

Job specific performance monitoring with ClusterCockpit

DKRZ TechTalk

Pay Giesselmann



Motivation:

Performance increases Throughput
Performant means Energy Efficient

<https://clustercockpit.dkrz.de>



The login form is centered on a dark blue background. It features the DKRZ logo at the top left, which consists of a stylized blue and white graphic followed by the text 'DKRZ DEUTSCHES KLIMARECHENZENTRUM'. Below the logo, the text 'Welcome to ClusterCockpit' is displayed. The form contains two input fields: 'Username' and 'Password', both with white text and white input boxes. A blue 'Log in' button is positioned at the bottom right of the form.

[Imprint](#) | [Privacy Policy](#) | Version 1.2.2 (a5b29f4) | Built 2023-11-21T15:15:23


Agenda

- Guided Tour
- Performance Metrics
- Use Cases

Available Monitoring

- `squeue -u <uid>`
- `scontrol show job <jobid>`
- `sacct -u <uid>`
- log files ...

<https://monitoring.dkrz.de>





Available Monitoring


- `queue -u <uid>`
- `scontrol show job <jobid>`
- `sacct -u <uid>`
- log files ...

- **But what happened to my jobs during runtime?**
 - Job specific?
 - Performance?
 - Resolved in time?

<https://monitoring.dkrz.de>



Guided Tour

ClusterCockpit  My Jobs Job Search Tags

Search 'type:<query>' ...

[Logout k202166](#)

Clusters

Name	Running Jobs	Total Jobs
levante	0 jobs	1 jobs

Guided Tour - My Jobs

ClusterCockpit My Jobs Job Search Tags

Search 'type:<query>' ...



[Logout k202166](#)



Sorting

Metrics


Filters

Reload


No periodic reload

Username	k202166
Name	Pay Giesselmann
Total Jobs	1
Short Jobs	0
Total Walltime	2
Total Core Hours	710

Duration Distribution




Number of Nodes Distribution




Job Info

[8382826](#) (levante)
 spawner-jupyterhub
[k202166](#) (Pay Giesselmann)
[k20200](#)
 l20047, 256 @
 compute
 Start: 11/01/2024, 14:05:14
 Duration: 2:46:25 cancelled
 Walltime: 4:00:00


cpu_load (%)




flops_any (GFlops/s)




ipc (IPC)



mem_used (GB)



mem_bw (GB/s)



Jobs per page: 50

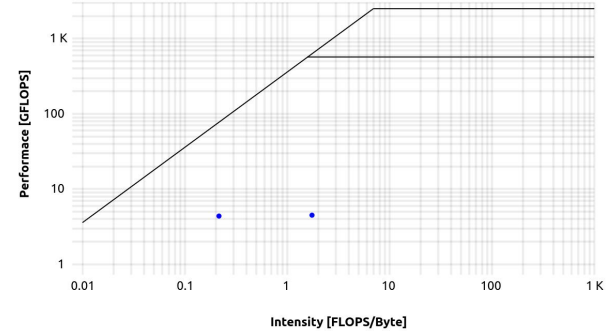

0 - 1 of 1 Jobs

Guided Tour - Job Details



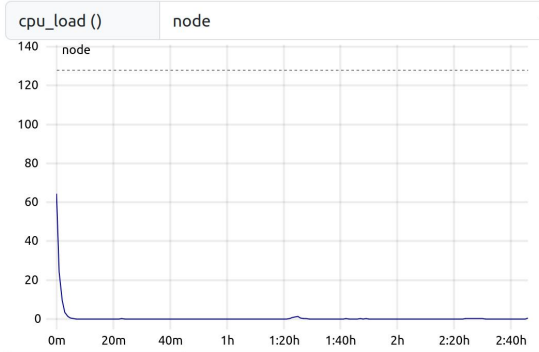
ClusterCockpit My Jobs Job Search Tags

Search 'type:<query>' ... [Logout k202166](#)

