

Frank Winkler (frank.winkler@tu-dresden.de)

# PIKA: Continuous Job Performance Monitoring



NHR Monitoring Workshop 2022  
07/18/2022 to 07/19/2022, Erlangen

# Motivation

- Performance is crucial for HPC
- Therefore measurement, analysis and validation also important
- Numerous established tools available
- Problem solved?
- ... for those who are aware of the problem
- Active preparation needed
- Not continuous for all jobs of all users

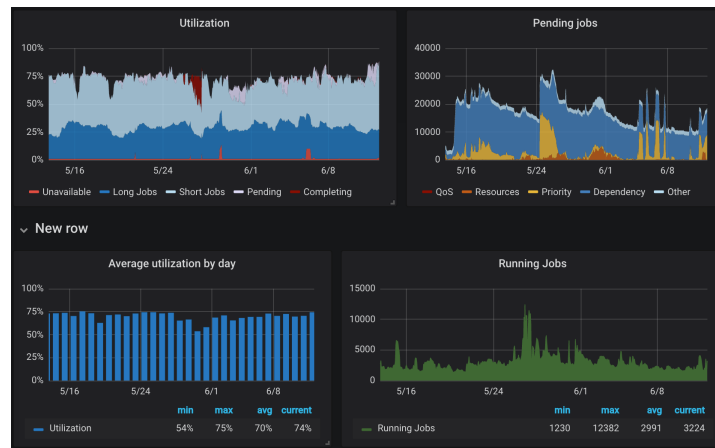


Frank Winkler  
NHR Monitoring Workshop 2022, Erlangen



# Motivation

- Several commercial, community and site-specific monitoring solutions already exist
- Most systems provide a system view
- A few are job-aware and track performance data like Flop/s (only applies for site-specific solutions)
- Existing site-specific monitoring systems are complex and difficult to adapt to other sites




Visualization of cluster utilization in Grafana

# PIKA: Continuous Job Performance Monitoring

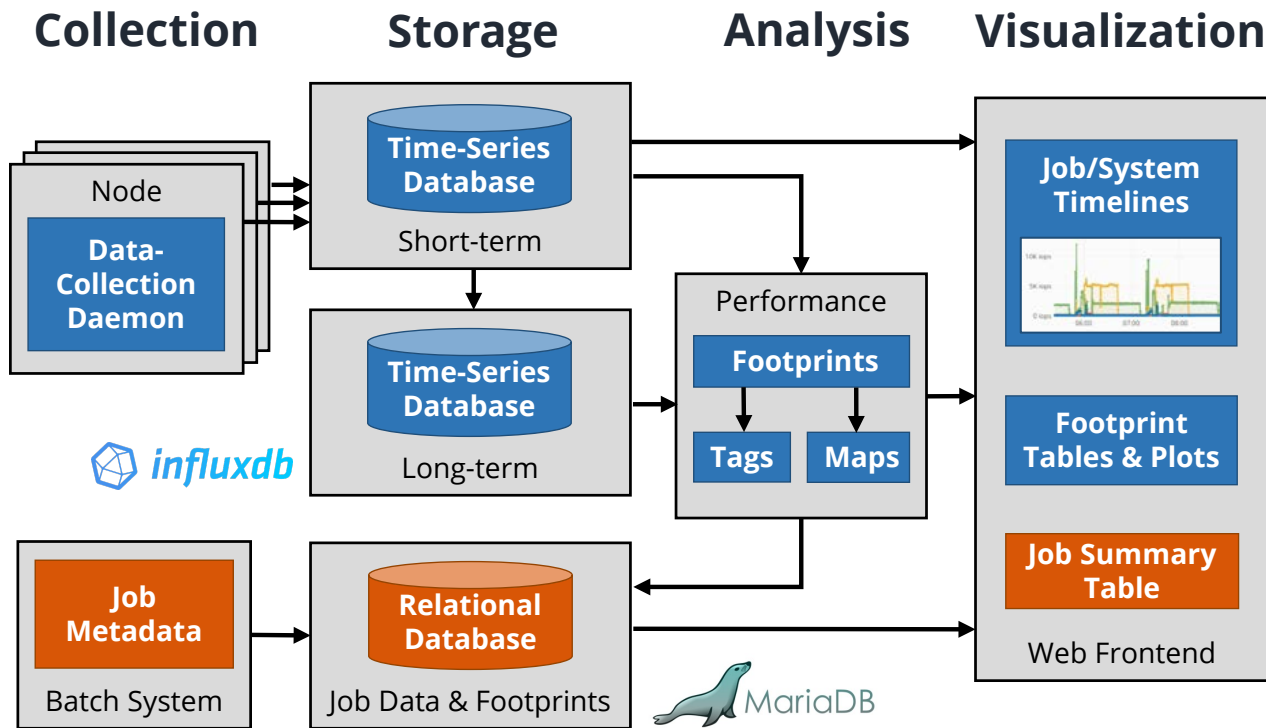
- Non-intrusive **data acquisition** on all cluster nodes
- Continuous **data collection**
- Web frontend for live and post-mortem **visualization**
- Detection of pathological jobs
- Automatic **job analysis and classification**
- Long-term **data storage**

