

CSD3 (Cambridge Service for Data-Driven Discovery)

Stuart Rankin

`sjr20@cam.ac.uk`

Research Computing Services (<http://www.hpc.cam.ac.uk/>)

University Information Services (<http://www.uis.cam.ac.uk/>)

- ▶ **Computer as a research tool** (~ telescope or microscope)
- ▶ Tier 2, multi-institutional research facility for data-intensive simulation and advanced data analytics.
- ▶ Tier 0 = European, Tier 1 = National
- ▶ £9M co-investment of UCAM, STFC, EPSRC, co-designed with Dell EMC, Intel & NVIDIA.
- ▶ Users from UCAM, other UK Universities (STFC & EPSRC) and industry.
- ▶ Production service commenced in November 2017, replaced Darwin service May 2018.
- ▶ Upgraded in November 2018, and Easter 2020.

- ▶ Computer as a research tool (\sim telescope or microscope)
- ▶ Tier 2, multi-institutional research facility for data-intensive simulation and advanced data analytics.
- ▶ Tier 0 = European, Tier 1 = National
 - ▶ £9M co-investment of UCAM, STFC, EPSRC, co-designed with Dell EMC, Intel & NVIDIA.
 - ▶ Users from UCAM, other UK Universities (STFC & EPSRC) and industry.
 - ▶ Production service commenced in November 2017, replaced Darwin service May 2018.
 - ▶ Upgraded in November 2018, and Easter 2020.

- ▶ Computer as a research tool (\sim telescope or microscope)
- ▶ Tier 2, multi-institutional research facility for data-intensive simulation and advanced data analytics.
- ▶ Tier 0 = European, Tier 1 = National
- ▶ £9M co-investment of UCAM, STFC, EPSRC, co-designed with Dell EMC, Intel & NVIDIA.
- ▶ Users from UCAM, other UK Universities (STFC & EPSRC) and industry.
- ▶ Production service commenced in November 2017, replaced Darwin service May 2018.
- ▶ Upgraded in November 2018, and [Easter 2020](#).

▶ Peta4

- * 1152x PowerEdge C6400 (2x Intel Xeon Gold 6142 @2.60GHz, 32 cores) [36864 skylake cores]
- * 342x PowerEdge C6320p (1x Intel Xeon Phi CPU 7210 @ 1.30GHz, 64 cores x4HT) [21888 KNL cores]
- * Intel Omni-Path interconnect (100Gbit/s)

▶ Wilkes2

- * 90x PowerEdge C4130 (1x Intel Xeon CPU E5-2650 v4 @2.20GHz, 12 cores)
+ 4x NVIDIA Tesla P100-PCIE-16GB [360 GPUs]
- * Mellanox EDR IB interconnect (100Gbit/s)

▶ Peta4

- * 1152x PowerEdge C6400 (2x Intel Xeon Gold 6142 @2.60GHz, 32 cores) [36864 skylake cores]
- * 342x PowerEdge C6320p (1x Intel Xeon Phi CPU 7210 @ 1.30GHz, 64 cores x4HT) [21888 KNL cores]
- * Intel Omni-Path interconnect (100Gbit/s)

▶ Wilkes2

- * 90x PowerEdge C4130 (1x Intel Xeon CPU E5-2650 v4 @2.20GHz, 12 cores)
+ 4x NVIDIA Tesla P100-PCIE-16GB [360 GPUs]
- * Mellanox EDR IB interconnect (100Gbit/s)

► Peta4

- * 1152x PowerEdge C6400 (2x Intel Xeon Gold 6142 @2.60GHz, 32 cores) [36864 skylake cores]
- * 342x PowerEdge C6320p (1x Intel Xeon Phi CPU 7210 @ 1.30GHz, 64 cores x4HT) [21888 KNL cores]
- * Intel Omni-Path interconnect (100Gbit/s)

► Wilkes2

- * 90x PowerEdge C4130 (1x Intel Xeon CPU E5-2650 v4 @2.20GHz, 12 cores)
+ 4x NVIDIA Tesla P100-PCIE-16GB [360 GPUs]
- * Mellanox EDR IB interconnect (100Gbit/s)

► Peta4

- * 1152x PowerEdge C6400 (2x Intel Xeon Gold 6142 @2.60GHz, 32 cores) [36864 skylake cores]
- * 342x PowerEdge C6320p (1x Intel Xeon Phi CPU 7210 @ 1.30GHz, 64 cores x4HT) [21888 KNL cores]
- * Intel Omni-Path interconnect (100Gbit/s)

► Wilkes2

- * 90x PowerEdge C4130 (1x Intel Xeon CPU E5-2650 v4 @2.20GHz, 12 cores)
+ 4x NVIDIA Tesla P100-PCIE-16GB [360 GPUs]
- * Mellanox EDR IB interconnect (100Gbit/s)

Showing off a bit . . .

From the June 2019 Top500 list (<http://www.top500.org>):

- ▶ CPU: 2.271 Pflop/s
- ▶ GPU: 1.193 Pflop/s

Currently the top two entries in the list of UK, academic segment machines.

Showing off a bit ...

