DUNE CPU usage estimates in 2024

Andrew McNab
University of Manchester

DUNE CCB
January 2025

Calculating CPU usage (1)

- During 2024 we switched from the general FIFE HTCondor pool to the dedicated DUNE Global Pool
- We lost access to FIFE usage monitoring at first
- In particular, the dedicated HTCondor schedds used by justIN were not able to log usage to FIFE/Landscape until late in the year
  - These were mostly production workflows
- To calculate 2024 CPU usage, I've taken non-justIN job totals from FIFEMON, and justIN job totals from justIN internal jobs database.
- Both work in wall time CPU.seconds and do not record overheard of pilots jobs
  - **From the point of view of sites these are underestimates by ~minutes per job**

Calculating CPU usage (2)

- Usage is recorded as seconds occupying a single processor jobslot within a Glidein pilot job
  - So a second in an 8 processor job slot would count as 8 jobslot.seconds
  - This is “core.seconds” or “CPU.seconds” or “processor.seconds”
- I worked out per country totals and then added the justIN and non-justIN numbers to get overall per country goals
- All of these figures are from 1st Jan to 31st Dec 2024
- This is the first time I’ve done this and it involves bits of SQL and sed filtering and working with spreadsheets
  - **It needs double checking to be finalised**
  - Hopefully we won’t need to do this again as everything is now in FIFEMON

Per site non-justIN usage

- Total is 1058 CPU.years
- This is all jobs recorded in FIFEMON for DUNE in 2024
  - excluding NERSC
  - but including user work like MARS jobs
- Completely dominated by use of Fermilab

