ARCHER Service
2015 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	2016-01-13	Inputting initial information	Anne Whiting, Jo Beech-Brandt, Andy Turner, Mike Brown, Lorna Smith
0.2	2016-01-20	Incorporating individual team input	Anne Whiting, Jo Beech-Brandt, Andy Turner, Mike Brown, Lorna Smith
0.3	2016-01-22	Incorporating Cray CofE and service reports	Steve Jordan, Felipe Popovics, Harvey Richardson
0.4	2016-01-22	Updated highlights, reviewed and commented	Andy Turner
0.5	2016-01-22	Added utilization graphs	Jo Beech-Brandt
0.6	2016-01-26	Internal review	Alan Simpson
0.7	2016-01-27	Updates post review	Anne Whiting, Harvey Richardson, Andy Turner, Mike Brown
1.0	2016-01-28	Version sent to EPSRC	Alan Simpson, Anne Whiting
1.1	2016-05-05	Minor change for website	Jo Beech-Brandt, 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
6.1 Best Practice for Data Management on ARCHER 11
6.2 The ARCHER Driving Test: encouraging new users on to the ARCHER service 11
6.3 Wee ARCHIE: a Raspberry Pi cluster to educate the next generation of HPC users 12
6.4 Women in HPC 13
6.5 Competitive eCSE Programme 13
7. Cray Service Group 15
7.1 Summary of Performance and Service Enhancements 15
7.2 Reliability and Performance 15
7.3 Service Failures 15
8. Cray Centre of Excellence (CoE) 17
8.1 CoE Project Highlights 17
8.2 Filesystem and I/O 17
8.3 Training and Workshops 18
8.4 ARCHER Queries and Software 18
8.5 eCSE Meetings 18

1. Introduction

This annual report covers the period from 1 Jan 2015 to 31 Dec 2015.

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 highlight some of their key projects from the year. They describe the work with the Consortia Contacts, the eCSE Programme, Women in HPC and the distributed training activity. The ARCHER Image Competition is also described.

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/2015.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.

- Broadening Access to HPC: Increasing the applicability of, and broadening access to, advanced computing for UK research is of key strategic importance. The ARCHER service has piloted a number of initiatives aimed particularly at increasing the diversity of research on the system. The ARCHER Driving Test provides a simple way for researchers new to HPC to gain initial training in the field and a modest amount of HPC resource to explore the possibilities of HPC in their research. The eCSE programme has included software development effort prioritising research communities that are new to HPC to boost their software to a level where it can exploit facilities such as ARCHER.

- Business Case for Future Systems: The service has worked closely with the research councils to gather information to support the business case for future HPC investment. The SAFE system has provided an invaluable resource allowing us to analyse how ARCHER (and previous systems) is used and by whom.

- ARCHER Outreach Project: This project was funded by EPSRC in 2015 to promote engagement and diversity in UK HPC, demonstrate impact from ARCHER, and enhance outreach activities. In *engagement and diversity,* 2015 has seen the initiation of the ARCHER Champions HPC peer support initiative, expansion of the Women in HPC network, and development of a Faces of HPC diversity website. In *impact*, a number of case studies have been developed and published. In *outreach*: Wee ARCHIE has been developed, a portable HPC cluster designed to promote the national HPC service to young people and the general public at a series of UK-wide outreach activities.

- Major Incident Management: Technical and management staff from all service partners collaborated effectively to resolve the issues arising from the Lustre file system issues, and provided successful and innovative solutions to minimise the impact on users and their work. Many of these solutions have been incorporated as ongoing service delivery improvements providing a more robust service for the future. All service partners are participating in further initiatives to make further improvements to coordination and information sharing as well as improving the joint Major Incident Plan.

- Utilisation of the service has remained very high and has grown steadily through 2015. Although this has been a challenging year for the service due to issues with the Lustre file systems, positive collaboration between all service partners has minimised the impact on the users, maintaining a high utilisation level. The majority of the the compute cycles have been expended on jobs exploiting hundreds or thousands of cores, which are difficult to run on smaller HPC systems..