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





# PIKA Architecture Overview



# PIKA Metadata Collection

**Slurm PrEp Plugin** to capture the following 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
  - GPUs per node
  - Job script



# PIKA Metadata Storage

Job metadata is stored in MariaDB (relational database) on a separate service node.

- Several indexes to improve query performance
- Statistics (September 2018 until today)
  - About 33 million jobs
  - About 12 GB of disk memory (six indices)
  - Job-specific query takes about 1ms






# PIKA Runtime Data Collection

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

## Collection daemon **collected**

- One collector/plugin for each metric source
- CPU counters are collected with LIKWID



# PIKA Runtime Data Storage with InfluxDB

Each metric value is tagged with timestamp, hostname as well as CPU or GPU

- Short-term storage
  - Retention policy: data older than 4 weeks is deleted
  - Shard length: 7 days
  - Complete shards are backed up and imported in long-term database
  - Requires SSD to handle more than 2K nodes
  - About 36GB of disk memory
- Long-term storage
  - No retention policy
  - Shard length: 7 days
  - Ceph filesystem
  - About 9GB per week → 0.5TB per year



# PIKA Job Visualization – Tables

Live

Project

User

Job

Footprint

Search

03/12/2019 01:00 - 08/12/2019 01:00

admin

Total Projects: 140

Project	Number of Jobs	Max Nodes	Max Cores	Overall Core Time	Max Pending	Overall Runtime	#Footprints
> pa	1	1	8	0000y 000d 01:07h	00d 00:00:01h	0000y 000d 00:08h	1
> pb	4	1	24	0000y 000d 05:12h	00d 00:14:10h	0000y 000d 00:22h	4
> pc	93	1	8	0000y 017d 17:07h	00d 02:27:01h	0000y 002d 05:08h	92
> pd	1	1	20	0000y 000d 01:46h	00d 00:07:01h	0000y 000d 00:05h	1
> pe	11		24	0000y 001d 12:06h	00d 00:03:12h	0000y 000d 01:30h	11

1

2

3

4

5

5

Jobs of 140 projects have been recorded for the selected time interval (top right).

# PIKA Job Visualization – Tables

Live

Project

User

Job

Footprint

Search

03/12/2019 01:00 - 08/12/2019 01:00

admin

Total Projects: 140

Project	Number of Jobs	Max Nodes	Max Cores	Overall Core Time	Max Pending	Overall Runtime	#Footprints
> pz	62	256	3072	0007y 194d 13:57h	00d 07:12:52h	0000y 002d 10:59h	62
> pb	3	80	1920	0000y 034d 04:00h	00d 03:41:07h	0000y 000d 08:49h	3
> pg	20	128	1536	0000y 035d 16:22h	00d 17:53:27h	0000y 000d 02:19h	20
> ps		48	1152	0000y 226d 07:41h	00d 01:52:03h	0000y 003d 10:48h	6
> pe		30	720	0004y 065d 20:50h	00d 07:21:18h	0000y 003d 17:25h	15

Unfolding

1

2

5

5

Unfolding

Get project with the highest number of cores.

