TECHNISCHE UNIVERSITÄT DRESDEN

DRESDEN concept

Job Monitoring using PIKA

Frank Winkler

44th VI-HPS Tuning Workshop
Dresden, 26 February 2024

HPC Performance Analysis

- Performance is crucial for HPC
- Hence, measurement, analysis, and validation are also important
- Numerous established tools are available
- Problem solved?

- … for those who are aware of the problem
- Active preparation is necessary
- Not continuous for all jobs of all users

HPCToolkit

VAMPIR

READEX
Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

MAQAO

LIKWID TOOLS

scalasca

Score-P
Scalable performance measurement infrastructure for parallel codes

ZIH

Continuous HPC Monitoring: Goals

Cluster-wide administrative view

— Existing generic solutions

— Pure hardware perspective

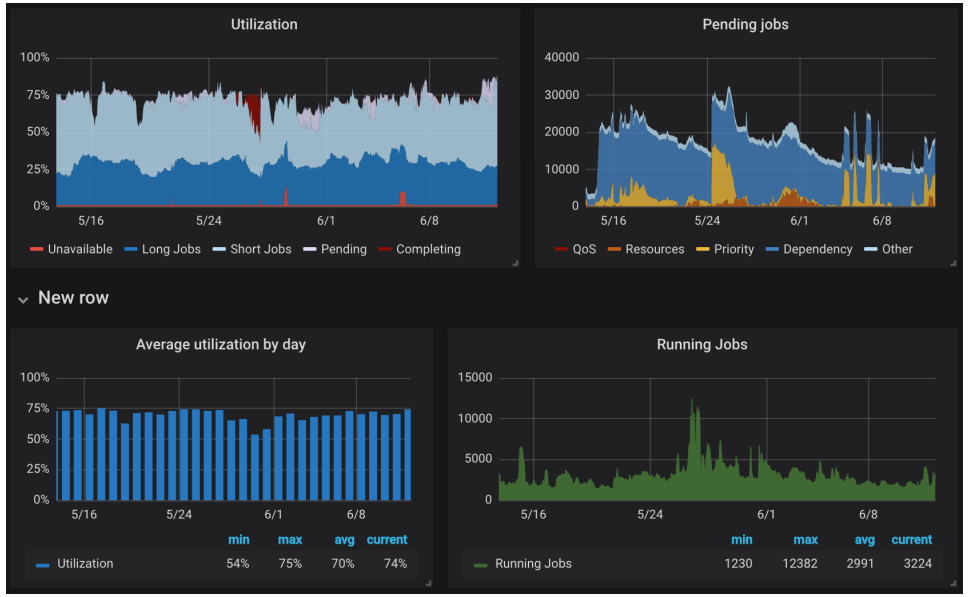

Grafana: Generic visualization of cluster utilization

Job-specific view for users and admins

— Representation for projects, users, jobs

Continuous

— For all jobs, without explicit activation/preparation

— Minimal overhead, coarse granularity

— Performance history available on demand (retrospectively)

PIKA: Continuous HPC Job Monitoring

- Non-intrusive **data acquisition** on all cluster nodes
- Continuous **data collection**
- Web frontend for live and post-mortem **visualization**
- Detection of pathological jobs
- Automatic **job analysis and classification**
- Long-term **data storage**

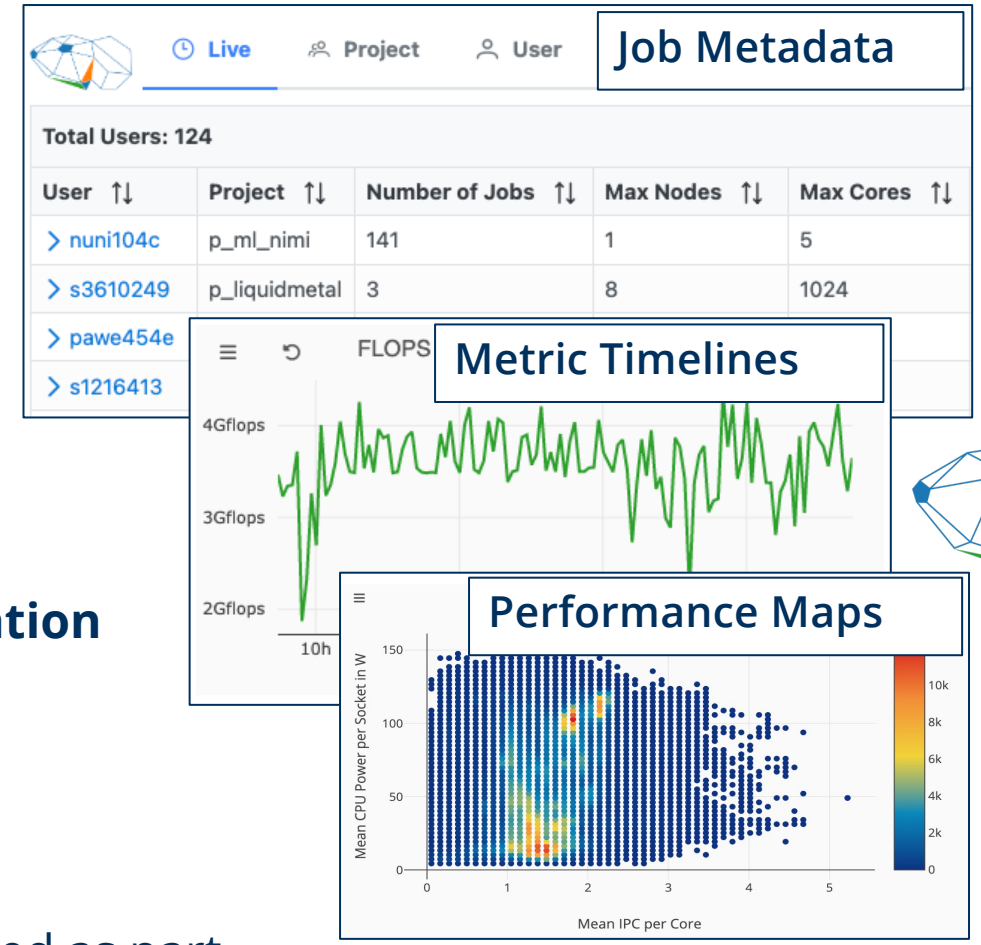

Funded by the DFG project ProPE, continued as part of NHR@TUD at ZIH.

PIKA: Continuous HPC Job Monitoring

“Pikas prefer rocky slopes and graze on a range of plants, mostly grasses, flowers and young stems. In the autumn, they pull hay, soft twigs and other stores of food into their burrows to eat during the long cold winter.”

Source: Walters, Martin (2005). *Encyclopedia of animals*. Parragon. p. 203. ISBN 978-1-40545-669-2.

