SH in REANA

06.10.2025

AKh

Extracting our homeland galaxy properties

The Gaia mission has provided an invaluable wealth of astrometric data for almost two billions of stars in our Galaxy.

The synergy with the Bayesian isochrone-fitting code StarHorse, allowed to derive additional stellar properties, like distance and dust extinction (Anders+2021, Queiroz+2023).

Then a regression technique, **xgboost**, trained with spectroscopically derived data from a variety of stellar surveys, was used to determine additional stellar parameters: **extinction, effective temperature, metallicity, surface gravity, and mass** (Khalatyan+2024).

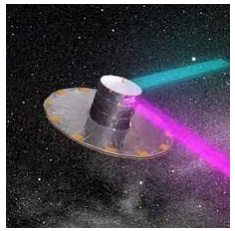

1.7Bln

>G18.5

400Mln

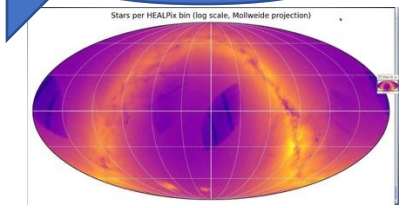

Starhorse

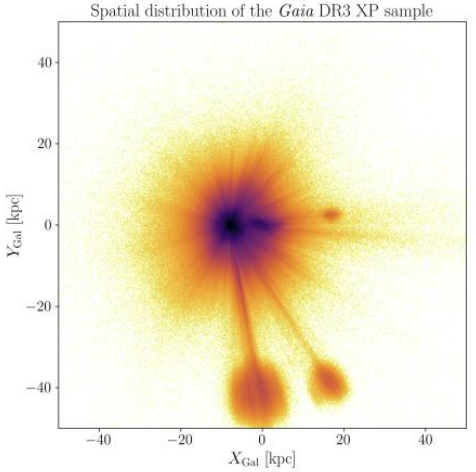

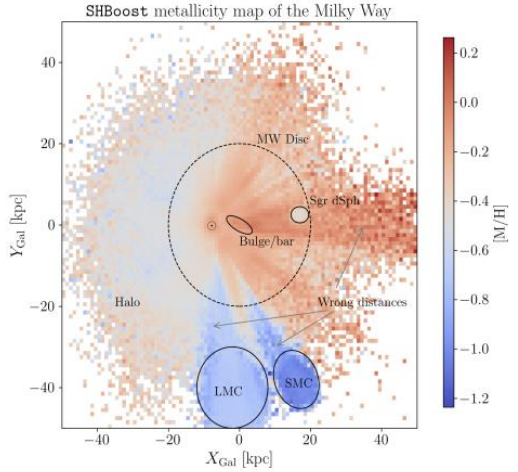

SH big run

- 12288x2 files to analyze
  - Parquet or fits files
  - Located in S3
  - Total 50GB
  - Single file 1MB-100MB
- Per file 2min-50 hours depend on Galaxy region(Bulge is most expensive)
- Cache file for warmup/run ~2GB*2
  - Located on S3
- RUN HW requirements:
  - With 6 cores 16GB RAM is minimum

REANA

- bash run_job.sh ????
  - ????: 0-12288 x2

Snakemake?


```
JOB_IDS = range(0, 4)  # Change to range(0, 12288) for full scale
rule all:
    input:
        expand("results{hp}.tar.gz", hp=JOB_IDS)
rule prepare_and_run:
    output:
        "results{hp}.tar.gz"
    container:
        "astro-ml:latest"
    resources:
     c4p_additional_requirements='regexp("gridka",Target.TardisDroneUuid)',
        compute_backend="compute4punch",
        c4p_memory_limit="40G",
        c4p_cpu_cores="16"
    shell:
        """
        touch results{wildcards.hp}.tar.gz
        sync
        sleep 5

        """
rule check_all:
    input:
        expand("results{hp}.tar.gz", hp=JOB_IDS)

    shell:
        """
        echo "========================================"
        echo "Job completed successfully"
        echo "Output size: $(du -h results*.tar.gz | cut -f1)"
        echo "Time: $(date)"
        echo "========================================"
        """
```

Problems: snakemake

- Jobs was running
- Environment is not propagated to c4p backend
- If we define on Snakefile multiple jobs: do they run in parallel?
- How to handle failed jobs?

S4P is very slow

```
[arm2arm@c4p-login workflows]$ du -bsh ./*
333M    ./be37eff7-24ca-4315-97c5-93fd54c39991
388M    ./d07a5359-090e-40a7-b51b-0883ad1f0104
86M     ./d0dda27d-6efd-476e-8453-24424ff14e23
68K     ./d39485c7-6dd7-4c62-b657-d4d39ca2c859
480M    ./e1d629eb-f2e0-4c46-b5a7-d4e0b63275a6
29M     ./e1dc1a92-d2f6-47a9-8bec-347914b733e8
388M    ./efe2be28-c945-4f70-9058-30cc418849a8
125K    ./fb26efd3-595e-40d5-9f2f-cefe25f31eb3
[arm2arm@c4p-login workflows]$ time rm -fr *
rm: cannot remove 'd0dda27d-6efd-476e-8453-24424ff14e23': Directory not empty

real    0m25.066s
user    0m0.025s
sys     0m0.880s
[arm2arm@c4p-login workflows]$
[arm2arm@c4p-login workflows]$ time rm -fr *
rm: cannot remove 'd0dda27d-6efd-476e-8453-24424ff14e23/.nfs00000000017e3c1700000001': Device or resource busy

real    0m0.004s
user    0m0.002s
sys     0m0.001s
```

WakeUp!

```
condor_master                    condor_ssh_start                    condor_wait
[arm2arm@c4p-login ~]$ condor_submit wakeup.sh
Submitting job(s)

Hello arm2arm! You are going to submit your HTCondor job(s) using the issuer(s) helmholtz.

Defining your credential for the issuer helmholtz.

A valid credential has been found with a remaining life time of 23 hours 27 minutes 47 seconds.

Its remaining life time is smaller than 24 hours!

Do you want to renew it? Please answer yes or no: yes

Please visit the following url in order to generate your credential: https://mytoken.data.kit.edu/c/ZQaPZVLb

Starting polling and waiting for your approval ......
```

