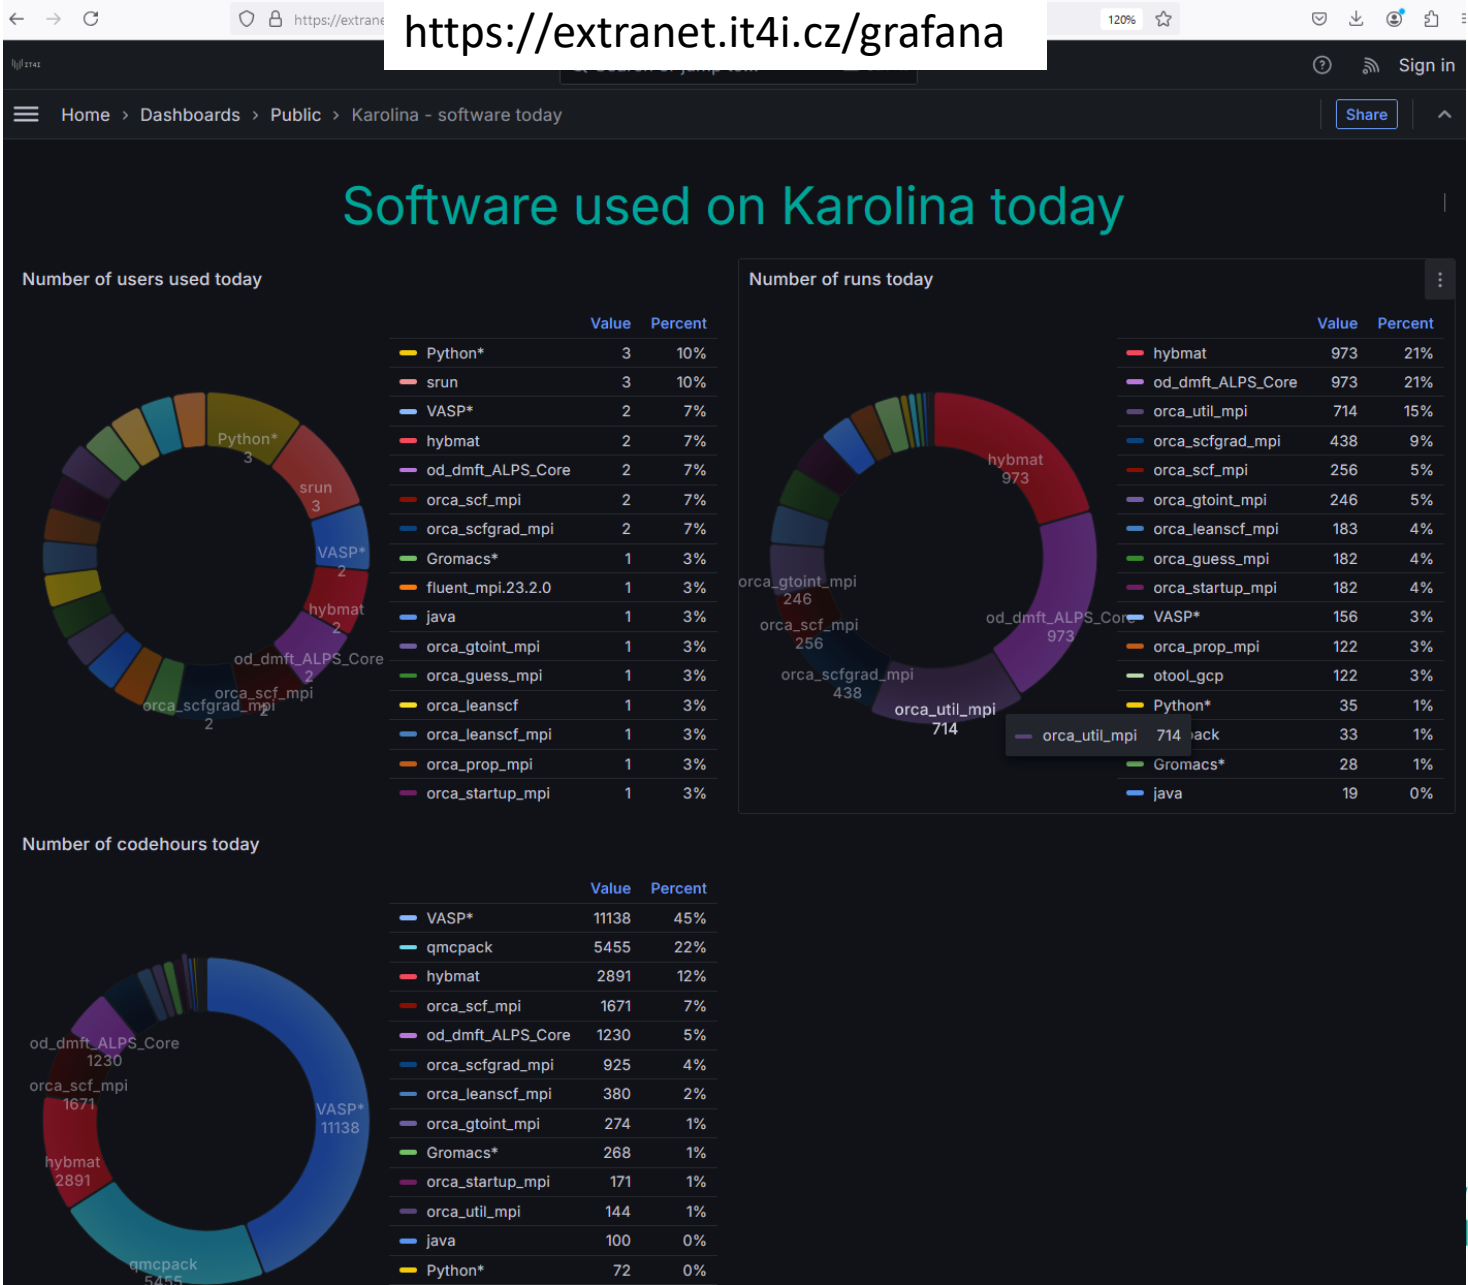



<https://extranet.it4i.cz/grafana>




Select Dashboard

DASHBOARDS AND XALT



\$ml --force rm XALT



HOST YOUR CODE AT CODE.IT4I.CZ




← → ↻ https://code.it4i.cz/jansik/mandelbrot/-/blob/master/README.md?ref_type=heads 170% ☆ 📧 ⬇️ 🔄 📄

📁 Branislav Jansik / mandelbrot

🔗 master ▾ mandelbrot / README.md Find file Blame History Permalink

 **Update README.md** e713a2e1 
Branislav Jansik authored 2 months ago

☰ ▾ M+ README.md  17.65 KIB </> 📄 Blame Edit ▾ Replace Delete 🔄 📄 ⬇️

Mandelbrot benchmark

Processor benchmark. Measures pure floating point performance of the processor (x86 CPU, ARM CPU, Power CPU, NVidia GPU and AMD GPU).
Code by Branislav Jansik, IT4Innovations


Intro

The Mandelbrot benchmark measures pure floating point performance of the processor (x86 CPU, ARM CPU, Power CPU, NVidia GPU and AMD GPU). All calculations are on registers only, no memory or disk access. **Very close to peak floating point performance is sustained**, nominal peak may be exceeded. The code puts extreme load on processor (CPU or GPU), forcing operation at TDP. Thermal throttling may be observed.

The code is implemented in assembly language. Fused multiply add (FMA) vector instructions or warp matrix-multiply-add (WMMA) instructions are executed. Instruction sets supported : x86 (SSE), x86_64(SSE, AVX, AVX2, AVX-512), ARM(AArch64 NEON, SVE), cuda(PTX), amd(RDNA), Power (ppc64/ppc64le VSX).