Outline

- PIKA infrastructure
- PIKA web frontend
- Hands-on exercise
- Discussions

PIKA Architecture Overview

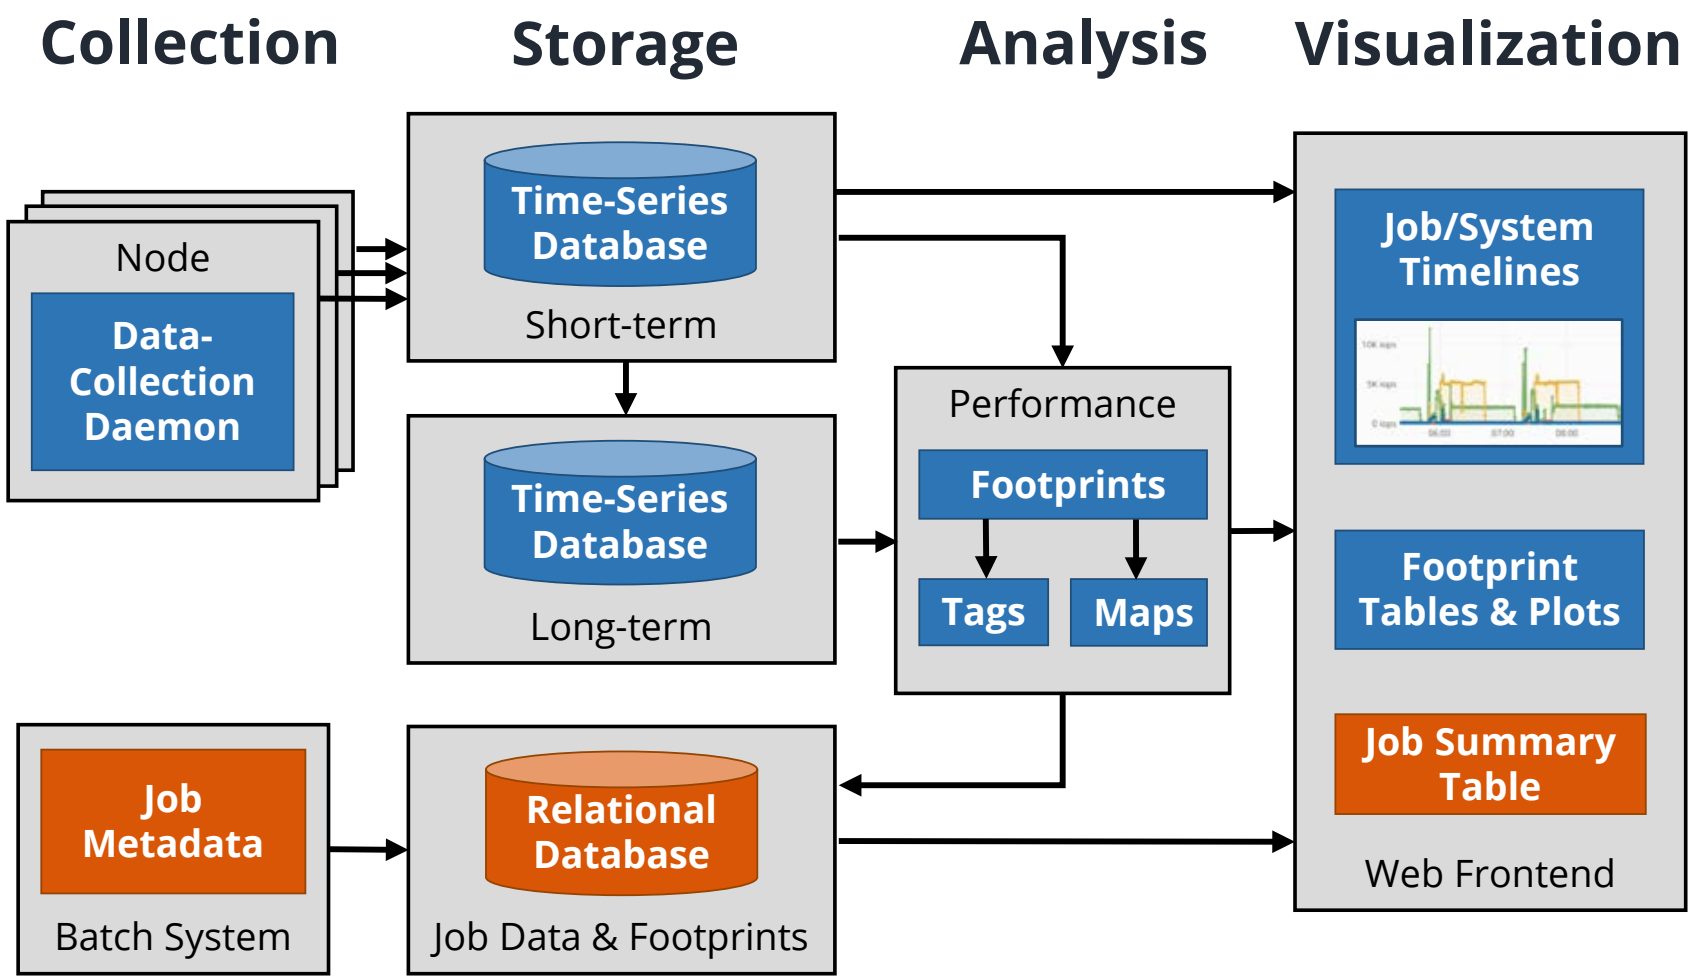

PIKA Architecture at TUD

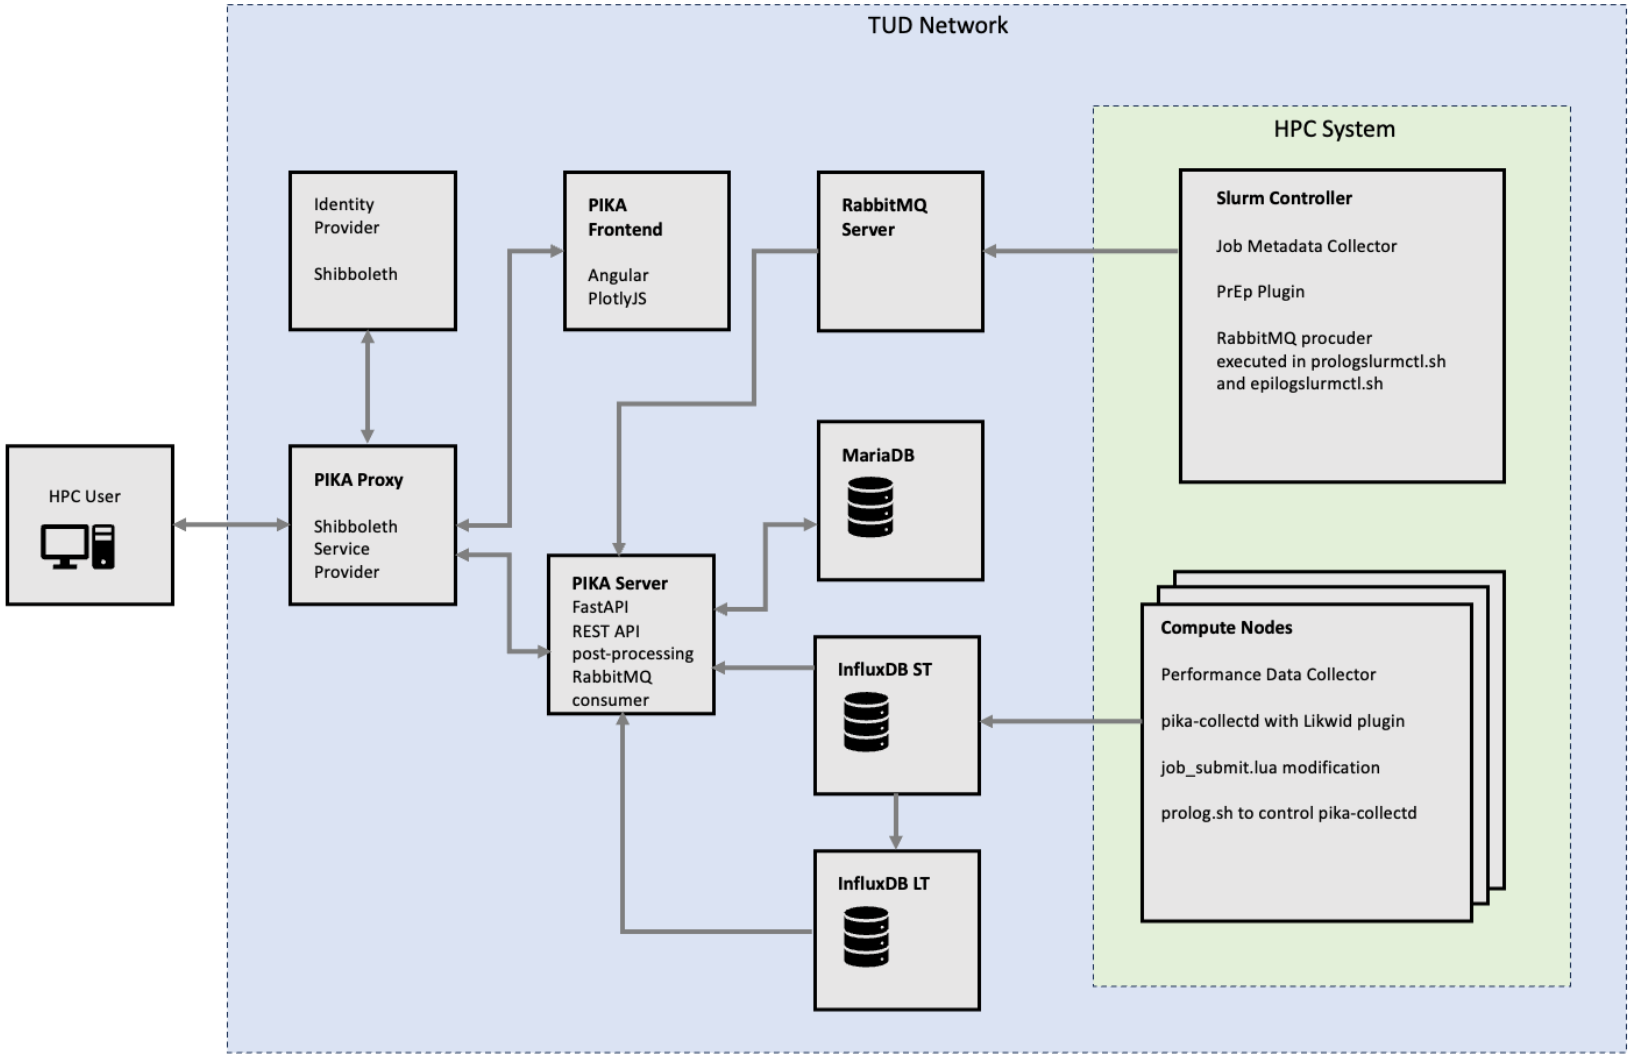

PIKA Metadata Collection

Slurm PrEp Plugin to capture job metadata:

- Unique job identifier, ArrayID
- Project, user, job name
- Start and end time, walltime
- Status (running, completed, timeout, failed, OOM, cancelled)
- Requested resources
  - Partition
  - Allocated compute nodes
  - Allocated CPUs on each node
  - Exclusive nodes
  - Main memory
  - GPUs per node

PIKA Runtime Data Collection

Monitored Metrics	Data Source	Hardware Unit
Instructions per Cycle (IPC)	LIKWID	Hardware Thread
FLOPS (SP Normalized)		Hardware Thread
Main Memory Bandwidth		CPU/Socket
Power Consumption		CPU/Socket
CPU Usage	proc & sysfs	Hardware Thread
Main Memory Utilization		Node
Network Bandwidth		Node
File I/O Bandwidth & Metadata	Local disk, Filesystems (Lustre)	Disk, Lustre Instance
GPU Usage	NVML	GPU
GPU Memory Utilization		
GPU Power Consumption		
GPU Temperature		

Collection daemon **collectd**

- One collector/plugin for each metric source
- CPU counters are collected with LIKWID
- Hardware thread metrics are summarized to the physical CPU core