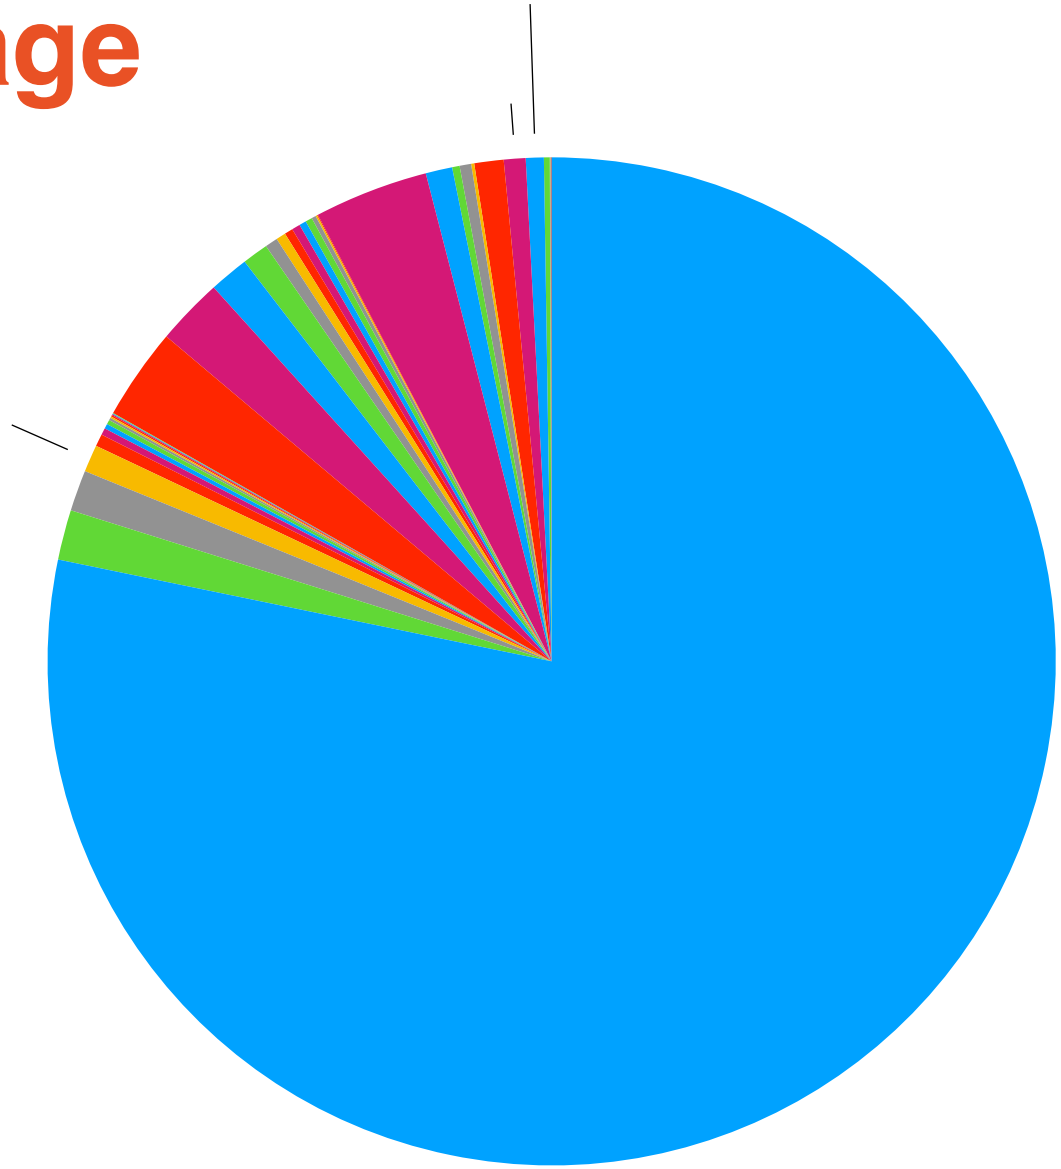

Per country non-justIN usage

- Total is 1058 CPU.years
- This is all jobs recorded in FIFEMON for DUNE in 2024
  - excluding NERSC
  - but including user work like MARS jobs
- Completely dominated by use of Fermilab