Optimized to perfectly fill the instruction pipeline, up to two instructions are retired every clock cycle on x86

KAROLINA SW UPGRADE



OS upgrade on Karolina done in April


From system perspective

- OS upgraded from CentOS 7.9 to Rocky Linux 8.9
- The SCRATCH storage backend was upgraded from ClusterStor 4.5 to ClusterStor 6.6, enhancing storage efficiency.
- The HPCM software was upgraded from version 1.4 to version 1.11

From a user perspective

- updated kernel, libraries, and the GPU and MOFED drivers
- performance and stability of the SCRATCH repository have been improved.
- New Slurm (ver 23.11.6)
- The **LibSci BLAS** library is now available, and the GPU and MOFED drivers

KAROLINA SW UPGRADE



BLAS library performance benchmark

<https://code.it4i.cz/jansik/lorenz>

<https://docs.it4i.cz/> (Pending!)

```
ml PrgEnv-intel
```

```
ml cray-pmi/6.1.14
```

```
export LD_LIBRARY_PATH=
```

```
$LD_LIBRARY_PATH:$CRAY_LD_LIBRARY_PATH:$CRAY_LIBSCI_PREFIX_DIR/lib:/opt/cray/pals/1.3.2/lib
```

```
mpiicx -qopenmp -L$CRAY_LIBSCI_PREFIX_DIR/lib -I$CRAY_LIBSCI_PREFIX_DIR/include -o lorenz-blas.x lorenz-blas.c  
-lsci_intel_mp -lsci_intel_mpi_mp
```


```
# MP ONLY VERSION
```

```
OMP_NUM_THREADS=128 OMP_PROC_BIND=true CRAY_OMP_CHECK_AFFINITY=TRUE ./lorenz-blas.x 10000
```


```
# MPI + MP VERSION
```

```
OMP_NUM_THREADS=64 OMP_PROC_BIND=true CRAY_OMP_CHECK_AFFINITY=TRUE mpirun -n 2 ./lorenz-blas.x  
10000
```

