ARCHER2 Capability Days

Andy Turner, EPCC, The University of Edinburgh

a.turner@epcc.ed.ac.uk

www.archer2.ac.uk

archer2

epcc

Outline

1st Capability Day: 0900 Thu 14 March – 0900 Fri 15 March

- Background
- Description and use
- Tips
- Summary

ARCHER2 Partners

epcc

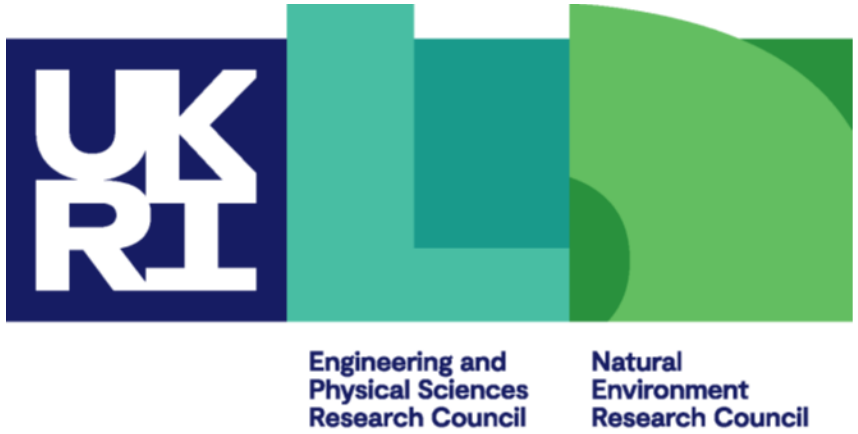

epcc

THE UNIVERSITY of EDINBURGH

Hewlett Packard Enterprise


Capability Day background
archer2
epcc

Motivation

- Enhancing world-leading science from ARCHER2 by enabling modelling and simulation at scales that are not otherwise possible.
- Enabling capability use cases that are not possible on other UK HPC services.
- Providing a facility that can be used to test scaling to help prepare software and communities for future exascale resources.

- All of these rely on enabling users to run very large, capability jobs on ARCHER2

Barriers

- Lack of available budget and storage resources - users have limited budgets and storage quotas and capability jobs use a lot of compute time and often require large amounts of storage resource.
- Ability of software to be used at scale - some modelling and simulation software are not able to be used at scale but development is difficult without access to be able to run at capability scale.
- Potential impact on day-to-day use of ARCHER2 – scheduling a larger number of capability jobs within the standard day-to-day operation of ARCHER2 has the potential to reduce the overall service utilisation efficiency.
- Some technology issues with running at scale – instabilities in the Slingshot interconnect that have made running at scale with consistent and expected performance difficult.
- Lack of user skills/experience of running at capability scale – as for the software development point, this is difficult to address without providing opportunities for users to gain experience by being able to run capability jobs.

Addressing barriers

- Lack of available budget and storage resources
  - Capability Days allows large jobs to run free of charge
  - Availability of scratch solid state storage allows users to go above usual quota
- Potential impact on day-to-day use of ARCHER2
  - Capability Days make dedicated resource available to reduce disruption to usual day-to-day use of ARCHER2
  - Some standard resource remains available during Capability Day period
- Some technology issues with running at scale
  - Situation substantially improved since major upgrade in first half of 2023
  - Large jobs are much more reliable

Capability Day description and use

epcc

Setup

- 0900 Thu 14 March – 0900 Fri 15 March
- 4352 nodes reserved (17 full cabinets) for capability jobs:
  - Should be chosen as a coherent block of nodes.
  - Leaves a maximum of 1508 nodes available during capability day for standard jobs.
- Jobs that run in the capability section will be uncharged.
- Users can queue up work beforehand.
- Maximum of 24h long – ends as soon as all pending work is complete.
  - Cancelled if no pending work 12h before start of session.
- Ends automatically if no work run for 3h period
- Any running jobs at end of period will be killed

Job limits

- Minimum job size: 512 nodes
  - Individual jobs steps (i.e. “srun” commands) within job scripts should also be a minimum of 512 nodes
  - Jobs that do not stick to these limits will be killed
- Maximum walltime: 3 hours
- Job numbers: 8 jobs maximum per user in the QoS
  - 2 jobs maximum running per user
- Users must have a valid, positive CU budget to be able to run jobs during Capability Days

- Submit to “capabilityday” QoS, “standard” partition

Example job submission script

```
#!/bin/bash
#SBATCH --job-name=capability_job
#SBATCH --nodes=1024
#SBATCH --ntasks-per-node=8
#SBATCH --cpus-per-task=16
#SBATCH --time=1:0:0
#SBATCH --partition=standard
#SBATCH --qos=capabilityday
#SBATCH --account=t01

export OMP_NUM_THREADS=16
export OMP_PLACES=cores

export SRUN_CPUS_PER_TASK=$SLURM_CPUS_PER_TASK

module load xthi
srun --hint=multithread --distribution=block:block xthi > placement-${SLURM_JOBID}.out

srun --hint=multithread --distribution=block:block my_app.x
```

