

# Verifying GPU Performance with the RWTH Job Monitoring

aiXcelerate 2024

## CLAIX-2023



# Agenda

---

- Accessing the Monitoring System
  - The Webportal
  - Joblist Overview
  - Job Details
  - Caveats and Gotchas
- Detecting Misconfiguration
  - Wrong Core Count
  - Single GPU Restriction
  - Underutilization of Resources
- Call to Action

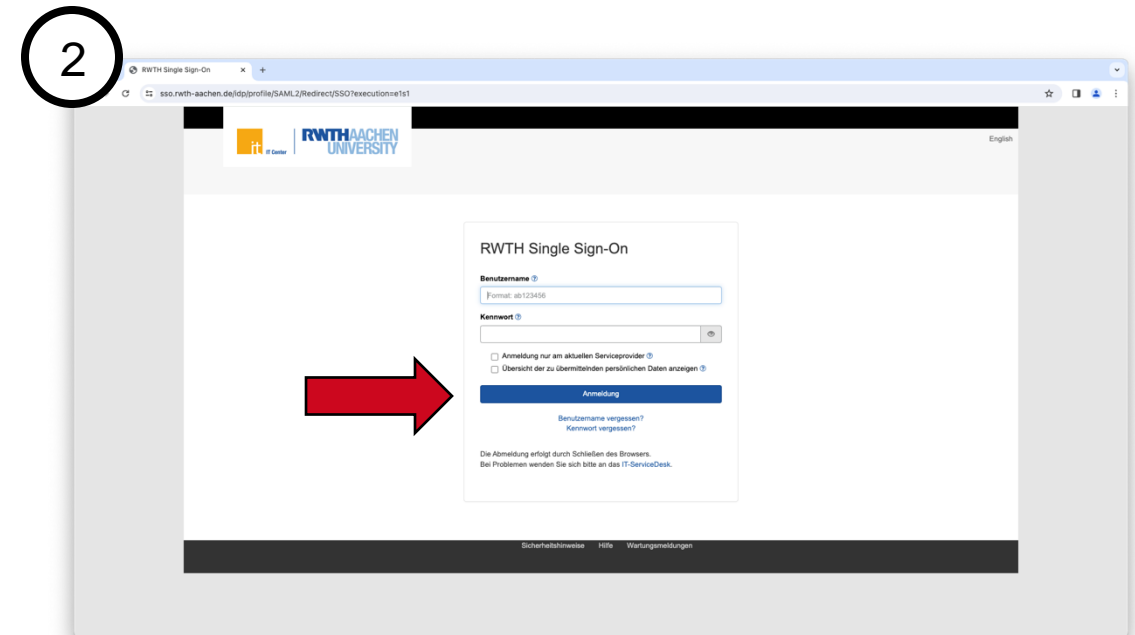
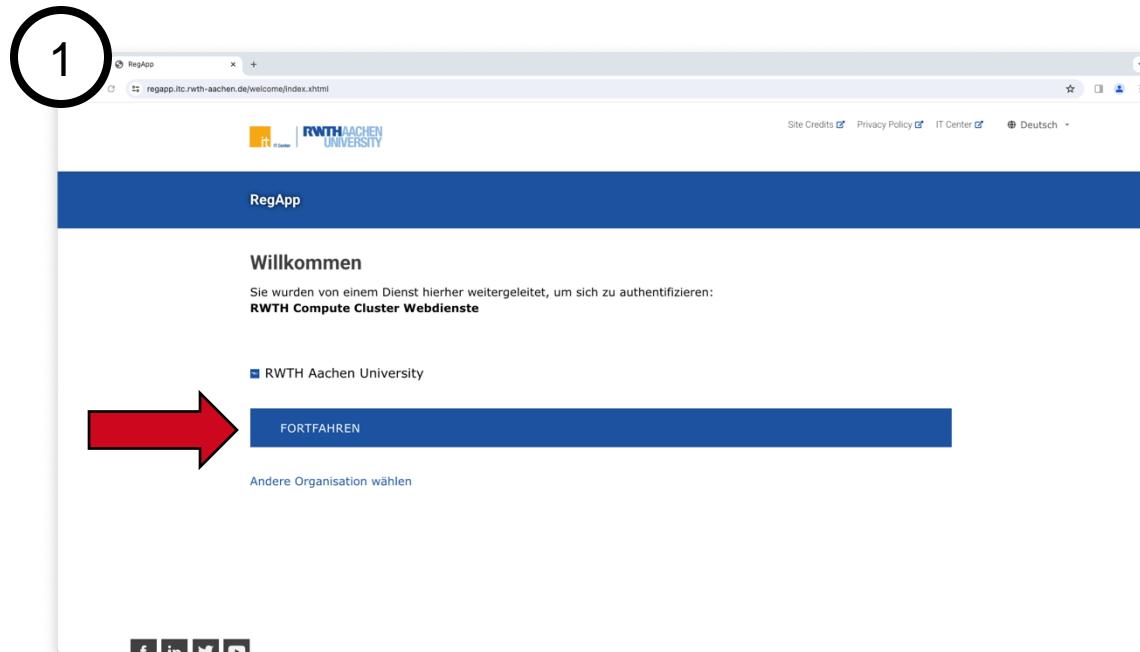
# Agenda

---

- **Accessing the Monitoring System**
  - **The Webportal**
  - **Joblist Overview**
  - **Job Details**
  - **Caveats and Gotchas**
- Detecting Misconfiguration
  - Wrong Core Count
  - Single GPU Restriction
  - Underutilization of Resources
- Call to Action

# Accessing the Monitoring System – The Webportal

- Online Documentation @ [help.itc](https://help.itc.rwth-aachen.de)
- Visit Grafana @ [perfmon.hpc.itc.rwth-aachen.de](https://perfmon.hpc.itc.rwth-aachen.de)



⚠ Single Sign-On Credentials (not HPC-Account)

