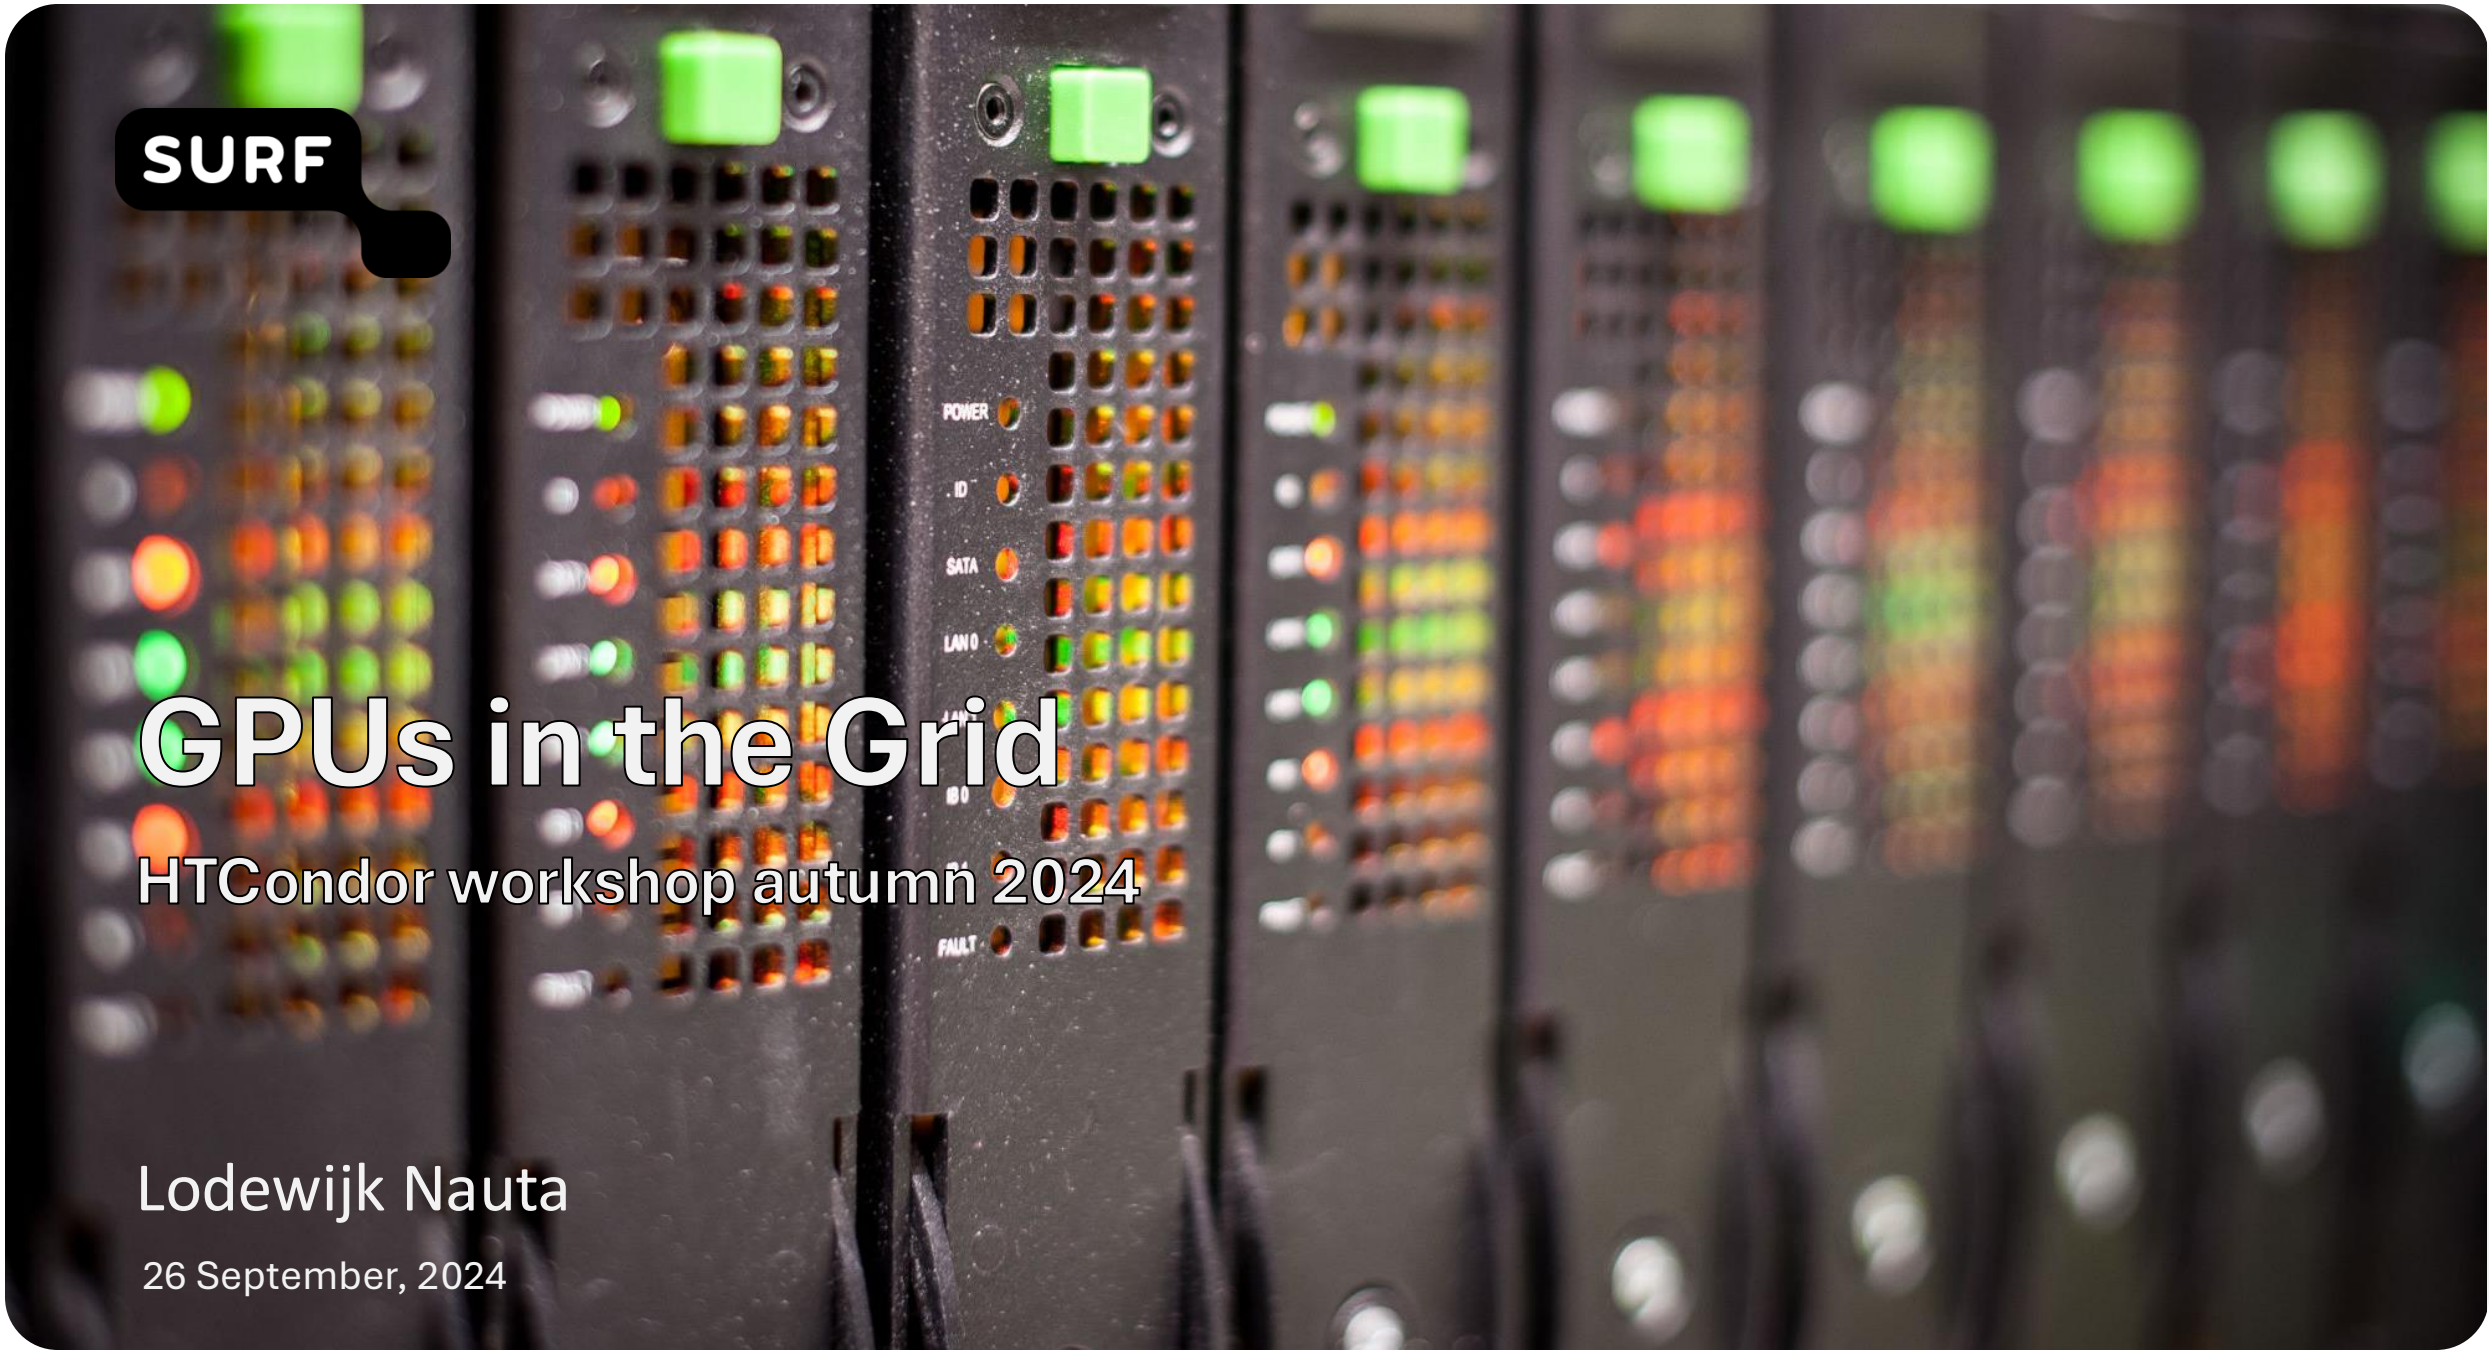

GPUs in the Grid

HTCondor workshop autumn 2024

Lodewijk Nauta

26 September, 2024

Outlook

- Grid cluster setup
- Performance
- User adoption

Grid site: Amsterdam

- Two providers: Nikhef & SURF (sara-matrix)
- SURF runs the local GinA cluster (‘Grid in Amsterdam’)
- Middleware: ARC / DIRAC
- GinA:
  - 135 machines in cluster with some ~10600 cores for users
  - Recently added 41 nodes with with 128 cores each (8 for overhead)
  - Virtualized
  - SLURM cluster

