



ARCHER Service 2016 Annual Report



Document Information and Version History

Version:	1.0
Status	Final
Author(s):	Alan Simpson, Anne Whiting, Stephen Booth, Andy Turner, Felipe Popovics, Steve Jordan, Harvey Richardson, Mike Brown, Lorna Smith
Reviewer(s)	Alan Simpson, Lorna Smith, Steve Jordan

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2017-01-05	Inputting initial information	Anne Whiting
0.2	2017-01-06	Added CSE report	Anne Whiting
0.3	2017-01-11	Added graphs	Jo Beech-Brandt
0.4	2017-01-13	Added updated CSE and 2 Cray reports	Anne Whiting
0.5	2017-01-13	Review	Andy Turner
0.6	2017-01-13	Review	Alan Simpson
1.0	2017-01-13	Version for EPSRC	Alan Simpson

Table of Contents

Document Information and Version History	2
1. Introduction.....	4
2. Executive Summary.....	5
3. Service Utilisation	6
3.1 Overall Utilisation.....	6
3.2 Utilisation by Funding Body.....	6
3.3 Additional Usage Graph	7
4. User Support and Liaison (USL)	8
4.1 Helpdesk Metrics	8
4.2 USL Service Highlights.....	8
5. Operations and Systems Group (OSG)	10
5.1 Service Failures	10
5.2 OSG Service Activities	10
6. Computational Science and Engineering (CSE).....	11
7. Cray Service Group	14
7.1 Summary of Performance and Service Enhancements	14
7.2 Reliability and Performance	14
7.3 Service Failures	14
8. Cray Centre of Excellence (CoE)	15

1. Introduction

This annual report covers the period from 1 Jan 2016 to 31 Dec 2016.

The report has contributions from all of the teams responsible for the operation of ARCHER;

- Service Provider (SP) containing both the User Support and Liaison (USL) Team and the Operations and Systems Group (OSG);
- Computational Science and Engineering Team (CSE);
- Cray, including contributions from the Cray Service Group and the Cray Centre of Excellence.

The next section of this report contains an Executive Summary for the year.

Section 3 provides a summary of the service utilisation. Section 4 provides a summary of the year for the USL team, detailing the Helpdesk Metrics and outlining some of the highlights for the year. The OSG report in Section 5 describes their four main areas of responsibility; maintaining day-to-day operational support; planning service enhancements in a near to medium timeframe; planning major service enhancements; and supporting and developing associated services that underpin the main external operational service.

In Section 6 the CSE team describe a number of highlights of the work in 2016. These include the work from the centralised team on parallel I/O performance; the training provided to support the KNL system; the Wee ARCHIE Raspberry Pi Supercomputer at the Big Bang Fair; the ARCHER Champions initiative; and Women in HPC.

In Sections 7 and 8, the Cray Service team and Cray Centre of Excellence give a summary of their year's activities, respectively.

This report and the additional SAFE reports are available to view online at <http://www.archer.ac.uk/about-us/reports/annual/2016.php>

2. Executive Summary


The sections from the various teams describe highlights of their activities. This section gives a brief summary of highlights from the first year of the overall ARCHER service. More details are provided in the appropriate section of the document.

- Work was carried out jointly between SP, CSE, and Cray to deliver the experimental Cray 12-node XC40 KNL system in October 2016. The CSE service created and delivered training courses to support user adoption of the new technology, and SAFE functionality has been introduced to manage and support KNL usage. In the first quarter of use 188 user accounts have been created, 3589 jobs were submitted using 3540 kAUs and the KNL utilisation was 47% for this period.
- Utilisation of the service has remained very high with a mean percentage utilization of 95% for 2016. Whilst this is positive, reflecting the popularity and usage of the service, it has presented challenges to the user community, in particular around job queuing times. The SP Service performed a detailed analysis of queue times that led to adjustments in the job priority formula in the ARCHER scheduling system. Analysis following these changes showed that they made queue times more equitable across different job sizes. There was both a dramatic reduction in the number of jobs that queued for very long times and a balancing of queue times across different job sizes.
- ARCHER has been instrumental in setting up and driving forwards the Women in HPC initiative, and this year has seen many highlights. The most noticeable was perhaps the recognition and involvement of Women in HPC at SC16. On 14 November 2016, WHPC was again recognised in the annual HPCWire Readers' and Editors' Choice Awards, receiving three prestigious awards. Getting recognition in this way highlights the impact that WHPC and the diversity activities put forward under the ARCHER Outreach programme are having.
- A programme of work was delivered by the CSE team to investigate parallel I/O performance on ARCHER and to formulate concrete advice to users and developers on how to measure, understand and optimise the I/O performance of their applications. The results of this work have been documented in the ARCHER Best Practice Guide (<http://www.archer.ac.uk/documentation/best-practice-guide/io.php>), incorporated into training material and will be used to produce a white paper and a webinar in early 2017.
- In total, the Service dealt with 7426 queries during 2016, meeting all query targets. Resolving user queries promptly allows users to maximise the research impact of the service. This level of support is only possible due to close and effective collaboration between all service partners.
- Responses received to the ARCHER Service annual user survey for 2015 were very positive, with the mean satisfaction score for the service of 4.3 out of 5. The highest rated aspect of the ARCHER service continues to be the helpdesk with a mean score of over 4.5 out of 5.

