ARCHER Service
2014 Annual Report

Document Information and Version History

Version:	1.0
Status	FINAL
Author(s):	Alan Simpson, Jo Beech-Brandt, Stephen Booth, Andy Turner, Jason Beech-Brandt, Steve Jordan, Mike Brown, David Henty, Lorna Smith
Reviewer(s)	Alan Simpson, Lorna Smith, Steve Jordan

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2014-11-30	Initial Template	Liz Sim, Alan Simpson, Jo Beech-Brandt, Stephen Booth, Andy Turner, Mike Brown, Steve Jordan, Jason Beech-Brandt
0.2	2015-01-16	Updates and additions	Jo Beech-Brandt
0.3	2015-01-20	Initial review of format	Alan Simpson
0.4	2015-01-26	Completed draft	Jo Beech-Brandt
0.5	2015-01-27	Internal review	Alan Simpson, Lorna Smith
0.6	2015-01-27	Draft version for EPSRC	Alan Simpson, Jo-Beech-Brandt
1.0	2015-02-25	Final Version for EPSRC	Alan Simpson, Jo-Beech-Brandt

Table of Contents

1. Introduction

This annual report covers the period from 1 Jan 2014 at 08:00 to 1 Jan 2015 at 08:00.

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

Section 3 provides a summary of the service utilisation. Section 4 details a summary of the year for the USL team and details the Helpdesk Metrics and outlines some of the highlights for the year. The OSG team report within Section 5 details 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 Programme, eCSE Programme, Women in HPC and the distributed training programme. The ARCHER Image Competition is also described.

In Sections 7 and 8 the Cray Service team and Cray Centre of Excellence team 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/2014.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.

- This has been a very positive first year of the ARCHER Service. Utilisation on ARCHER has been good throughout the year and has gradually increased. Most of the cycles are used on jobs exploiting hundreds or thousands of cores. These are the sorts of jobs that are generally hard to run on other smaller systems.

- In total, the Service dealt with around 8,500 queries during 2014, meeting all query targets. Dealing promptly with queries requires all the elements of the Service to work closely together.

- ARCHER successfully underwent a major upgrade to Phase 2 during 2014, which increased its capability by more than 60%. Phase 2 was completed on schedule and included a period of free access to users where they were able to test out their codes and stress test the system. The successful completion of this upgrade involved careful planning and close collaboration between all the elements of the ARCHER service, particularly Cray Service Group and the Operations and Systems Group.

- The various elements of the ARCHER service collaborate effectively to provide a high quality service to users. These collaborations benefit from both informal links provided by co-location and the more formal mechanism of the ARCHER Operations meeting that brings together key staff from all of the service elements to discuss any issues and to plan for future improvements to the service.

- The ARCHER Service will continue to be an innovative service that engages effectively with all users. Highlights in the first year include: the successful Women in HPC initiative; training at a wide variety of sites reflecting the geographic distribution of users; named contacts for all Consortia; and promoting the science outputs via the ARCHER image competition and associated calendar.

- The eCSE programme is the primary mechanism for improving the application codes on ARCHER and it is very positive that it has been so popular with record-breaking numbers of proposals. Both technical staff and PIs have come from a broad range of institutions through the UK.

- At the end of the year, we carried out a survey of ARCHER users and this demonstrated that users are very happy with the service. The survey showed that users overall satisfaction with the service is very high and that they particularly valued their interactions with service staff via the Helpdesk.

3. Service Utilisation

3.1 Overall Utilisation

Utilisation over the year was 71% (or 88% of optimum). This utilisation has seen an increase as the year has progressed. The Phase 2 upgrade took place in November and so there was a dip in percentage use but the machine had increased in capability by more than 60%.