First light with serial jobs

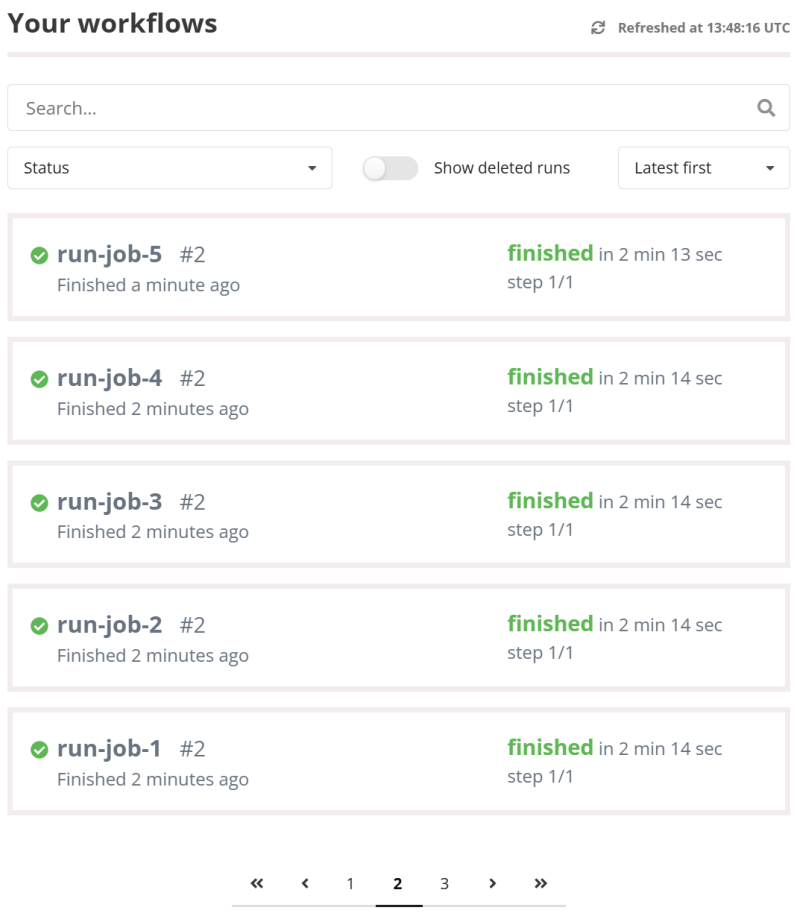

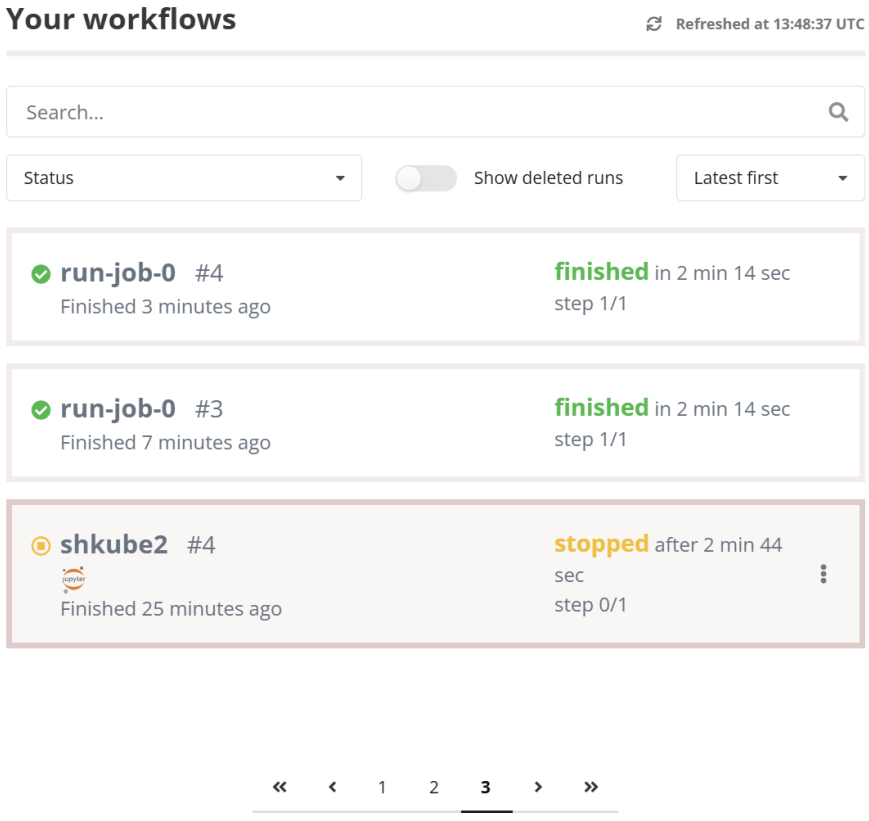

The data curation: S3->c4p->S3 works!

```
Added `mci` successfully.

Mode: none

Path mode: WITHOUT_SHBOOST

Checking S3 file existence: s3:starhorsereana/arman/results/WITHOUT_SHBOOST/results10.tar.gz

File found on S3

Results already exist on S3: s3:starhorsereana/arman/results/WITHOUT_SHBOOST/results10.tar.gz
```