3. Service Utilisation


3.1 Overall Utilisation

Utilisation over the year was 94%, up from 87% in 2015.



3.2 Utilisation by Funding Body


The utilisation by funding body relative to their allocation can be seen below.



This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER.

3.3 Additional Usage Graph

The following heatmap provides a view of the distribution of job sizes on ARCHER in 2016.



The heatmap shows that most of the kAUs are spent on jobs between 192 cores and 12,288 cores (8 to 512 nodes). The number of kAUs used is closely related to money and shows how the investment in the system is utilised.

4. User Support and Liaison (USL)

4.1 Helpdesk Metrics

Query Closure

It was a busy year on the helpdesk but all Service level agreements were met. A total of 7426 queries were answered by the Service Provider, and over 99.4% were resolved within 2 days. In addition to this, the Service Provider passed on 222 in-depth queries to CSE and Cray.

	15Q1	15Q2	15Q3	16Q4	TOTAL
Self-Service Admin	1722	1172	775	1693	5288
Admin	654	616	408	497	1869
Technical	118	91	67	83	269
Total Queries	2494	1879	1250	2273	7426

Other Queries

In addition to the Admin and Technical Queries detailed above, the Helpdesk also dealt with Phone queries, Change Requests, internal requests and User Registration.

	16Q1	16Q2	16Q3	16Q4	TOTAL
Phone Calls Received	82 (25)	81 (21)	56 (17)	80 (16)	299 (79)
Change Requests	2	10	4	7	23
User Registration Requests	338	264	264	218	1084

The numbers shown in brackets for the phone calls received are the calls resulting in new or updated queries. It is worth noting that the volume of telephone calls was low throughout the year. Of the 299 calls received in total, only 79 (26%) were actual ARCHER user calls that resulted in queries. All phone calls were answered within 2 minutes, as required.

4.2 USL Service Highlights

Work on Scheduler Prioritisation Formula

The mean system utilisation for ARCHER in 2016 was 95%. Whilst this reflects the extensive use made of the service, and supports the case for the future investment in HPC, it also presents challenges to the user community, primarily around the queuing times for smaller jobs. The SP Service performed a detailed analysis of queue times that led, after consultation with users and the Research Councils, to adjustments in the job priority formula in the ARCHER scheduling system. Analysis following these changes showed that they made queue times more equitable across different job sizes. There was both a dramatic reduction in the number of jobs that queue for very long times and a balancing of queue times across different job sizes. The volume of users registering their concerns with queuing times has greatly reduced since this change has been made.

KNL Support

User support has been extended to provide support to users of the KNL system including additional SAFE functionality and user documentation and assistance videos.

SAFE Improvements

Extensive work has been carried out to improve both the usability and functionality of SAFE this year. Changes include:

- A new version of SAFE was rolled out providing an improved and more user-friendly interface Updated training material was produced which includes documentation and training videos. The new version has received positive feedback from the user community and has been rolled out across the DIRAC and EPCC SAFE instances to provide a consistent user experience.
- SAFE functionality has been introduced to allow users to register publications, making it easier for them to record publications in Researchfish. This functionality will also provide a useful input to the service benefits realisation data collected.
- New functionality has been introduced to the SAFE to facilitate the easy sign up of users to applications such as VASP and CASTEP. It provides the option of linking automating access to particular licensed software for particular user groups e.g. all new NCAS users could be given automatic access to the Unified Model when their user account is set up.
- New functionality to support sign-up and management of KNL resources.