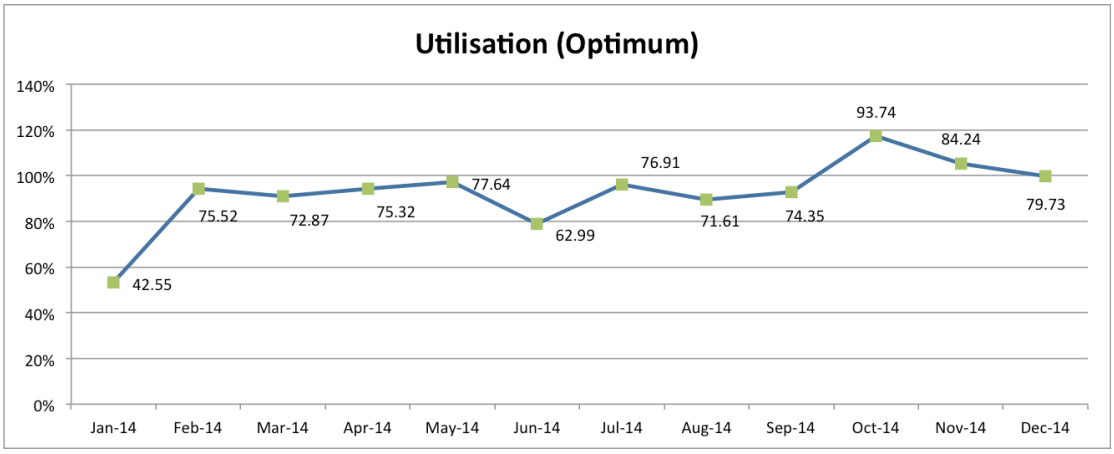

3.2 Utilisation by Funding Body

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

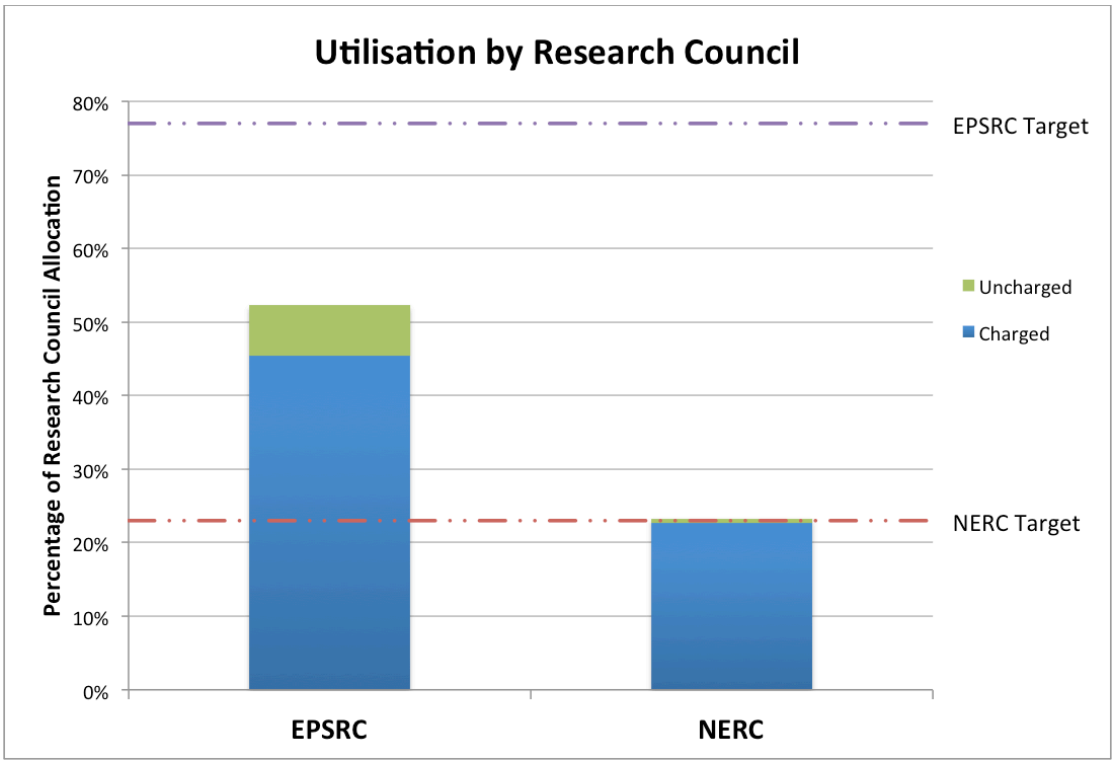

This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that NERC achieved their target allocation, whereas EPSRC underused their allocation.