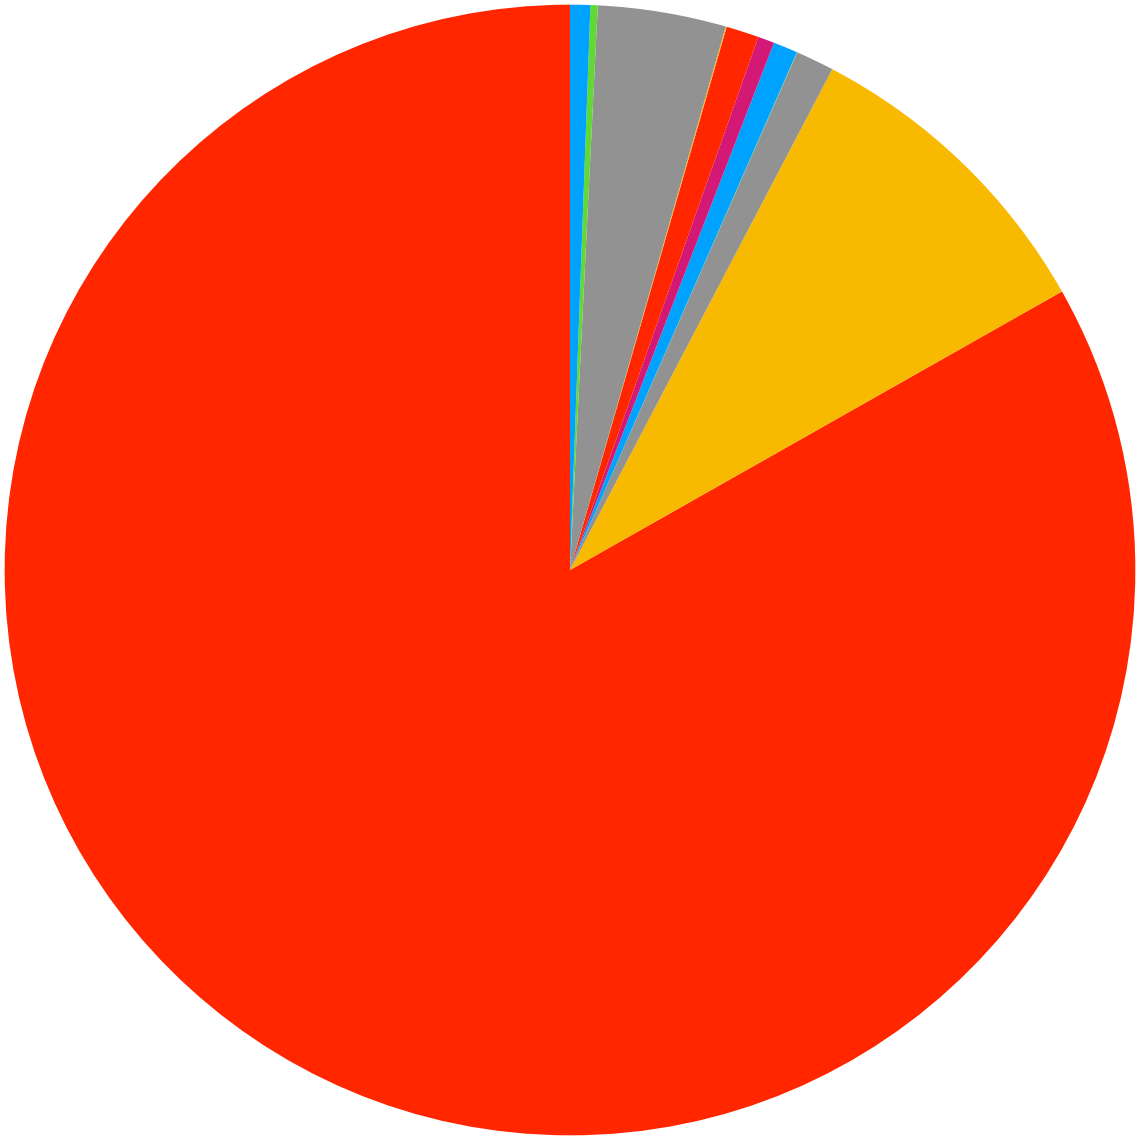

Per site justIN usage

- Total is 962 CPU.years
- Mostly production jobs
- Short Automated Workflow Testing monitoring jobs are excluded
- justIN targets jobs at sites near relevant data which helps pull jobs outside Fermilab

