

# Analysis Facilities Monitoring Discussion

Rob Gardner, Liz Sexton-Kennedy

with contributions from Alessandra Forti, Markus Schulz, Thomas Hartmann, Yves Kemp, James Lett, Giordon Stark

WLCG/HSF Workshop @ DESY  
16 May 2024

# Plan for our discussion

- Offer questions, define a scope for today
- Show some examples of existing monitoring of AFs and grid-based analysis
- Open discussion, your ideas

## Analysis Facilities Whitepaper

### 8 Monitoring and Metrics

To provide an extensive overview on how resources are used in order to guide infrastructure development and to allow users to make an informed decision about which infrastructure to use, key parameters should be published for each AF.


Metrics for AFs can be categorized into four areas

1. User experience metrics
2. User trend metrics
3. Performance metrics
4. Facility metrics


# What are the questions, and from whom?

- Users (experience)<sup>1</sup>
  - What resources are available?
  - How long will my jobs wait in queue? Why do they run so slow? Why is my notebook hanging? Why did my last few jobs not finish? Why are my jobs held? Why did my jobs fail? Why are they being held?
  - How do I access my data? Is it local? How do I get X software installed? How do I run with my container?
- Resource providers (trends, performance, facility metrics)
  - What resources (cpu, disk-capacity, disk-fast, network, gpu) are under-provisioned?
  - What are the performance bottlenecks?
  - What are the (unexpressed) requirements?
  - Managing the storage - scratch, precious, freeing up space, group storage
  - Scheduling bursty workflows & precious resources (GPUs, fast storage)
- A fifth category: metrics for **framework & platform developers**
  - Which data formats are physicists most often using and by which frameworks?
  - Are performance targets met? (e.g. X TB / Y minutes)
  - Where are the inefficiencies and user pain points?
  - What capabilities are missing?

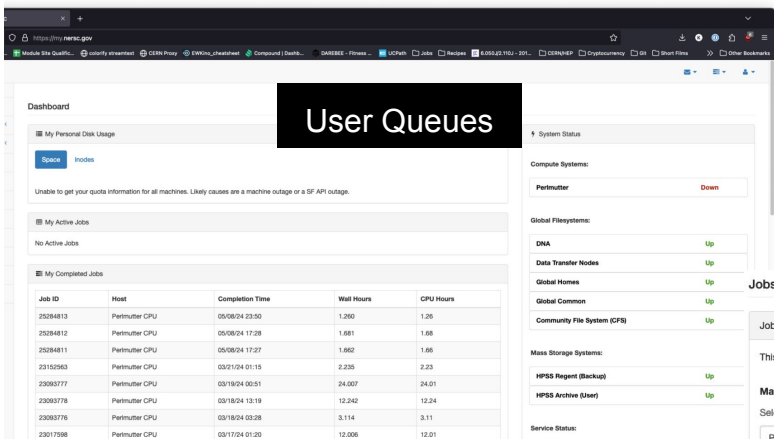
[link](#)




some  
examples



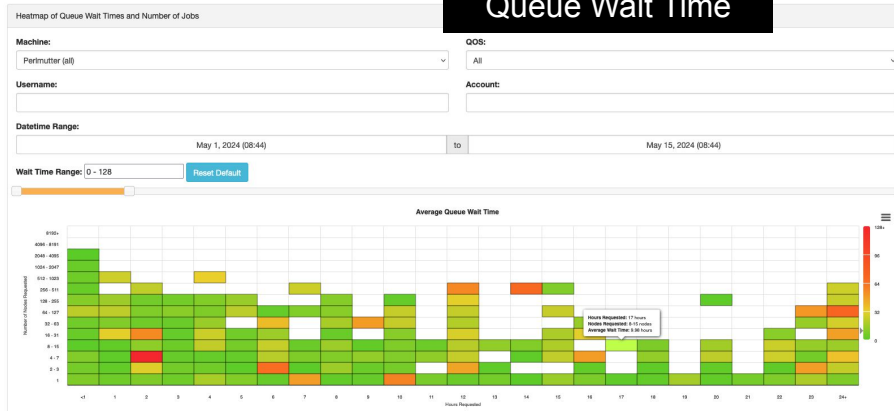
# NERSC (Perlmutter)



## Backlog History



## Queue Wait Times



## Jobscript Generator

### Job Information

This tool generates a batch script template which also realizes specific process and thread binding configurations.

### Machine

Select the machine on which you want to submit your job.

Perlmutter - GPU

### Application Name

Specify your application including the full path.

mctm

### Job Name

Specify a name for your job.

nnlo

### Email Address

Specify your email address to get notified when the job enters a certain state.

### Quality of Service

Select the QoS you request for your job.

regular

### Wallclock Time

Specify the duration of the job. The max walltime for the regular QoS is 24 hours.

10 hours 30 minutes 0 seconds

Number of Nodes

## Job Script Generator

```
#!/bin/bash
#SBATCH -N 128
#SBATCH -C gpu
#SBATCH -G 512
#SBATCH -q regular
#SBATCH -J nnlo
#SBATCH -l 10:30:00
#SBATCH -A m1256


#OpenMP settings:
export OMP_NUM_THREADS=1
export OMP_PLACES=threads
export OMP_PROC_BIND=spread

#run the application:
#Applications may perform better with --gpu-bind=none instead of --gpu-bind=single:1
lsrun -n 512 -c 32 --cpu_bind=cores -G 512 --gpu-bind=single:1 mctm
```




# DESY NAF

aggregate resource delivery

Running job slots on the NAF since 2014




NFS traffic to/from NAF, PiB/month



# NAF analysis resource profiles


Count of records

cpu eff




CpusUsage

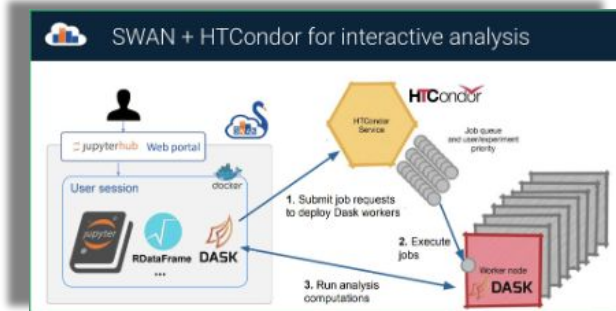
memory usage



cms





# CERN AF-Pilot




- The AF-pilot is linking and expanding existing services at CERN and connecting them to additional resource and load management systems
  - SWAN, CONDOR, Dask, RDataFrame, Coffea, CERN's shared GPUs, EOS...
  - There is quite detailed monitoring for all existing services and the infrastructure...
  - These have and will be combined to create a worthwhile AF specific monitoring
- The CERN AF-pilot is still in its infancy...
  - We have checked that in both systems we can track the use
  - But we haven't come to adding the specific attributes into our probes




# CERN AF-Pilot      Examples: SWAN



# CERN AF-Pilot Ideas for Condor:



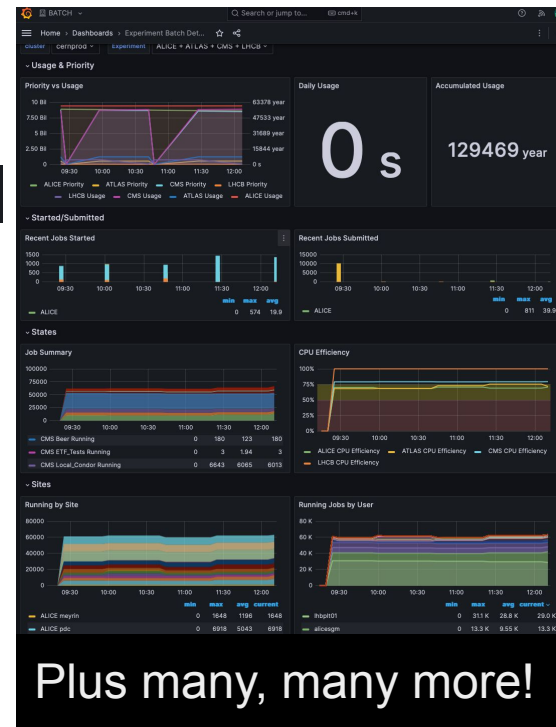
But with very different time ranges  $\ll 1h$




# CERN AF-Pilot Ideas for Condor-2

Home > Dashboards > Experiment Batch Details

- We have detailed and complex monitors for the experiments' activities
  - 50+ diagrams...
- To this we consider to add similar plots covering the “interactive” usage patterns
- The accounting data at CERN can be queried to extract the AF specific resource usage
  - We haven't looked at this in detail
  - There hasn't been a lot of usage yet

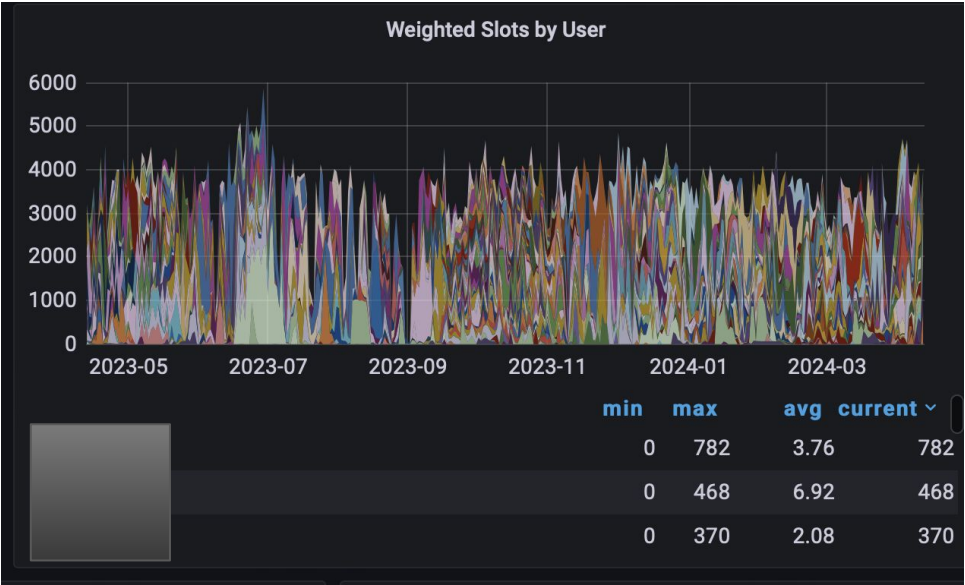


# CMS - grid and local

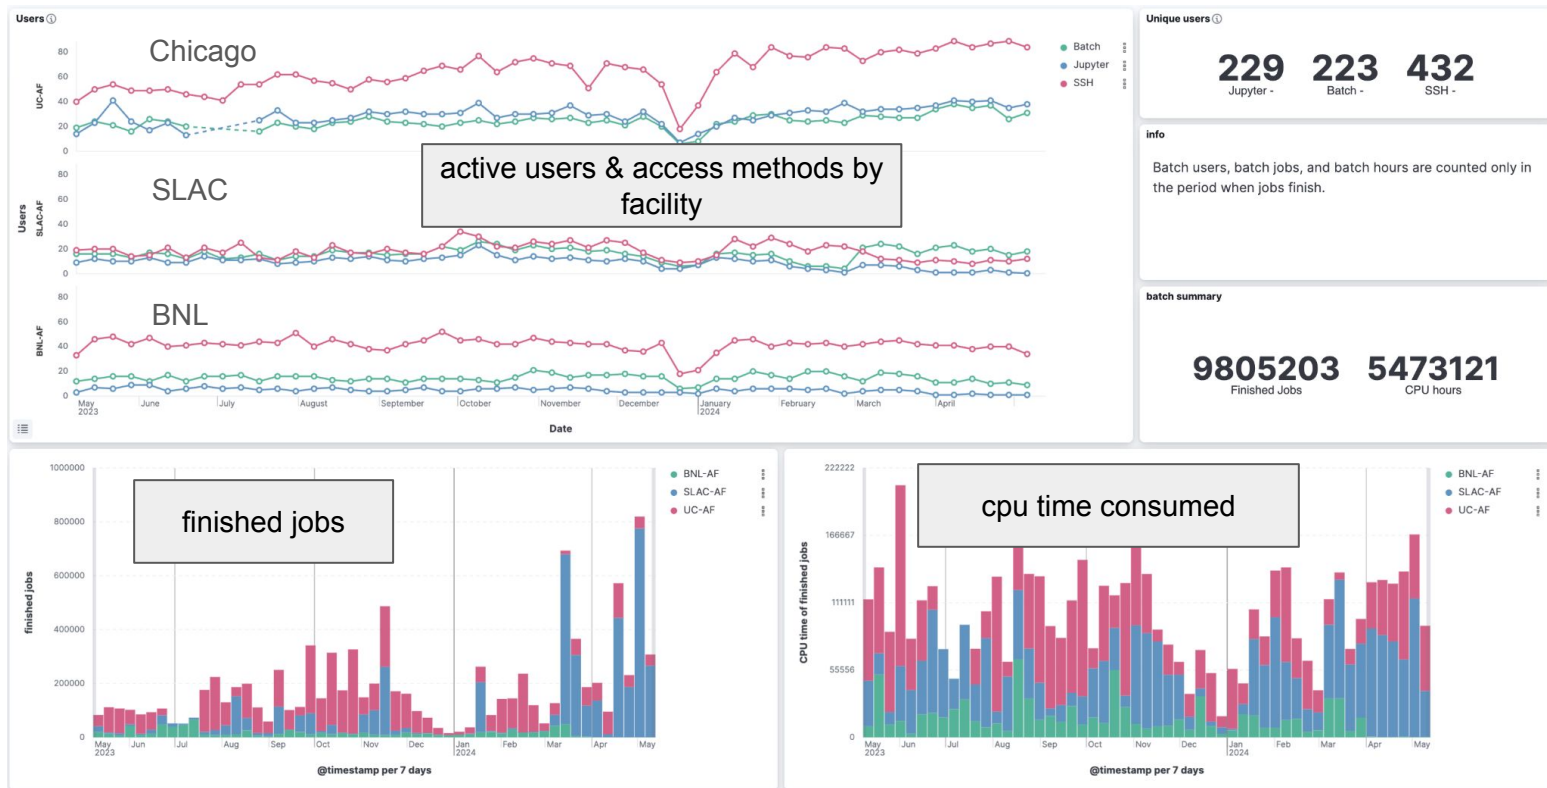


## CMS Users in the Global Pool of WLCG sites

## LPC users monitor over a year



# ATLAS US Shared Tier3s ([link](#))




The chart displays GPU utilization for various physical GPUs over time. The y-axis represents utilization percentage from 0 to 110. The x-axis shows time from 05/13 18:00 to 05/15 16:00. A pie chart inset shows the distribution of GPU types: 11 NVIDIA A100-SXM4-40GB (blue), 8 NVIDIA GeForce RTX 2080 Ti (orange), and 4 NVIDIA GeForce GTX 1080 Ti (green).




# ATLAS grid analysis ([link](#))


MK Analysis jobs per job type (users)




FL Analysis jobs per job type (walltime)




MK Analysis jobs per job type (counts)



MK Analysis jobs per job type (events)



# ATLAS grid analysis ([link](#))






N datasets

# ATLAS


#datasets  
accessed &  
format

Name	Total
DAOD_PHYS	599 K
DAOD_PHYSLITE	67.4 K
Other	370



N Unique Users


unique users per  
day & format




Data volume

data volume  
processed &  
format

Name	Total ~
DAOD_PHYS	619 PB
DAOD_PHYSLITE	221 PB
Other	203 TB



# discussion time!

[link](#)

submit ideas  
here