



# GridPP

UK Computing for Particle Physics



UNIVERSITY  
of  
GLASGOW

# ARM for WLCG

GDB  
CERN  
12<sup>th</sup> June 2024

David Britton  
GridPP Project Leader  
WLCG Deputy Leader



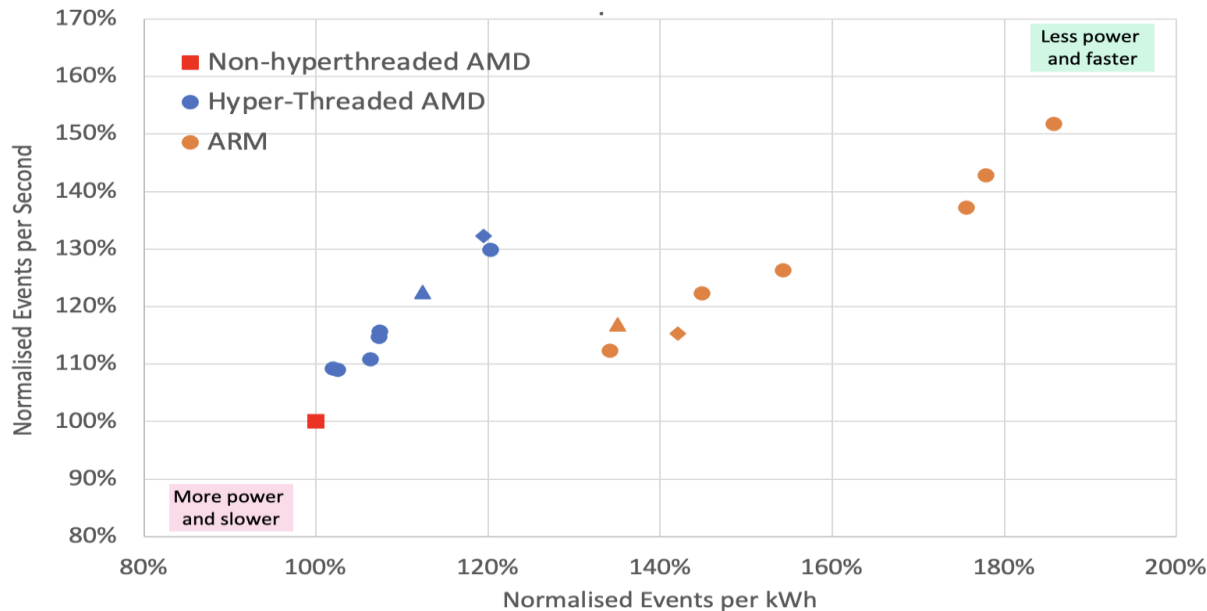
- We've spent 2 years investigating ARM for WLCG.
- Site interest: Three (non-orthogonal) axes:
  - a) Carbon - sites now often have NetZero aspirations/constraints
  - b) Energy - sites may have energy constraints (peak-time or total supply; or a budget cap)
  - c) Cost - sites see benefit from being able to widen the range of procurable hardware (choice, availability, and driving costs with competition).
- Experiment interest: Two axes:
  - i) Carbon - sustainability is an increasingly important agenda item within the experiment collaborations.
  - ii) Resources levels - the ability to use ARM can enable additional resources (opportunistic, pledged, HPC, cloud, Grid).



# Initial Tests (2022)

Compared two *same-price* (but different core-count) machines:  
AMD EPYC 7643 48C/96T @ 2.3GHz TDP 225W, 48-cores (96 HT).  
ARM Q80-30 80 core 210W TDP processor, 80-cores (no HT).

Comparison of 80-core ARM Altra vs 48 (96 H.T.)-core AMD 7003  
Events/Second vs Event/kWh plotted for 8 HEPsScore workloads, normalised to non-H.T. AMD.