# Accessing the Monitoring System – Joblist Overview

- History of Finished & Running Jobs

3

Filter by Cluster Partition

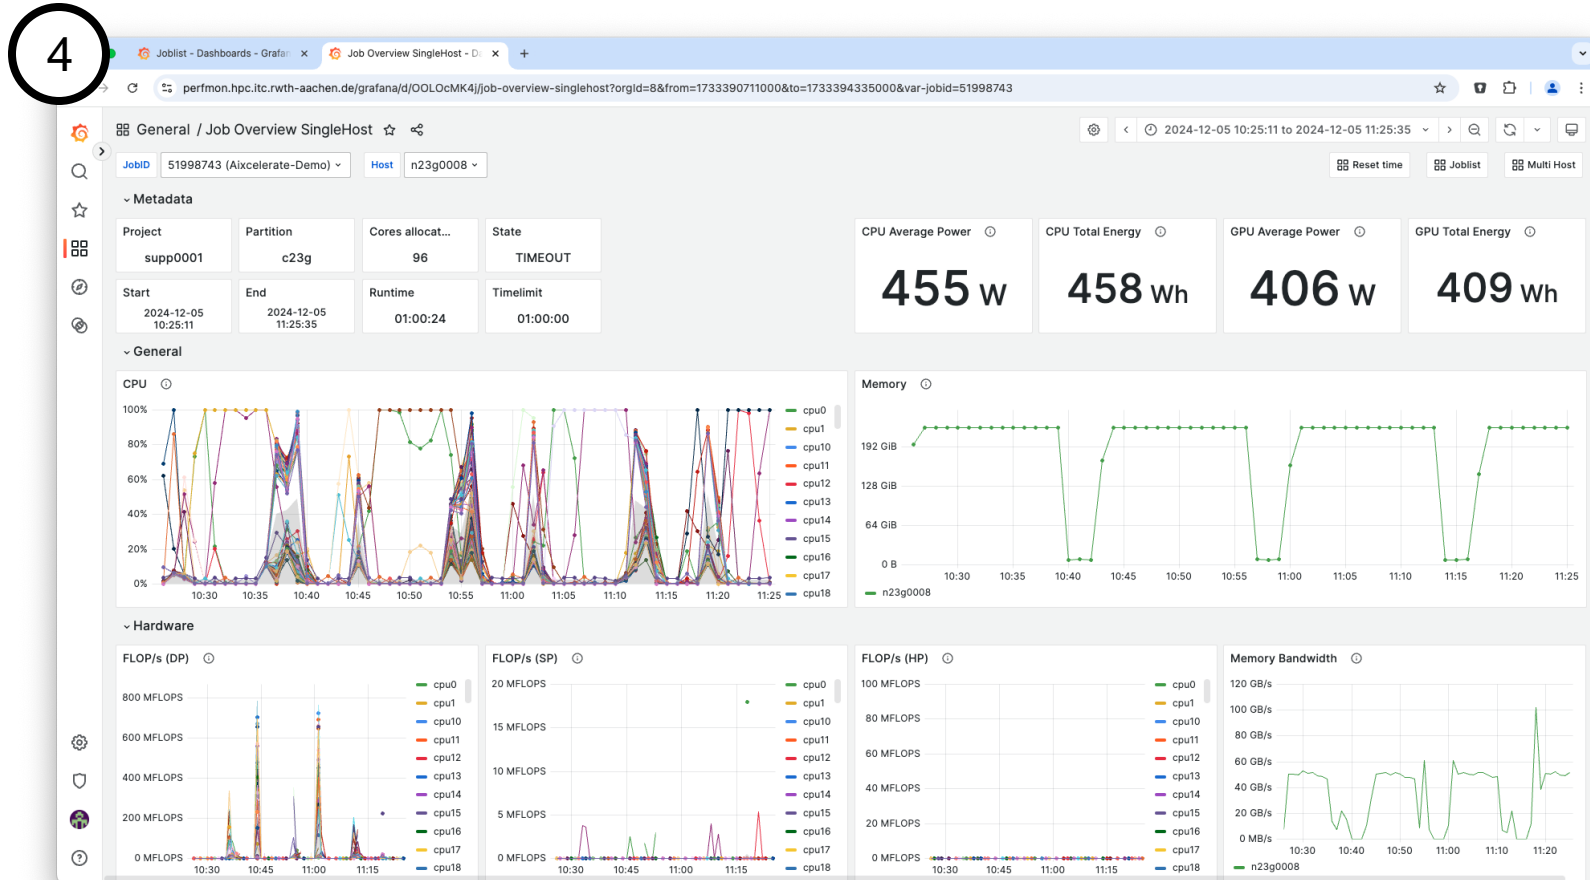
Sort by Runtime

Increase Time Range

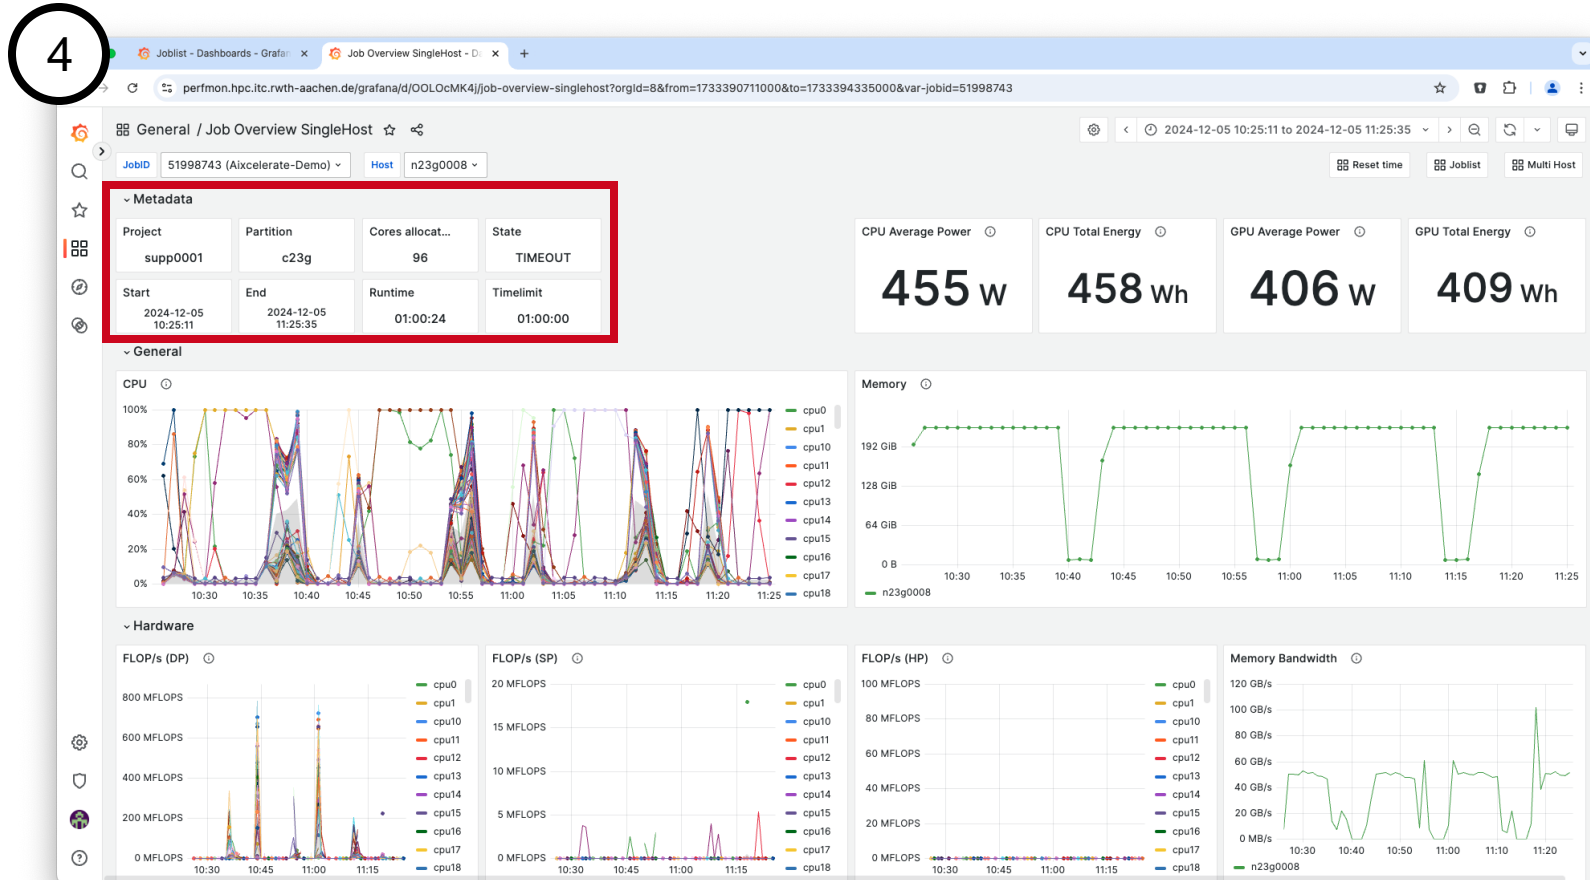
Access Job Details

JobID	JobName	Partition	Start	End	Runtime	JobState	NumNodes	Nodelist
52052438	Aixcelerate-Demo	c23g	2024-12-06 14:38:43	2024-12-06 15:38:43	1 hour	RUNNING	1	r23g0002
52052437	Aixcelerate-Demo	c23mm	2024-12-06 14:34:07	2024-12-06 15:34:07	1 hour	RUNNING	1	n23m0218
52052436	Aixcelerate-Demo	c23mm	2024-12-06 14:34:07	2024-12-06 15:34:07	1 hour	RUNNING	1	n23m0160
52052435	Aixcelerate-Demo	c23g	2024-12-06 14:38:11	2024-12-06 15:38:11	1 hour	RUNNING	1	n23g0015