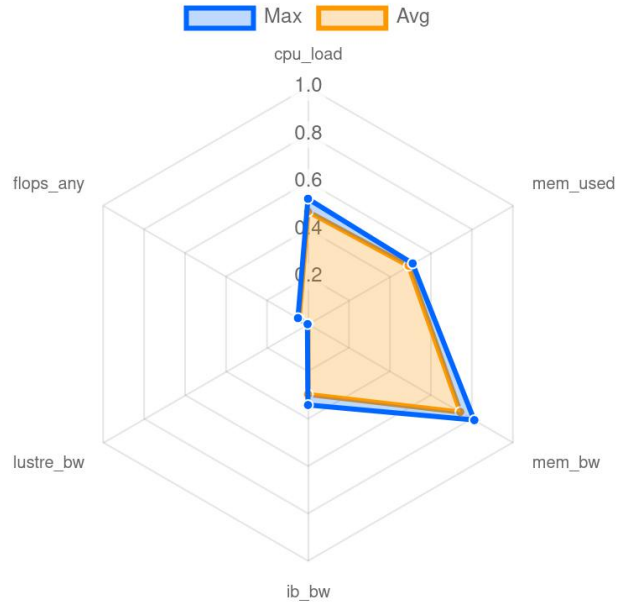
8382826 (levante)
 spawner-jupyterhub
[k202166](#) (Pay Giesselmann)
[k20200](#)
 l20047 , 256
 compute
 Start: 11/01/2024, 14:05:14
 Duration: 2:46:25 cancelled
 Walltime: 4:00:00



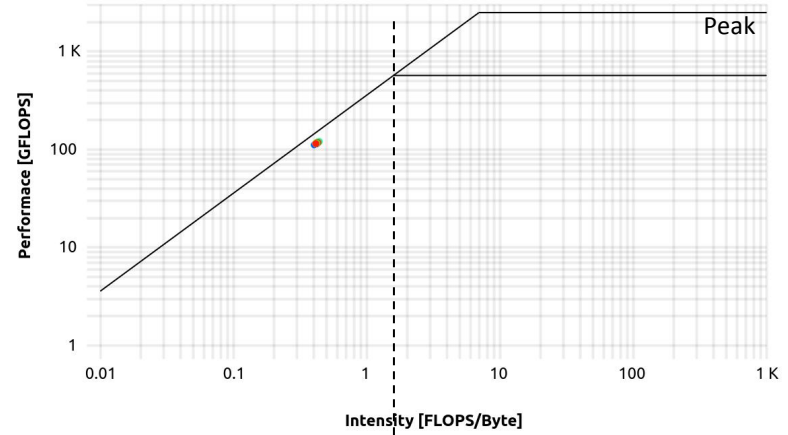
[Manage Tags](#)