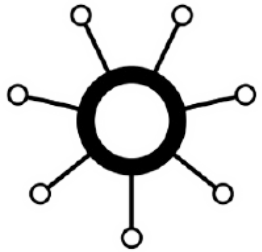

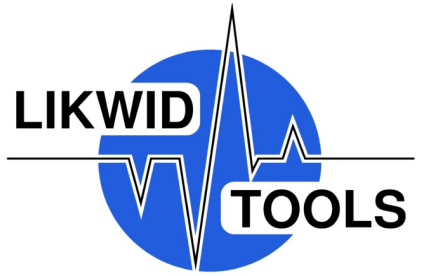

PIKA Job Visualization – Tables

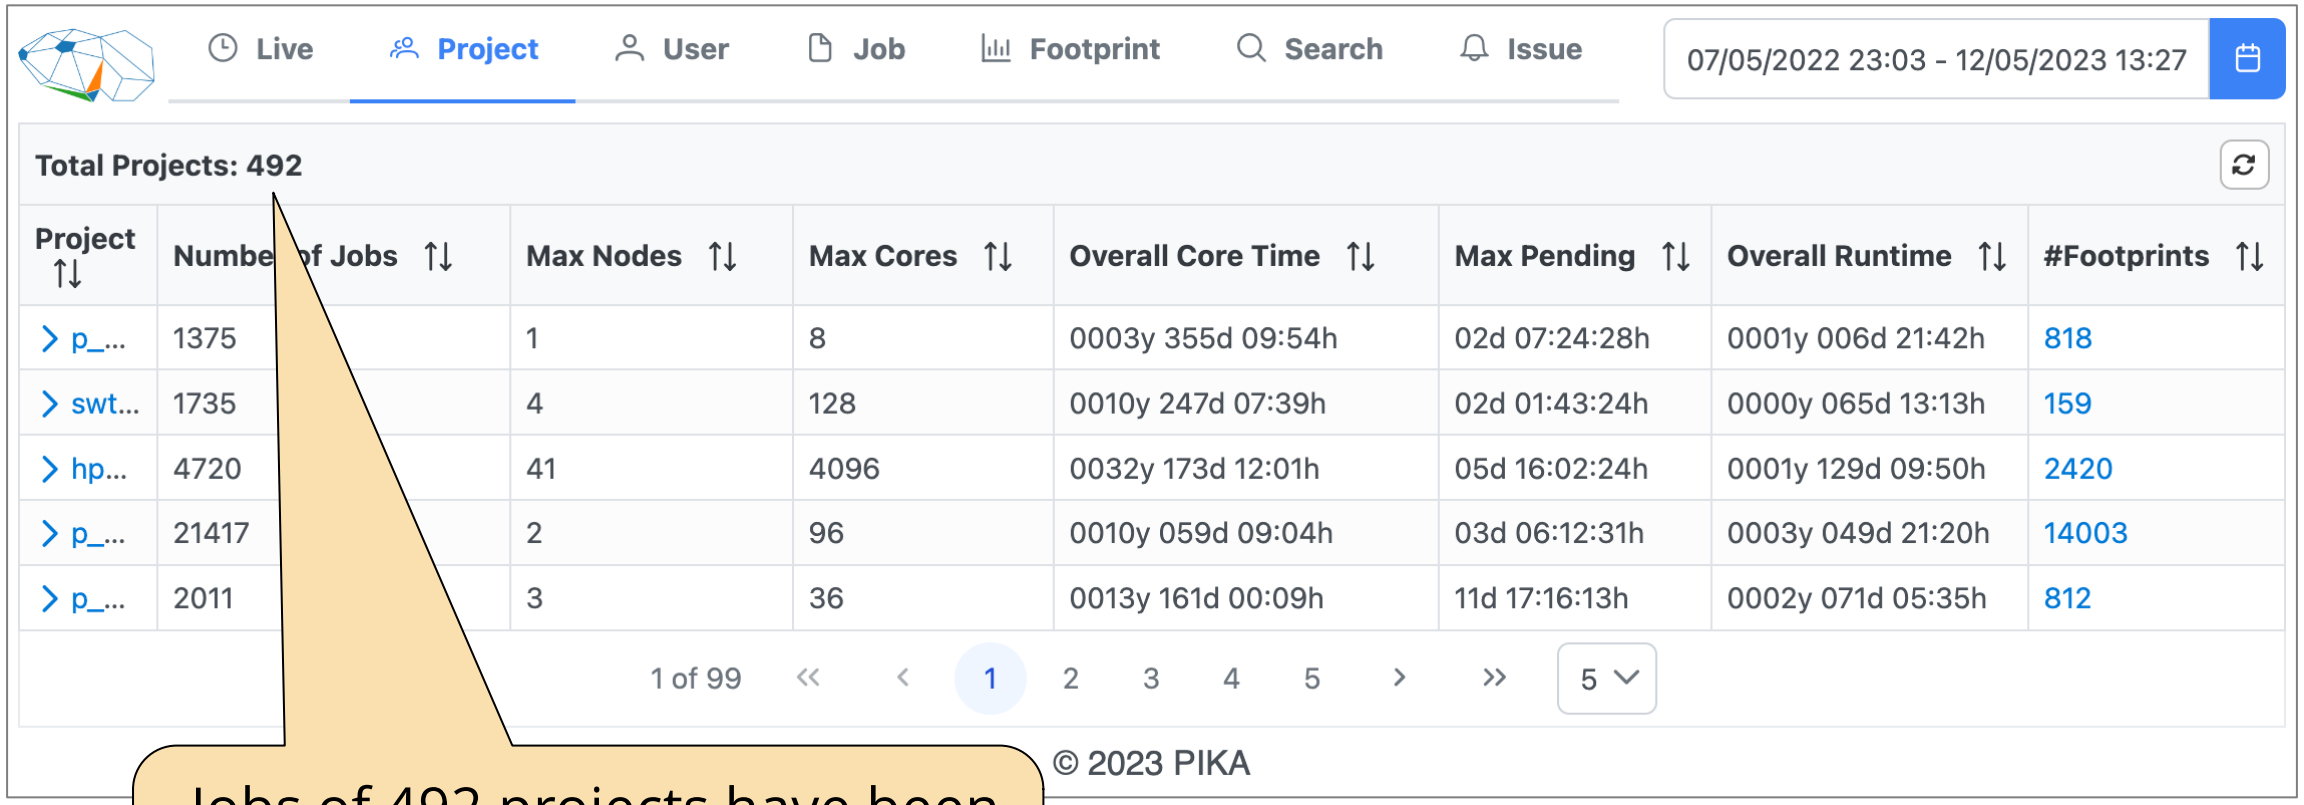

Jobs of 492 projects have been recorded for the selected time interval (top right).

PIKA Job Visualization – Tables

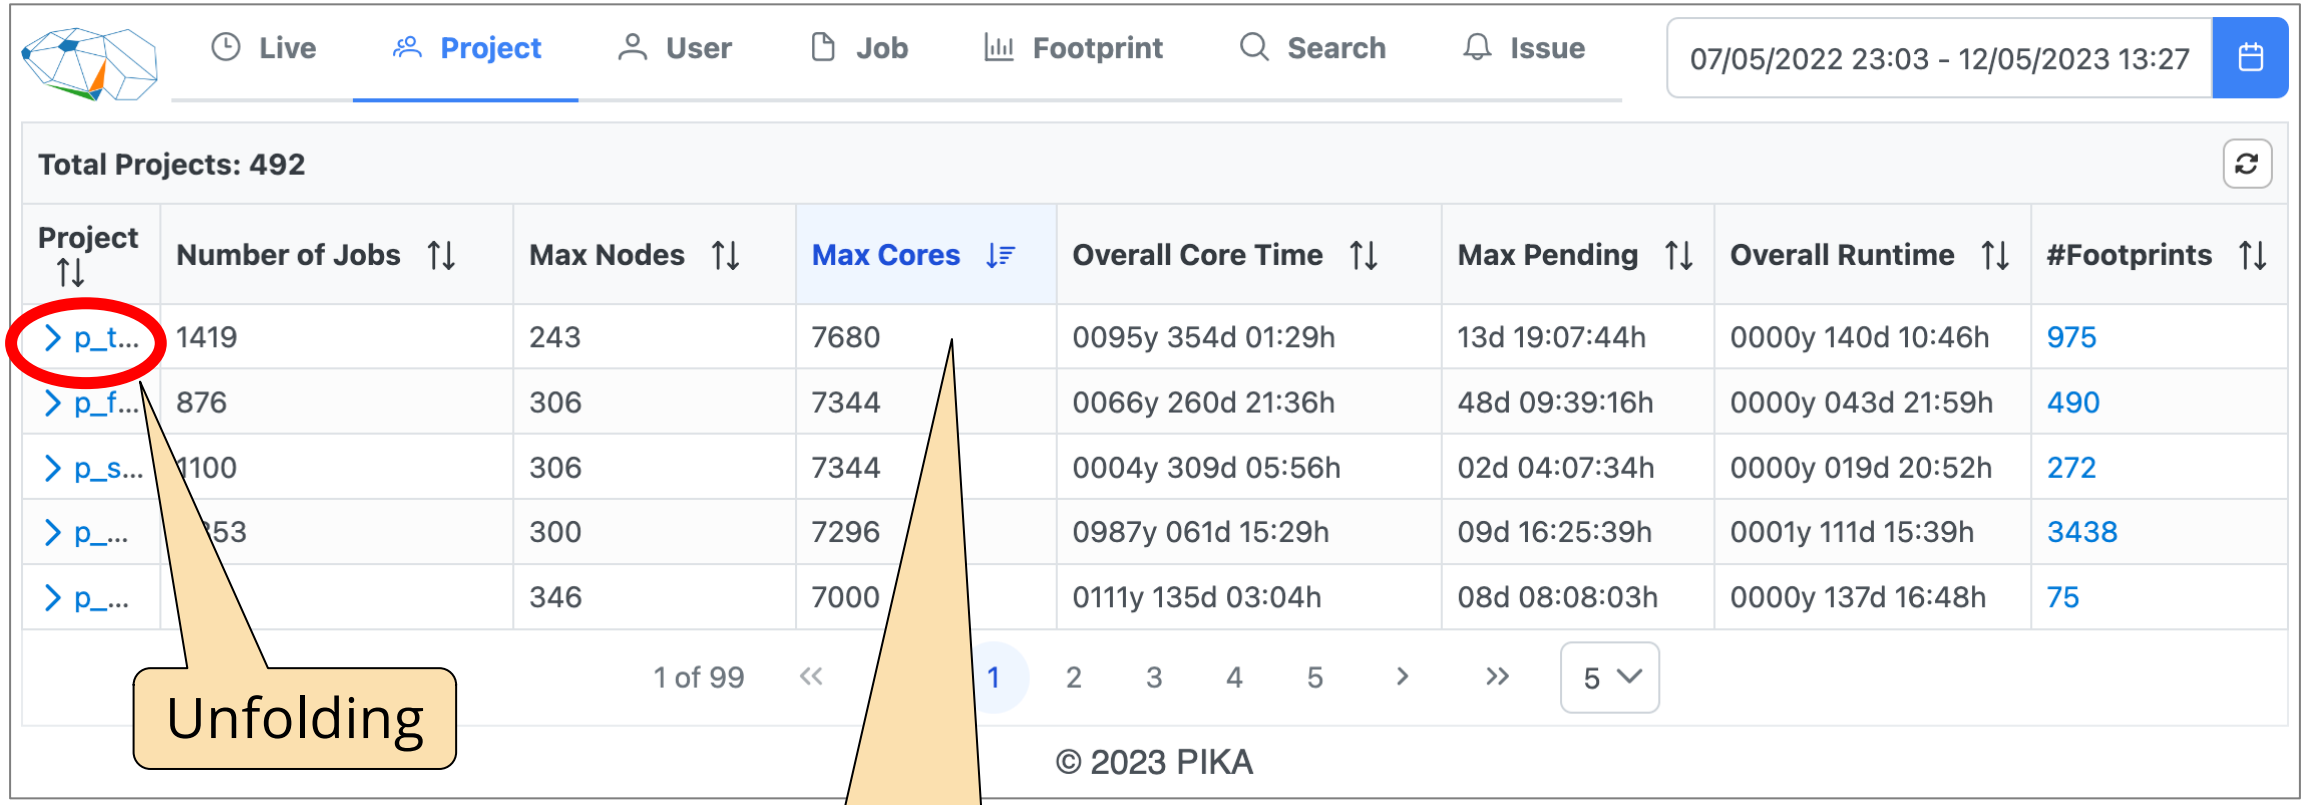

Unfolding

Get project with the highest number of cores.

