

ClusterCockpit Monitoring Service

Jan Eitzinger, 12.07.2022

ClusterCockpit 

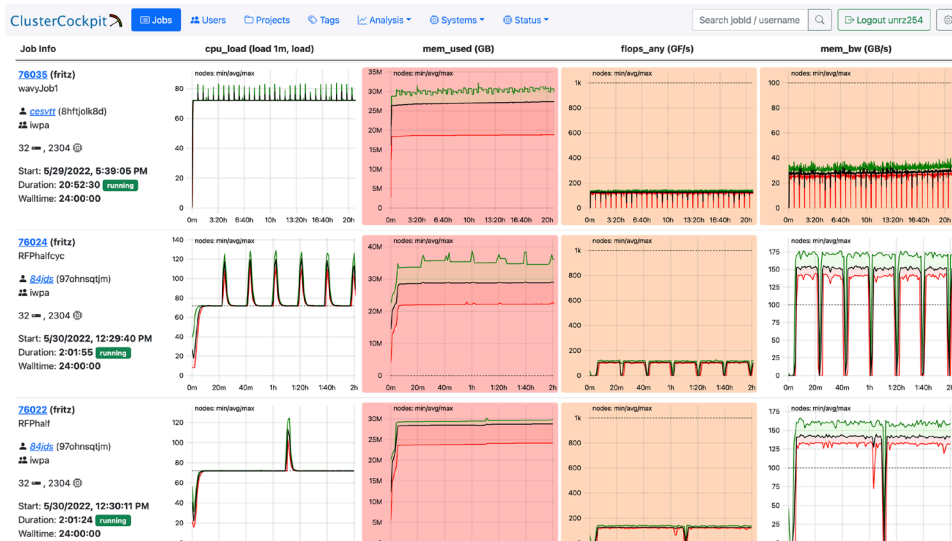


Overview

- Web-based job-specific performance monitoring
- Rudimentary job accounting and access to job performance metric data
- User Authentication with IDM HPC accounts
- **HPC-Portal** accounts cannot login yet
- ClusterCockpit Framework developed at NHR@FAU



<https://github.com/ClusterCockpit/>

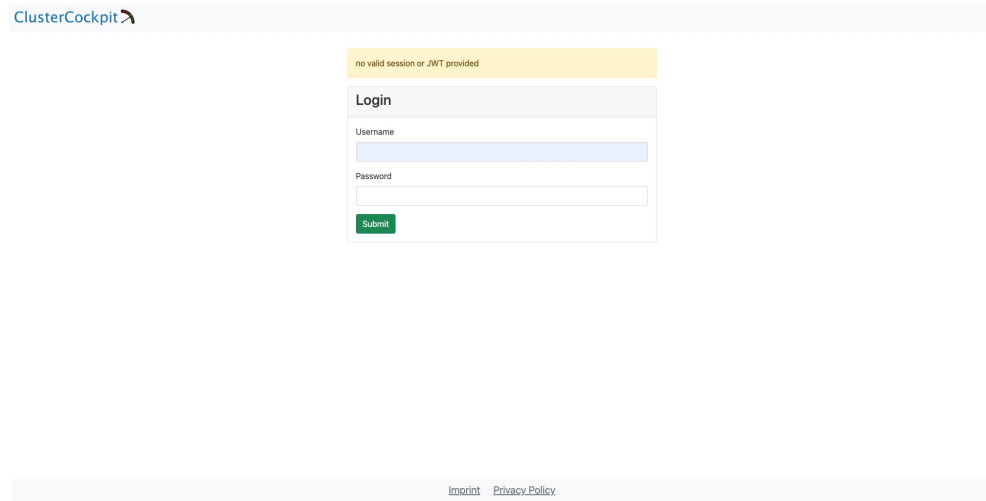


Service still
in BETA!