LDAP/connection random failure

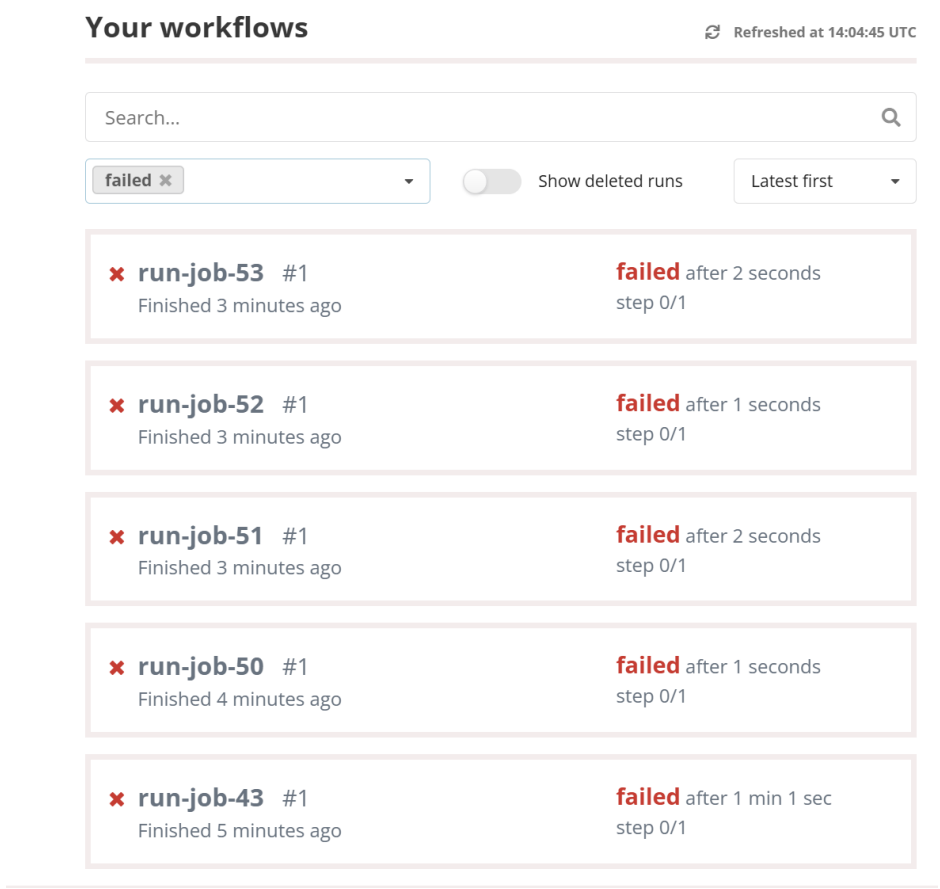

steps.sh

chmod +x *.sh

env

bash prepare.sh

bash **run_job.sh ${1}**

bash cleanup.sh

REANA Script: array.yaml

```
inputs:
  parameters:
    hp_index: 0
  files:
    - steps.sh
    - run_job.sh
    - prepare.sh
    - cleanup.sh
    - git-clone.sh
    - templates/param_template.ini
    - templates/config_template.ini
    - templates/config_withshboost_template.ini
    - templates/param_withshboost_template.ini
workflow:
  type: serial
  specification:
    steps:
      - name: starhorse
        environment: 'astro-ml:latest'
        compute_backend: "compute4punch"
        c4p_memory_limit: "40G"
        c4p_cpu_cores: "16"
        c4p_additional_requirements:  'regexp("gridka",Target.TardisDroneUuid)'
        commands:
          - bash steps.sh ${hp_index}
```

Pushing limits

- reana-client run --skip-validation -w run-job-$1 -f array.yaml -p hp_index=$1
- seq 151 1000 | xargs **-P16** -I{} ./many.sh {}