Improved reporting – scheduling coefficient

Functionality was added to the live status page on the website with the scheduling coefficient matrix and usage matrix for various periods to allow users to plan their use of ARCHER more effectively. Historic data for these two plots has also been made available. These updates have been well received by the user community.

ISO 9001:2015 certification

Preparatory work was carried out across the year to improve and consolidate processes and documentation primarily to introduce a consistent framework to improve the service provided to our users but also to achieve ISO 9001:2015 certification. Benefits are already being seen in process improvement and the first stage external audit was successfully carried out in December 2016. The full external audit is taking place in February 2017.

5. Operations and Systems Group (OSG)

5.1 Service Failures

There were no SEV1 Service Failures in the period as defined in the metric.

5.2 OSG Service Activities

Principal activities undertaken (in addition to day-to-day operational cover) included:

- (1) Operating system and applications software support:
 - a. Planning for CLE 5.2 UP04 upgrade on the XC30;
 - b. Installing regular compiler and programming development upgrades;
 - c. Supporting OS enhancements to external login nodes.
- (2) Login nodes issue:
 - a. Substantial investigation and testing associated with serious issue regarding multiple recurrent failures of login nodes
- (3) KNL installation
 - a. Close cooperation with Cray and the CSE service on installation of KNL test and development platform
- (4) ISO9001:
 - a. Assessment and analysis of appropriate operational processes undertaken for ISO9001 accreditation
- (5) System administration:
 - a. Further development and expansion of automated ticket handling;
 - b. Refinement of locally-developed systems administration tools;
 - c. Increase in the short queue hours from 0900 – 1700 Monday to Friday to be 0800 – 2000 Monday to Friday;
 - d. Modification of maintenance schedule with the approval of the Research Councils, reducing the number of full maintenance days to one a month to minimise user impact;
- (6) Supporting Cray hardware operations:
 - a. Providing additional on-site support for Cray personnel during major hardware upgrade operations (such as the optical cable re-work).
- (7) Security:
 - a. Implementing enhancements to security monitoring;
 - b. Installing Cray-supplied security field notices;
 - c. Providing additional hardening of security measures – specific details are not available for obvious reasons.
 - d. Successful external security audit undertaken with no outstanding issues as a result
- (8) Outreach:
 - a. Attendance at two UKCSF meetings: - Daresbury (March) and ECMWF (September)
 - b. Establishment of regular 3-way meeting with EPSRC and Cray to discuss operational issues

6. Computational Science and Engineering (CSE)

Parallel I/O Performance Studies

At the start of 2016, in collaboration with user groups and the other Service partners, the CSE service identified a number of priority areas to invest technical effort from the centralised CSE team. One of the key areas identified was gaining a better practical understanding of parallel I/O performance on ARCHER and how it compares to other systems. In general, the HPC community often poorly understand parallel I/O performance in terms of:

- what “good” performance actually is on a particular file system;
- what a particular application does;
- what benchmarks illustrate.

We designed a programme of work for the CSE team for 2016 to address these problems by providing useful data on parallel I/O performance on ARCHER, and concrete advice to users and developers on how to measure, understand and optimise the I/O performance of their applications.

This work has led to a number of positive impacts for the ARCHER community (and the wider HPC community):

- Understanding of what information can be gained from different synthetic parallel I/O benchmarks and how well they model real HPC applications.
- Quantifying the maximum performance available from the ARCHER (and other) parallel file systems in production. This gives users useful values to compare current performance to, to assess if the applications are doing well or badly in terms of performance.
- Understanding the performance and scaling characteristics of different common parallel I/O patterns. This allows users to assess what the best approach to take is for I/O in their application and allows ARCHER service partners to configure ARCHER file systems to best meet user requirements.
- Producing statistical data on the variability of parallel I/O performance on ARCHER to aid users in assessing if their performance variations are within normal bounds or not and to aid service partners in best configuring ARCHER.