Guided Tour - Job Details



Footprint



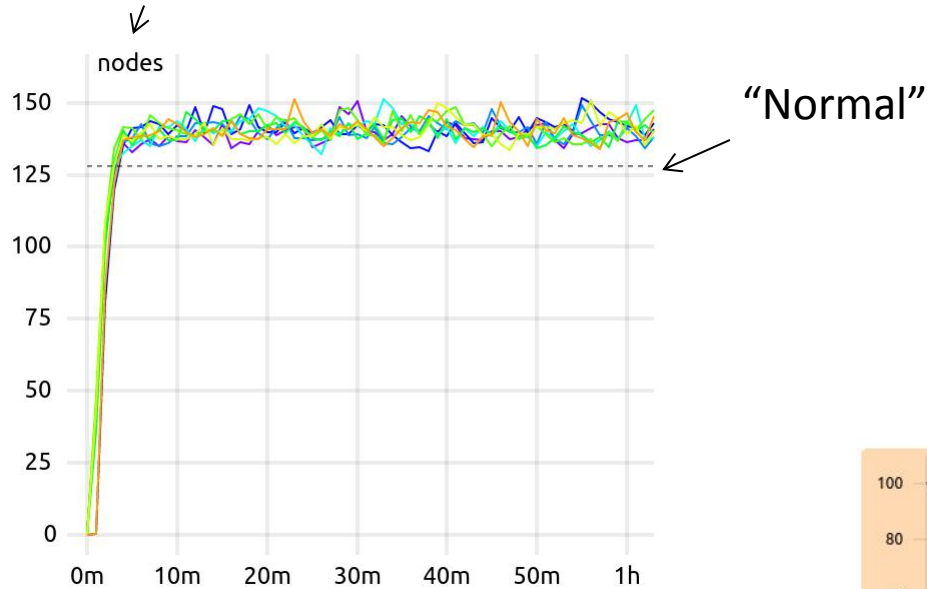
Memory bound



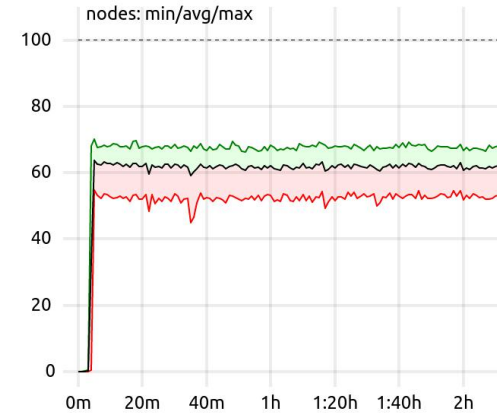
Compute bound

Performance Metrics


Metric Scope




Large Jobs > 15 nodes



FLOPS (Caution)



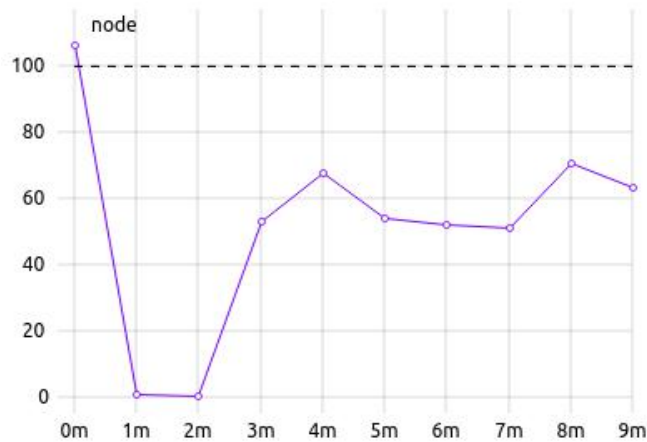


CPU Load (Alert)



Sampling Artefacts and Aliasing

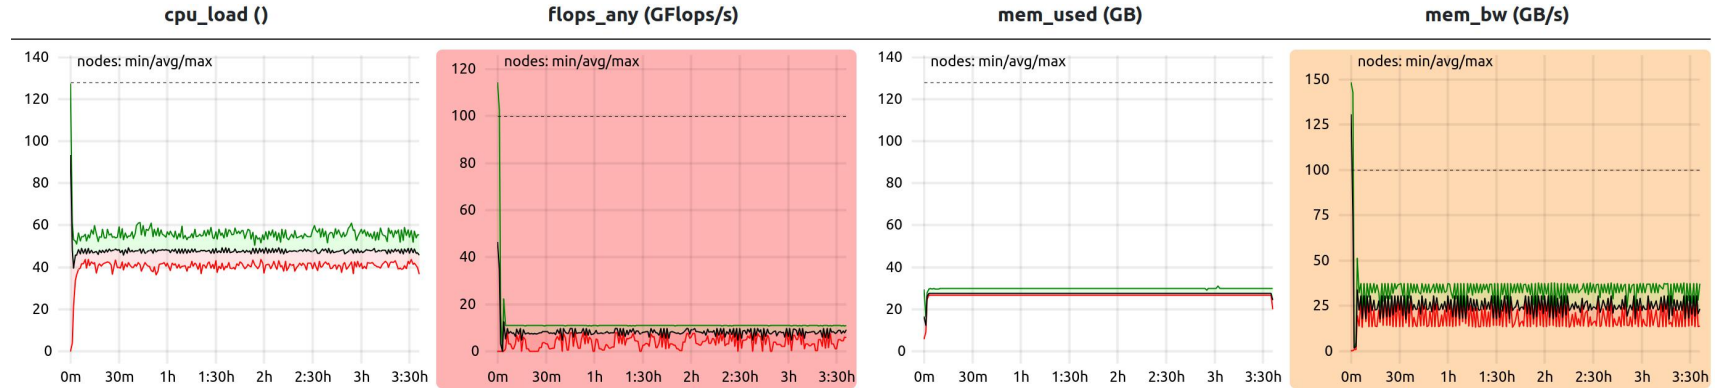
- Fixed one minute sampling
- Async job scheduling