51998743	Aixcelerate-Demo	c23g	2024-12-05 10:25:11	2024-12-05 11:25:35	1.01 hour	TIMEOUT	1	n23g0008
51998104	Aixcelerate-Demo	c23g	2024-12-05 10:19:00	2024-12-05 10:54:36	35.6 min	COMPLETED	1	r23g0001
51997558	Aixcelerate-Demo	c23g	2024-12-05 10:10:16	2024-12-05 10:13:29	3.22 min	COMPLETED	1	n23g0004
51997287	Aixcelerate-Demo	c23g	2024-12-05 10:06:41	2024-12-05 10:23:13	16.5 min	COMPLETED	1	n23g0008
51997206	Aixcelerate-Demo	c23g	2024-12-05 10:05:39	2024-12-05 10:06:45	1.10 min	CANCELLED	1	n23g0004
51996913	Aixcelerate-Demo	c23g	2024-12-05 10:01:58	2024-12-05 10:05:13	3.25 min	COMPLETED	1	n23g0004
51996783	Aixcelerate-Demo	c23g	2024-12-05 09:59:53	2024-12-05 10:01:38	1.75 min	CANCELLED	1	n23g0004
51996310	Aixcelerate-Demo	c23g	2024-12-05 09:55:14	2024-12-05 09:58:22	3.13 min	COMPLETED	1	n23g0004