Virtualized setup

- SURF has a private cloud running on OpenStack in the AMS DC
- ~500 machines
- The services these machines make up:
  - HPC Cloud: ‘a HPC machine in a cloud setup’
  - Spider (HTC): local batch cluster for data heavy users in NL
  - GinA (HTC): local cluster that is part of NL-GRID
  - MS4: cloud machines to support Spider/GinA user setups
  - Npuls (restructure educational infrastructure)
  - (innovation cluster)

GPU nodes in GinA

- GPU nodes:
  - 4 machines with GPUs
    - Intel Xeon 6342, 48 cores per machine (4 for overhead)
    - 350GB memory / node → ~7.3GB / core
    - 4 Nvidia A10's (24GB) per node: 16 GPUs total in GinA
  - Slurm GPU Partition:
    - Job time limit set to 4 days
    - 1 GPU / job (default)
    - 11 cores / job (default)
  - ARC / DIRAC lands the jobs on the machines
- Virtual setup: you have to expose the GPU directly

SURF

Performance

- GPU Accounting is WIP (on our side and on EGI side)
- As we only track GPU-hours used
  (or: CPU-hours used assuming user is always using a GPU)
- We don’t have in-depth performance metrics (unfortunately)

- We do have some *quick&dirty* metrics:
  - Data wrangling on SLURM-db entries
  - logging in to a worker node (WN) and looking 👀