PIKA Job Visualization – Tables

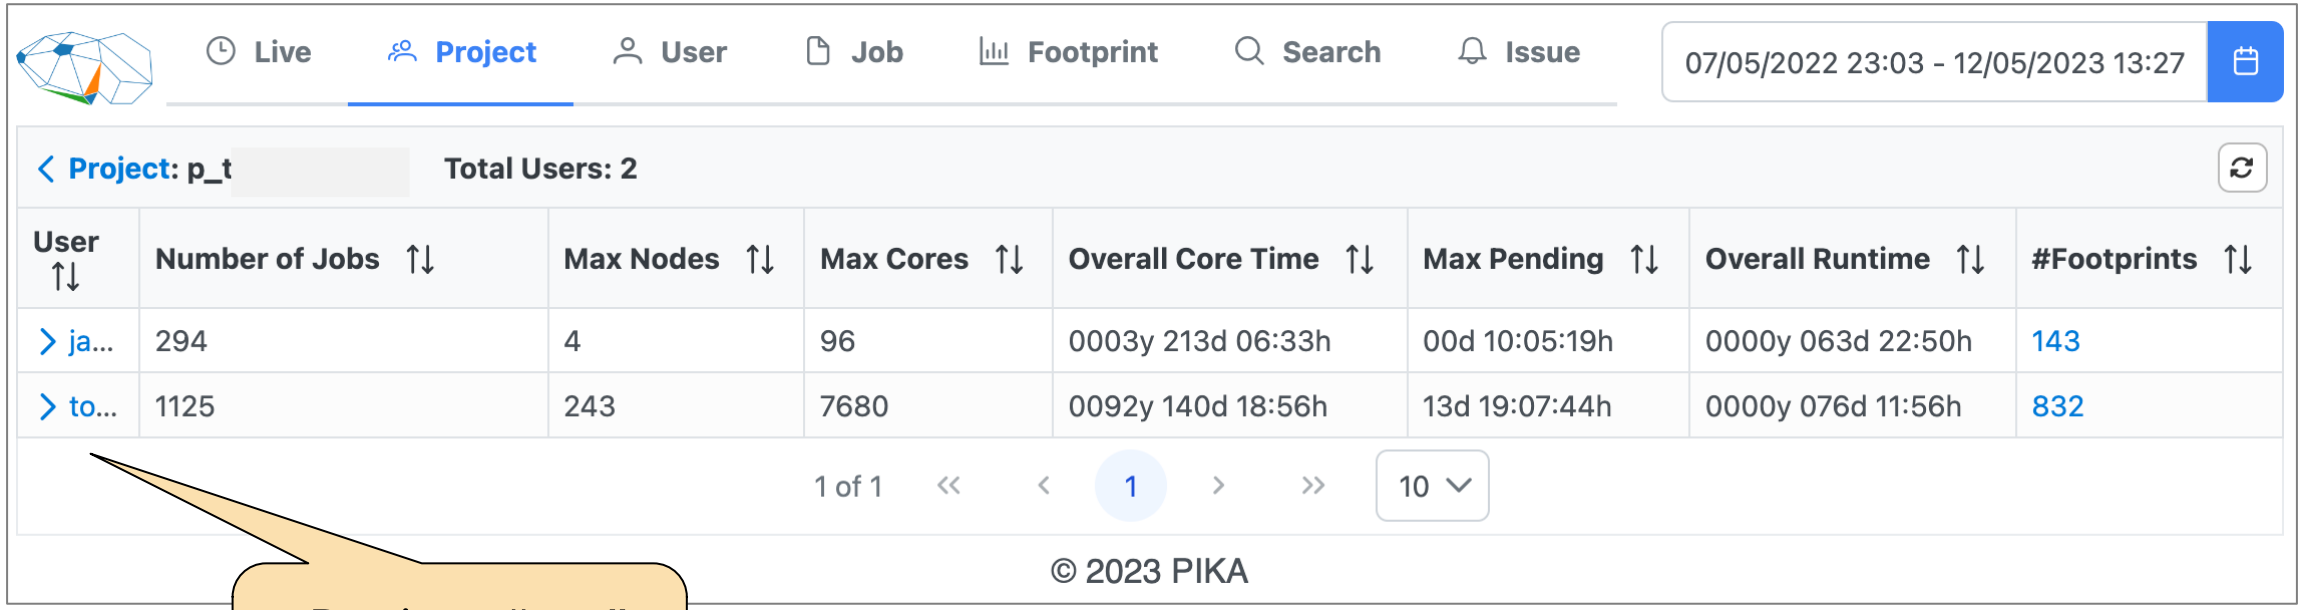

Project “p_t” has two users.

PIKA Job Visualization – Metadata & Timelines

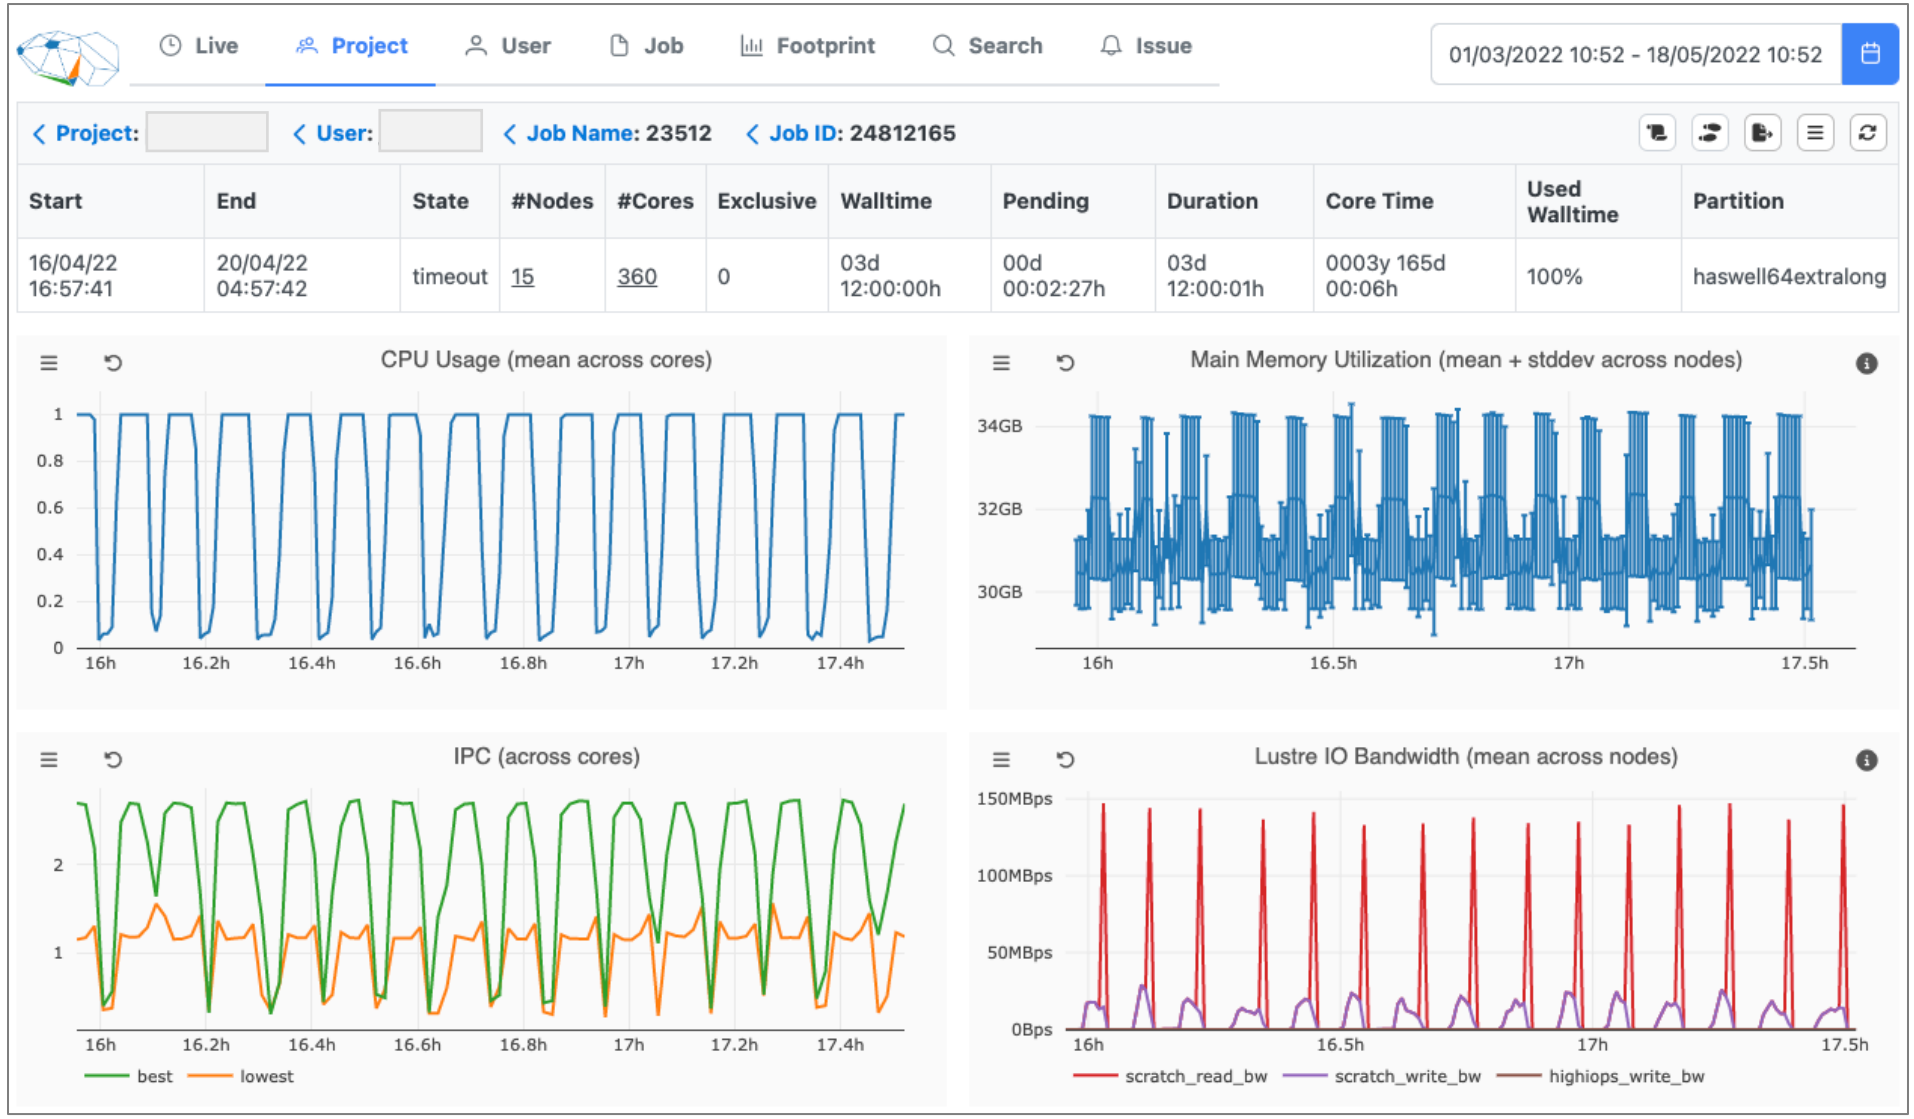

PIKA Post Processing

Job characterization via tagging

- **Footprints** based on summarized runtime data
  - **Average** (CPU and GPU usage, IPC, FLOPS, main memory bandwidth, CPU and GPU power, InfiniBand traffic)
  - **Total** (file IO read/write)
  - **Maximum** (host and GPU memory usage)
- Job tags based on formulas and thresholds

Tag Name	Formula and Threshold
unrestrained	-
memory-bound	$\frac{\texttt{memory bandwidth (measured)}}{\texttt{memory bandwidth (maximum)}} > 80\%$
compute-bound	$\frac{\texttt{FLOP/s (measured)}}{\texttt{FLOP/s (maximum)}} > 70\%$ or $\frac{\texttt{IPC (measured)}}{\texttt{IPC (optimal)}} > 60\%$
GPU-bound	$\texttt{GPU utilization} > 70\%$ or $\texttt{GPU utilization} > \texttt{CPU utilization}$
IO-heavy	$\frac{\texttt{IO bandwidth (measured)}}{\texttt{IO bandwidth (maximum)}} > 60\%$
network-heavy	$\frac{\texttt{network bandwidth (measured)}}{\texttt{network bandwidth (maximum)}} > 60\%$