Load… is ok

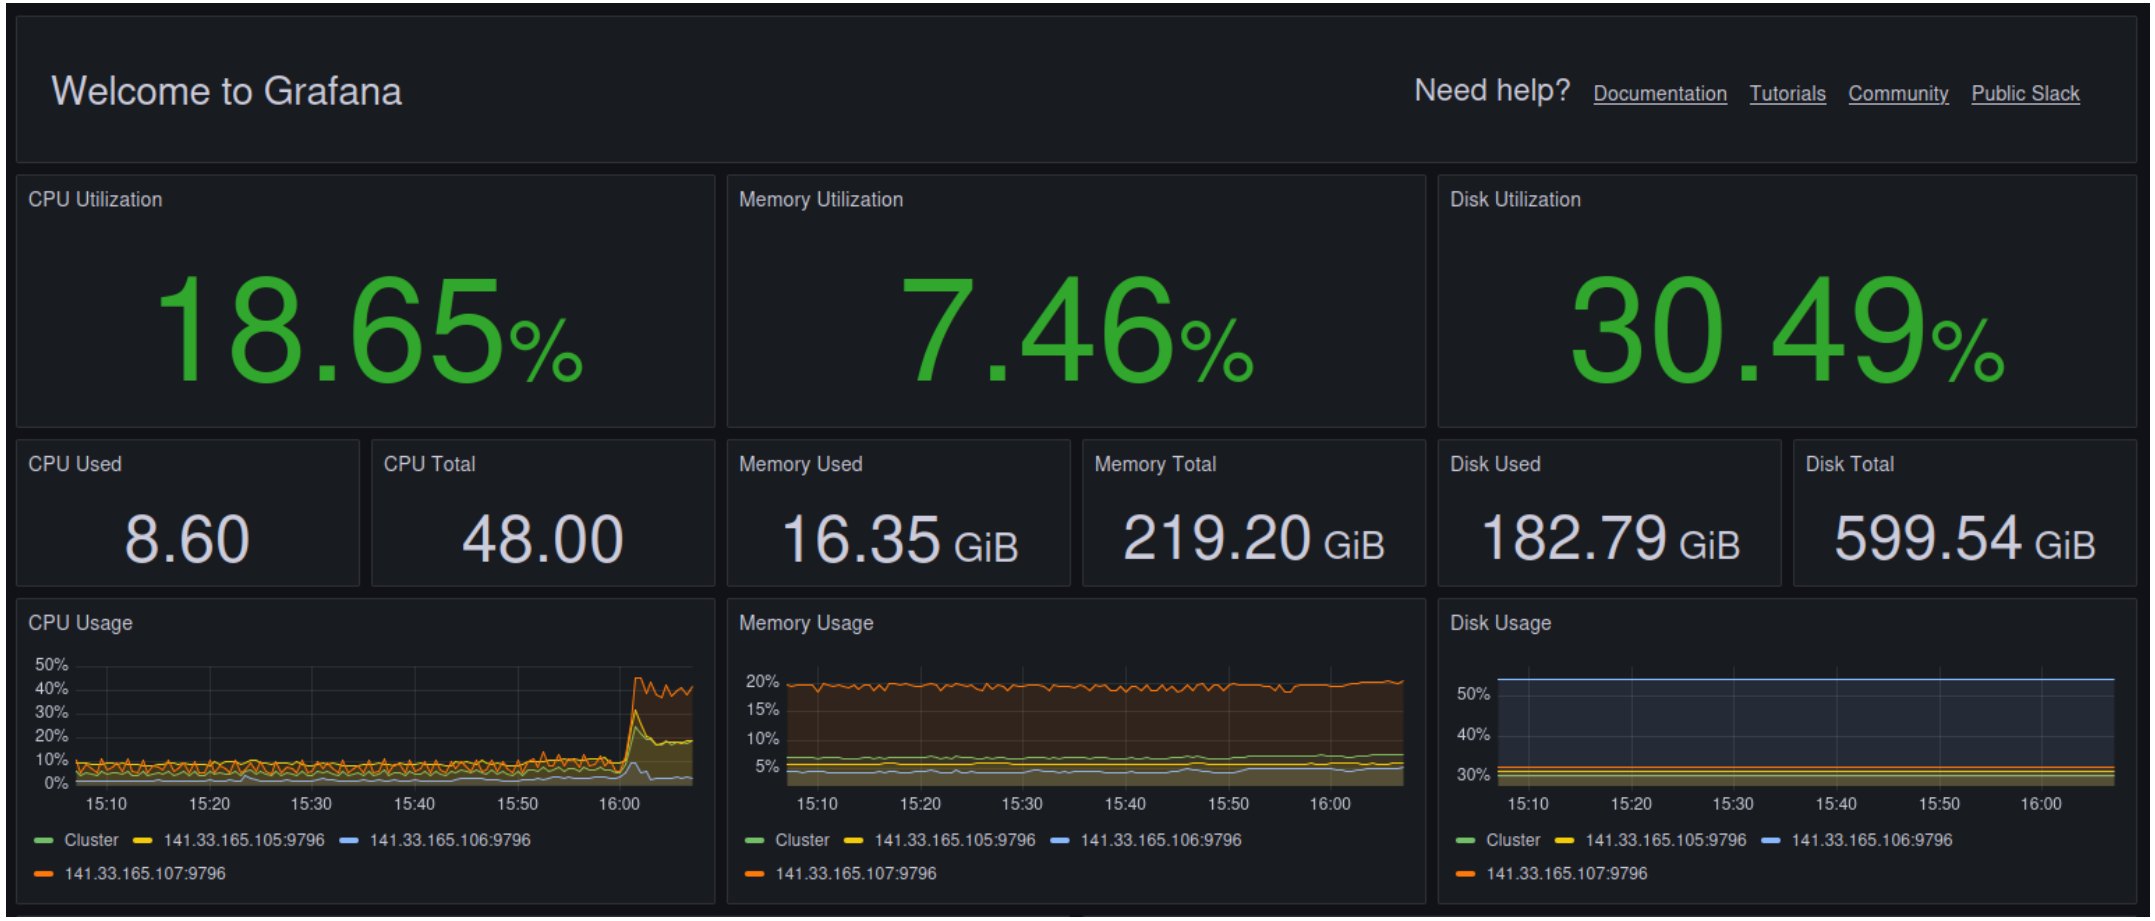
Welcome to Grafana
Need help? Documentation Tutorials Community Public Slack
CPU Utilization
18.65%
Memory Utilization
7.46%
Disk Utilization
30.49%
CPU Used
8.60
CPU Total
48.00
Memory Used
16.35 GiB
Memory Total
219.20 GiB
Disk Used
182.79 GiB
Disk Total
599.54 GiB
CPU Usage
50%
40%
30%
20%
10%
0%
15:10
15:20
15:30
15:40
15:50
16:00
Cluster
141.33.165.105:9796
141.33.165.106:9796
141.33.165.107:9796
Memory Usage
20%
15%
10%
5%
15:10
15:20
15:30
15:40
15:50
16:00
Cluster
141.33.165.105:9796
141.33.165.106:9796
141.33.165.107:9796
Disk Usage
50%
40%
30%
15:10
15:20
15:30
15:40
15:50
16:00
Cluster
141.33.165.105:9796
141.33.165.106:9796
141.33.165.107:9796