# PIKA Job Visualization – Tables

Live

Project

User

Job

Footprint

Search

03/12/2019 01:00 - 08/12/2019 01:00

admin

< Projects

pz

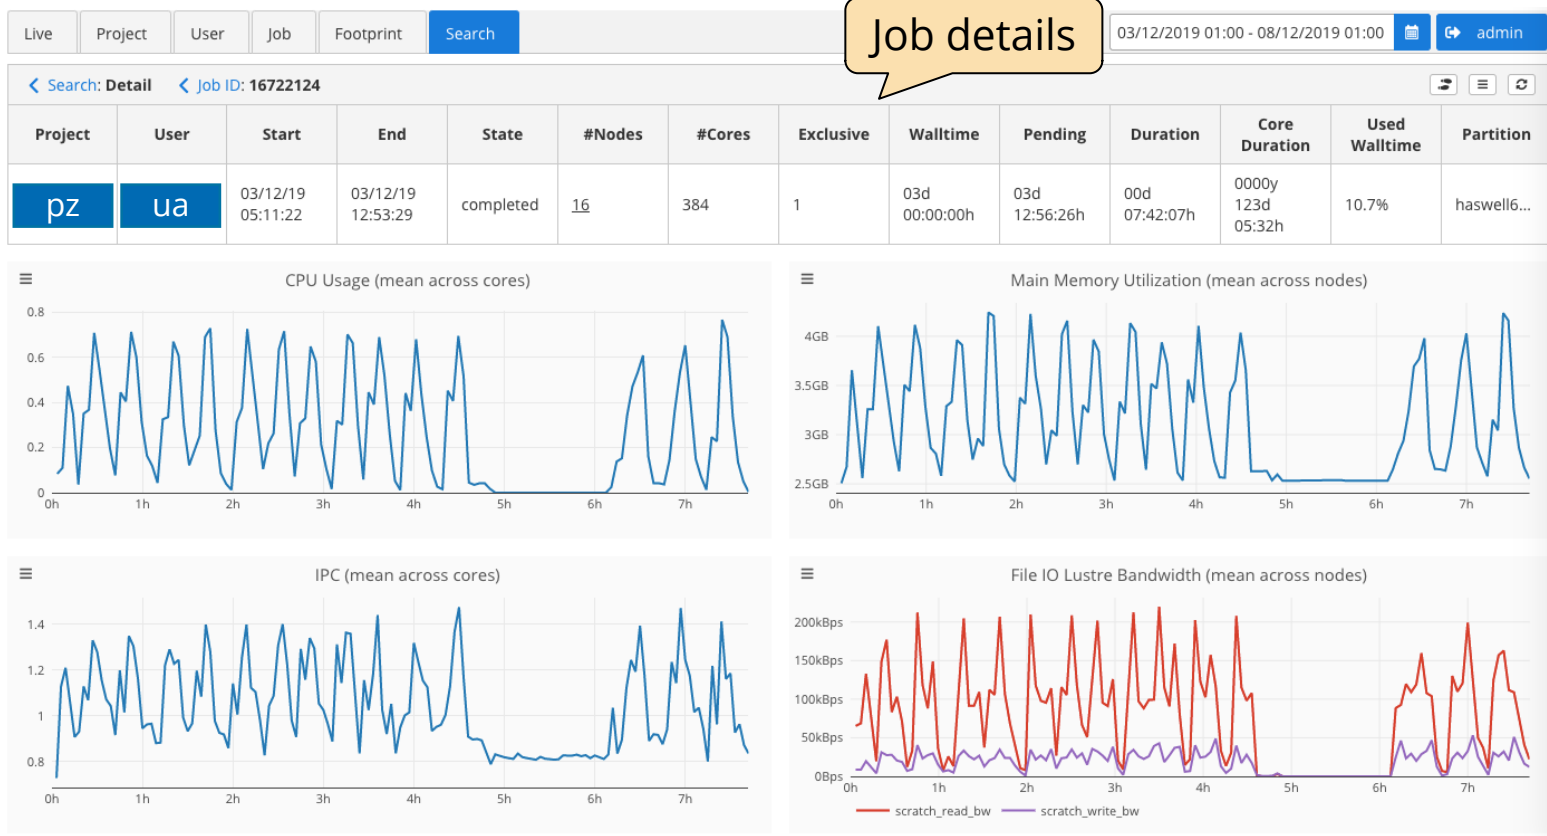
Total Users: 4

User ▾	Number of Jobs ▾	Max Nodes ▾	Max Cores ▾	Overall Core Time ▾	Max Pending ▾	Overall Runtime ▾	#Footprints ▾
> ua	9	45	1080	0000y 361d 19:21h	00d 00:00:03h	0000y 000d 12:54h	9
> ua	4	1	24	0000y 000d 09:52h	00d 00:00:02h	0000y 000d 02:17h	4
> ua	47	256	3072	0006y 194d 09:06h	00d 07:12:52h	0000y 001d 19:34h	47
> ua	2	40	960	0000y 002d 23:36h	00d 00:00:01h	0000y 000d 00:12h	2

</

Project "pz" has  
4 users and a  
total of 62 jobs.