Must be 512 or larger

Specify the correct QoS

Propagate cpus-per-task setting to srun

Check process/thread placement


Capability jobs tips
archer2
epcc

Tips

- OFI communications protocol seems to work more reliably at capability scale than UCX protocol
  - UCX often sees memory/timeout errors
- All-to-all collective patterns do not generally scale well to large MPI process counts, particularly when there are high MPI process counts per node
  - c.f. On the Frontier exascale system there are typically a maximum of 8 MPI processes per node (1 per GPU). 9,408 compute nodes gives a maximum of 75,264 MPI processes for a whole system job.
  - 4096 ARCHER2 compute nodes, 1 MPI process per core is 524,488 MPI processes!
- MPI-IO does not generally scale well to high process counts unless the IO pattern is very simple
  - Same for IO libraries based on MPI-IO: parallel HDF5, NetCDF
  - Consider a different parallel IO approach, e.g. ADIOS2
  - Follow IO performance tips: https://docs.archer2.ac.uk/user-guide/io/#achieving-efficient-io
- Make use of the scratch, solid state file system so you do not hit unexpected storage quota issues
- With very high MPI process counts, you may see long MPI startup times, take this into account in wall times in your job scripts

Practical considerations

- Test job submission scripts thoroughly before submitting
- If possible, test scaling (even on very short tests) up to 256/512 nodes before submitting Capability Day jobs
- Consider using the solid-state scratch file system to avoid quota issues
  - Access needs to be requested ahead of time if your project does not already have access
  - Check with “ls /mnt/lustre/a2fs-nvme/work” – if your project ID is there, you already have access
- If possible, submit jobs before Capability Day starts

Dr Sam Azadi

Summary

Capability Day: 14 Mar 2024

- No charge for jobs in “capabilityday” QoS
- Limits:
  - Minimum of 512 compute nodes (for both the job and individual job steps within the job)
  - Maximum walltime of 3 hours
- Use
  - Can submit in advance
  - Use “standard” partition and “capabilityday” QoS
- Plan to run more in future – frequency and format based on feedback from this first Capability Day

https://docs.archer2.ac.uk/user-guide/scheduler/#capability-days

GPU eCSE Programme

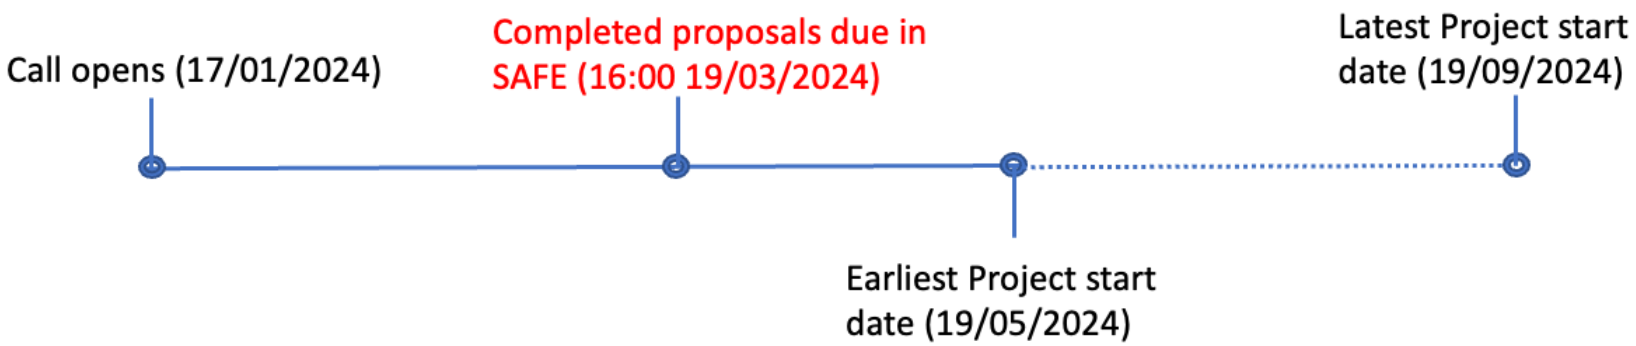

- Programme of GPU eCSE software development calls
  - Expect to run 3 calls this year
  - 1st call open now
- Open to proposals to support research across all of UKRI remit
- Up to 36 person months of effort available per call
  - Max duration 2 years
  - Flexibility of how effort spend (e.g. 1 person 50% for 2 years, 2 people 100% for 18 months, etc.)
  - Funding can be for RSE at PI's institution, RSE at third-party institution, member of ARCHER2 CSE team – or combination of the above

https://www.archer2.ac.uk/ecse/calls/