PIKA Post Processing

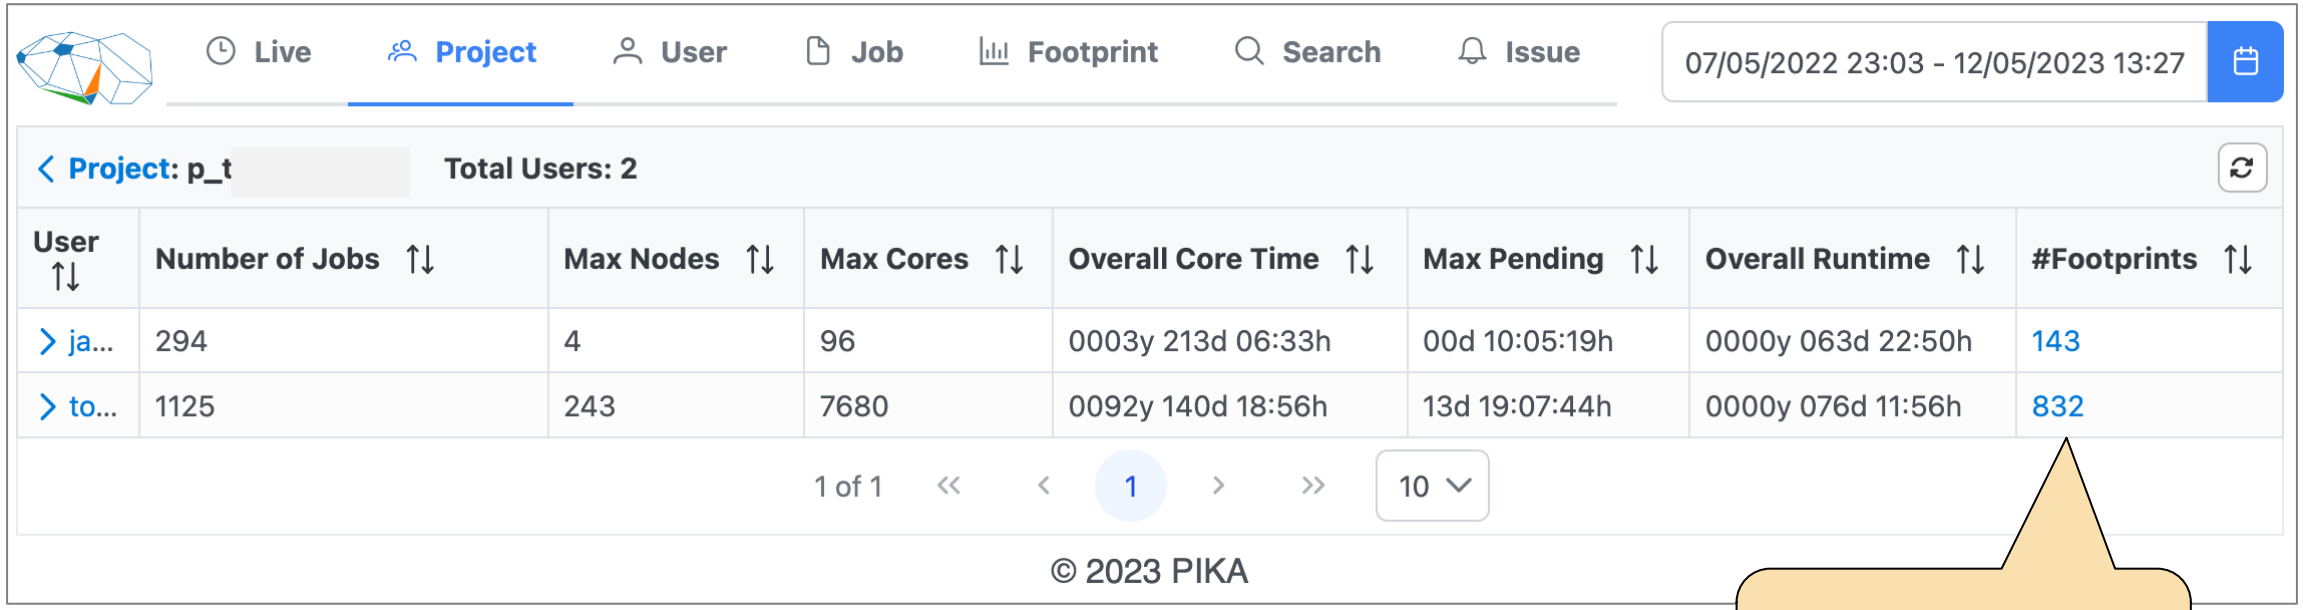

832 user jobs are tagged.

PIKA Job Visualization – Footprints

PIKA – Search for Suboptimal Jobs

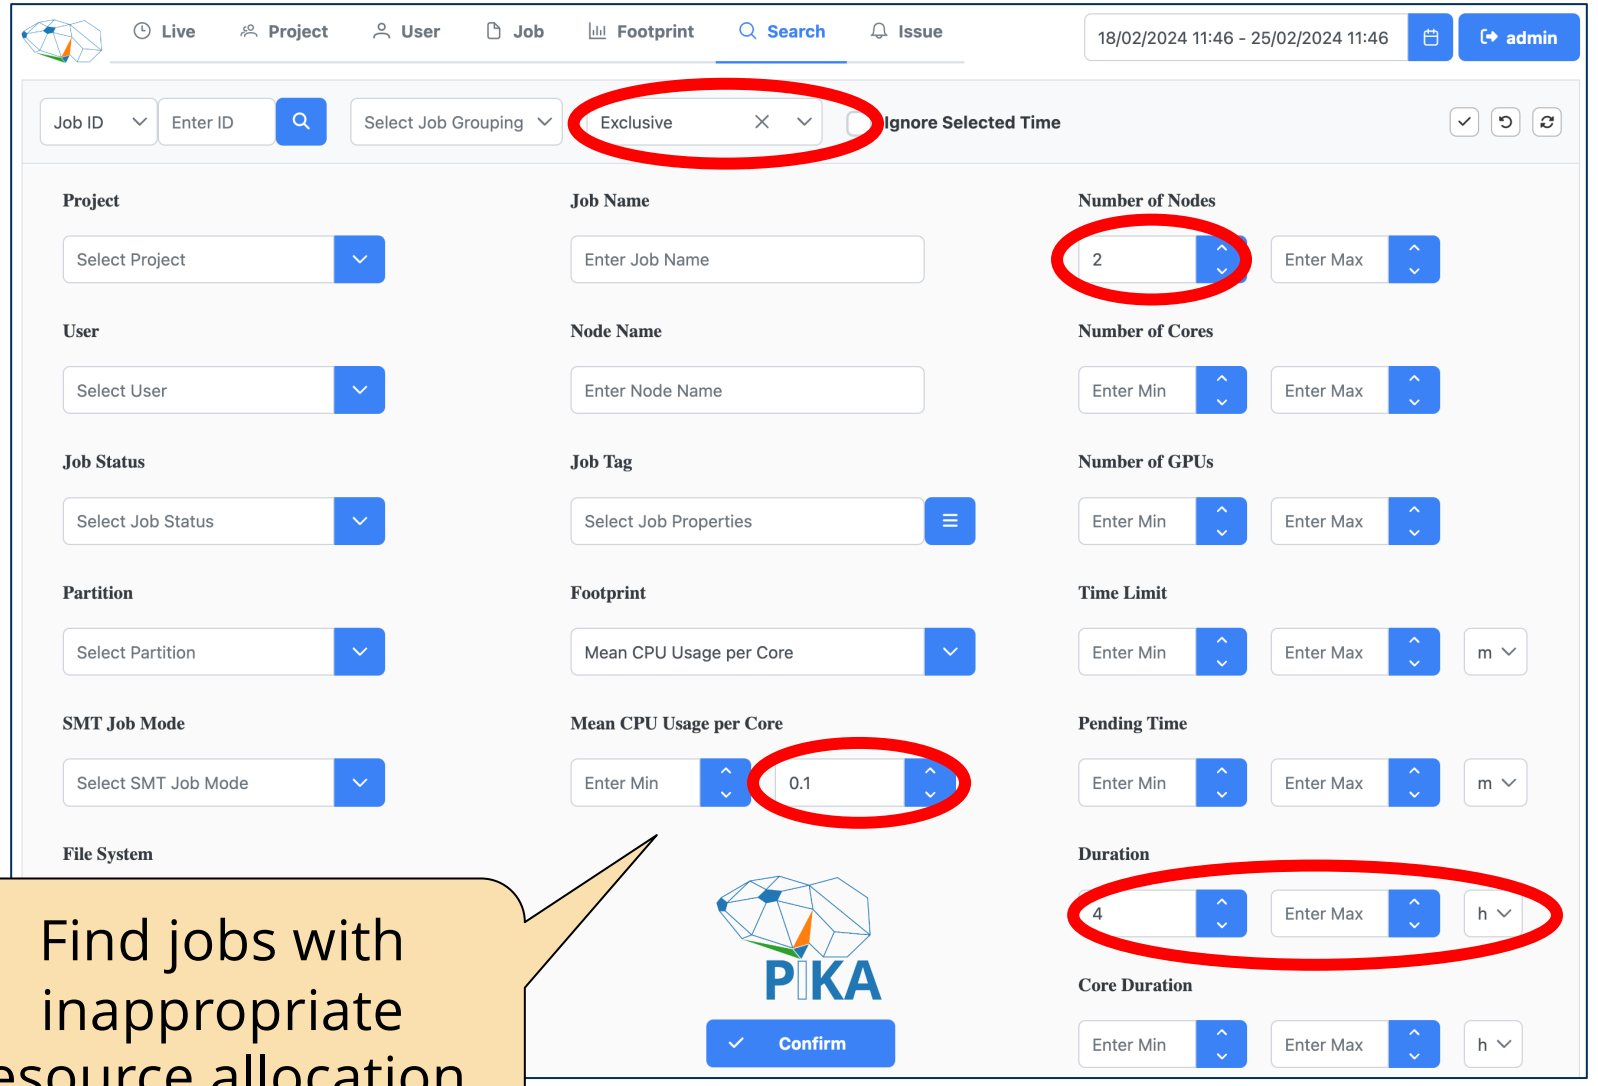

Find jobs with inappropriate resource allocation.