# PIKA Job Visualization – Timelines



# PIKA Job Visualization – Timelines



# PIKA Job Visualization – Timelines





# PIKA Post Processing

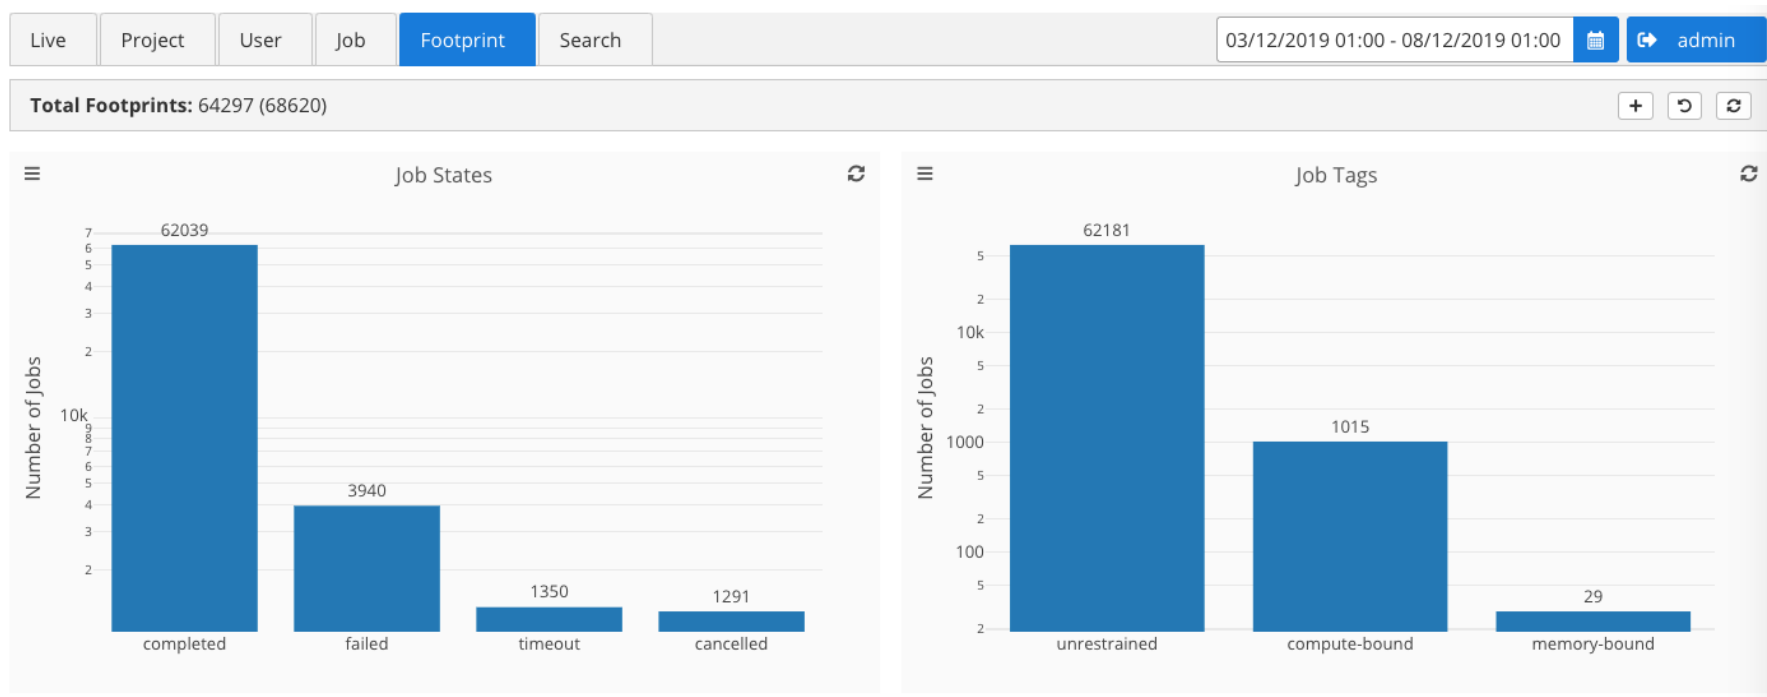
## Job characterization via tagging

- **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

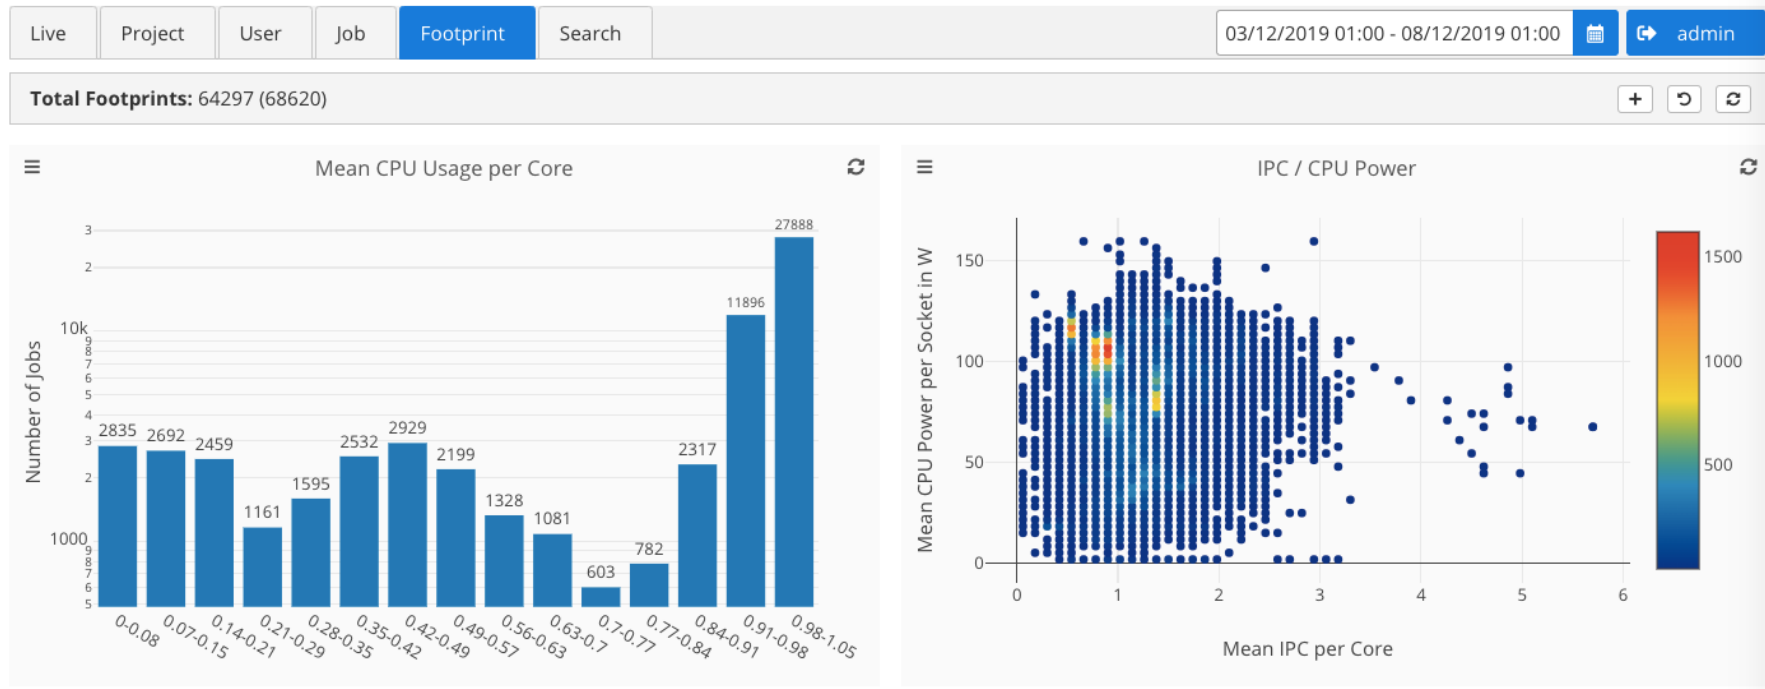
Tag Name	Formula and Threshold
unrestrained	-
memory-bound	$\frac{\text{memory bandwidth (measured)}}{\text{memory bandwidth (maximum)}} > 80\%$
compute-bound	$\frac{\text{FLOP/s (measured)}}{\text{FLOP/s (maximum)}} > 70\%$ or $\frac{\text{IPC (measured)}}{\text{IPC (optimal)}} > 60\%$
GPU-bound	GPU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

**Automatic detection of job performance issues** based on metadata (requested resources) and raw runtime data analysis (still in development)

# PIKA Job Visualization – Footprints



# PIKA Job Visualization – Footprints



# PIKA Filter Mask

Live Project User Job Footprint **Search** Validation 01/09/2020 20:38 - 06/09/2021 11:02 admin

Job ID Enter ID Select Search Option **Exclusive** Ignore Selected Time Submit

<b>Project</b> Select Project	<b>Job Name</b> Enter Job Name	<b>Number of Nodes</b> 2 Enter Max
<b>User</b> Select User	<b>Node Name</b> Enter Node Name	<b>Number of Cores</b> Enter Min Enter Max
<b>Job Status</b> Select Job Status	<b>Job Tag</b> Select Job Properties	<b>Time Limit</b> Enter Min Enter Max m
<b>Partition</b> Select Partition	<b>Footprint</b> Mean CPU Usage per Core	<b>Pending Time</b> Enter Min Enter Max m
	<b>Mean CPU Usage per Core</b> Enter Min 0.01	<b>Duration</b> 4 Enter Max h
		<b>Core Duration</b> Enter Min Enter Max h



Find jobs with inappropriate resource allocation.