What is it good for?

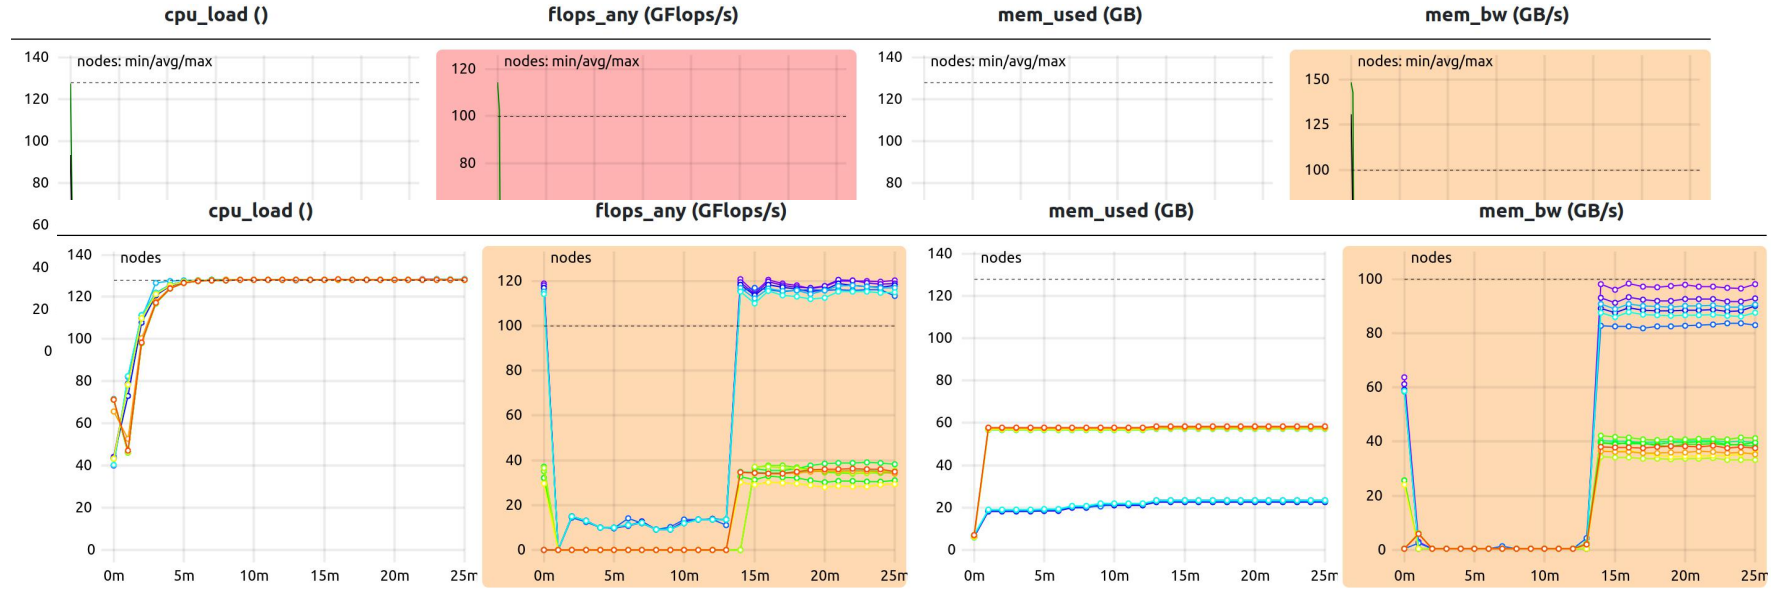
Intuitive Overview

101 nodes




Intuitive Overview

101 nodes



Intuitive Overview

101 nodes



80 nodes, OMP_NUM_THREADS=4, missing libiomp5.so

Performance Goals

Jobs:

> **5min** with

State **completed**


From 01.01 to 07.01

Levante **compute** partition

Plots:

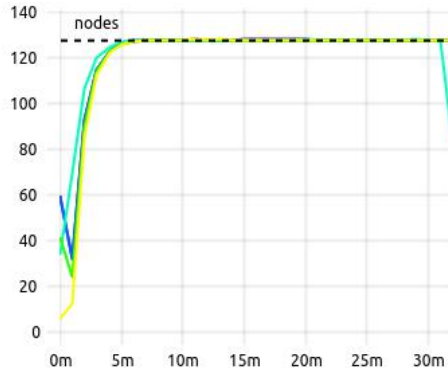
Distribution of job averages

Category	Jobs	Node Hours
<5 nodes	11.329	19.477
5-10 nodes	5.617	35.087
>10 nodes	1.948	225.773

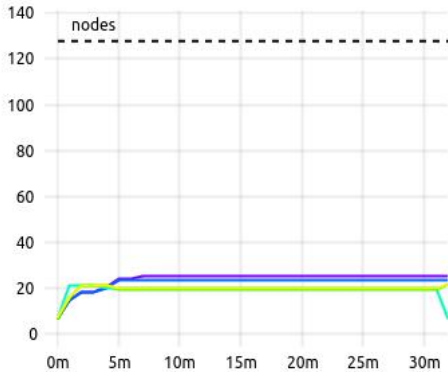



Performance Goals - On Point

128 Cores




256 HW-Threads




Levante:
 256 GB (2670 nodes)
 512 GB (294 nodes)
 1024 GB (18 nodes)

variable = cpu_load