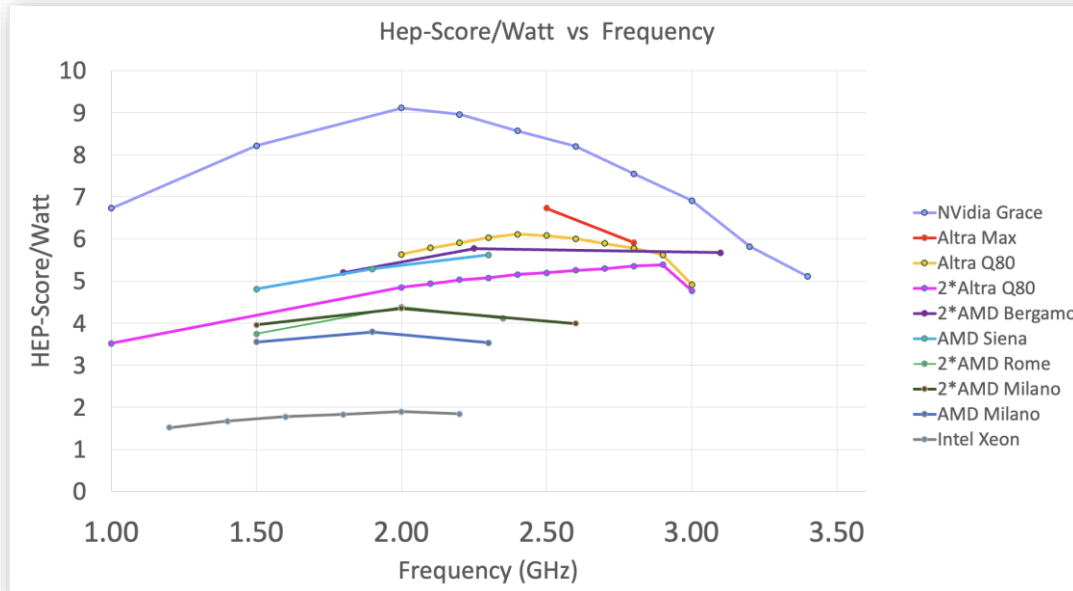
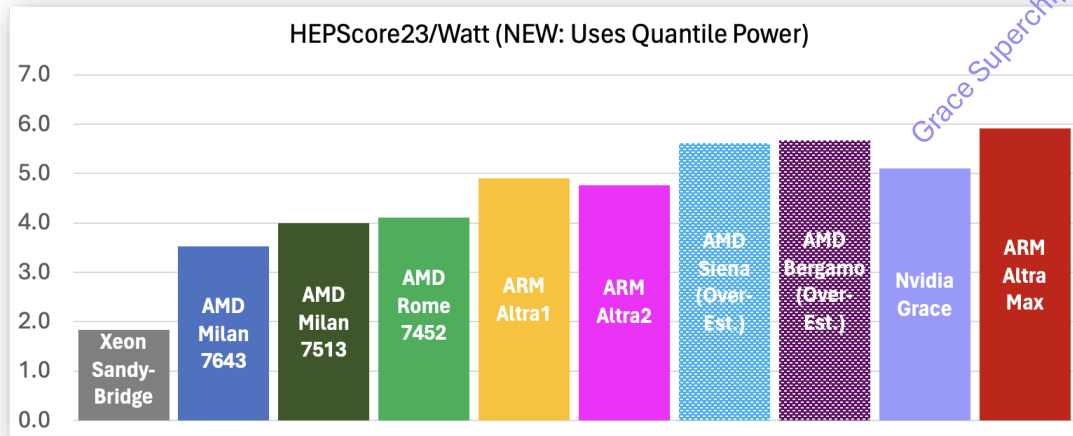


Ran 8 candidate HEPsScore workloads:

- On AMD, HT improves speed (efficiency) by up to 30% (20%) depending on workload.
- On ARM, average speed is a little quicker, but average efficiency is notably better.