PIKA – Search for Suboptimal Jobs

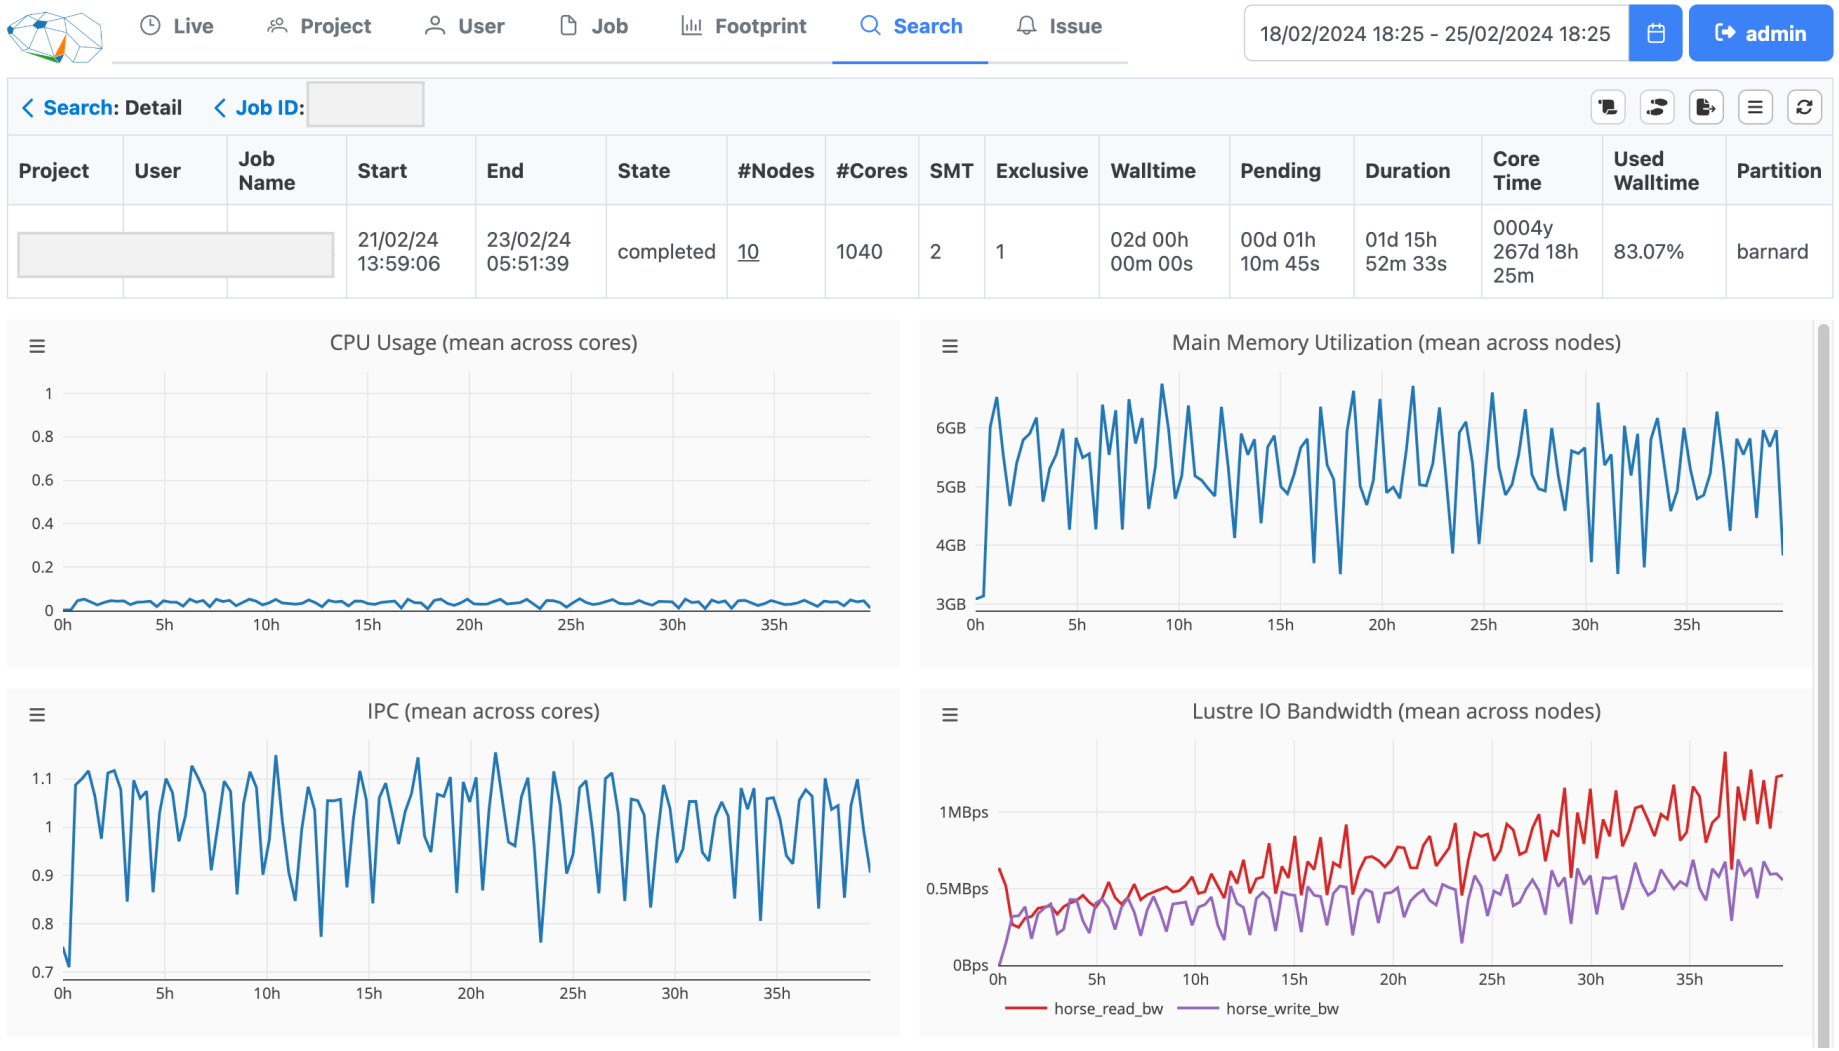

PIKA Issue Analysis

Automatic detection of job performance issues on eligible jobs

- **Prerequisite:**
  - Duration >= 1 hour
  - Number of physical cores > 1
  - Slurm Status: completed, out of memory, timeout
  - Metric timeline vectors*: CPU/GPU load, memory usage, I/O bandwidths and I/O metadata operations
- **Heuristics** to detect inefficient jobs
- **Criteria for efficient usage**
  - Shortest possible runtimes (compared to similar jobs)
  - High utilization of the hardware
  - Even distribution of computational workloads across processing units

* Sampled every 30 seconds

PIKA Issue Analysis – Summarized User View

Possible performance issues with the inefficient HPC jobs of a user

Performance Issue	Description
Idle CPU/GPU Time (**ICT/IGT**)	Summed time intervals of all CPUs/GPUs across all jobs in which the load was close to zero.
Idle CPU/GPU Ratio (**ICR/IGR**)	Quotient of “Idle CPU/GPU Time” and “Total CPU/GPU Time” across all jobs.
Maximum Unused CPU/GPU Ratio (**Max UCR/UGR**)	Maximum ratio of “unused” to “used” CPUs/GPUs across all jobs.
Maximum CPU/GPU Load Imbalance (**Max CLI/GLI**)	Maximum of the average standard deviation of CPU/GPU load across all jobs.
Maximum I/O Congestion (**Max IOC**)	Maximum rate of metadata operations at a measuring point across all jobs. The attribution per job starts with 40 operations.
Maximum I/O Blocking Phases (**Max IOB**)	Maximum periodic number of phases with an inverse correlation between CPU load and I/O metrics across all jobs. The attribution per job starts with 10 periodic phases.
Maximum Synchronous Offloading (**Max SO**)	Maximum periodic number of phases with an inverse correlation between CPU and GPU load across all jobs. The attribution per job starts with 10 periodic phases.
Maximum Memory Leak (**Max ML**)	Maximum of the linear increase of memory usage over time across all jobs.

PIKA Issue Analysis – Issue Table

User jobs are sorted by idle CPU time

PIKA Issue Analysis – Issue Table

User jobs are sorted by idle GPU time

User	Project	#Runs	ICT	ICR	Max UCR	Max CLI	Max IOB	Max IOC	Max ML	IGT	IGR	Max UGR	Max GLI	Max SO	Max S
sek...	p_sca...	406	0002y 096d 03:59h	0.05	0.83	0.45	0	39226	0.86	96d 05:14:30h	0.04	1	0.6	0	0
s12...	p_da...	79	0000y 099d 19:02h	0.18	0.54	0.22	0	0	0.1	91d 07:31:30h	0.98	1	0	0	0
s9...	p_sca...	8781	0078y 345d 18:15h	0.84	1	0.85	0	0	0	919d 09:55:30h	0.98	1	0	0	0
s6...	zihfor...	27	0000y 341d 19:04h	0.41	0.6	0.63	0	336	0.01	90d 04:47:30h	0.65	0.88	0.46	0	0
s3...	zihfor...	413	0005y 337d 04:12h	0.39	1	0.38	0	19464	0.04	893d 10:59:30h	0.61	1	0.54	0	0