AUTH/Ldap, Striking again ☹

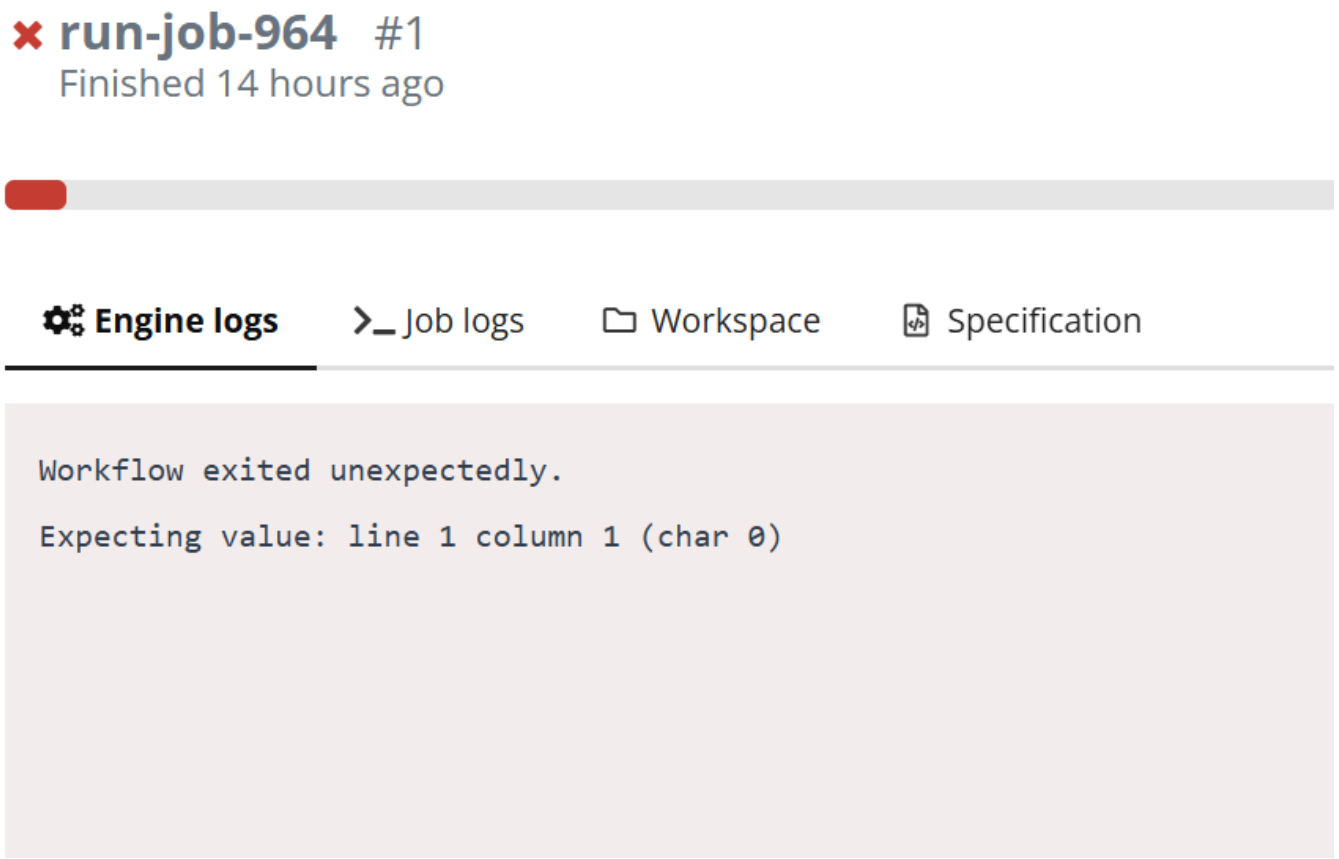

REANA limits

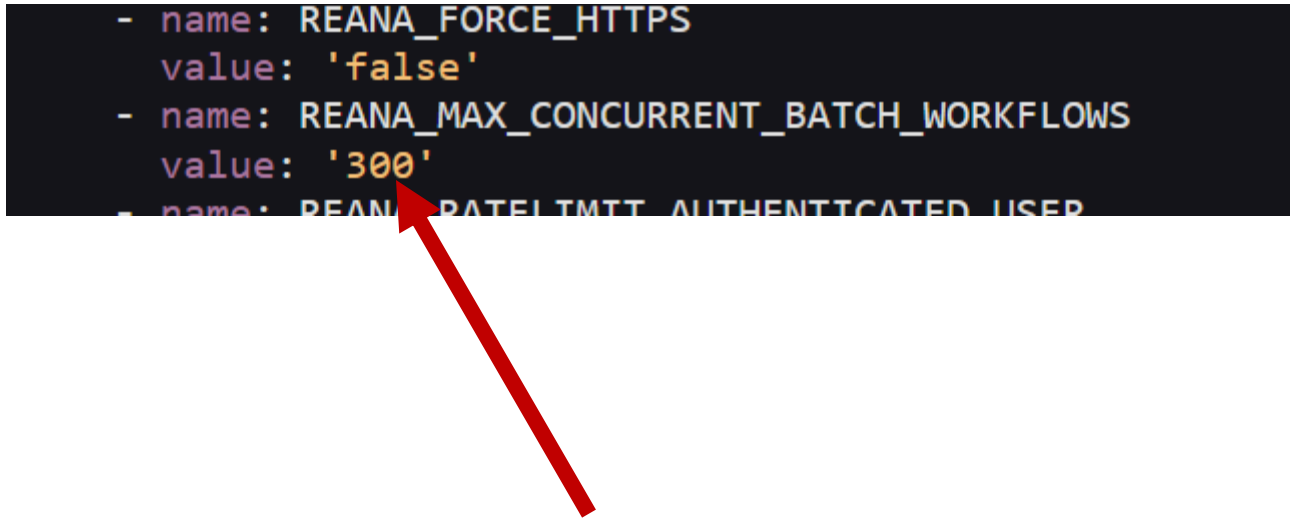

More fair share policy from GridKa for C4P

1,824 vCORES, 4.5TB RAM

```
Total for query: 246 jobs; 0 completed, 0 removed, 132 idle, 114 running, 0 held, 0 suspended
Total for arm2arm: 246 jobs; 0 completed, 0 removed, 132 idle, 114 running, 0 held, 0 suspended
Total for all users: 246 jobs; 0 completed, 0 removed, 132 idle, 114 running, 0 held, 0 suspende
d

[arm2arm@c4p-login ~]$
```

Scaling jobs.