Access to ClusterCockpit

URL: <https://monitoring.nhr.fau.de/>

- Login with **IDM HPC account**
- **HPC Portal users cannot login (yet)**



ClusterCockpit ↗

no valid session or JWT provided

Login

Username

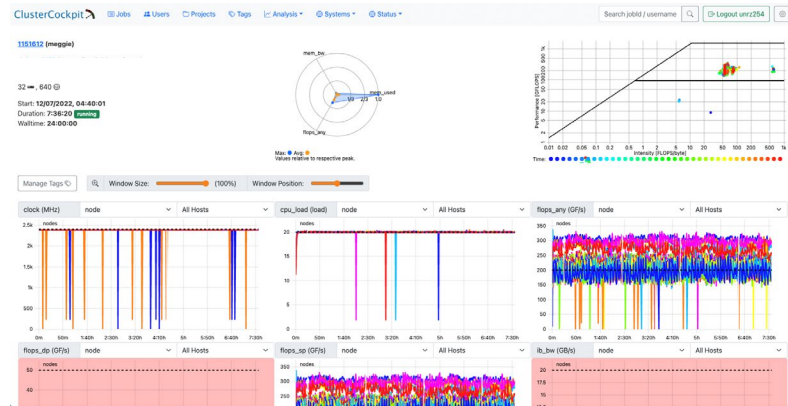
Password

[Imprint](#) [Privacy Policy](#)

- You can see **your running** and **completed jobs**
- Future: Project managers can see jobs from all users in a project

Purpose and use cases

- **Access to job performance metrics**
 - **Feedback** about job performance
 - Identify **pathological jobs** early on
 - **Monitor** and **classify** your job performance
- **Simple job monitoring and accounting**
 - **Overview** about running and completed jobs
 - **Histograms** about walltime and node count
 - Powerful **filter** and **search** features
 - **Tag jobs** to group and organize your jobs



Empirical roofline plot

How fast can tasks be processed? P [flop/s]


The bottleneck is either

- The execution of work: P_{peak} [flop/s]
- The data path: $I \cdot b_S$ [flop/byte x byte/s]

This is the “Naïve Roofline Model”


- High intensity: P limited by execution
- Low intensity: P limited by data transfer
- “Knee” at $P_{\text{max}} = I \cdot b_S$:
Best use of resources

Measured using hardware performance counter data





How to detect bad jobs



- Load metric indicates allocation or placement issues



- Basic resource utilization (arithmetic and memory bandwidth) can be seen in the roofline and polar plots



Max: ● Avg: ●
Values relative to respective peak.



Max: ● Avg: ●
Values relative to respective peak.

Node statistic table to detect performance imbalance

Statistics Table [Job Script](#) [Slurm Info](#)

Metrics

Node	min [△]	avg [▲]	max [△]	min [△]	avg [△]	max [△]	min [△]	avg [△]	max [△]
m0445	0	28.9	34.6	0	41.8	55.3	2	20	22.1
m0709	0	26.5	32	0	35.3	46	2	20	21.3
m0705	0	26.4	32.6	0	35.4	45.8	2	20	21.3
m0706	0	26.4	32.5	0	36.3	47.4	2	20	21.4
m0710	0	26.2	32.5	0	33.3	45.6	2	20	21.6
m0711	0	26.2	32.4	0.1	35	47.9	20	20.1	21.2
m0668	0	26.1	32.5	0	37.9	47.8	2	20	21.2
m0661	0	25.8	32.4	0	37.8	49.6	20	20.1	21.3
m0269	0	17.3	27.3	0	29.1	42.4	2	20	21.7
m1160	0	17.2	26.3	0	30.8	42.7	2	20	20.8
m1133	0	16.9	27.3	0.1	32.5	43.5	20	20.1	21.4
m1135	0	16.9	25.3	0	33.7	45.5	2	20	21.5
m1131	0	16.8	25.8	0	33.3	45.3	2	20	20.8
m1158	0	16.6	25.9	0	31.7	44.3	2	20.1	21.3
m1141	0	16.5	24.2	0	34.6	46.4	20	20.1	20.5
m1137	0	16.4	24.1	0	33.5	44.1	2	20	20.3
m1139	0	16.4	25.3	0	32.2	42.7	2	20	20.9
m1138	0	16.3	24	0	32.7	43.8	2	20	20.6

Outlook

- **Ongoing work on:**
 - Login support for HPC Portal users
 - Support all Clusters
 - Intuitive Job performance visualization
 - Automatic Job classification
 - Automatic application tagging

- **Please open a ticket if you encounter a problem!**

Any remarks or questions?

DEMO