




Science and  
Technology  
Facilities Council


# X86 CPU

Tom Birkett



# Introduction

- I manage the Tier-1 Batch Farm.
  - ~50,000 core farm that processes data from the LHC experiments.
- The LHC workflows require High Throughput Computing.
  - Each individual collision within the LHC can be treated separately and processed by a <8 cores and <64GB memory in a few minutes.
  - We have billions of events to process.
- Other projects such as JASMIN, SCARF and the STFC cloud also purchase similar CPU nodes.
  - Where possible we try and combine our CPU procurements.



# Current Hardware

- In 2022, we purchased ~150 1U servers:

CPU	Dual AMD EPYC 7763
RAM	1024 GB
OS Disk	480 GB
Scratch Disk	6.4 TB NVMe with 3DWPD for 5 years Requirement = 25GB capacity and 0.5MB/s write per thread
NIC	25 Gb/s or 100 Gb/s Mellanox NIC's Network infrastructure is Mellanox

- We purchase a 3 year warranty and run the nodes for 5-6 years.
- We purchased AMD EPYC 7763 in 2021 and AMD EPYC 7452 in 2019 and 2020. Before that we had been Intel.

# Future Procurements

- The LHC experiments have switched to using a new benchmark called HEP Score23 (HS23).
  - Replaces the (32 bit) HEP SPEC06 benchmark (HS06).
  - HS23 uses multiple LHC workflows.
- We could make use of Liquid Cooling (direct to Chip) in this years procurement.
- Minimising the carbon footprint of the procurement is becoming a higher priority.
  - We need to monitoring the energy usage of our systems more.
- We are always looking for ways to optimize the amount of work the servers can do.

# STFC Cloud CPU

- The compute farm can "balloon" into the local STFC cloud.
- Along with aforementioned requirements:
  - Disks must be able to be spanned and mirrored using hardware raid or other bios functions (RAID10 equivalent)
  - Network interface adaptors must be able to act as routers without running anything within the OS
    - They must be able to run a network operating system purely on the NIC
    - They must support vxlan offload, ecmp and evpn
    - They must have out of band management access on a separate port on the same nic



Science and  
Technology  
Facilities Council

# Questions?