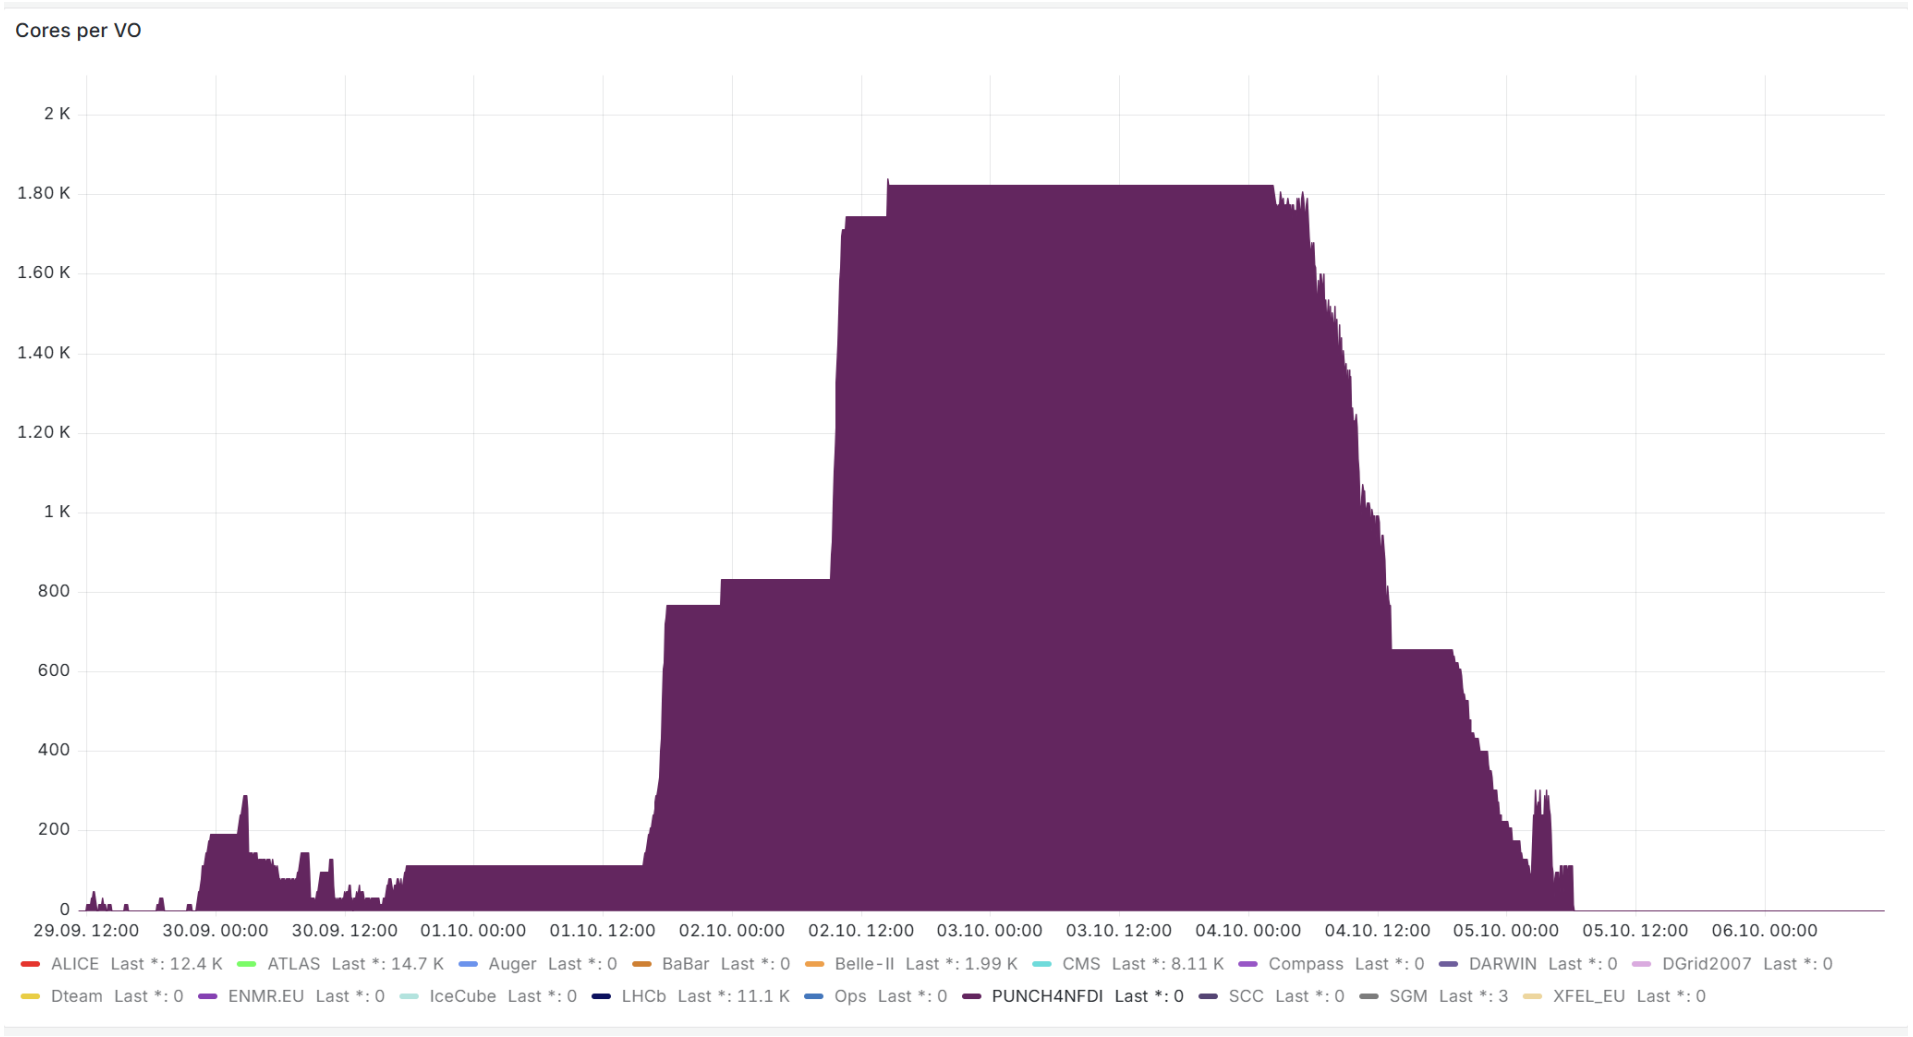

LIMITS for C4P

- REANA:
  - Max Concurrent Batch Workflows: 10000
- C4P
  - Type of nodes available:
    - 8 vCores 20GB RAM
    - 16 vCores 40GB RAM
      - compute_backend: "**compute4punch**“
      - c4p_additional_requirements: 'regexp("**gridka**",Target.TardisDroneUuid)'
  - Number of nodes available
    - Unknown, dynamic
    - Max RUNtime: <5d (at least at **gridka**)