PIKA Issue Analysis – Issue Table

User jobs are sorted by maximum I/O congestion

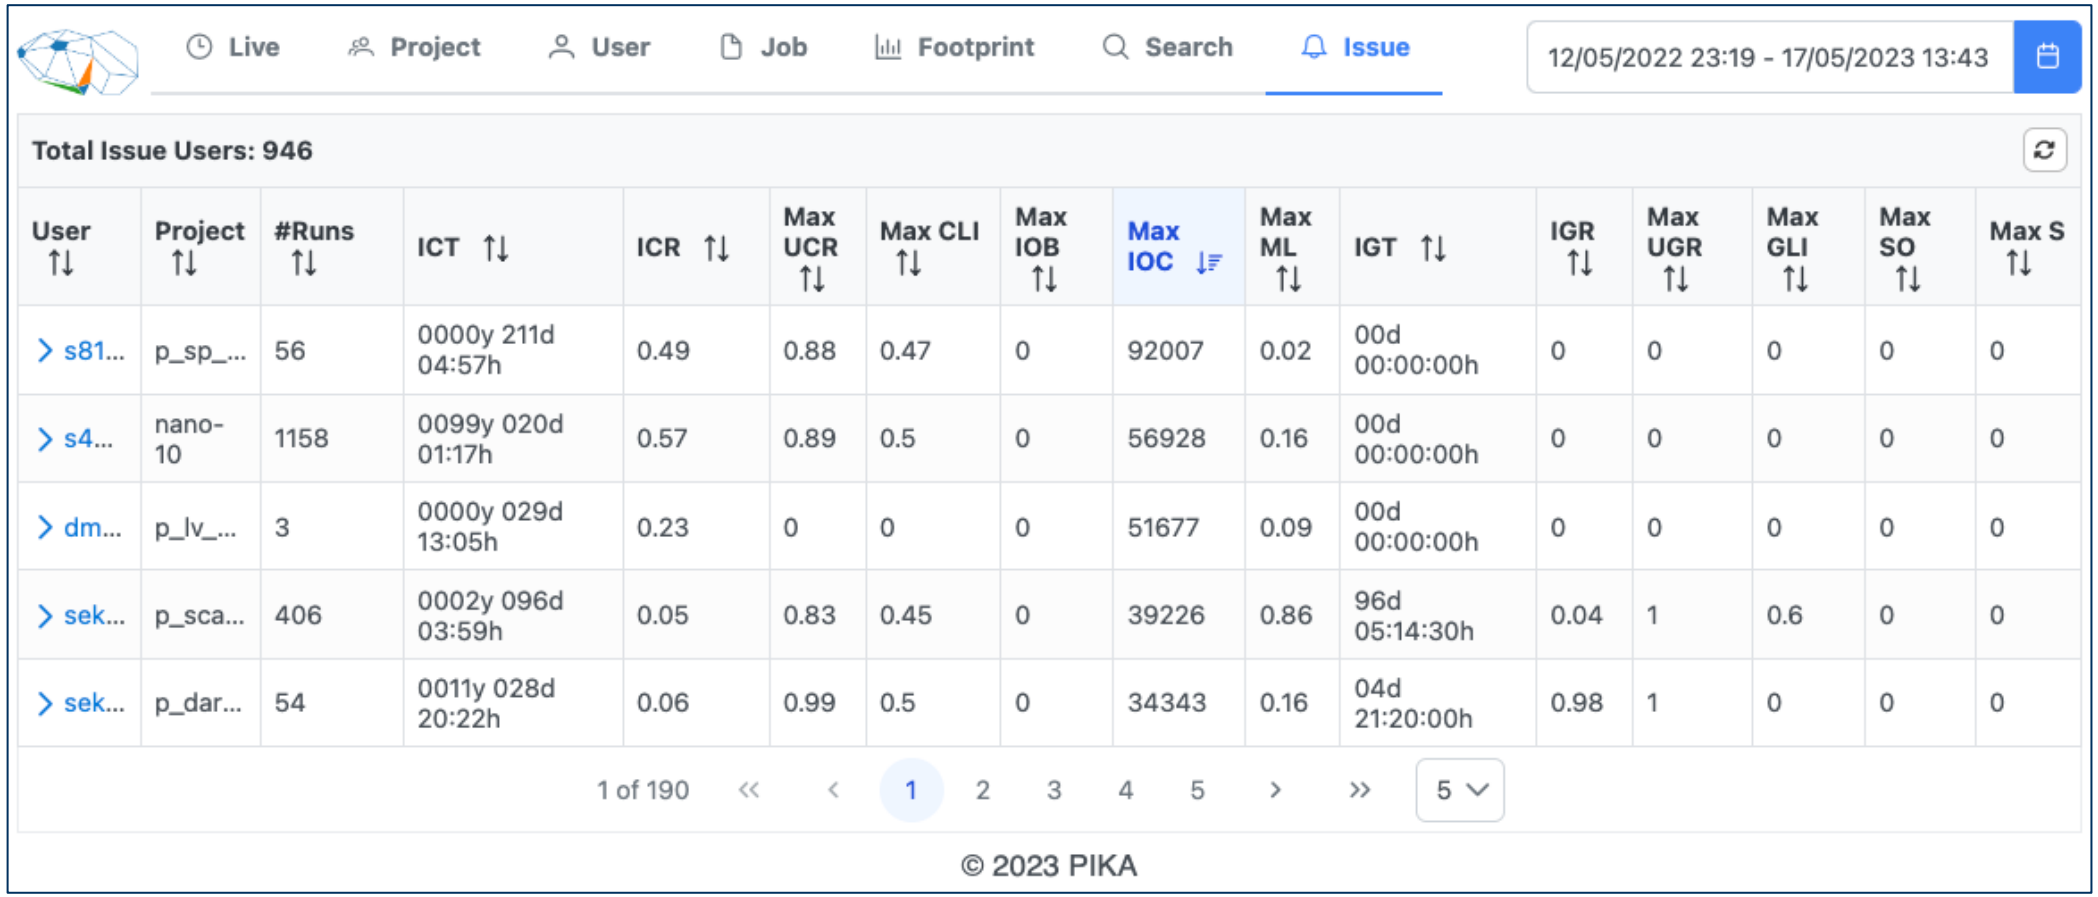

PIKA Issue Analysis – Issue Table

User jobs are sorted by maximum I/O blocking

PIKA Issue Analysis – Issue Table

User jobs with I/O blocking issues

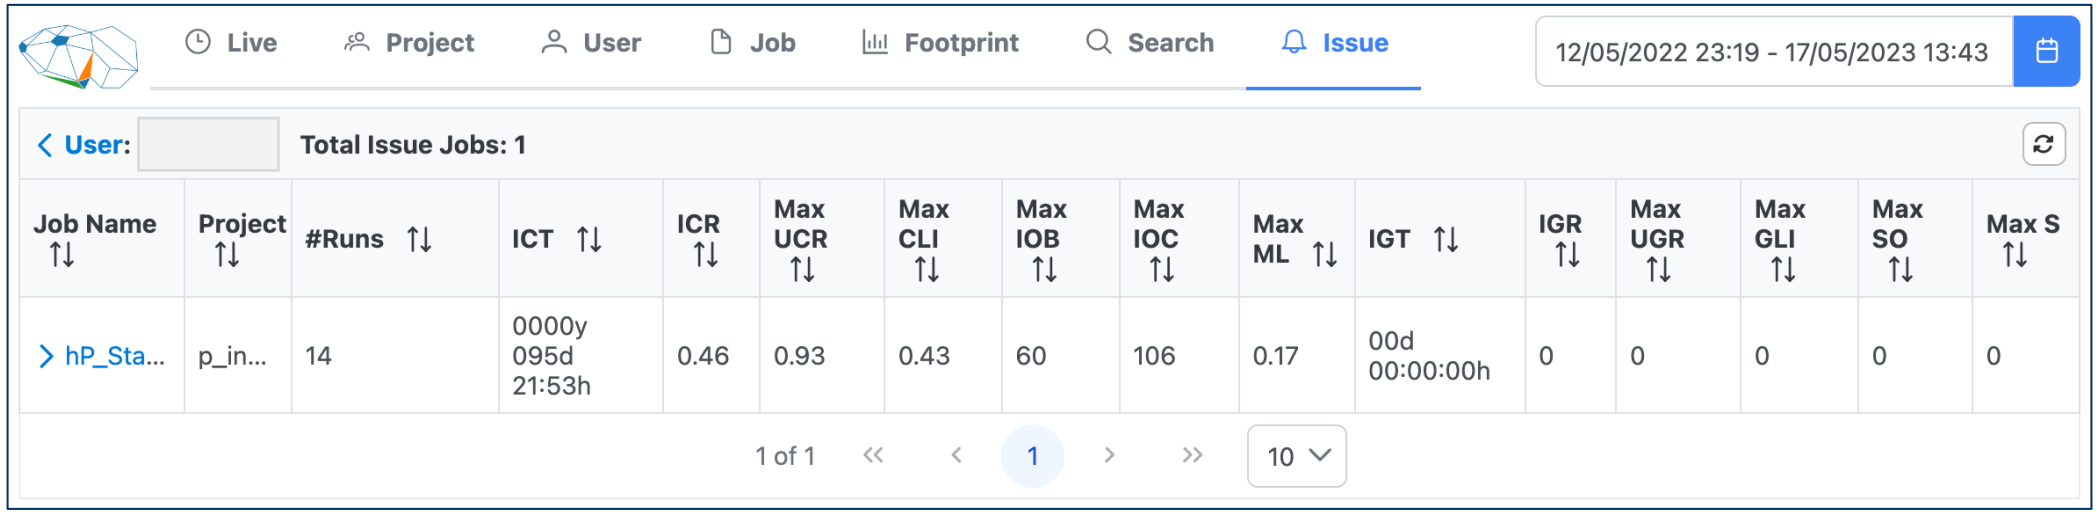

PIKA Issue Analysis – Issue Table

User jobs with I/O blocking issues

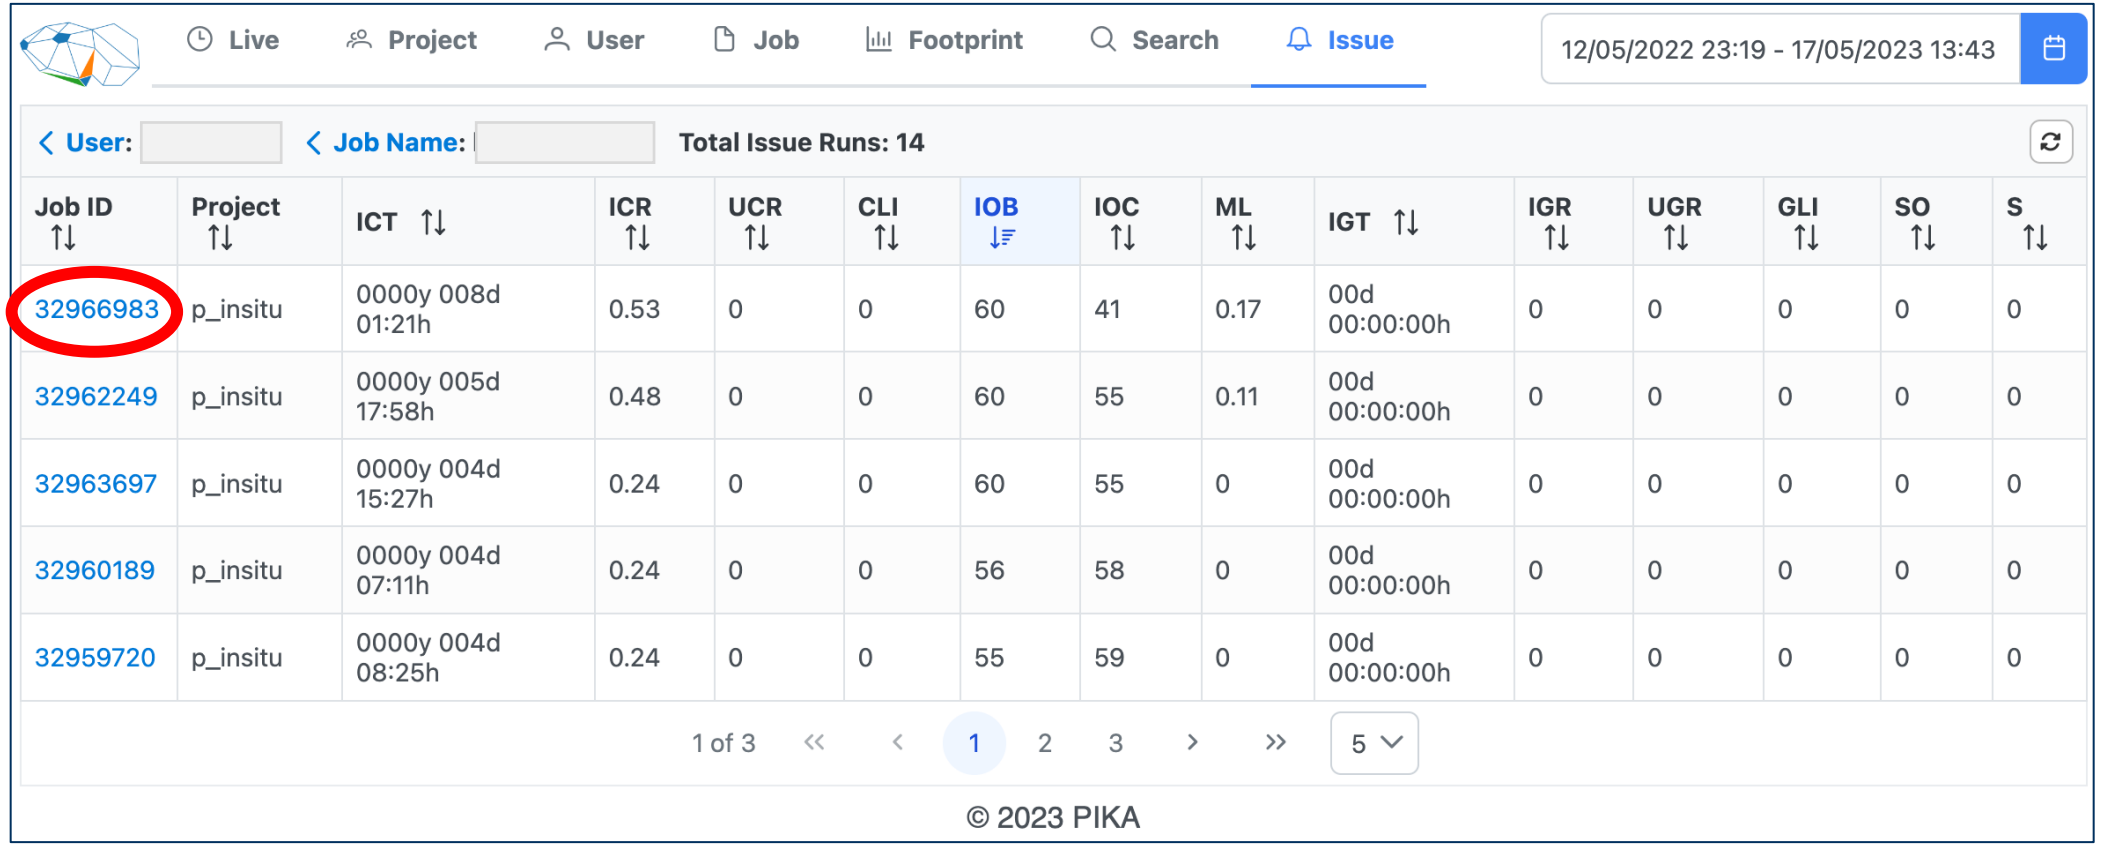

Job ID	Project	ICT	ICR	UCR	CLI	IOB	IOC	ML	IGT	IGR	UGR	GLI	SO	S
32966983	p_insitu	0000y 008d 01:21h	0.53	0	0	60	41	0.17	00d 00:00:00h	0	0	0	0	0
32962249	p_insitu	0000y 005d 17:58h	0.48	0	0	60	55	0.11	00d 00:00:00h	0	0	0	0	0
32963697	p_insitu	0000y 004d 15:27h	0.24	0	0	60	55	0	00d 00:00:00h	0	0	0	0	0
32960189	p_insitu	0000y 004d 07:11h	0.24	0	0	56	58	0	00d 00:00:00h	0	0	0	0	0
32959720	p_insitu	0000y 004d 08:25h	0.24	0	0	55	59	0	00d 00:00:00h	0	0	0	0	0

PIKA Issue Analysis – Metadata & Timelines

Pro	Start	End	State	#Nodes	#Cores	Exclusive	Walltime	Pending	Duration	Core Time	Used	Partition
p...	09/02/23 22:27:43	10/02/23 00:58:26	com...	6	144	0	00d 04:00	00d 01:57	00d 02:30:43h	0000y 015d 01:43h	62.8%	haswell...

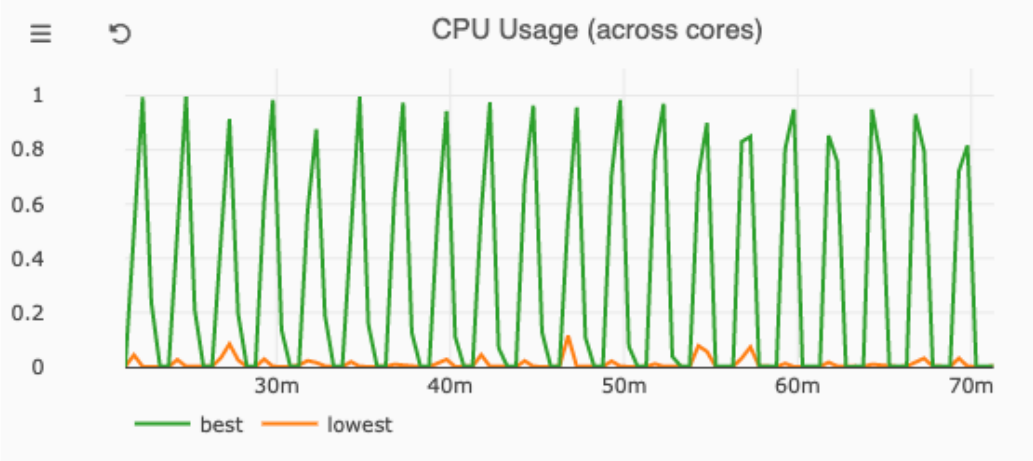

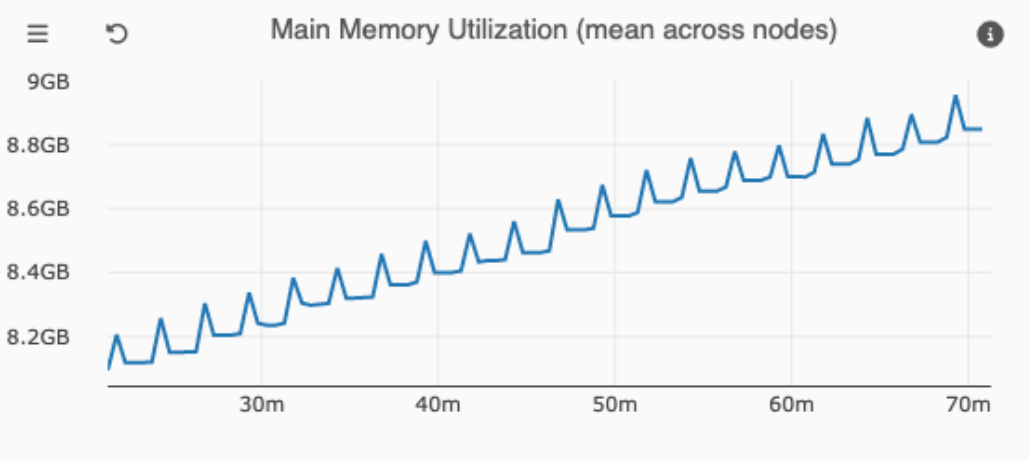

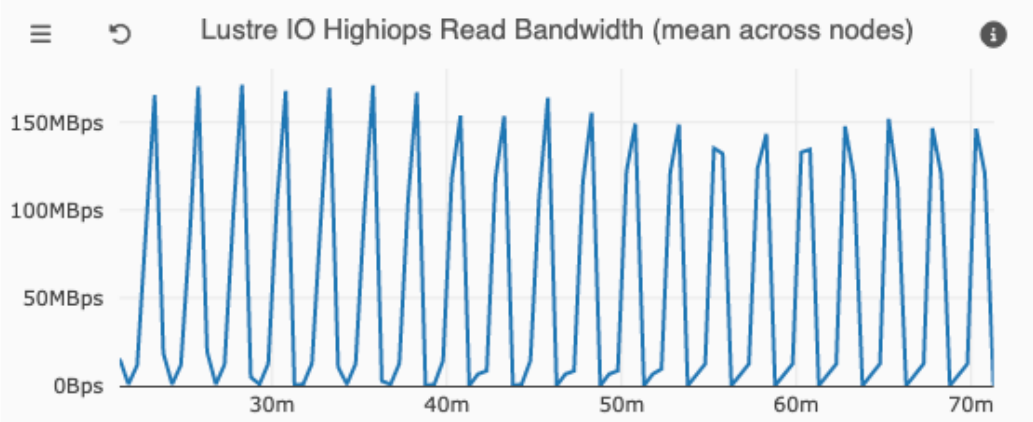

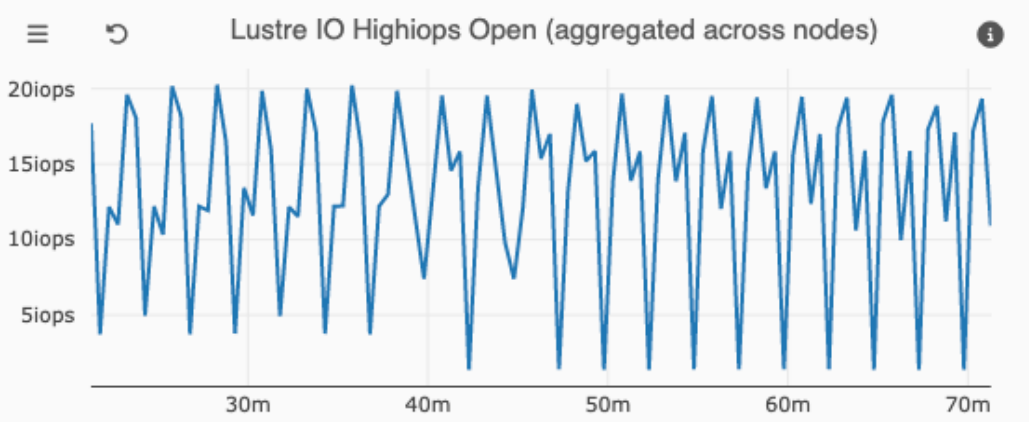

Conclusion