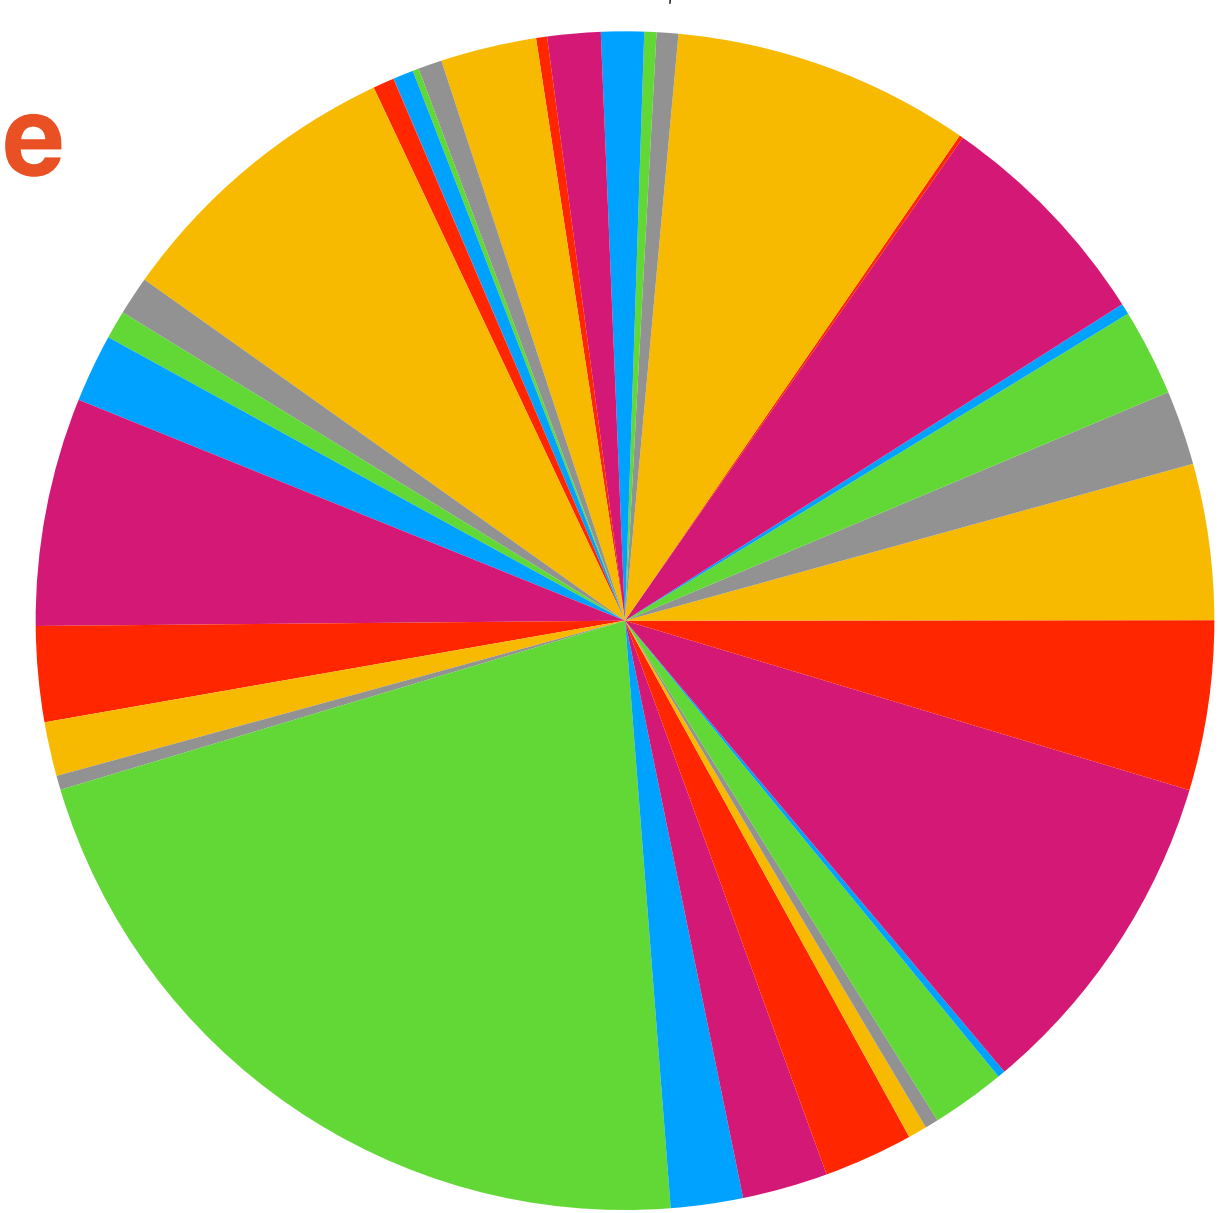

Per country justIN usage

- Total is 962 CPU.years
- Mostly production jobs
- Short Automated Workflow Testing monitoring jobs are excluded
- justIN targets jobs at sites near relevant data which helps pull jobs outside Fermilab