The data, analysis, and conclusions from the studies so far are being disseminated to ARCHER users and the wider HPC community through several mechanisms to ensure the work has the largest possible impact:

- Production of a new I/O chapter in the ARCHER Best Practice Guide containing practical advice and reference I/O performance data:
<http://www.archer.ac.uk/documentation/best-practice-guide/io.php>
- Production of an ARCHER white paper comparing parallel I/O performance across different systems in the UK. This is currently being reviewed and will be published in Q1 2017.
- Incorporation into ARCHER Training materials.
- Parallel I/O webinar in Q1 2017.

ARCHER KNL Training

The 12-node Knights Landing (KNL) manycore system was launched towards the end of October 2016, and a major challenge was to provide sufficient training to new users so that they could readily make use of it. The KNL processor itself is very new, but its integration into the standard Cray environment is even more recent. We worked closely with the EPCC Intel Parallel Computing Centre and the Cray Centre of Excellence, both of whom had early access to KNL

development systems, so that we could provide a wide range of training to users as soon as the system was open for user service.

The training delivered included:

- A virtual tutorial (i.e. interactive online webinar) on 14 September titled “The Intel Knights Landing Processor” to introduce users to basic KNL concepts and to advertise the upcoming eCSE call that included KNL development projects.
- A virtual tutorial on 21 September explaining how to apply to the upcoming eCSE call, with a focus on applications to use the KNL system.
- A virtual tutorial on 12 October titled “Using KNL on ARCHER” explaining how to use the KNL system in practice, including information on the local details such as compilers and queue structures on ARCHER.
- A 1-day hands-on course “Using Knights Landing Manycore Processors on ARCHER” that allowed users to run real programs on the ARCHER KNL nodes.

We have also launched a new driving test for the ARCHER KNL system which is backed up by training material at <http://www.archer.ac.uk/documentation/knl-guide/knl-training-resources.php>. This contains videos of the two technical KNL virtual tutorials described above, a new video describing how to request KNL access via SAFE, and links to all the material (slides and practical exercises) from the hands-on course.

We will be continuing the KNL training programme into 2017:

- The 3-day hands-on “Cray Optimization Workshop: ARCHER and Knights Landing”, being run in collaboration with Cray at their Bristol offices in late January, will integrate use of the ARCHER KNL system into existing advanced material on Cray hardware, environment, compilers and tools.
- The 2-day course “Programming the Manycore Knights Landing Processor”, to be run in London in 1Q2017, will extend the existing 1-day hands-on course to provide a complete introduction to developing and optimising parallel codes for the ARCHER KNL system.

Despite the challenging timescales, we have developed and delivered a wide range of traditional and online training on the ARCHER KNL system allowing users to make the best use of the unique opportunity provided by access to this system.

Wee Archie, Wee Archlet and the Big Bang Fair

Wee Archie is a suitcase-sized supercomputer, designed to let school children try their hand at computing and learn about the benefits of supercomputing. The Raspberry Pi 2 system has been created to be representative of the system design in massively parallel architectures. Each Raspberry Pi has an LED display that lights up when in use, providing a visual display that helps demonstrate how multiple processors work in parallel to solve complex tasks.

Developed in 2015, a key highlight of 2016 has been using Wee Archie as a tool to educate the next generation of HPC users. Wee Archie was the centrepiece of our booth at the Big Bang Fair. The Big Bang Fair is the UK’s largest STEM event, with 70,000 people attending over 4 days. Our booth had around 6000 people take part and Wee Archie was a considerable draw. Young people could use Wee Archie and see how Wee Archie works. This was the largest event we have participated in, by some margin, and was months in the planning. The event ensured we could showcase ARCHER to the next generation of scientists from across the UK. Wee Archie has proved so popular that we have built a second one to cope with demand.

2017 will see Wee Archie returning to the Big Bang Fair, as well as various other events. Coupled with this, we will be introducing Wee Archlet. This is the younger sibling of Wee Archie, an even smaller Raspberry Pi cluster. Wee Archlet is designed to be cheap and easy to build while still demonstrating the key concepts of parallel computing. On-line instructions will be available for download by schools and community groups wanting to build and configure a system themselves.

ARCHER Champions

2016 saw the creation and development of the ARCHER Champions (<http://www.archer.ac.uk/community/champions/>), a peer support network between staff members whose role involves advising users on access to local, regional and national HPC resources. The aim is to help to promote a coherent access structure to HPC resources across the UK, with coordination between tiers. There is also a focus on supporting and promoting activities designed to provide career development to research software engineers seeking a career in HPC. Activities are also designed to broaden the UK HPC user base to new disciplines and communities. There have been two successful ARCHER Champions Workshops this year, one in Edinburgh in March, and one in Oxford in September. Both events were very well attended and developed the ARCHER Champions network. 2017 will see the next workshop in Leeds, associated with the HPC-SIG meeting.