PIKA is a hardware performance monitoring stack in order to identify potentially inefficient jobs.

— Easy access to processed performance data for individual jobs, projects, and users

— Performance overview facilitates quick identification of pathological or suboptimal jobs

— Awareness for performance analysis among new HPC users

R. Dietrich, F. Winkler, A. Knüpfer and W. Nagel, “PIKA: Center-Wide and Job-Aware Cluster Monitoring,” 2020 IEEE International Conference on Cluster Computing (CLUSTER), Kobe, Japan, 2020, pp. 424-432.

https://gitlab.hrz.tu-chemnitz.de/pika

Hands-on Exercise

PIKA Hands-on (1)

Open the PIKA web interface and log in with the PIKA demo user:

https://pika.zih.tu-dresden.de

User / Password: TBA

PIKA Hands-on (2)

Connect to Barnard cluster via ssh:

```
% cp ${VIHPS_ROOT}/pika/run_nhr_pika_example.sh .

% sbatch run_nhr_pika_example.sh
  Submitted batch job 19349695
```

Open the PIKA web interface and log in with your own login credentials:

https://pika.zih.tu-dresden.de

- Click on the Live tab
- Update the timelines by clicking on the refresh button (data is updated about every minute)
- Examine the benchmark output of triad_avx and check whether FLOPS and Memory Bandwidth are displayed correctly in PIKA
- Do the IO bandwidths match the instructions in the job script?