# Accessing the Monitoring System – Job Details

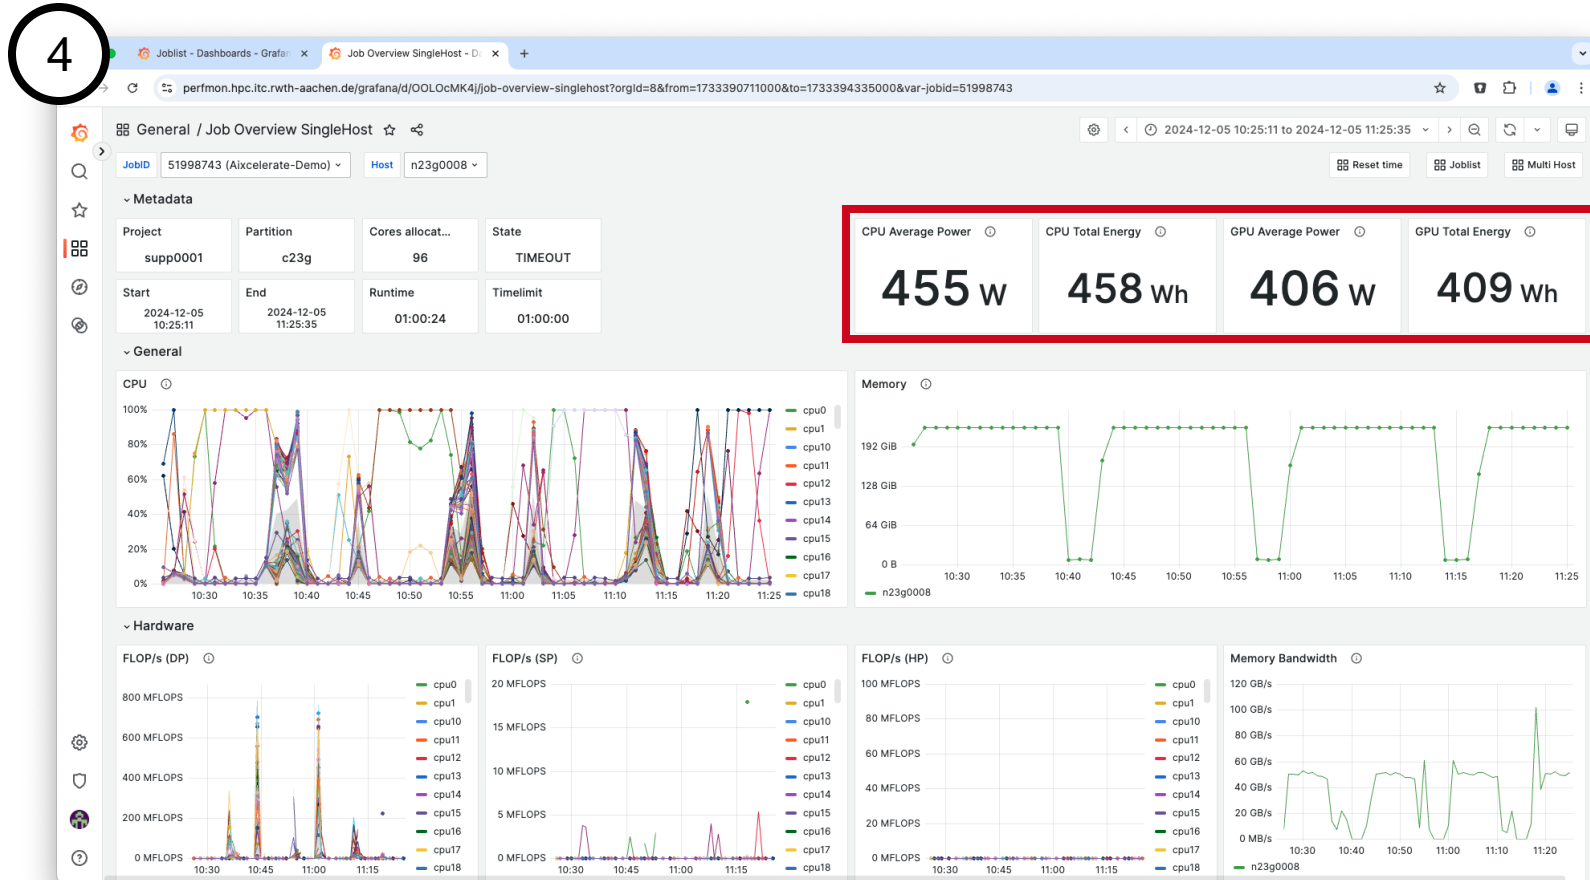


# Accessing the Monitoring System – Job Details





# Accessing the Monitoring System – Job Details



# Accessing the Monitoring System – Job Details

**4**

**Select Different Host**

**Switch Single-/Multi-Host View**

**Select Subset with CTRL/CMD+Click**

**Scroll Down for More Details**

Project	Partition	Cores allocat...	State
supp0001	c23g	96	TIMEOUT

Start	End	Runtime	Timelimit
2024-12-05 10:25:11	2024-12-05 11:25:35	01:00:24	01:00:00

**CPU Average Power** 455 w

**CPU Total Energy** 458 wh

**GPU Average Power** 406 w

**GPU Total Energy** 409 wh

**CPU**

**Memory**

**FLOPs (DP)**

**FLOPs (SP)**

**FLOPs (HP)**

**Memory Bandwidth**

# Accessing the Monitoring System – Caveats and Gotchas

---

- Sampling Frequency: 1 Sample / Minute
  - Almost no influence on application runtime
  - ⚠ Short application stages can be missed
- Metric-specific Sampling Quality:
  - Sampling length of 1 second for FLOP/s, power, memory bandwidth
    - ⚠ This data is missing with `#SBATCH --hwctr=OPTION` with `OPTION ∈ { vtune, vtuneperf, likwid, papi }`
  - Continuous measurement for fabric, Lustre, NFS
  - Collection of current value for CPU, memory, GPU
- Shared Node Usage
  - ⚠ Data may be influenced by behaviour of other jobs
  - Force exclusive node availability with `#SBATCH --exclusive`
- Power & Energy Measurements
  - For CPUs measured with Intel RAPL → for details see [here](#) and [here](#)
  - For GPUs measured with nvidia-smi → for details see [here](#) and [here](#)
  - ⚠ Data quality and availability limited by available tools