# PIKA Filter Mask

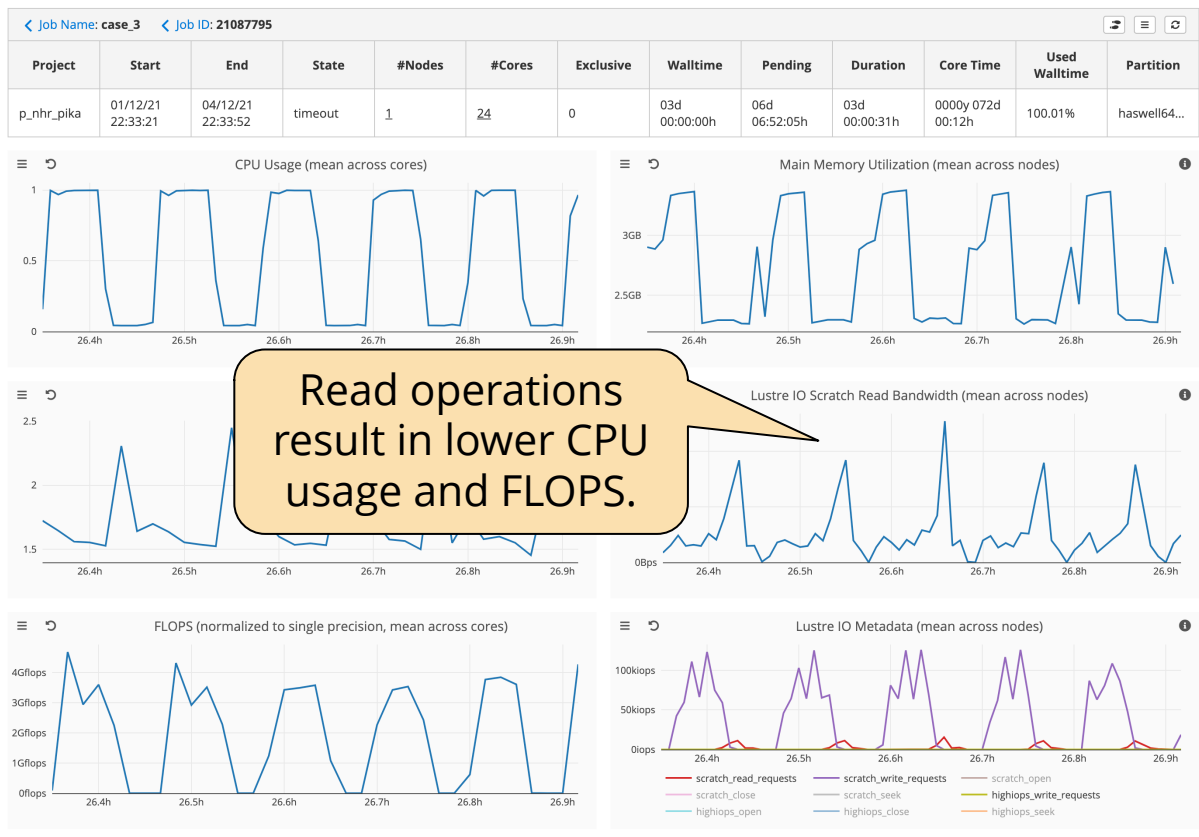
Live	Project	User	Job	Footprint	Search	Validation	01/09/2020 20:38 - 06/09/2021 11:02		admin
------	---------	------	-----	-----------	--------	------------	-------------------------------------	--	-------