From the June 2019 Top500 list (<http://www.top500.org>):

- ▶ CPU: 2.271 Pflop/s [107]
- ▶ GPU: 1.193 Pflop/s [368]

Currently the top two entries in the list of UK, academic segment machines.

What is it good for?

- ▶ **Skylake nodes are general purpose:**
 - ▶ Single CPU throughput jobs (R, python).
 - ▶ Shared memory parallel jobs (OpenMP, pthreads).
 - ▶ MPI parallel jobs (C++, C, Fortran).
 - ▶ Hybrid MPI/OpenMP jobs.
 - ▶ Peculiar jobs requiring containers (via [singularity](#)).
- ▶ GPU nodes are specialised for:
 - ▶ Highly vector codes written in CUDA or OpenCL.
 - ▶ Machine learning & AI.
- ▶ KNL nodes are specialised for:
 - ▶ Highly scalable codes.
 - ▶ Machine learning & AI.
- ▶ All nodes have fast connectivity to the RDS and RCS storage and support data intensive workload.
- ▶ 3D visualization nodes for manipulation of data sets.

What is it good for?

- ▶ **Skylake nodes are general purpose:**
 - ▶ Single CPU throughput jobs (R, python).
 - ▶ Shared memory parallel jobs (OpenMP, pthreads).
 - ▶ MPI parallel jobs (C++, C, Fortran).
 - ▶ Hybrid MPI/OpenMP jobs.
 - ▶ Peculiar jobs requiring containers (via [singularity](#)).
- ▶ **GPU nodes are specialised for:**
 - ▶ Highly vector codes written in CUDA or OpenCL.
 - ▶ Machine learning & AI.
- ▶ **KNL nodes are specialised for:**
 - ▶ Highly scalable codes.
 - ▶ Machine learning & AI.
- ▶ All nodes have fast connectivity to the RDS and RCS storage and support data intensive workload.
- ▶ 3D visualization nodes for manipulation of data sets.

What is it good for?

- ▶ **Skylake** nodes are general purpose:
 - ▶ Single CPU throughput jobs (R, python).
 - ▶ Shared memory parallel jobs (OpenMP, pthreads).
 - ▶ MPI parallel jobs (C++, C, Fortran).
 - ▶ Hybrid MPI/OpenMP jobs.
 - ▶ Peculiar jobs requiring containers (via [singularity](#)).
- ▶ **GPU** nodes are specialised for:
 - ▶ Highly vector codes written in CUDA or OpenCL.
 - ▶ Machine learning & AI.
- ▶ **KNL** nodes are specialised for:
 - ▶ Highly scalable codes.
 - ▶ Machine learning & AI.
- ▶ All nodes have fast connectivity to the RDS and RCS storage and support data intensive workload.
- ▶ 3D visualization nodes for manipulation of data sets.

What is it good for?

- ▶ **Skylake** nodes are general purpose:
 - ▶ Single CPU throughput jobs (R, python).
 - ▶ Shared memory parallel jobs (OpenMP, pthreads).
 - ▶ MPI parallel jobs (C++, C, Fortran).
 - ▶ Hybrid MPI/OpenMP jobs.
 - ▶ Peculiar jobs requiring containers (via [singularity](#)).
- ▶ **GPU** nodes are specialised for:
 - ▶ Highly vector codes written in CUDA or OpenCL.
 - ▶ Machine learning & AI.
- ▶ **KNL** nodes are specialised for:
 - ▶ Highly scalable codes.
 - ▶ Machine learning & AI.
- ▶ All nodes have fast connectivity to the RDS and RCS storage and support data intensive workload.
- ▶ 3D visualization nodes for manipulation of data sets.

Who is it good for?

- ▶ Any user with a large data set on RDS to compute with:
Life scientists, bioinformaticians, cryoEM, astronomers, physicists, engineers, image analysts, theoretical linguists, ...
- ▶ Any HPC user:
Physicists, chemists, material scientists, chemical engineers, astronomers, engineers, ...

Generous free usage tier!

Who is it good for?

- ▶ Any user with a large data set on RDS to compute with:
Life scientists, bioinformaticians, cryoEM, astronomers, physicists, engineers, image analysts, theoretical linguists, ...
- ▶ Any HPC user:
Physicists, chemists, material scientists, chemical engineers, astronomers, engineers, ...

Generous free usage tier!

- * ~ 600x 56-core Cascade Lake CPU nodes
- * ~ 360x NVIDIA GPUs
- ▶ 20% of the new resources earmarked for UCAM.
- ▶ Expecting service in Easter Term.
- ▶ Exact timing and size dependent on external factors (Br****).

- * ~ 600x 56-core Cascade Lake CPU nodes
- * ~ 360x NVIDIA GPUs
- ▶ 20% of the new resources earmarked for UCAM.
- ▶ Expecting service in Easter Term.
- ▶ Exact timing and size dependent on external factors (Br****).

- * ~ 600x 56-core Cascade Lake CPU nodes
- * ~ 360x NVIDIA GPUs
- ▶ 20% of the new resources earmarked for UCAM.
- ▶ Expecting service in Easter Term.
- ▶ Exact timing and size dependent on external factors (Br****).