

Discussion of Introduction of HepScore CPU Power units

DUNE CCB Meeting 13 December 2023 at CERN

Introduction of HepScore CPU Power units in DUNE for pledging

- Up to now we produce CPU requests in units of “cores” and ask for pledges in same.
- This was instigated at beginning for ultra-simplicity – as was appropriate at the time.
- Of course, requirements can only be estimated in units of CPU-Work. We knew/know that a “core” is not a well-defined unit of CPU-Power and hence cannot be used to provide a well-defined amount of CPU-Work.
- But we didn’t care as the main aim was to establish international contributions etc... and so getting the exact amount of CPU-Work pledge was not a priority, especially given the history of an evolving computing model and in practice underusing somewhat.
- These slides are to propose that we now introduce the proper units of CPU-Work/Power in parallel, to be used by those who prefer to do so.
- To be totally clear: We propose to retain an average “CPU-Power” → “core” translation for those who still prefer to pledge in cores.

- HEP CPU-Work units are formally now HepScore-23-years(HS23-years).
 - HepScore is the new benchmark suite
 - It is crystalised in a given year. Currently we have HS23
 - By agreement 1 HS23 is normalised to equal 1 HS06 unit
- CPU requirements in LHC experiments are estimated in HS23-years
- In practice, a default “DC use” is built in so that the request and provision is in “Deployed CPU-Power”
 - I.e. a Federation pledges to deploy XXX kHS23 of deployed CPU-Power by a certain date.
 - I.e. a Federation does not pledge to deliver XXX kHS23-years of CPU-Work
 - In HEP this is effectively equivalent as aggregation means data-centres can smooth out non-DC use
- In DUNE we estimate requirements in (in effect) HS23-years units anyway. We have always just converted this to cores using 1 core \sim 11 HS06 \Rightarrow 11 HS23.
- Most funding vehicles understand HS06/HS23
- The proposal is simply to publish both.
 - Sites can then pledge in HS23 if they prefer (typically any WLCG site already pledging to the LHC)
 - Sites who prefer can still pledge cores.
 - DUNE will understand that not all cores are equal,
 - Such sites will know their core power and can pledge an appropriate number.
 - In any case it’s not crucial to be exact at this moment

Comment on Memory Weighting

- This factorises off from this discussion somewhat
- One can as easily have MWHS23 as MWC
- Some of us do not want to pledge MW units as it would mean seriously reducing our pledge
 - So we need the non-MW numbers published as per the previous slide
 - This is easy as the raw calculation is for HS23-years - the MW bit is added after
- But we want to retain MW units as well this year for continuity ?

Summary

- The proposal is to publish in the requirements document
 - kHS23-years of Work needed
 - Its equivalent kHS23 deployed Power needed over 1 year (formally 1:1 equivalent in these units)
 - Its equivalent in cores at 11 HS23 per core (or whatever)
 - For continuity this year : the formal request to sites in kMWS23 and equivalent MWC
- The summary tables should be simple, and apportion the requirement between FNAL, CERN, and Others (as we normally do)
- Countries can then pledge what they wish as a fraction of the “Others” line

Final comment on memory per-core

- DUNE needs more than the canonical 2 GB per core.
- As we finalize our updated framework and production workflows over the next three years, we should have enough experience and knowledge to know if 4 GB is the correct threshold, or if it needs to be adjusted.
- We expect that come 2028, we will have a well-defined memory ask.
- Is it reasonable to specify DUNE PREFERS ≥ 4 GB per-core from this year ?
 - Who could not do this ? I.e. who has only 2 GB per-core machines ?
 - By saying we PREFER we are not excluding.