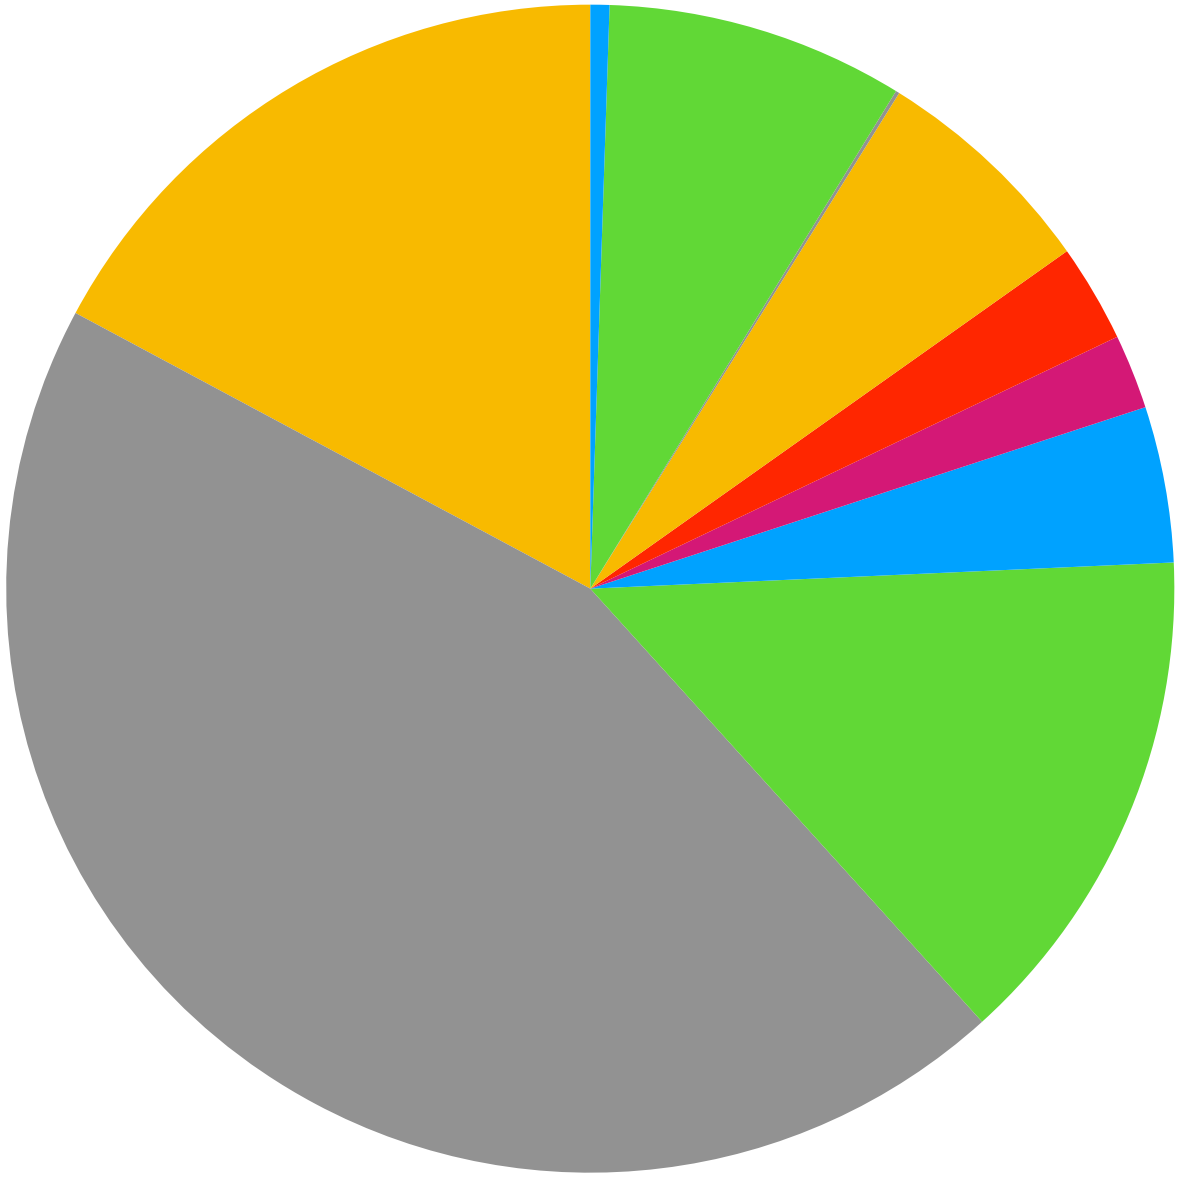

Overall 2024 per-country CPU usage

- Adding up per-site usage
- Assume conversion factor of 11 HS23 per CPU

Total for justIN and non justIN CPU

Country	Non justIN CPU.years	justIN CPU.years	Total CPU.years	Total HS23.years
BR	6.1	5.1	11.1	122.6
CA	2.2	0.0	2.2	24.3
CERN	38.7	79.3	118.0	1298.0
CH	0.3	1.0	1.3	13.9
CZ	9.9	60.8	70.7	778.2
ES	5.0	26.1	31.0	341.3
FR	7.4	19.9	27.3	300.0
IN	0.1	0.0	0.1	1.3
IT	0.0	41.5	41.5	456.5
NL	11.6	134.8	146.4	1609.9
UK	96.6	428.3	524.8	5773.3
US	880.6	165.2	1045.8	11504.3
Totals	**1058.5**	**961.8**	**2020.3**	**22223.6**

Summary

- Production CPU usage has shifted from jobsub and POMS to justIN
  - Slightly more than half is still done without justIN, mostly by individual users running at Fermilab
- We are seeing excellent justIN usage of participating sites around the world rather than just maxing out our share of Fermilab
- We can produce usage figures, charts etc despite lack of unified recording of usage until late 2024
  - Should not be necessary from now on though