# Deeper Studies



- Top plot is a comparison of different boxes running at max frequency (except Siena and Bergamo).
- *But this is not the full story! Efficiency (and, of course, HS) depends on clock speed.*
- Comparison and optimisation is complex!
- Need to consider:
  - Cost
  - Carbon (Scope 2 + 3)
  - Performance (HS23)
  - Efficiency (HS23/Watt)
- Optimisation will not always be the same but it's clear that both AMD and ARM are viable.




# Optimising Procurement



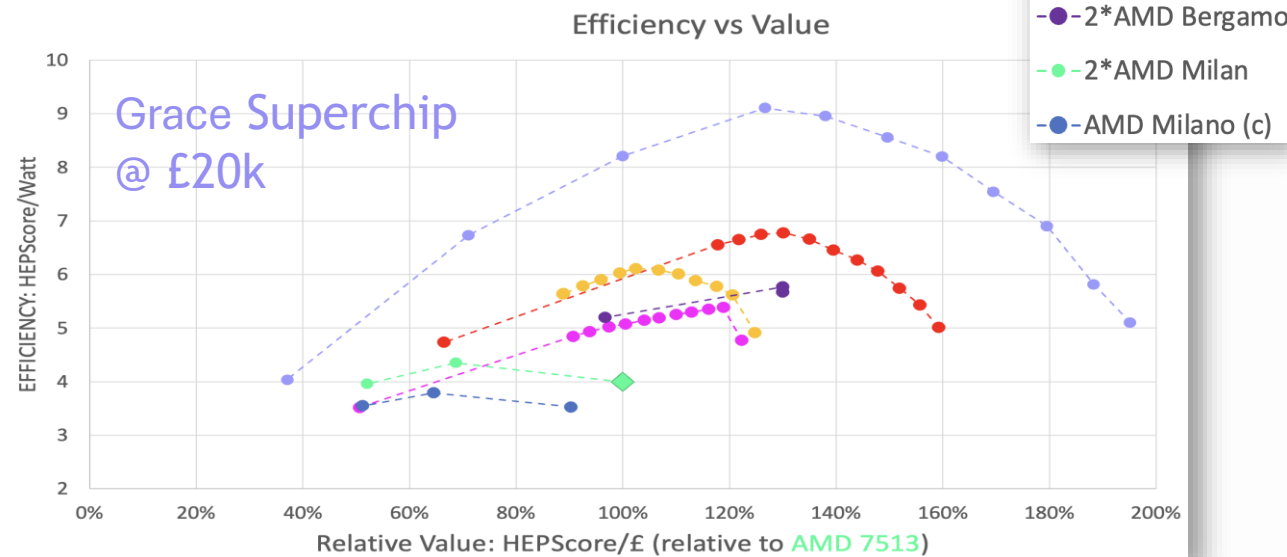
- These curves show effect of running hardware at different frequencies.
- Individual curves move left-right as prices change (frequently), but not up and down.



# Example Procurement



Choose Altra Max?



Choose Grace?



# Where do we go from here?

- To complete this work, we need to agree if / when ARM resources can be pledged:
  - Experiments need to have validated a sufficient number of workloads.
  - Experiments need to be willing (happy?) to do the work necessary to submit work to heterogeneous resources at sites.
  - Sites need to understand/accept the operational implications of providing heterogeneous resources.
  - Sites need to be assured that the resources can be pledged AND will be used.
  - WLCG needs to ensure there are no impediments to pledging, monitoring, or accounting\* of ARM resources. \* BTW, can accounting handle variable frequency (HS)?
  - We all need to recognise that accepting pledges mean at least a 4-5 year commitment to accept ARM resources, maintain software (experiment and things like Geant), even if market moves in other directions and ARM is less interesting.
  - What else?
- Purpose of this GDB session is to try to establish where we are, what still needs to be done, and when we can proceed.