# Agenda

---

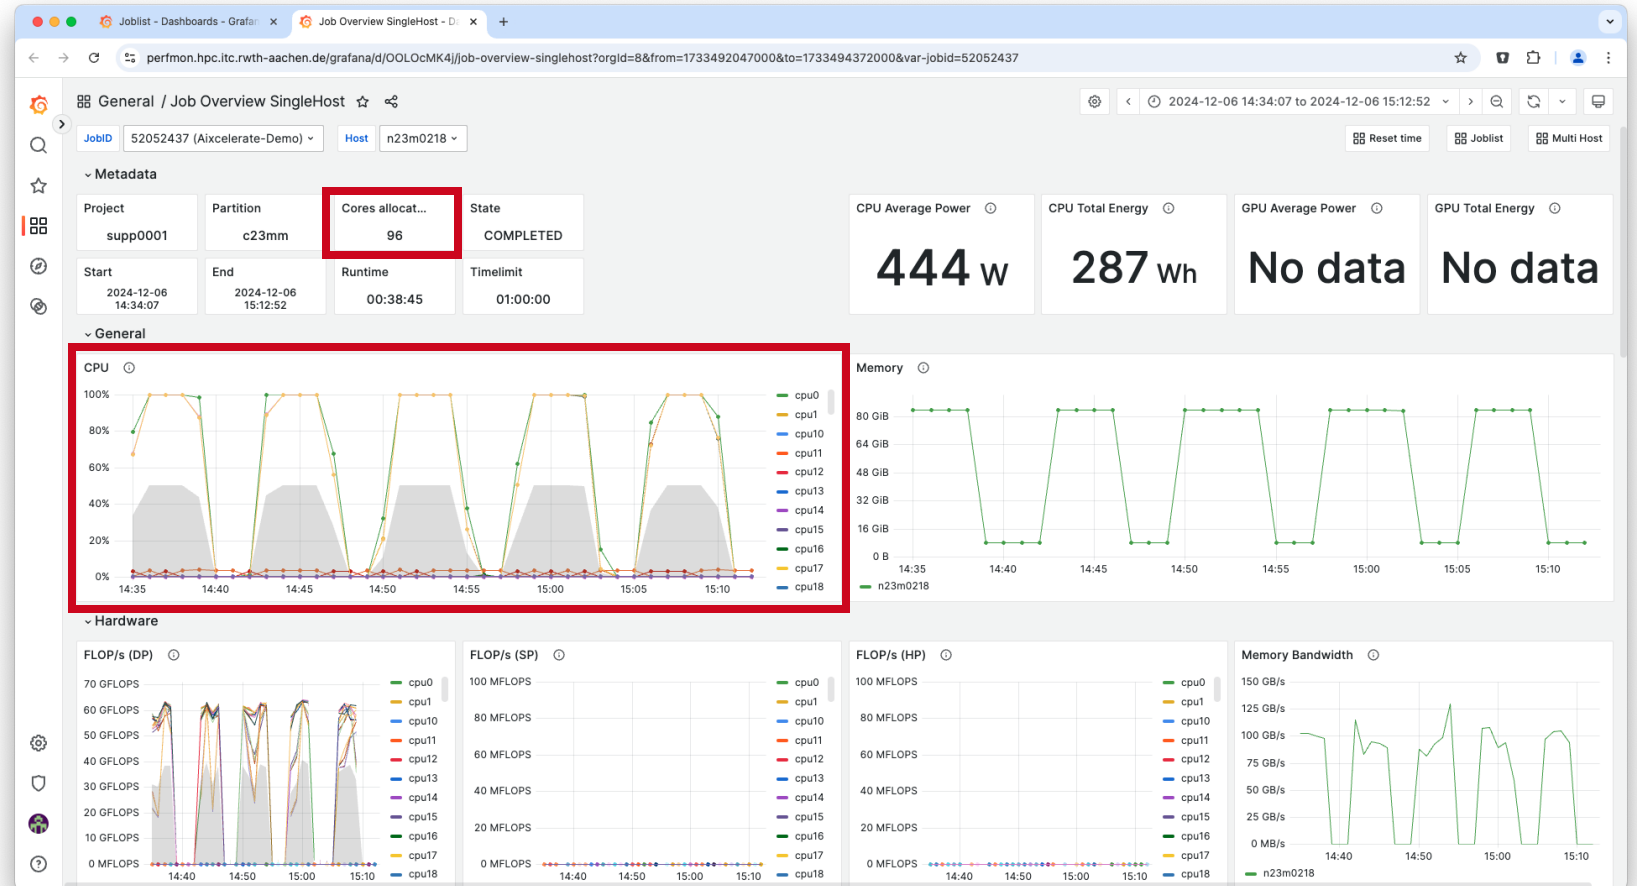
- Accessing the Monitoring System
  - The Webportal
  - Joblist Overview
  - Job Details
  - Caveats and Gotchas
- **Detecting Misconfiguration**
  - **Wrong Core Count**
  - **Single GPU Restriction**
  - **Underutilization of Resources**
- Call to Action

# Detecting Misconfiguration – Wrong Core Count

✓ 96 Cores Allocated

⚠ Only 50% Mean CPU Usage

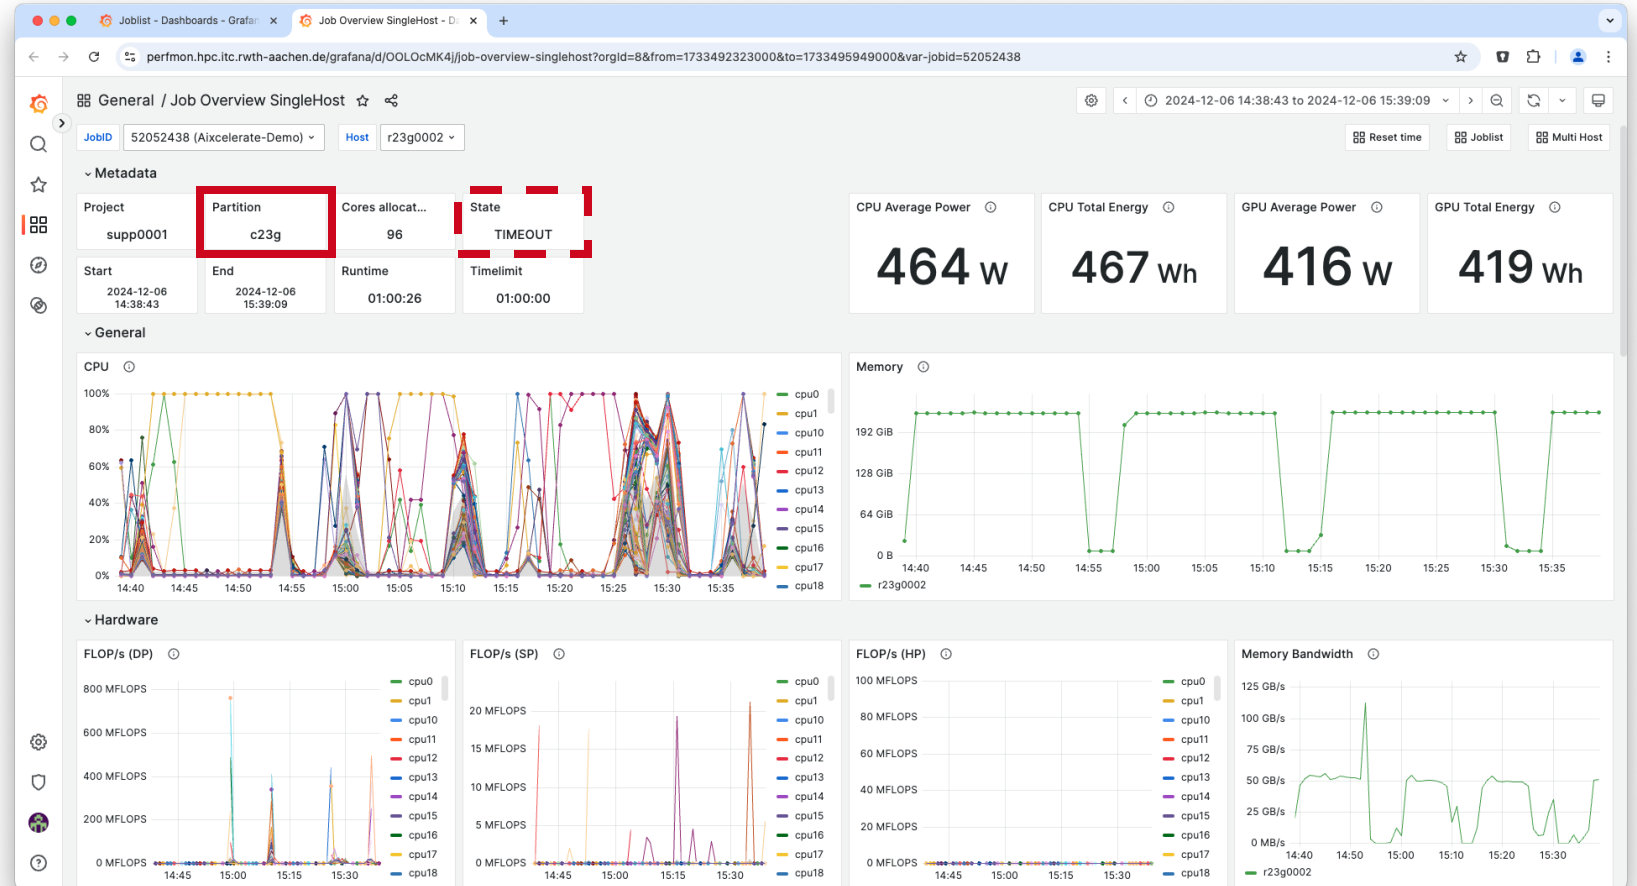
- Takeaways:
  - CLAIX-2018: 48 Cores per Node
  - CLAIX-2023: 96 Cores per Node
- Remove hard-coded values for
  - `OMP_NUM_THREADS`
  - `srun -n / mpirun -n / mpiexec -n`