Women in HPC Recognition

ARCHER has been instrumental in setting up and driving forwards the Women in HPC initiative and this year has seen many highlights. The most noticeable is perhaps the recognition and involvement of Women in HPC at SC16. On 14 November 2016, WHPC was recognised once again in the annual HPCWire Readers' and Editors' Choice Awards, receiving the following honours: Readers' Choice: Workforce Diversity Leadership Award; Editors' Choice: Workforce Diversity Leadership Award and Readers' Choice: Outstanding Leadership in HPC, for Toni Collis. The annual HPCwire Readers' and Editors' Choice Awards are determined through a nomination and voting process with the global HPCwire community, as well as selections from the HPCwire editors. The awards are an annual feature of the publication and constitute prestigious recognition from the HPC community. This is the second year that WHPC has been honored to receive the Reader's Choice Workforce Diversity Leadership Award. Receiving recognition from HPCwire's readers highlights the impact that WHPC and the diversity activities put forward under the ARCHER Outreach programme are having.

7. Cray Service Group

7.1 Summary of Performance and Service Enhancements

2016 was an excellent year for the ARCHER service with no unscheduled system outages caused by technology failures. The excellent level of reliability in conjunction with high utilisation of system resources has provided users with an efficient and stable national service for computational science.

A small Cray XC40 KNL Testbed system was installed in October 2016. This system was provided to enable early access to Intel Phi Knights Landing technology for the ARCHER user community within a familiar programming environment framework.

7.2 Reliability and Performance

Where aspects of technology provision have performed below the high standards expected, Cray has worked to resolve issues with a minimum of disruption to users. Occasionally, complex problems of an intermittent nature can be difficult to identify and resolve. In such cases Cray's regional and global support teams, Engineering group and service partners work together to investigate and implement solutions.

The most significant technology areas of the ARCHER service where issues were encountered in 2016 were:

- Memory fragmentation on job-launch service nodes. A workaround to prevent this issue impacting users was implemented in October 2016.
- Intermittent periods of instability affecting external login and pre/post-processing nodes. There have been two separate causes of instability on external login and pre/post-processing nodes in 2016:
 - A GPFS client related issue which was resolved in August 2016.
 - A lustre related issue first seen in November 2016 which will require a patch to resolve.

7.3 Service Failures

There were no unscheduled incidents classified as full service failures in 2016.

8. Cray Centre of Excellence (CoE)

At the start of the year the Cray Centres of Excellence moved into a new organisation, the Cray EMEA Research Lab (CERL) under the leadership of Adrian Tate, which made additional expertise available for CoE activities. Further information about the CERL is available on the Cray [blog post](#) and [announcement](#). We also started the year with various discussions to decide on the right focus areas for the CoE, in particular there was a wish to focus on projects that would be beneficial beyond individual applications and to impact communities rather than specific research groups.

Karthee Sivalingam joined the CoE during the year bringing expertise in elementary particle physics and Numerical Weather Prediction (NWP) and Climate applications.

Longer-Term Projects

I/O performance

The CoE is leading a broad investigation into I/O performance and optimisation that touches several areas. Several communities, in particular the UK NWP community have displayed an interest in I/O optimisation toolsets and abstraction layers. We are currently investigating how existing I/O technologies developed at Cray can be leveraged by these communities. As part of this broad investigation, we have also been experimenting with the ADIOS implementation as a platform to both provide a wide set of storage APIs for parallel I/O and as a platform for future development. Additionally, we have developed a tool that provides the user with a way to optimize data movement to provide many small files to an application and avoid possible contention in Lustre. This tool has been shown to be beneficial for OpenFOAM applications at scale. The tool was announced to users but there has been limited interest so we will attempt to contact users individually.

AUTO-TUNING

The aim of this project is to determine the usefulness of a simple-to-use auto-tuning tool. This builds on previous work undertaken by Cray under the EU CRESTA Exascale research project. There are two main aspects to this project – gaining more experience with applications both to determine the usefulness of the current mock-up implementation and then some effort to make improvements. We updated the documentation and created a summary presentation. Both EPCC and the CoE have started on the first part of the project with EPCC looking at using the tool to optimize configuration of VASP. The CoE has started a discussion with NCAS and hopes to apply the auto-tuning to the UM.