< Search: Detail < Job ID: 4792772

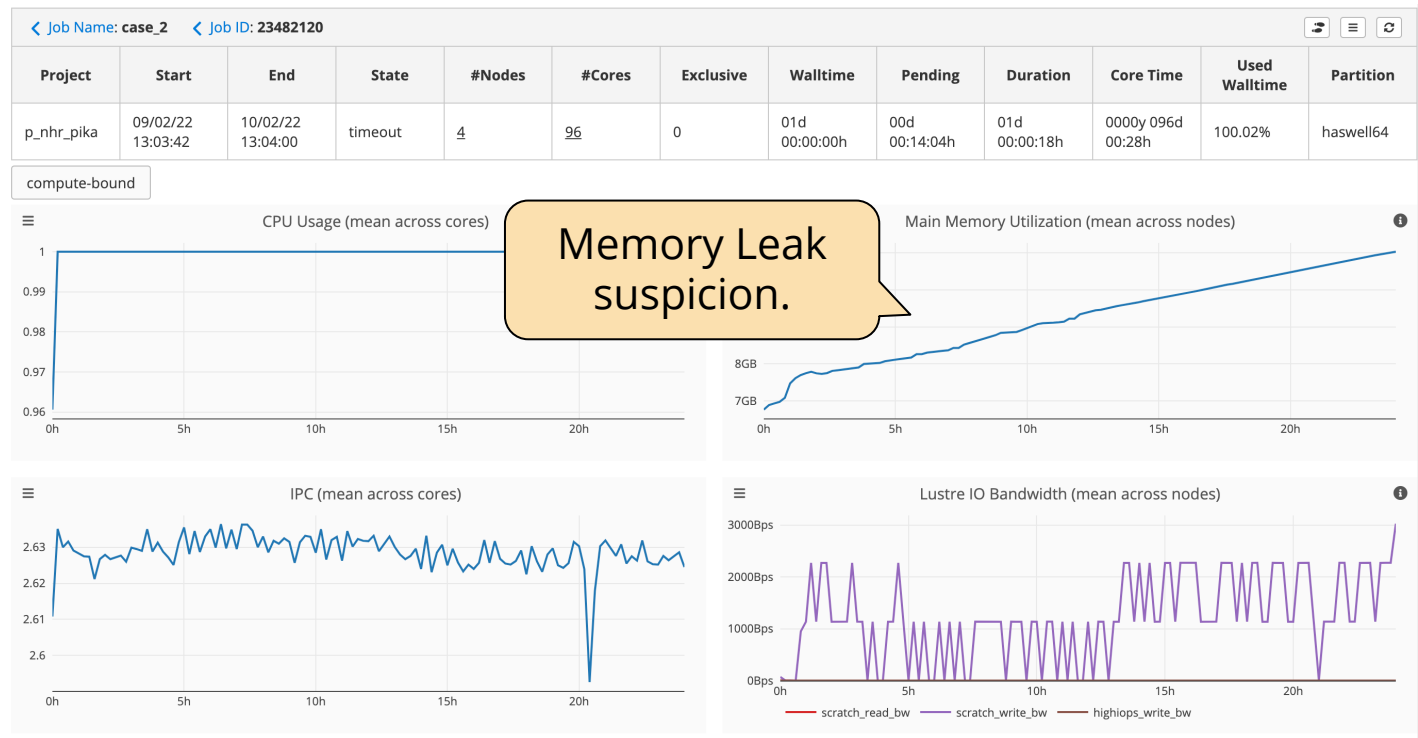
Project	User	Start	End	State	#Nodes	#Cores	Exclusive	Walltime	Pending	Duration	Core Duration	Used Walltime	Partition
		17/09/20 16:32:45	20/09/20 10:41:35	complet...	4	512	1	03d 00:00:00h	00d 22:46:58h	02d 18:08:50h	0007y 267d 06:45h	91.87%	romeo



# PIKA Job Performance Issues



# PIKA Job Performance Issues



# PIKA Demo

<https://hpcmon.zih.tu-dresden.de/nhr>

User / Password: TBA



# Conclusion


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

- **Metric Data Collector:** An extension of the collection daemon collectd to record metric data on each compute node
- **Job Metadata Collector:** Centralized capture of job metadata for both exclusive and shared jobs using a Slurm PrEp Plugin
- **Frontend:** Powerful interactive GUI with top-down approach
- **Post-processing:** Python analysis scripts for job tagging and automatic detection of job performance issues

PIKA uses **InfluxDB** to store metric data and **MariaDB** to store job metadata. Most of the above components are available as **docker containers** and **RPM packages**.

# Conclusion

- Continuous job performance monitoring is important
- Easy access to prepared performance data for projects, users and jobs
- Performance overview for admins, easy scan for pathological / suboptimal jobs
- Allows targeted consulting offers



Paper: 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.

PIKA available as open source software at  
 <https://gitlab.hrz.tu-chemnitz.de/pika>