# Detecting Misconfiguration – Single GPU Restriction

✓ GPU Partition

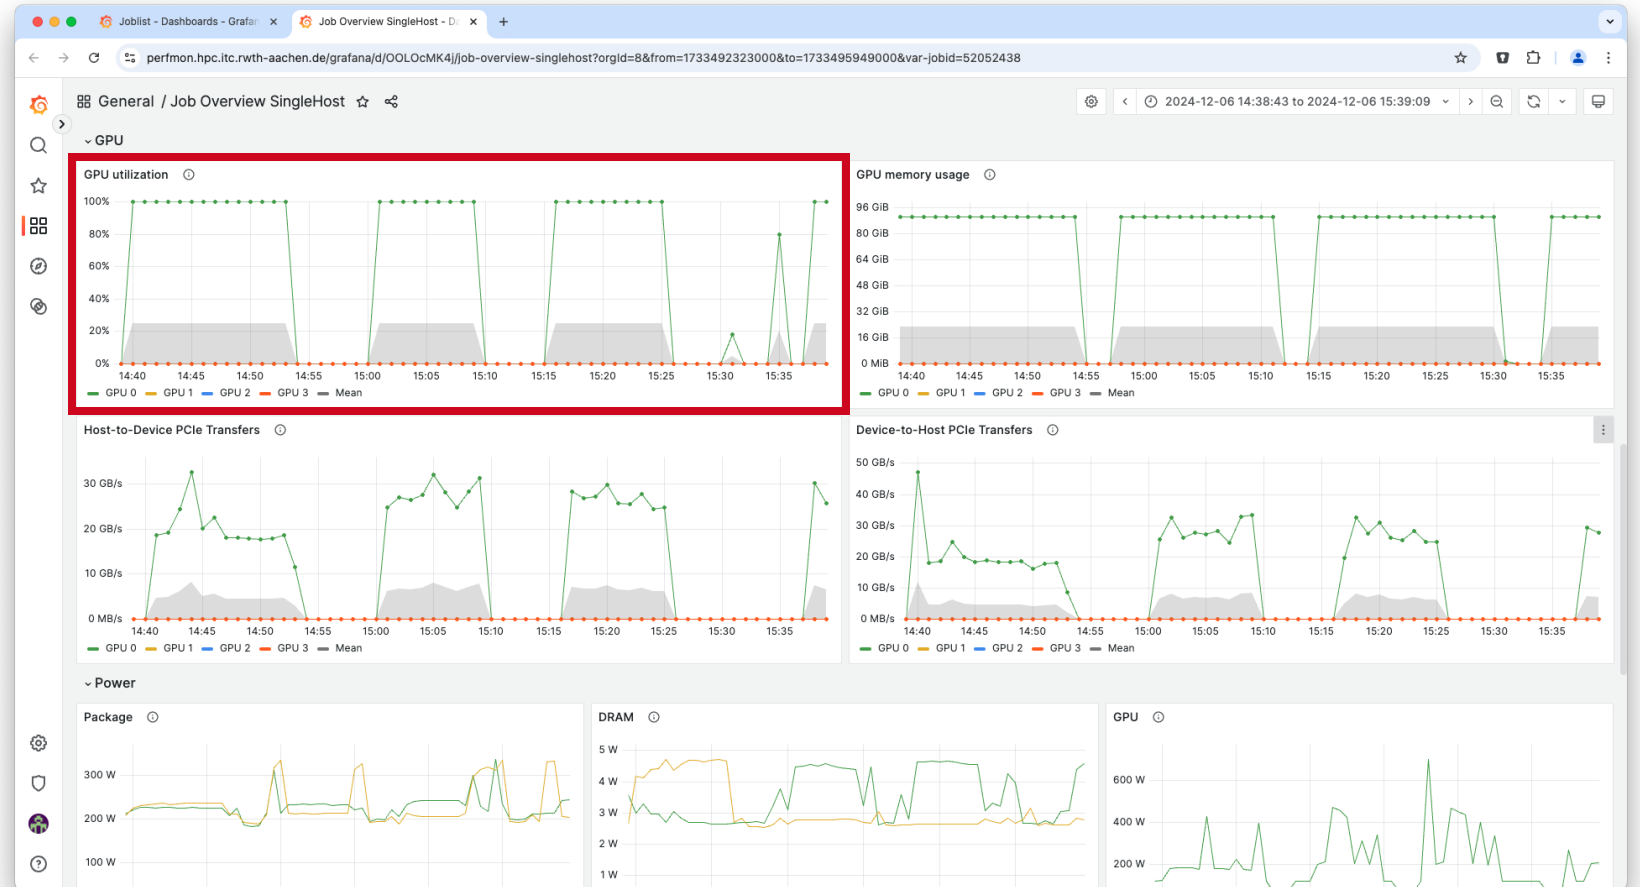
⚠ Job Timeout



# Detecting Misconfiguration – Single GPU Restriction

⚠️ Only 25% Mean GPU Usage


- Takeaways:
  - Check GPU-ID restrictions
  - Might differ from interactive usage
- Consult documentation for
  - `CUDA_VISIBLE_DEVICES`