<https://www.youtube.com/watch?v=uHrIXC7Tr6U>




BARBORA END OF LIFE



- Deployed 10/2019
- End of vendor support 31.12.2024
- Phase out 2025



COMPLEMENTARY SYSTEMS



Clusters

Karolina >

Barbora >

NVIDIA DGX-2 >

Complementary Systems >

Introduction

Accessing CS

Specification

Complementary System Job Scheduling

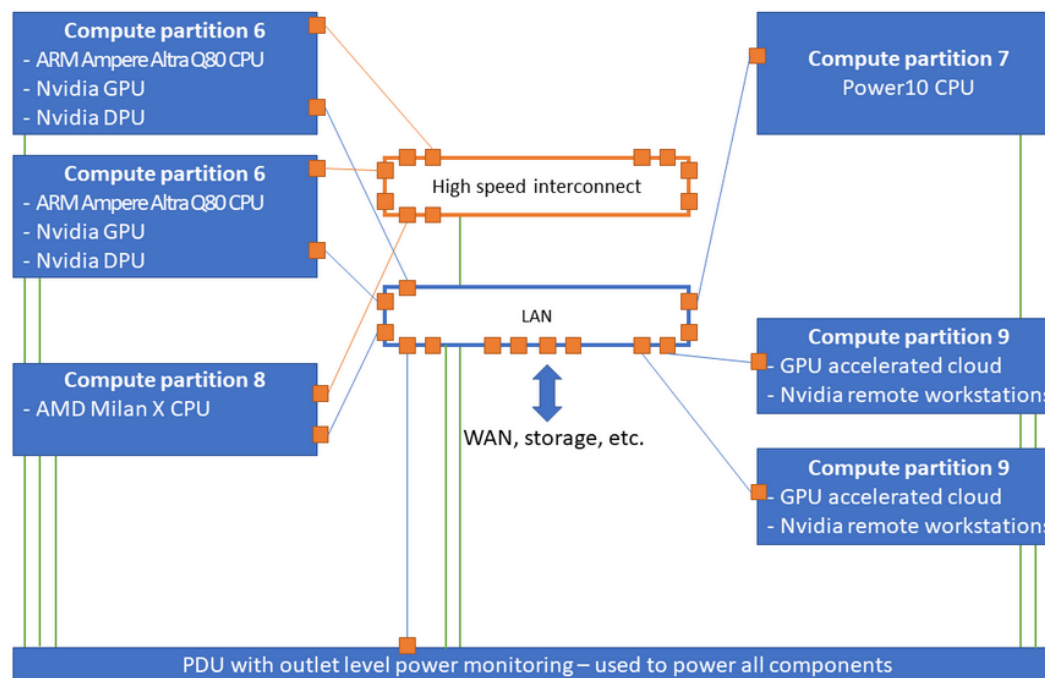
Guides >

Archive >

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



Ta


Cc

Cc

M

Av

COMPLEMENTARY SYSTEMS



Partition 10 - Sapphire Rapids-HBM Server

HBM memory on the x86 processor on the performance of the user applications.


- 2x Intel® Xeon® CPU Max 9468 48 cores base 2.1GHz, max 3.5Ghz
- 16x 16GB DDR5 4800Mhz + **8x 16GB HBM2**
- 2x Intel D3 S4520 960GB SATA 6Gb/s
- FP64 Peak of 7.8TFLOPS

Partition 11 - NVIDIA Grace CPU Superchip


The [NVIDIA Grace CPU Superchip](#) ,144 Arm® Neoverse V2 cores and 1TB/s of memory bandwidth. Runs all NVIDIA software stacks and platforms, including NVIDIA RTX™, NVIDIA HPC SDK, NVIDIA AI, and NVIDIA Omniverse™

- Superchip design with up to **144 Arm Neoverse V2** CPU cores with Scalable Vector Extensions (**SVE2**)
- World's first LPDDR5X with error-correcting code (ECC) memory, 1TB/s total bandwidth
- 900GB/s coherent interface, 7X faster than PCIe Gen 5
- NVIDIA Scalable Coherency Fabric with 3.2TB/s of aggregate bisectional bandwidth
- 2X the packaging density of DIMM-based solutions
- 2X the performance per watt of today's leading CPU
- FP64 Peak of 7.1TFLOPS

NEW SYSTEMS, MCIII





Compute		cca 140 nodes
Cores		cca 27000 cores
Memory		768 GB
Freq		2.1 GHz
Throughput		Over 1.2 TB/s Over 8000 MT/s
Net	IB NDR	400 Gbps
Ridge point		Below 8 F/B



Contract signature pending!