Current Situation

- 2 groups of runs:
  - 12K with SHBoost 200M
  - 12K without SHBoost 200M
- 12000 – with SHBoost
  - 5178 done
- 12000 – without SHBoost
  - 1792 done
  - **1000 on C4P**

Next steps

- Making more robust against failures and network outages
  - Check Git pull networks outages
  - Check S3 outages
- Make jobs smaller:
  - 20K Stars – 5h
  - 8Cores 20GB Ram
- Try to go more

Kubernetes limits for scheduler(Benoit)

- The maximum jos are not going more than 200-300
- Every job is one pod for monitoring

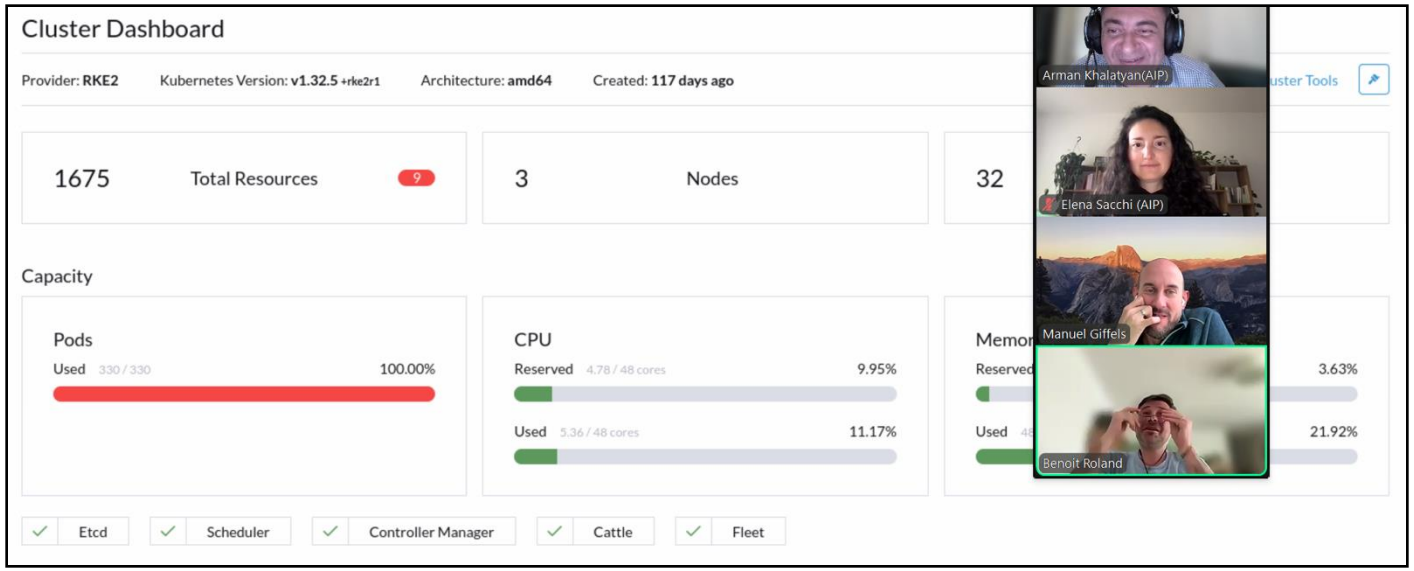

Status: 13.10.2025

- **Good news**
  - Still running
  - Many bugs are fixed
- We hit Max: **2290 Cores**
- Black areas still calculating