  - Check allocated CPUs / GPUs: `scontrol show job ${SLURM_JOB_ID} --details`



# Detecting Misconfiguration – Underutilization of Resources

- Reference Run on CLAIX-2018

- ⚠️ 60 GB/s vs. 200 GB/s

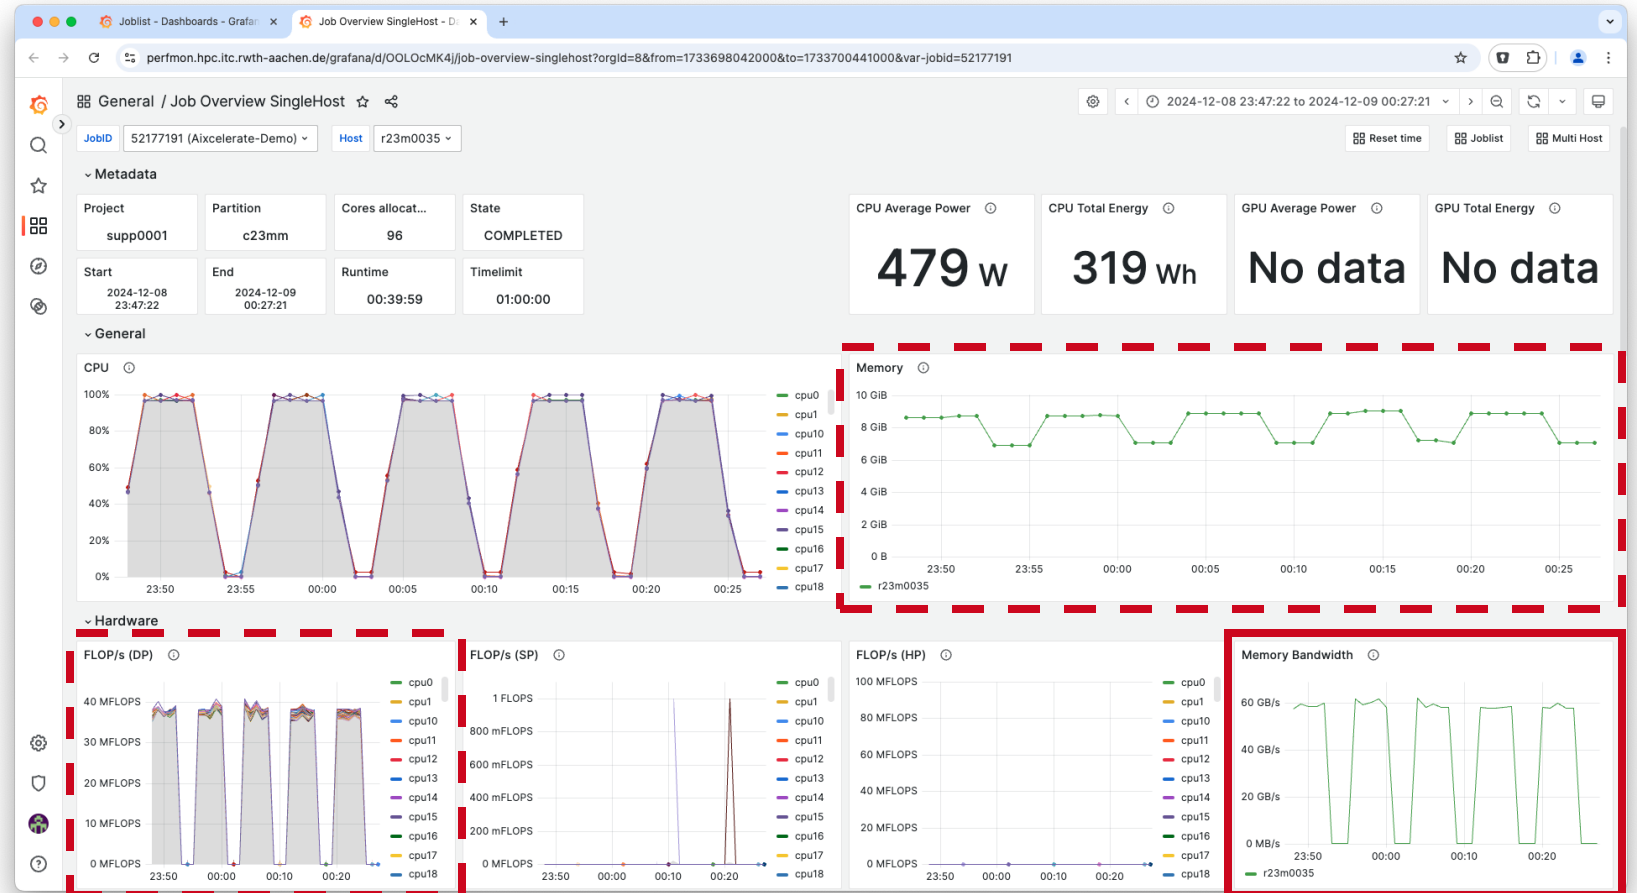


- Previous Experience with Job:

- “Usually 90% of Peak Bandwidth”
- “Before > 200 MFLOP/s occurred”

- Talk to us about:

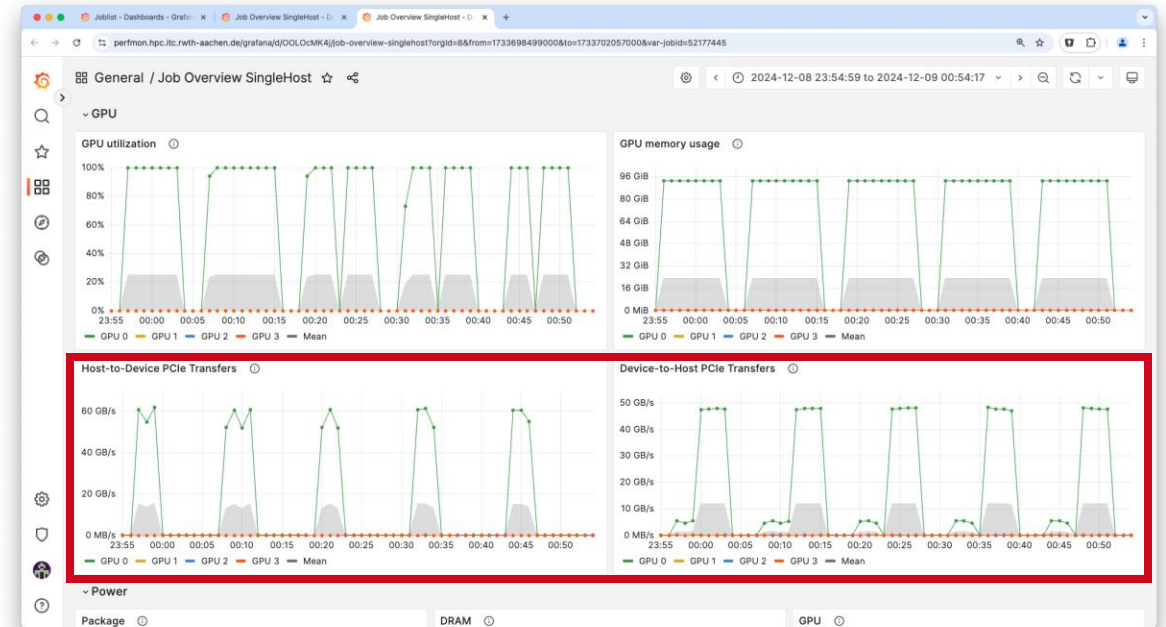
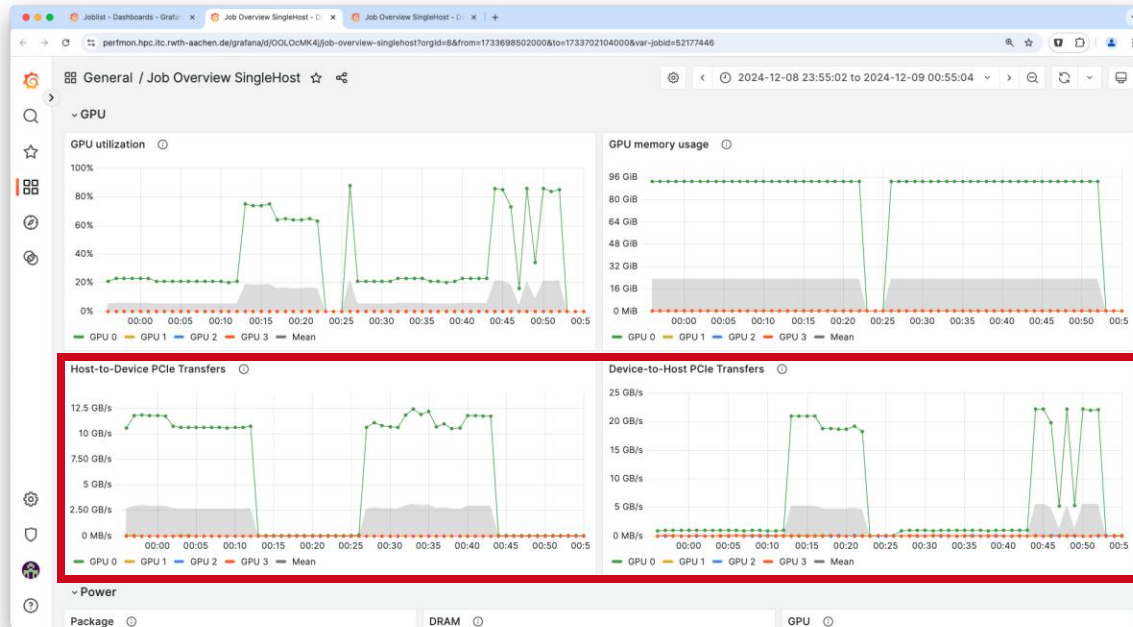
- Expectations for CLAIX-2023
- Deviations from previous results





# Detecting Misconfiguration – Underutilization of Resources

- Applies similarly to other GPU-specific execution characteristics like
  - PCIe transfer speeds: 10-25 GB/s for pageable memory vs. 45-60 GB/s for pinned memory



- Also compare your expectations for: GPU memory usage, CPU runtime portion, network / storage utilization

# Agenda

---

- Accessing the Monitoring System
  - The Webportal
  - Joblist Overview
  - Job Details
  - Caveats and Gotchas
- Detecting Misconfiguration
  - Wrong Core Count
  - Single GPU Restriction
  - Underutilization of Resources
- **Call to Action**

## Call to Action


---

Click icon to visit Grafana:



Or go to: [perfmon.hpc.itc.rwth-aachen.de](https://perfmon.hpc.itc.rwth-aachen.de)

Missing something?  
New ideas are welcome!



Let's discuss your  
results together!

# References

---

- Images & Icons:
  - Meeting by Hermine Blanquart from <https://thenounproject.com/icon/meeting-4220250> (CC BY 3.0)
  - Ideas by Ian Ransley from <https://thenounproject.com/icon/ideas-1879490> (CC BY 3.0)
  - Grafana from [https://github.com/grafana/grafana/blob/main/public/img/grafana\\_icon.svg](https://github.com/grafana/grafana/blob/main/public/img/grafana_icon.svg)
  - Gauge High from <https://fontawesome.com/icons/gauge-high>
  - Plug Circle Bolt from <https://fontawesome.com/icons/plug-circle-bolt>
  - File Code from <https://fontawesome.com/icons/file-code>