LUMI-Q QUANTUM COMPUTER

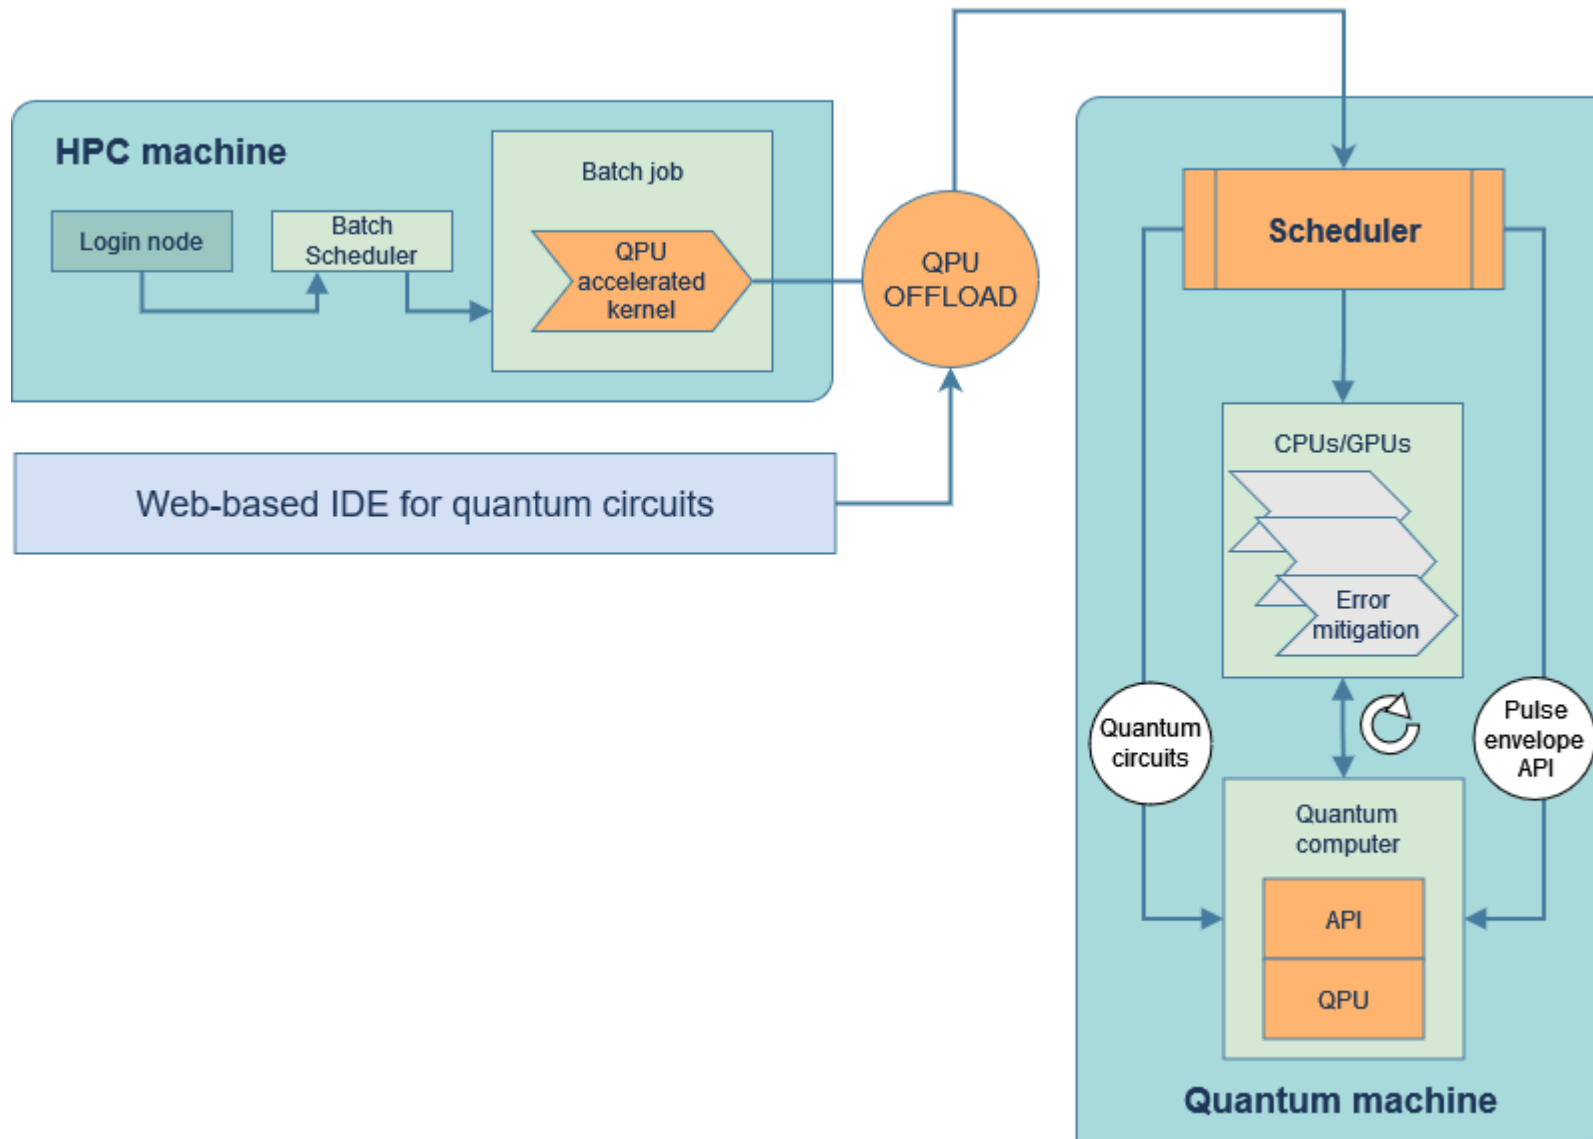




- Star-shape qubit topology, one-to-all qubit connectivity,
- Contract signed 25.9.2024
- Implementation ongoing
- Delivery 25.9.2025

- Upgrade discussion
- 6 qubit star (Twin star)
- Delivery April 2025

Metric	Value
Qubits	24 (22)
Qubit connectivity	one-to-all, star-shape
T1 relaxation time	typically $\sim 40 \mu\text{s}$ minimum for all qubits: $15 \mu\text{s}$
T2 dephasing time	typically $\sim 20 \mu\text{s}$ minimum for all qubits: $15 \mu\text{s}$
1-qubit gate time	$< 50 \text{ ns}$
2-qubit gate time	$< 120 \text{ ns}$
1-qubit gate fidelity	> 0.997
2-qubit gate fidelity	> 0.99
Readout fidelity	> 0.95

LUMI-Q EUROHPCQS-INTEGRATION






LUMI-Q RENDERING


Block Layout

Framing and wires not shown, sized placeholders used



Window views





<https://events.it4i.cz/>


Nov 07


[ONLINE] Introduction to NVIDIA CUDA Quantum

Nov 26-27

[HYBRID] Moving your AI training jobs to LUMI: A Hands-On Workshop (EuroCC)







REGISTER YOUR PUBLICATIONS




Agendas ▾ Requests ▾ Extranet ▾ <https://scs.it4i.cz/> 

Training
Publications