- In total, the Service dealt with around 8,100 queries during 2015, meeting all query targets. Resolving user queries promptly so that the resolution allows users to maximise their research on the service is only possible due to close and effective collaboration between all service partners.

3. Service Utilisation

3.1 Overall Utilisation

Utilisation over the year was 87% which is similar to the percentage utilisation for 2014> However, following the Phase 2 upgrade, which took place in late 2014, the capacity of ARCHER was increased by 60%.

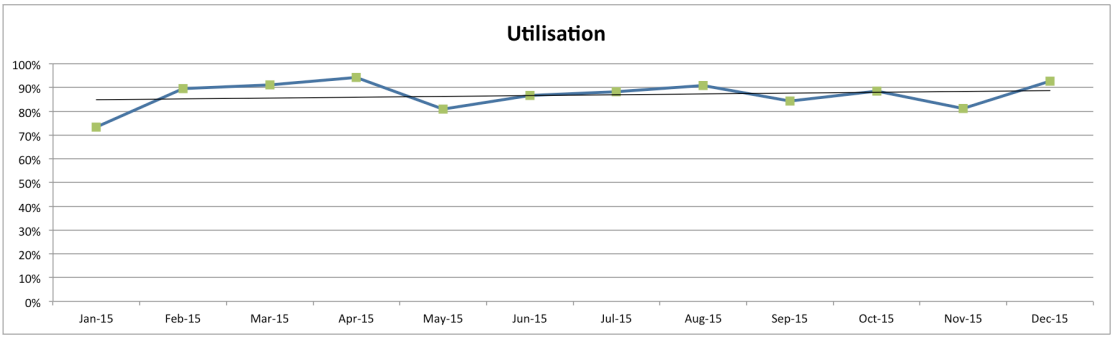

3.2 Utilisation by Funding Body

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

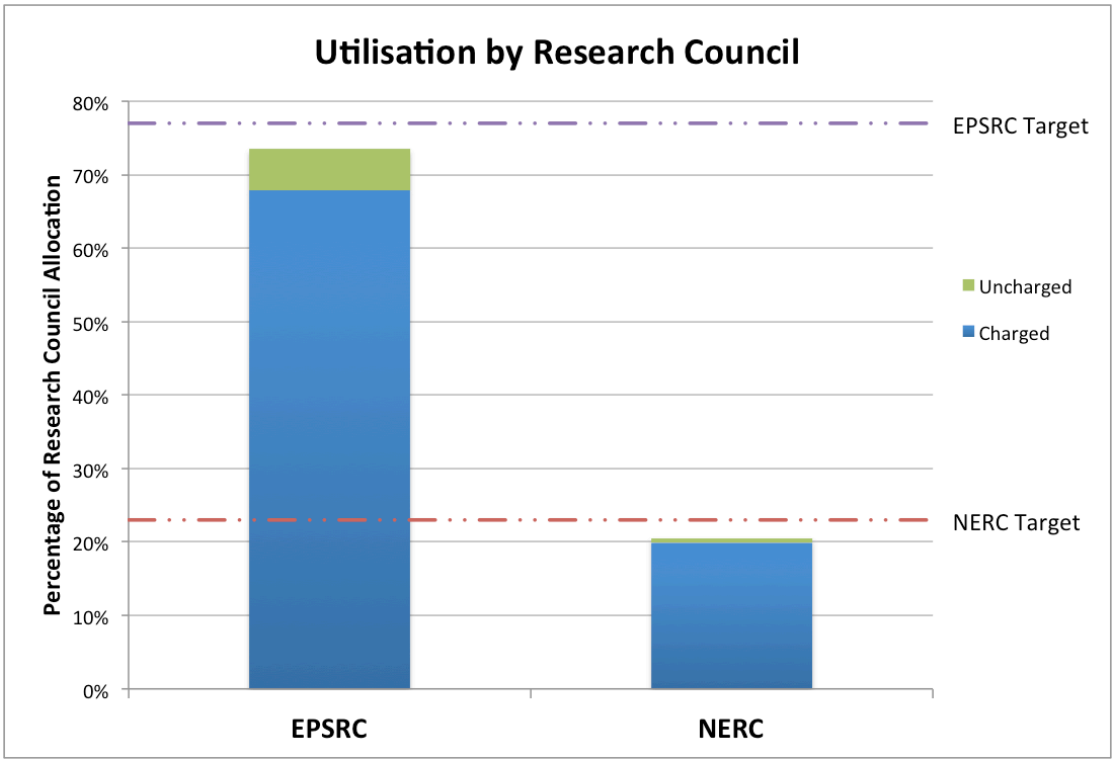

This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. The uncharged proportion for EPSRC includes the temporary project v01 that was put in place during the filesystem issues.