ONETEP

We were previously spending a small effort (via other UK Application Staff) supporting a Poisson-Boltzmann Equation solver for ONETEP (in collaboration with the University of Southampton). This effort is now continuing as part of a project funded under eCSE-07 which started in June. The most recent work has involved design, implementation and testing work for transferring the higher order correction from ONETEP to the DL_MG multigrid solver.

Filesystem and I/O Issues

From time to time some users have reported concerns over I/O performance. Occasionally this proves to be due to a system problem, but more often the filesystem is simply busy. Improving and understanding I/O performance is a focus area for the CoE and we have been working on I/O performance benchmarking (as has EPCC); in particular the project mentioned above which we hope will deliver some new options for users to help optimize I/O. We are developing an experimental framework that will address many of the concerns of the ARCHER community, while being accessible to users in as transparent a way as possible and also looks forward to future capabilities, for example storage-class memory devices and SSDs.

Training and Workshops

The CoE assisted with various workshops during the year. Examples were the Porting and Optimization workshop, run around the time of CUG 2016, and the ARCHER Advanced OpenMP course run at Cray's EMEA HQ in Bristol in August.

CoE staff presented an [ARCHER webinar](#) outlining new features in recent updates of the Programming Environment on ARCHER.

The CoE was able to engage with ARCHER users at various events including the Computing Insight UK meeting in Manchester, the ARCHER Champions meeting in Oxford, the ECMWF HPC workshop, and events held by the EPSRC Centre for Doctoral Training in Pervasive Parallelism. Cray sponsored the EuroMPI event which this year was held in Edinburgh,

We have started planning for an Optimization Workshop specifically covering the newly installed ARCHER KNL system to be held at Cray's EMEA HQ from 31st January to 2 February 2017.

The KNL Addition to ARCHER

The CoE was heavily involved with the introduction of the new Intel Knights Landing (KNL) system to ARCHER. While the system was being installed the CoE provided guidance on the KNL hardware modes and helped EPCC and the ACF team decide on an appropriate configuration for start of service. The CoE also ran some sanity checks and helped resolve teething issues as the system went live. The KNL talk given by the CoE at the ARCHER Champions meeting in September was augmented and given to EPCC along with other information to assist with preparation of the user webinars and documentation. The CoE worked with EPCC on the ARCHER-KNL web documentation and the CoE provided a tool and associated advice to assist users with proper binding of hybrid applications (also relevant to ARCHER).

Case Studies and ARCHER Promotion

The following Cray Application Brief was published:

Removing Bottlenecks to Large-Scale Genetic and Genomic Data Analysis with DISSECT and the Cray® XC™ Supercomputer ([pdf](#)), which showcases work of the [Roslin Institute](#) using the [DISSECT](#) genomic and epidemiologic analysis tool on ARCHER.

ARCHER Queries and Software

The CoE helps resolve a range of issues that come in from users via the helpdesk, some of which require significant effort and need interaction with Cray R&D experts.

Of particular note was an issue relating to a research project which was utilising "Director's Time" on ARCHER. The project is trying to integrate software that uses DMAPP and uGNI and this has raised some subtle problems with memory registration. We were able to explain how to set up the software correctly to enable coexistence of both APIs.

Late in the year there was a weekend when users complained of slow filesystem performance and we had to work with the systems team and one user in order to obtain a detailed system profile of an application that was significantly contributing to the load on the filesystem. As a result we noticed that NEMO can cause stress on the filesystem by opening the same files many times per process. This is still under investigation.

eCSE panel meetings

The CoE completed technical assessments and final reviews for the three eCSE calls during the year and staff attended the project planning meetings. Advice was also provided in advance on technical concerns over projects prior to panel meetings.

2017

For 2017 the Head of CERL, Adrian Tate, is keen to ensure that all the Cray CoEs continue to provide excellent service and results to customers, in particular making use of the unique skills of individual staff in the team as well as wider Cray resources. The overall aim is to assist ambitious users in the effective use of all the ARCHER resources for world-class science and research. As part of this continual review we are likely to seek further guidance from EPSRC and from key members of the ARCHER community.