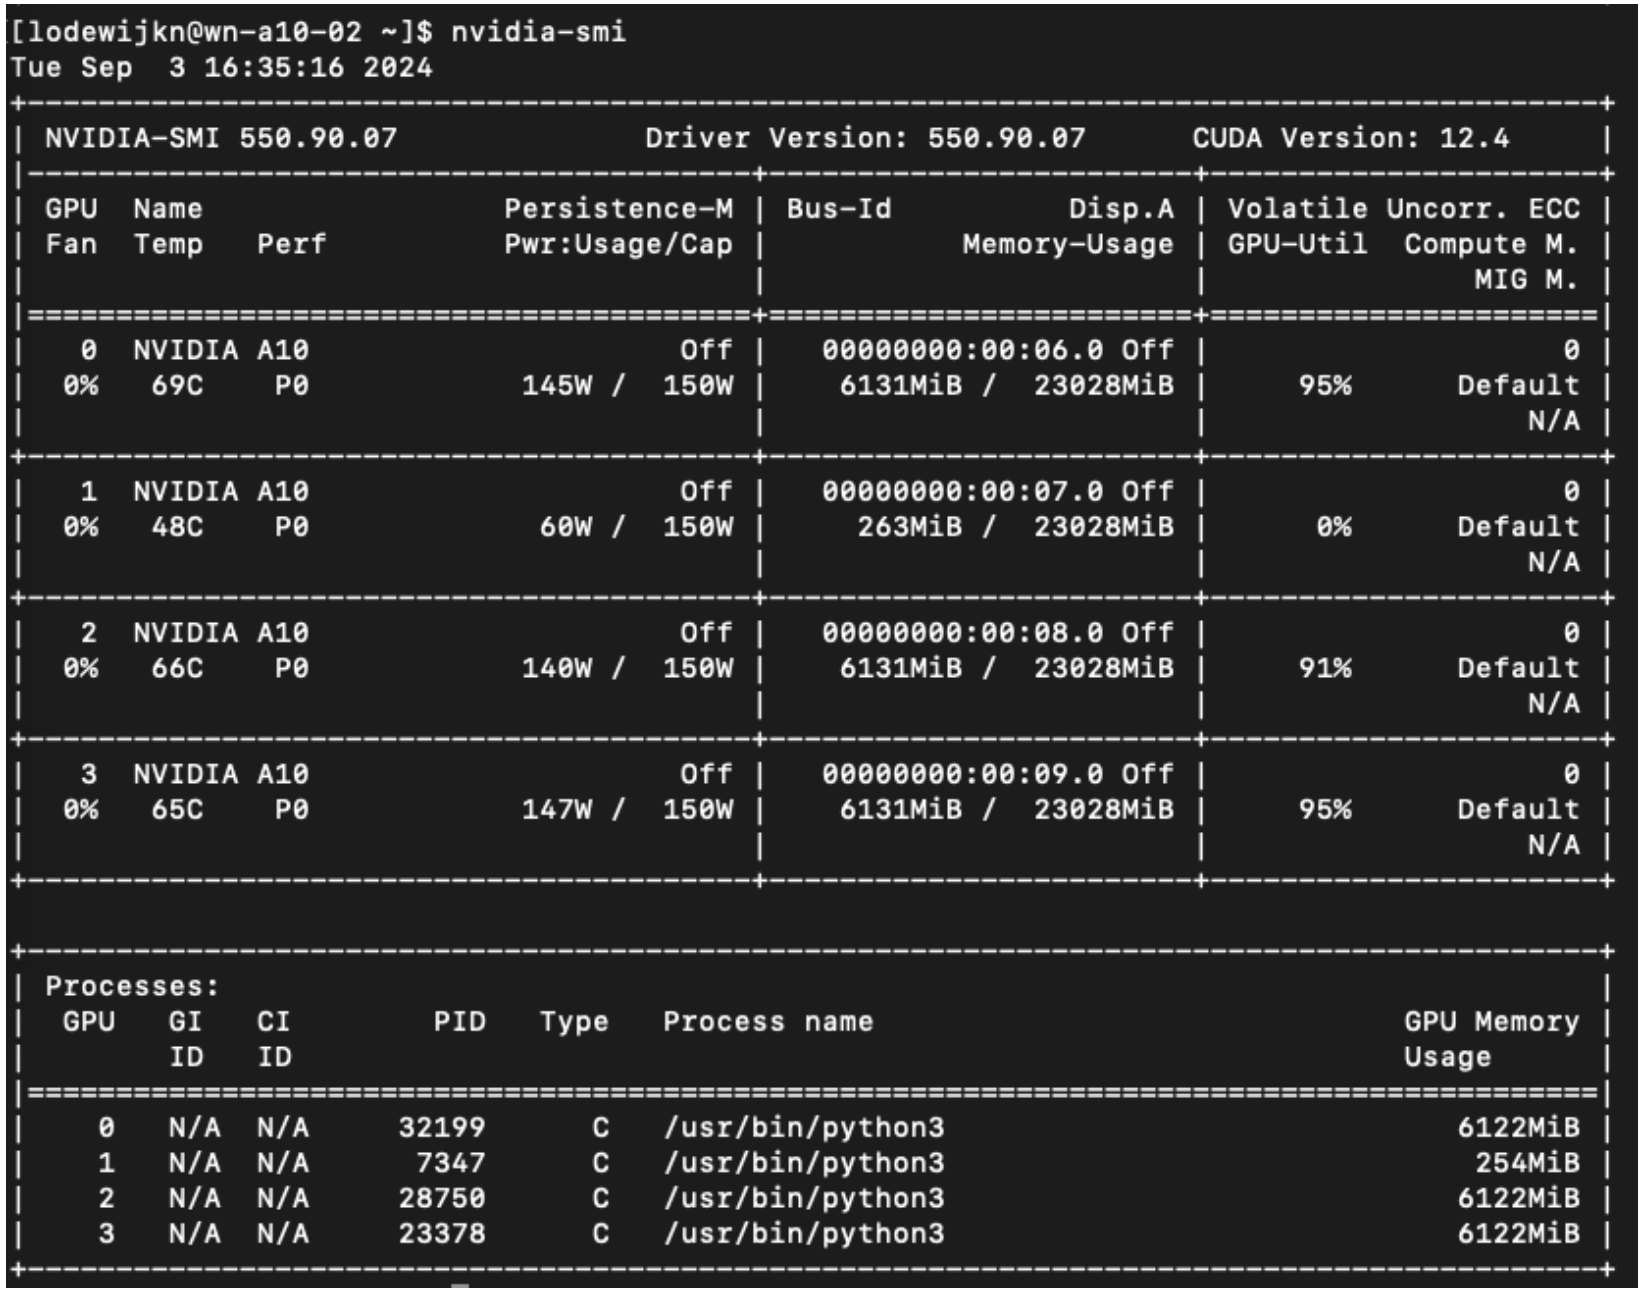

```
[lodewijkn@wn-a10-02 ~]$ nvidia-smi
Tue Sep  3 16:35:16 2024
+-----------------------------------------------------------------------------------------+
| NVIDIA-SMI 550.90.07              Driver Version: 550.90.07      CUDA Version: 12.4     |
|-----------------------------------------+------------------------+----------------------+
| GPU  Name                 Persistence-M | Bus-Id          Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf          Pwr:Usage/Cap |           Memory-Usage | GPU-Util  Compute M. |
|                                         |                        |               MIG M. |
|=========================================+========================+======================|
|   0  NVIDIA A10                     Off |   00000000:00:06.0 Off |                    0 |
|  0%   69C    P0            145W /  150W |    6131MiB /  23028MiB |     95%      Default |
|                                         |                        |                  N/A |
+-----------------------------------------+------------------------+----------------------+
|   1  NVIDIA A10                     Off |   00000000:00:07.0 Off |                    0 |
|  0%   48C    P0             60W /  150W |     263MiB /  23028MiB |      0%      Default |
|                                         |                        |                  N/A |
+-----------------------------------------+------------------------+----------------------+
|   2  NVIDIA A10                     Off |   00000000:00:08.0 Off |                    0 |
|  0%   66C    P0            140W /  150W |    6131MiB /  23028MiB |     91%      Default |
|                                         |                        |                  N/A |
+-----------------------------------------+------------------------+----------------------+
|   3  NVIDIA A10                     Off |   00000000:00:09.0 Off |                    0 |
|  0%   65C    P0            147W /  150W |    6131MiB /  23028MiB |     95%      Default |
|                                         |                        |                  N/A |
+-----------------------------------------+------------------------+----------------------+

+-----------------------------------------------------------------------------------------+
| Processes:                                                                              |
|  GPU   GI   CI        PID   Type   Process name                              GPU Memory |
|        ID   ID                                                               Usage      |
|=========================================================================================|
|    0   N/A  N/A     32199      C   /usr/bin/python3                             6122MiB |
|    1   N/A  N/A      7347      C   /usr/bin/python3                              254MiB |
|    2   N/A  N/A     28750      C   /usr/bin/python3                             6122MiB |
|    3   N/A  N/A     23378      C   /usr/bin/python3                             6122MiB |
+-----------------------------------------------------------------------------------------+
```

GPU utilisation

- GPU utilisation: > 90% (unless starting/stopping)
- GPU memory usage: ~6 GB per job
  - Either optimized for ~6-8GB
  - Or not memory intensive (or both?)
- Power usage: 140-150W

- Job length distribution:
  - 94% of jobs < 3 hours
  - These jobs are 55% of runtime

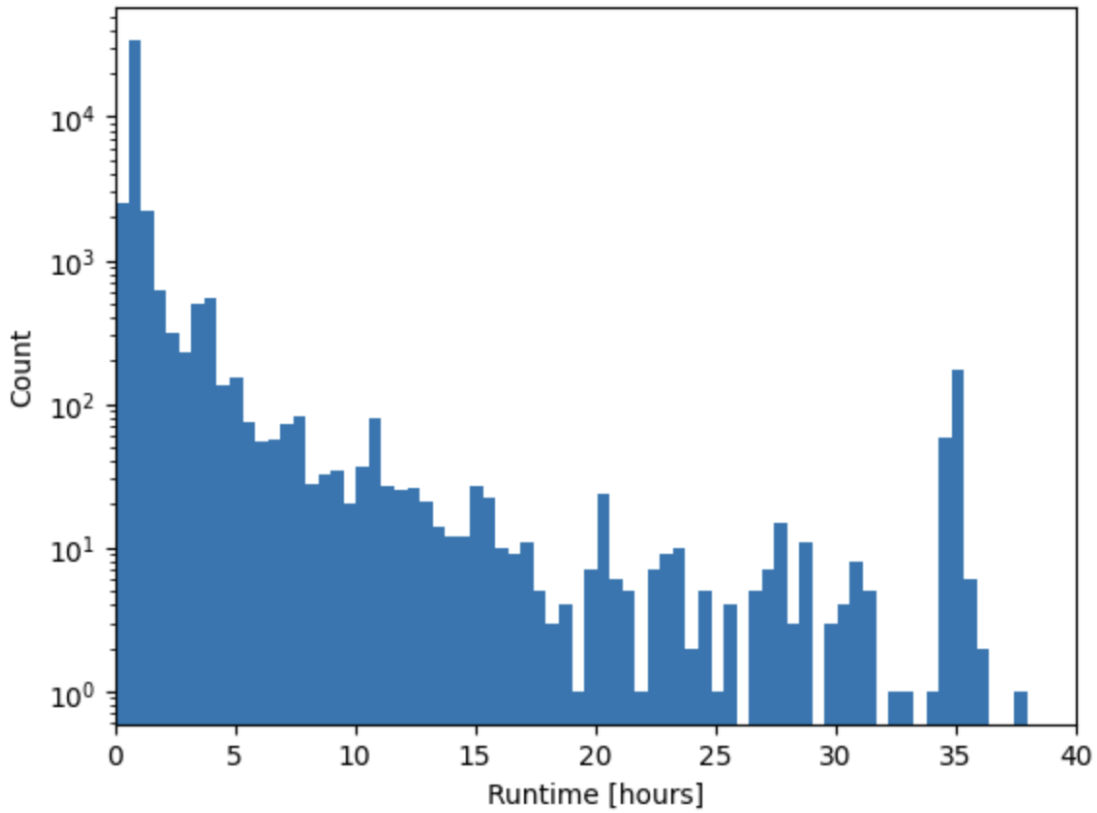

GPU processes

- Do a sneaky: check running processes

```
28750 virgo003  20   0   14.9g   1.6g 315948 R 100.3  0.5  15:25.77 integrate_likel
32199 virgo003  20   0   14.8g   1.5g 315916 R 100.3  0.4  14:13.62 integrate_likel
 7347 virgo003  20   0   14.4g   1.5g 315896 R 100.0  0.4   1:22.91 integrate_likel
 8567 virgo003  20   0 8067660 594188 207628 R 100.0  0.2   0:24.04 integrate_likel
```

- Apptainer binary: taken from CVMFS
- Apptainer image: taken from CVMFS
- Used drivers in container: CUDA 11.x (2022)
- Our drivers: CUDA 12.4 (March 2024)

SURF

User adoption

- We have 1 active user group: VIRGO (gravitational waves)
- Usage is excellent:
  ‘constant’ usage of 16 GPUs, some 100 jobs in queue at peak times

- Capacity: **max** 16*24*31 = 11.904 GPU-hours / month
- Usage: 7.000-10.000 GPU-hours / month

- Due to fair-share and only 1 user group:
  - No competition, can use all GPUs 100% of time in principle
  - Once a second group arrives, fairshare ‘kicks in’: you get less resources
  - No discovery of queues and machines in Grid

SURF

Wrapping up

- GinA slurm cluster contains 16 A10 GPUs in a GPU partition
- DIRAC sends GPU jobs to this partition
- 1 GPU per job

- Performance:
  - Good GPU utilisation: >90% when in use
  - 6GB VRAM in use
  - 140-150W usage
  - However no reliable statistics available

- Accounting for GPUs is Work In Progress (for SURF *and* EGI)

- Excellent usage due to fair-share and no competition!
  NB: This changes when a second group joins the GPU partition