3.3 Additional Usage Graph

The following graph provides a view of the distribution of job sizes on ARCHER.

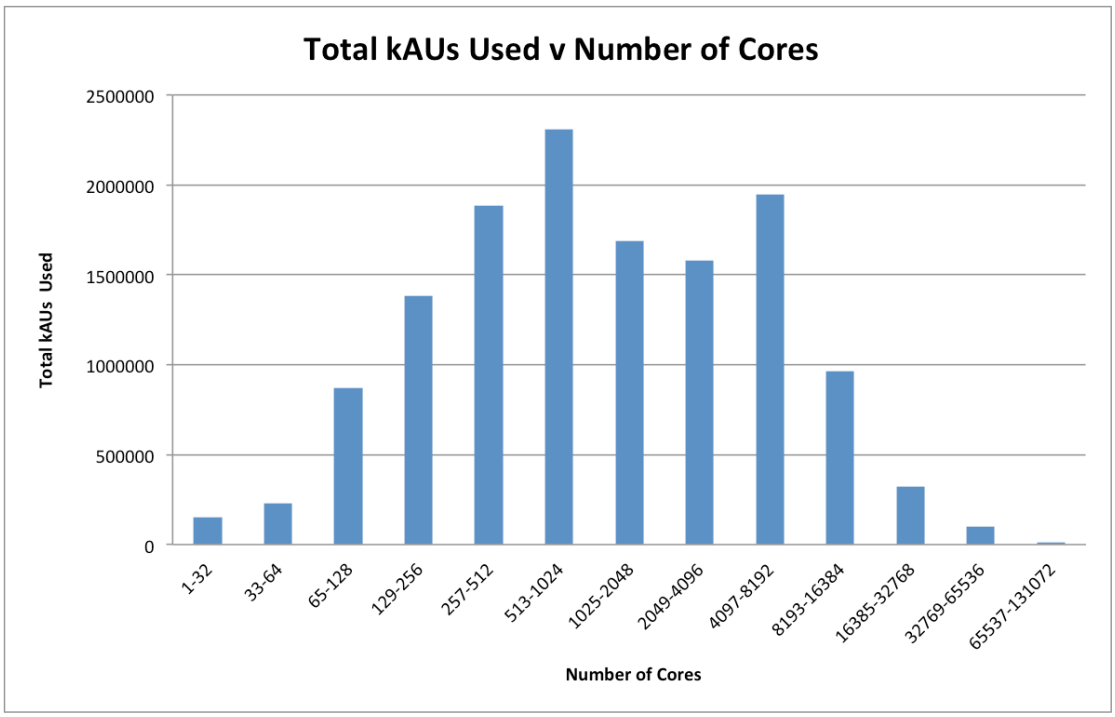

The graph shows that most of the kAUs are spent on jobs between 257 cores and 8192 cores. The number of kAUs used is closely related to money and shows how the investment in the system is utilised.

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 7874 queries were answered by the Service Provider, and over 98.5% were resolved within 2 days. In addition to this, the Service Provider passed on 296 in-depth queries to CSE and Cray.

	15Q1	15Q2	15Q3	15Q4	TOTAL
Self-Service Admin	1722	1172	775	1564	5233
Admin	654	616	408	601	2278
Technical	118	91	67	87	363
Total Queries	***2494***	***1879***	***1250***	***2252***	***7874***

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.