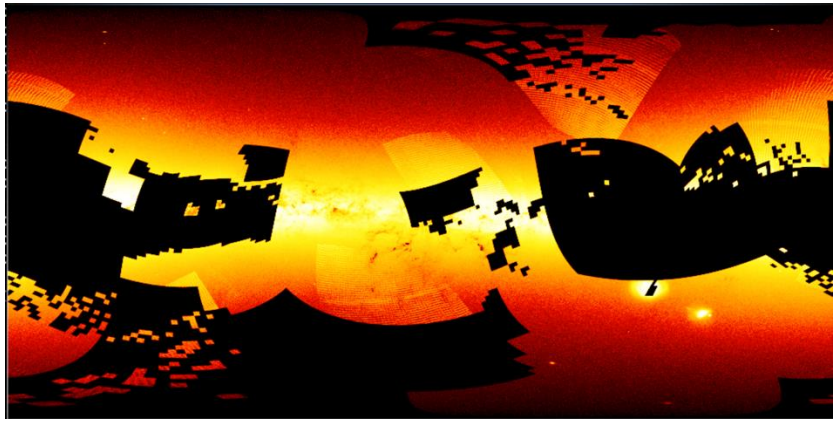

Last 7days

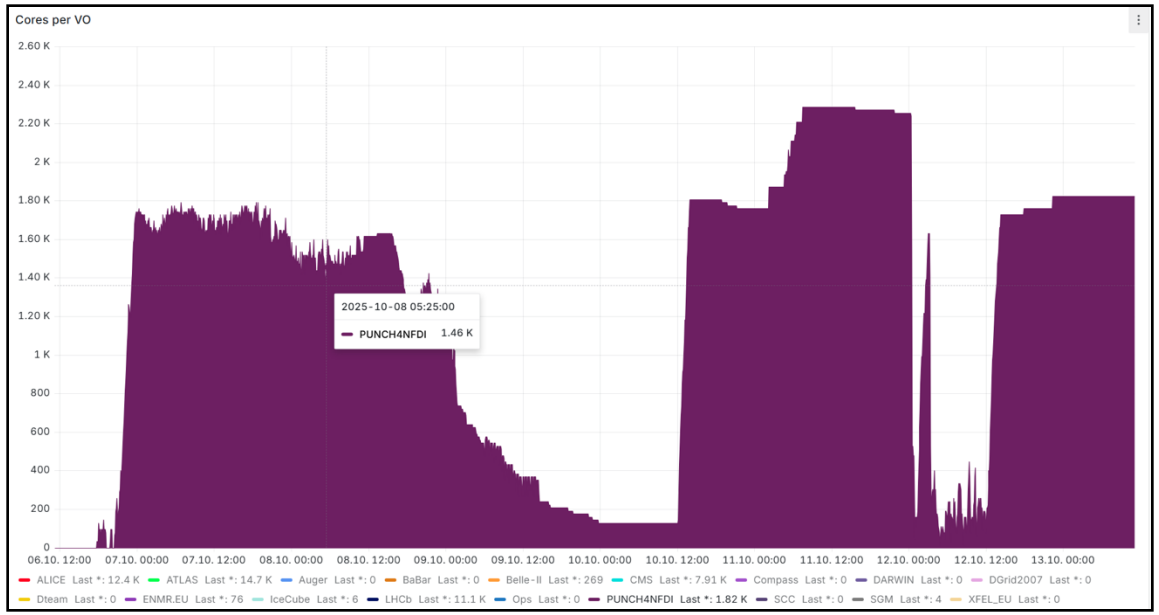

Open questions

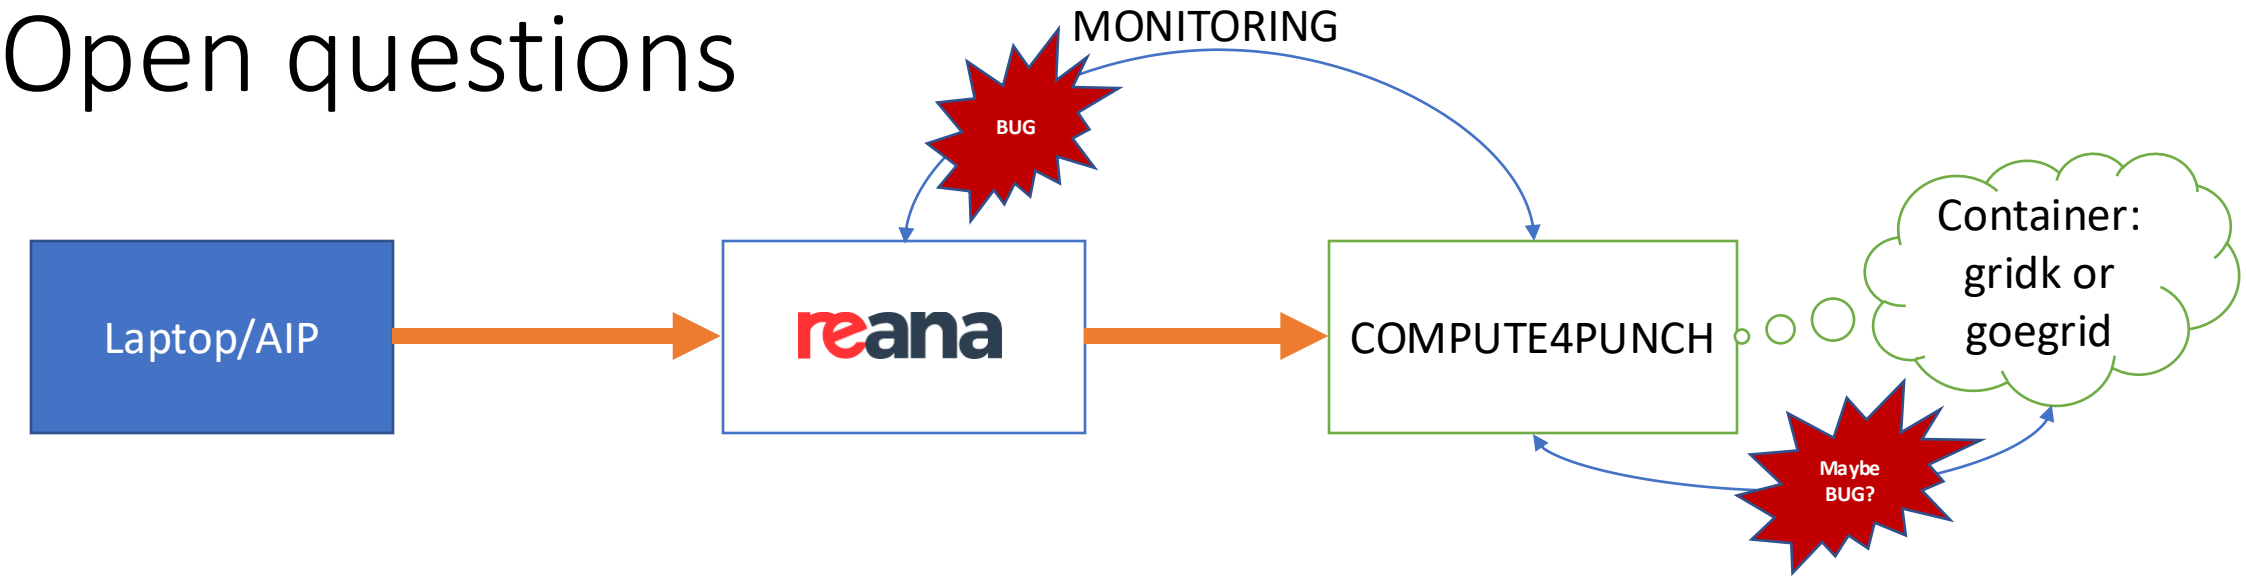

- Babysitting 25K jobs is hard, many cli scripts to find the really failed jobs
- The jobs are running/hold/idle : REANA thinks it is still RUNING
- Some jobs are failed on REANA-FAILED but condor_q- running/hold/idle
- Monitoring on c4p-login node: here is no feedback on jobs logs
  - User need to login to running container and show status
- ~~BIG BUG on login failed with "LDAP" -~~ SOLVED?(more questions @Benoit)