Final reports
Feedbacks

isik) are signed in as Karel Voprsalek (owner). [Back to admin](#) ✕.

 **Karel Voprs**      Back

Login: **Last Login:** 2014-12-28 (about 8 hours spent)
Name: Karel
Surname: Voprsalek
Email: kvop999@gmail.com
Status: active
Telephone:
Country: CZ
Uid/Gid Number: 1025 | 2015

Last Job: 237781.dm2 (2014-08-27 18:20:32)
Dn: uid=vop999,ou=users,ou=master,dc=it4i,dc=cz
Home: /home/vop999
Login Shell: /bin/bash
Group Policy:
Token: UsUVSwa2ZYojkB1SmzcE
Description: description
SSH Public Key: ssh-rsa
AAAAB3NzaC1yc2...SFQqWCaFPmqhSHoTix02b07d2Lo
TXDV0VU69iAKZMhB9Q== vop999 

Groups

Group	Status	GID	Grouping	Type
vop999	active	2015	posix	user

Total 1

Projects Summarize: ALL(4) DD(4)

PID	Status	Project	Allocation	WCHs project	NCHs project	WCHs user	NCHs user				
DD-13-5 	open	IT4I Employee Access Project	0	N/A%	7258181	N/A%	7251453	N/A%	0.00	N/A%	0.00
DD-14-17 	expired	OMICS Installation	10000	900%	89983	900%	89983	0.0%	0.00	0.0%	0.00



Branislav Jansík
branislav.jansik@vsb.cz

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

VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER