Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CMS Feedback

Hartmut Stadie, Johannes Lange

10 December 2024

Silides from DESY CMS group

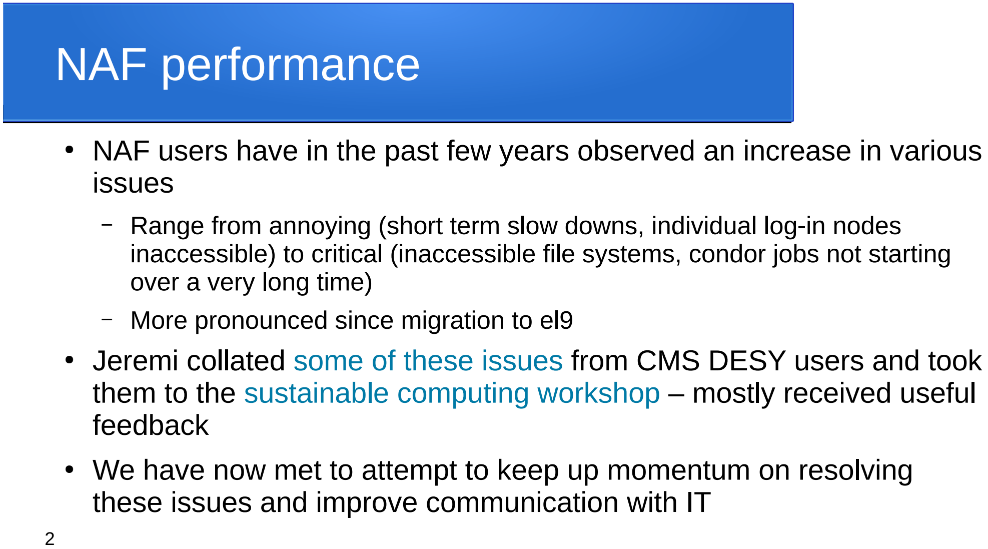

NAF performance

- NAF users have in the past few years observed an increase in various issues
  - Range from annoying (short term slow downs, individual log-in nodes inaccessible) to critical (inaccessible file systems, condor jobs not starting over a very long time)
  - More pronounced since migration to el9
- Jeremi collated some of these issues from CMS DESY users and took them to the sustainable computing workshop – mostly received useful feedback
- We have now met to attempt to keep up momentum on resolving these issues and improve communication with IT

Silides from DESY CMS group

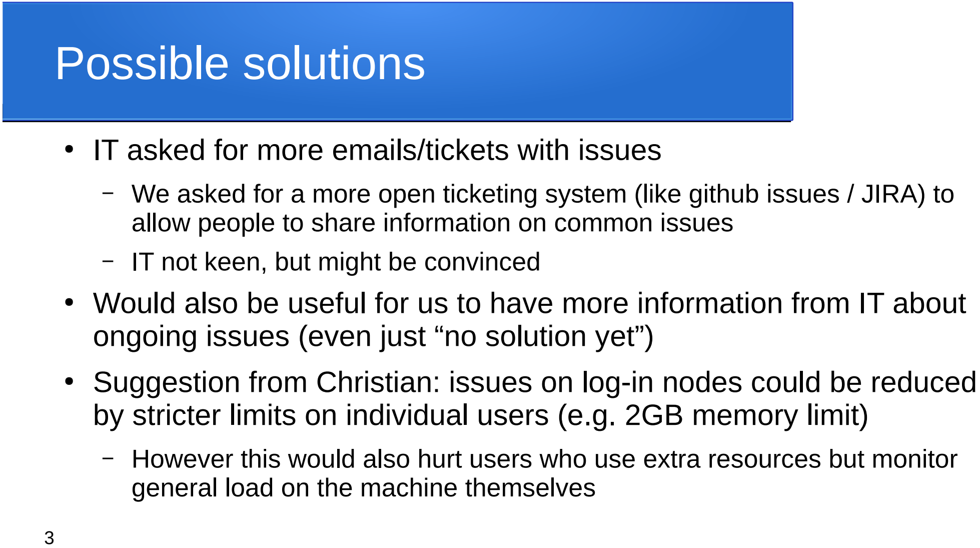

Possible solutions

- IT asked for more emails/tickets with issues
  - We asked for a more open ticketing system (like github issues / JIRA) to allow people to share information on common issues
  - IT not keen, but might be convinced
- Would also be useful for us to have more information from IT about ongoing issues (even just “no solution yet”)
- Suggestion from Christian: issues on log-in nodes could be reduced by stricter limits on individual users (e.g. 2GB memory limit)
  - However this would also hurt users who use extra resources but monitor general load on the machine themselves

3

Silides from DESY CMS group

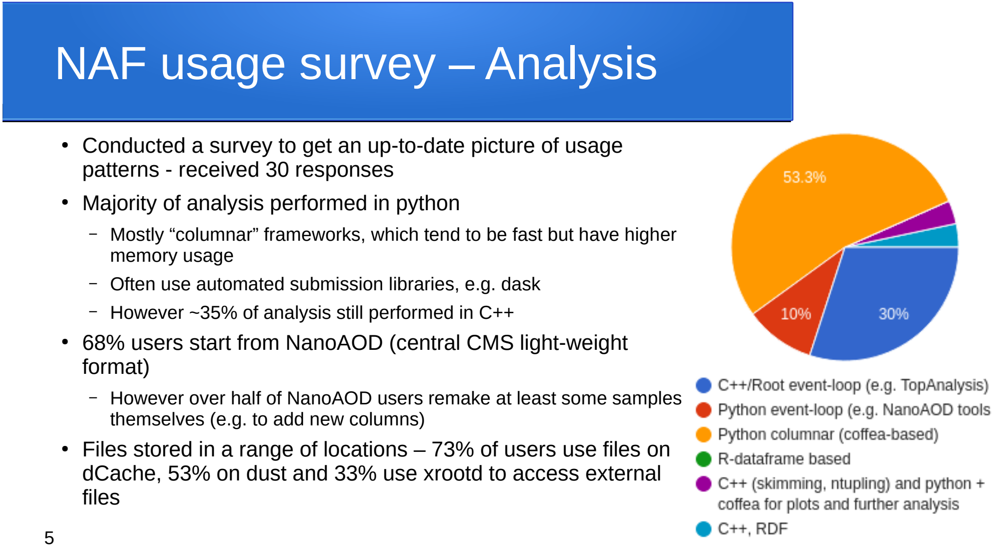

NAF usage survey – Analysis

- Conducted a survey to get an up-to-date picture of usage patterns - received 30 responses
- Majority of analysis performed in python
  - Mostly “columnar” frameworks, which tend to be fast but have higher memory usage
  - Often use automated submission libraries, e.g. dask
  - However ~35% of analysis still performed in C++
- 68% users start from NanoAOD (central CMS light-weight format)
  - However over half of NanoAOD users remake at least some samples themselves (e.g. to add new columns)
- Files stored in a range of locations – 73% of users use files on dCache, 53% on dust and 33% use xrootd to access external files

53.3%
10%
30%
● C++/Root event-loop (e.g. TopAnalysis)
● Python event-loop (e.g. NanoAOD tools
● Python columnar (coffea-based)
● R-dataframe based
● C++ (skimming, ntupling) and python + coffea for plots and further analysis
● C++, RDF

5

Silides from DESY CMS group

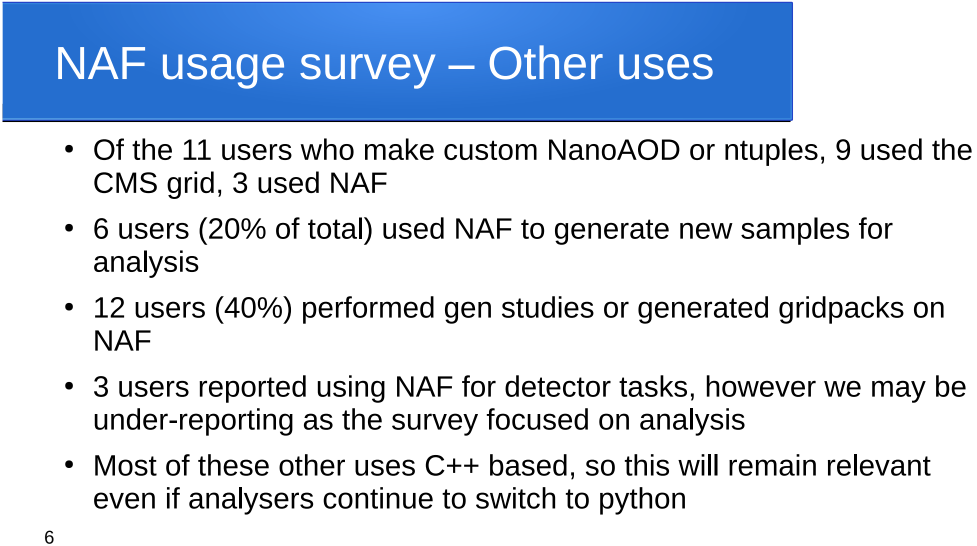

NAF usage survey – Other uses

- Of the 11 users who make custom NanoAOD or ntuples, 9 used the CMS grid, 3 used NAF
- 6 users (20% of total) used NAF to generate new samples for analysis
- 12 users (40%) performed gen studies or generated gridpacks on NAF
- 3 users reported using NAF for detector tasks, however we may be under-reporting as the survey focused on analysis
- Most of these other uses C++ based, so this will remain relevant even if analysers continue to switch to python

6

Silides from DESY CMS group

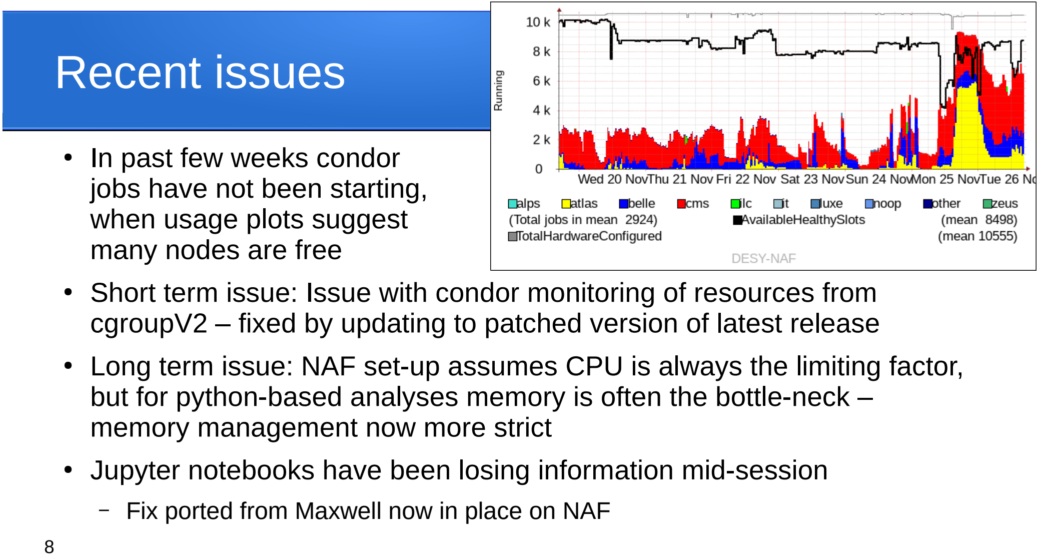

Recent issues

- In past few weeks condor jobs have not been starting, when usage plots suggest many nodes are free
- Short term issue: Issue with condor monitoring of resources from cgroupV2 – fixed by updating to patched version of latest release
- Long term issue: NAF set-up assumes CPU is always the limiting factor, but for python-based analyses memory is often the bottle-neck – memory management now more strict
- Jupyter notebooks have been losing information mid-session
  - Fix ported from Maxwell now in place on NAF

8

Ideas/Suggestions

User support/Feedback loop

- reestablish monthly NUC meetings
- improve handling of problems that affect many users
  problem how to deal with multiple tickets
  - give better feedback to tickets, do not wait until it is fixed if this takes days
  - use DESY-wide known issues or new page
  - gitlab issues?

Change in usage patterns

- much more usage of columnar analysis/python frameworks
- sometimes with many intermediate steps/data
- memory more important than CPU for scheduling
- less production like slimming/skimming of official data sets
- modern analysis tools need more memory