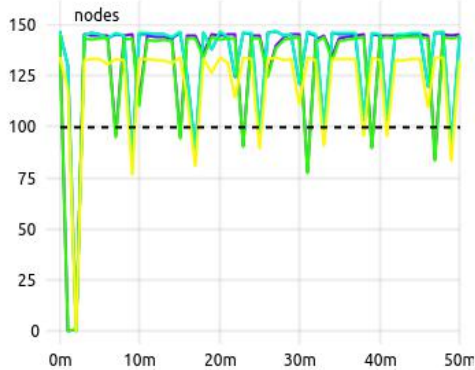


variable = mem_used




Performance Goals - Let It Burn

NEMO 4.2 @5 nodes





LIKWID Bench:
570 GFLOPS (Scalar)
2500 GFLOPS (SIMD)







ICON @8 nodes




STREAM:
360 GB/s



Performance Goals - As Much As Needed



Infiniband:
100 GBit/s Full Duplex



Limitations

- Node level metrics on shared/interactive
- Jobs > 2min imported every 5min
- 6 Months data retention

- No GPU FLOPS & memory bandwidth
- GPU Power disabled (aliasing)

- Statistics filter only on some metrics

Summary

- ClusterCockpit job monitoring platform
- Quick overview of running and completed jobs
- Reading job specific performance metrics
- Performance goals and limits