	15Q1	15Q2	15Q3	15Q4	TOTAL
Phone Calls Received	135 (41)	100 (22)	104 (20)	92 (14)	**431 (97)**
Change Requests	8	7	5	5	**25**
User Registration Requests	313	214	220	302	**1049**

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 431 calls received in total, only 97 (22.5%) were actual ARCHER user calls that resulted in queries. The trend through the year has been a falling number of actual ARCHER calls resulting in a query. All phone calls were answered within 2 minutes, as required.

4.2 USL Service Highlights

File system issues and improvements arising

Major service disruption was experienced in May and June due to Sonexion file system issues. In conjunction with work to resolve the issues, successful measures were put in places to minimize the user impact. Collaborative working between all service partners and carefully constructed and targeted user communication were key to this. The success of the measures could be seen in the 83% utilisation maintained during May and June with the temporary file space utilisation accounting for 46% of the utilization for the period. There were minimal user complaints received during the period of disruption and an appreciation of the effort taken to keep the service running from the user community.

Many of the measures devised and implemented to minimize user impact and downtime are now included as standard processes and functionality. Recommendations from the lessons learned reports are also being implemented. The measures implemented included:

- SAFE functionality to be able to lock job submission on a per-file system basis (this prevents users from wasting resources when their file space is not available for running jobs).
- Provision of temporary project space when a particular file system is unavailable to allow users who are affected by file system issues to keep running calculations if possible.
- Move to resilient package installation across all file systems to enable users to access packages independently of any particular file system being unavailable
- An improved coordinated Major Incident Procedure

Period allocations for consortia and large research groups were successfully implemented and then staggered

In Q1 of 2015, under the direction of EPSRC, 6 monthly period allocations were introduced. This was done to help ensure that projects used their allocations more evenly over the lifetime of the project. This change has had a positive effect, though the simultaneous ending of a large number of both EPSRC and NERC allocations in March 2015 caused the machine to be very busy. Since then the EPSRC project allocations have been staggered throughout the year to avoid a recurrence of this issue.

The impact of these changes have been measured using the Scheduling Coefficient report. These reports show no recurrence of the problems from March 2015.

SAFE changes

Changes have been made to SAFE this year to support service improvements. These include:

- The implementation of sub-project management allowing the PIs to devolve management of parts of a project to project managers;
- The move to an improved reporting engine to speed up the creation of user reports;
- The addition of career stage monitoring in particular to allow EPSRC to track the number of early career stage researchers;
- And, the implementation of automatic tweeting of user mailings to increase the mailing delivery options.

UK-Federation authentication to SAFE implemented

UK-Federation authentication to the SAFE was implemented allowing users to authenticate with the same credentials as for their home institution. The impact of this was to reduce the number of credentials that the user needs to remember and troubleshoot. 215 users have signed up to use this functionality to date.

5. Operations and Systems Group (OSG)

5.1 Service failures

There were no 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 and implementing CLE 5.2 upgrade on the XC30;
   b. installing regular compiler and programming development upgrades;
   c. supporting OS enhancements to external login nodes.

(2) Resource management:
   a. PBS queue enhancements such as the SHORT development queue and further support for creation of advanced reservations;
   b. assessing and monitoring problems with the job scheduling cycle.

(3) Storage:
   a. significant involvement in the handling of major storage problems encountered during the year;
   b. upgrade of Sonexion (lustre) filesystem software;
   c. further integration of the RDF into the operational environment.

(4) System monitoring:
   a. further enhancement of use of external monitoring tools such as Nagios and OMD;
   b. expansion of internal system health checks.

(5) System administration:
   a. development and expansion of automated ticket handling;
   b. refinement of locally-developed systems administration tools;
   c. integration of the RDF data-analysis cluster into the wider operational configuration.

(6) Communications:
   a. installation and configuration of multiple 40G connections to JANET core network;
   b. further hardening of internal ACF networks that underpin both external operational and internal secure management services.

(7) Service support systems:
   a. further development of automated failover of hypervisor-based virtual servers that provide resilient services such as SAFE, website and wiki.

(8) 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).

(9) 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.

6. Computational Science and Engineering (CSE)

These are selected highlights from the CSE Service during 2015. Full operational details on the CSE service (including metrics) can be found in the quarterly reports on the web.

6.1 Best Practice for Data Management on ARCHER

The amount of data required and produced by modeling and simulation is increasing year on year. This is reflected in the fact that data management and file system (IO) performance are now major concerns for many ARCHER users. Until relatively recently, neither of these were issues that concerned the majority of users. There is a lack of generally-available material on these topics for HPC users and it also tends to be an area where many HPC users have little expertise or experience.

In the second half of 2015, the CSE service focused on providing a set of practical resources for ARCHER users with the aim improving their data management and/or IO performance on ARCHER and the RDF. We provided both general advice, and advice targeted at specific application users where we are aware of particular issues with data management. In particular, we have produced:

- Data Management Guide on the ARCHER website, covering:
  - Archiving data to the RDF
  - Data transfer between ARCHER and the RDF
  - Data transfer to/from external sites to ARCHER and the RDF
  - Different ARCHER and RDF file systems and their use
- White Paper on Performance of Parallel IO on ARCHER:
  - Initially covering MPI-IO performance on ARCHER Lustre file systems
  - Currently expanding work to look at NetCDF and HDF5 performance
  - Working with the DiRAC facility to comparing performance across different file system and vendor architectures
- Webinars:
  - Data Management: best practice in using tools to manage data on ARCHER and the RDF, including how to efficiently move data between the different file systems.
  - Using OpenFOAM on ARCHER: the popular Open Source CFD software OpenFOAM has particular issues with the numbers of files it can produce when run in parallel. This webinar raised awareness of these issues in the user community and provided advice for how to deal with the problems.
  - Lustre and IO Tuning: provided a description of the ARCHER Lustre file systems, where users may see issues with performance, and tips for getting best performance out of the file systems depending on your usage pattern.
- Training:
  - Data management and IO performance best practice has been built into our Introductory face-to-face courses and the online ARCHER Driving Test.
  - Advanced material on parallel IO performance has been used as the basis of the Efficient Parallel IO on ARCHER course run in December 2015 in Oxford.

6.2 The ARCHER Driving Test: encouraging new users on to the ARCHER service

The ARCHER driving test was launched at the start of the year to give a mechanism for new users easily to gain access to the service whilst also ensuring that they had enough knowledge of HPC to make use of their ARCHER account. The test at https://www.archer.ac.uk/training/course-material/online/driving_test.php comprises 20 questions chosen randomly from a bank of 60, distributed to ensure coverage of all aspects of the system:

Category	# questions
Hardware	2
I/O	3
Programming	4
Compiling	3
PBS	1
Running jobs	3
Random category	4
Total	**20**

It is also supported by online training material including slides and video lectures addressing all the areas covered by the test. The test is promoted at ARCHER training courses and also mentioned on the email sent to all attendees after the course is finished where we encourage them to fill in the feedback form.

In the first year, the test was successfully completed by 122 people, 82 of whom have gone on to obtain accounts on ARCHER; those passing the test are sent a certificate of completion. After an initial burst of interest, take-up has remained very consistent throughout the year:

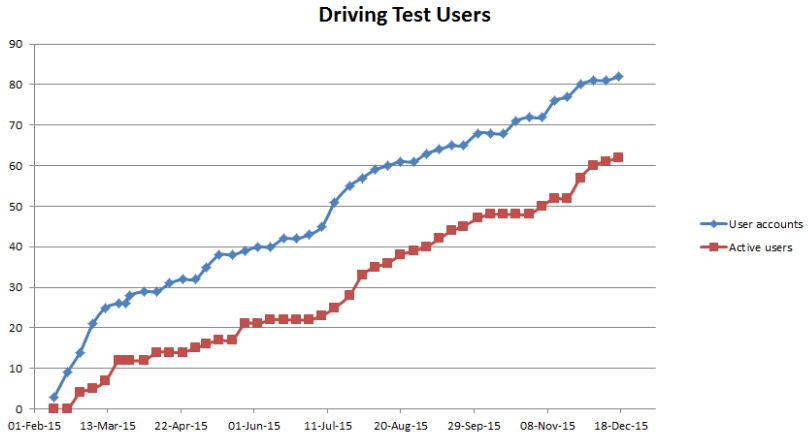

It is interesting to note that, from Q2 onwards, almost all new users have become active users (i.e. have submitted compute jobs). In total, some 34,600 kAUs have been spent by these 62 users, an average usage of around 560 kAUs; a typical active user is therefore spending almost half of their total allocation of 1,200 kAUs.

The driving test has been a great success and shows every sign of continuing to attract new users for the remainder of the ARCHER service.

6.3 Wee ARCHIE: a Raspberry Pi cluster to educate the next generation of HPC users

The ARCHER Outreach project aims to engage new communities and the next generation to take advantage of HPC technologies. However, one common problem in reaching out to these communities is helping them to understand the relationship between everyday computing, be it through a tablet, laptop, or smart phone, and supercomputing or 'high performance computing'. The ARCHER team appreciates the importance of helping everyone to understand how HPC can improve their science, allow the community to do complete more science, and also to ensure that the next generation understands HPC is a tool for all, not just the few lucky enough to work at an institution with an HPC resource. So we developed Wee ARCHIE.

Wee ARCHIE has been designed and built to help explain what HPC is, the difficulties in using such technologies but also the possibilities available when using HPC platforms. While Wee

ARCHIE is only a model of a real HPC cluster, it has all the key components. The cluster consists of 18 Raspberry Pi 2s, each of which has four cores, simulating the idea of a node, each with four cores. The cluster has been designed to enable explanation of the hardware and how the components are connected and interact. The processors, switches, power supply units and networking cables are all visible through a Perspex case, which is designed to be highly portable enabling us to take it to outreach events around the UK. Each Raspberry Pi has also been fitted with an LED array to allow us to show when the Pi is active and future code development will allow us to show the load on each 'node', enabling us to teach people about the importance of load balancing.

The design plans for the Wee ARCHIE cluster will be made available online in 2016, enabling anyone to purchase and build their own cluster. We will also be developing a range of software to highlight the advantages and also the difficulties of using HPC and the importance of using the right tool for your problem.

Wee ARCHIE will be taken to a series of outreach events in 2016 including the Big Bang Fair at the NEC, Birmingham in March 2016.

6.4 Women in HPC

Women in HPC, started in 2013, with the official launch in April 2014, and has become an internationally recognised name in the last year. In 2015, the Women in HPC initiative went from the two events held in 2014 to seven different events in 2015, the launch of a new website, the signing of our first international Women in HPC partner organisation and winning the HPCWire Readers Choice Award for Diversity.

During the last year, Women in HPC has participated in three international conferences: PraceDays15, Dublin, Ireland; ISC 2015, Frankfurt, Germany and Supercomputing 2015, Austin, USA. At each conference we have had an array of events, including 'Bird's of a Feather' discussions, workshops, training sessions and networking receptions. In September 2015, Women in HPC ran the first ever Women in HPC careers event in collaboration with BCSWomen, in London, bringing together leading women working with HPC in the UK to discuss career opportunities with early career women interested in a career in HPC or looking for a new direction to follow within the HPC community. The day culminated with a speed networking session, which despite many being apprehensive of, was the best-received activity of 2015.

At ISC 2015, Women in HPC signed an agreement with Compute Canada as the first international partner to work with Women in HPC. The partnership enables the establishment of a Canadian Women in HPC chapter organisation which will run training events and networking sessions aimed at the Canadian HPC community, and sharing information and ideas with Women in HPC. This is a model that Women in HPC is in the process of developing, with the plan to establish chapters and partnerships around the globe providing the opportunity for women in the HPC community to network internationally as well as the opportunity to encourage other women to move into a career within the HPC community.

In 2016, Women in HPC is set to expand, signing up additional international and regional partners to establish best practice in broadening participation in the HPC community around the world. We will also be working with a variety of conferences as well as expanding our dissemination activities.

6.5 Competitive eCSE Programme

The embedded CSE (eCSE) programme provides funding for 14 FTEs embedded directly into the scientific community through a series of competitive, peer-reviewed calls. 2015 saw a high demand from the community for funding, resulting in a very high quality threshold.
Over the course of the six calls, 54 projects have been funded. These projects have made a significant impact on the quality and performance of the software suite on ARCHER – more than

ten of the most heavily used codes on the system have benefitted from eCSE investment effort. This in turn has facilitated greater scientific output and impact, allowing previously untenable science.

The programme has a focus on early career researchers and on developing the UK software skills base. The distributed and embedded nature of the programme allows for this skills development to be spread across the whole of the UK, and a key highlight of the programme has to be the fact that staff from ~30 institutions from a wide geographical distribution have benefitted from eCSE investment. Coupled with this, we will have early career researchers observing at future panel meetings. The aim is to give them a better insight into the mechanism of selection to assist in their future preparation of funding proposals.

A final highlight is the successful new communities programme that encourages proposals from new communities, looking to enhance the diversity of science being carried out on ARCHER. Over the three eCSE calls that have included this initiative we have received thirteen new community applications.

7. Cray Service Group

7.1 Summary of Performance and Service Enhancements

2015 has been another strong year for the ARCHER service. Overall system reliability and utilisation of resources have continued to be at a high level. Where technology areas have performed below the high standards expected, corrective action has been taken to resolve issues with the minimum amount of disruption and after careful consultation with service partners. More details of specific technology failures can be found in the table and associated descriptions below.

7.2 Reliability and Performance

The performance and reliability of the hardware and software technologies underpinning the ARCHER service continues to be of a high standard. New versions of software that provide feature enhancements and bug fixes to the user community are continually under development and are then being implemented on the ARCHER service following periods of evaluation on appropriate test platforms.

Large, complex HPC systems such as ARCHER are not immune from technology failures but under most circumstances those failures can be managed by utilising well-designed resiliency features and robust configurations.

On occasions, technology failures do result in impact upon the user community. The most significant technology area of the ARCHER service where issues were encountered in 2015 was in the parallel lustre filesystem and associated storage components. Acknowledging that improvements could be made in both hardware and software areas of the storage subsystem, these improvements were forthcoming and integrated with a minimum of disruption to the user community.

7.3 Service Failures

Seven unscheduled incidents classified as full service failures were encountered during 2015. As can be seen, six of these failures occurred in the first half of the year with a much-improved performance and only a single service failure in the second half of the year.

Incident	Date	Description
1	08-Jan-15	System reboot required following storage controller failure
2	06-May-15	Storage failure on lustre filesystem /fs3.
3	07-May-15	Storage failure on lustre filesystem /fs2
4	13-May-15	System reboot following PBSPro batch system server failure
5	10-Jun-15	System reboot following a failure in the system bootraid device.
6	30-Jun-15	Running user work lost following PBSPro batch system becoming unresponsive
7	06-Oct-15	System reboot following unintended initialization of system components

The details of these seven technology service failures were:

- One service failure due to a lustre storage controller fault requiring a system reboot to clear.

- Two service failures occurred due to multiple storage component failures affecting two different lustre parallel filesystems.
- Two service failures due to problems related to the PBS Pro batch subsystem which caused the loss of running user work.
- One service outage due to a controller failure in a bootraid device, which houses the operating system filesystems for the ARCHER service.
- One service failure was caused by the accidental use of an initialisation command on system components.

8. Cray Centre of Excellence (CoE)

Michael Neff joined the CoE and brings specific expertise in computational chemistry to the CoE.

8.1 CoE Project Highlights

HIPSTAR

A case study was produced on some previous work done by the CoE on the HiPSTAR code from the University of Southampton. In a previous CoE project, OpenMP was added to HiPSTAR improving the code scalability considerably. This OpenMP work was then used as a basis for an OpenACC port of the application (done in conjunction with the users by the ARCHER and ORNL CoEs). The OpenACC port of the application allowed the users to run HiPSTAR to very large scale on the Titan system at ORNL, and form a collaboration with GE in this work. This work has been documented in a Cray case study - http://www.cray.com/sites/default/files/XC30-ARCHER-HiPSTAR-0315.pdf.

ONETEP

We expended a small amount of effort during the year (via other UK Applications Staff) supporting a Poisson-Boltzmann Equation solver for ONETEP (in collaboration with The University of Southampton). An eCSE proposal for further work with ONETEP and CASTEP is about to be submitted and it is our intention that continuing support from the CoE would be provided if the project was to be funded.

HADOOP/Spark

The CoE was involved with a project with users from the University of Nottingham to look into the potential of analyzing data generated by molecular dynamics applications with Hadoop. The general aim here was really a proof-of-concept study to understand what can be done with Hadoop technologies in processing of HPC data. For this project, the CoE brought in experts in MapReduce and Spark from Cray's Data Analytics division. In April, work started on a basic Hadoop application with subsequent initial testing. This work was highlighted in a Cray Case Study.

PDNS3D

The CoE investigated a performance problem with the PDNS3D code. The coarray implementation was not performing well relative to the MPI implementation. The CoE used an expert from the US Cray Performance team to continue investigating this. He found a performance issue with the Cray compiler that disadvantages the coarray version of the code at all scales. This issue was resolved and subsequent analysis showed that the halo-swap communication patterns are implemented differently in the MPI and coarray versions. The results were communicated at the end of the year and we hope to discuss them further in the near future.

8.2 Filesystem and I/O

The CoE was engaged to understand the adverse effects reported by users as a result of filesystem rebuilds. Working with the Cray team onsite, we performed a detailed investigation of individual storage unit (OST) performance and were able to determine the cause of the adverse performance. As part of this work, we were also able to show that the tuning of both raid-check and the rebuild process was working to the extent that applications would be less impacted when these operations were throttled. Filesystem tasks are now much less intrusive as a result of software improvements and configuration of raid-check and rebuild operations.

As part of this effort the CoE also started to engage with the NCAS community, were given access to the NCAS Puma service, and have been able to run a representative UM job on ARCHER via that service. Initial investigations concentrated on the I/O server configuration and this is a topic that we hope to revisit in the future. The direct investigation of NCAS UM jobs became less relevant

once the filesystem performance issues were understood, and the new raid-check regime was put in place.

The question of how to optimize I/O comes up often so we are considering how we can do more (beyond existing material we have presented in optimisation workshops and in the ARCHER tuning guide) to get appropriate information to users.

8.3 Training and Workshops

The CoE assisted with various workshops during the year. Particular examples were the Porting and Optimisation workshop run at the time of EASC 2015, the 1st Euro OpenACC Hackathon, and the ARCHER serial Optimisation course run at Cray's EMEA HQ in Bristol in December.

CoE staff presented a seminar on modern Fortran, as well as a talk on coarrays and ARCHER projects at a joint meeting of the British Computer Society and Institute of Physics.

The CoE was able to engage with ARCHER users at various events including the Insight UK meeting in Coventry, the UK Turbulence Consortium Annual review meeting, and the 24th Discrete Simulation of Fluid Dynamics (DSFD) conference in Edinburgh.

The ARCHER CoE organised a mini-symposium at the PARCO 2015 conference with a focus on programming for manycore nodes (including GPUs, multicore CPUs, and Intel Xeon Phi). This (along with other ParCo events) was a useful way to interact with users and the wider community on concerns and requirements for programming models as we look at current and future architectures.

8.4 ARCHER Queries and Software

Of particular note this year was an issue with suboptimal performance of NWChem. NWChem was not performing optimally on ARCHER and, for some cases, was slower than HECToR. The problem was difficult to diagnose due to large runtimes, but was found to be due to a new GA implementation. Cray CoE, Cray USA developers, EPCC and the user were all involved in working on this. An early fix did not work due to a race condition but, as of the end of the year, a new ARMCI communication model has resolved the performance problem.

Updates of CLE and the Programming Environment towards the end of the year on ARCHER mean that new features are available and we will produce a seminar in 2016 to outline these.

8.5 eCSE Meetings

The CoE completed technical assessments for the two eCSE calls during the year, and staff attended the project planning meetings.