3.3 Additional Usage Graph

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

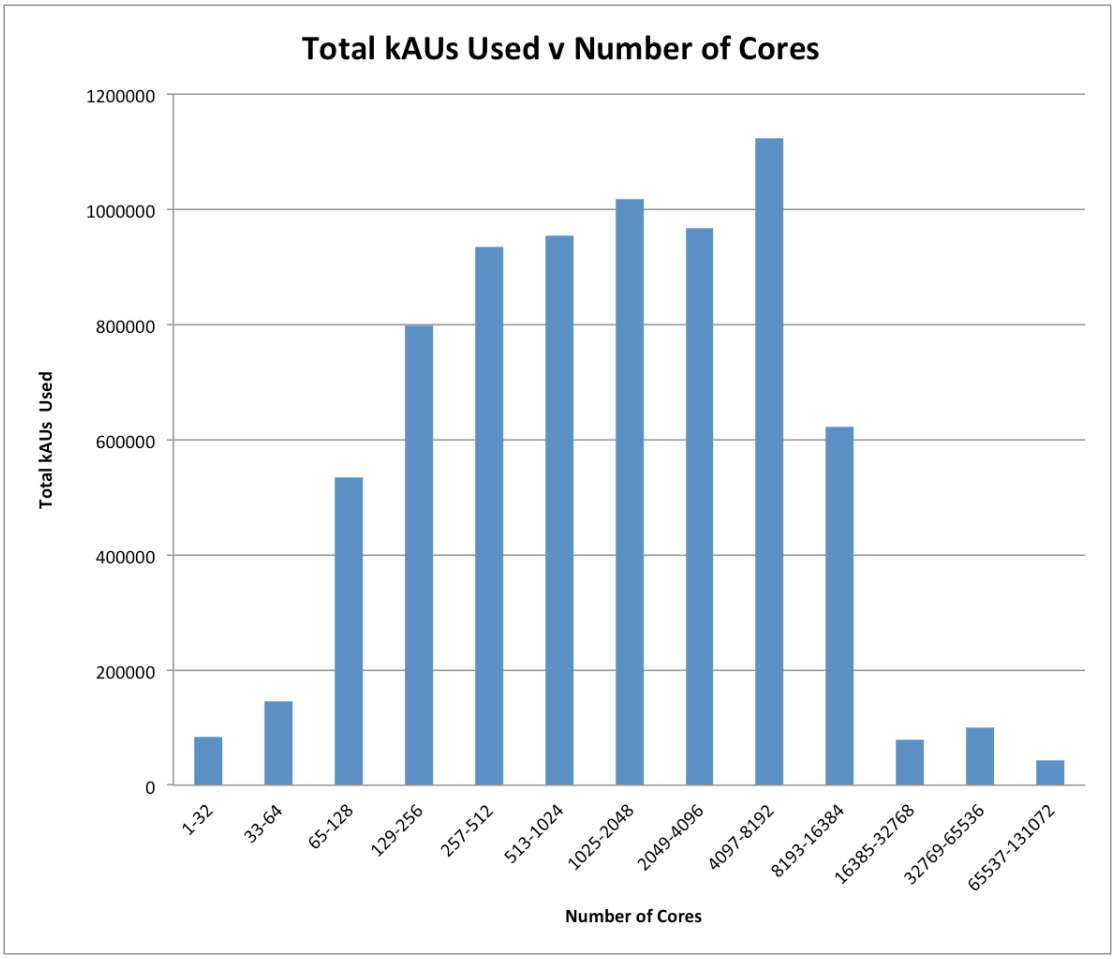

The graph shows that most of the kAUs are spent on jobs between 257 cores and 8192 cores. The number of kAUs used is closely related to money and shows how the investment in the system is utilised.

4. User Support and Liaison (USL)

4.1 Helpdesk Metrics

Query Closure

It was a busy first year on the helpdesk but all Service level agreements were met. A total of 8161 queries were answered by the Service Provider, and over 98% were resolved within 2 days. In addition to this, the Service Provider passed on more than 300 in-depth queries to CSE and Cray.

	14Q1	14Q2	14Q3	14Q4	TOTAL
Self-Service Admin	2651	1085	1042	971	5749
Admin	527	534	360	452	1916
Technical	144	166	112	117	2455
Total Queries	***3322***	***1785***	***1514***	***1540***	***8161***

Other Queries

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

	14Q1	14Q2	14Q3	14Q4	TOTAL
Phone Calls Received	162(40)	218(28)	122 (28)	97 (21)	**599(117)**
Change Requests	1	21	4	12	**38**
User Registration Requests	65	340	198	248	**851**

It is worth noting that the volume of telephone calls was low throughout the year . Of the 599 calls received in total, only 117 (19.5%) were actual ARCHER user calls that resulted in queries. All phone calls were answered within 2 minutes.

4.2 USL Service Highlights

Transition of Users from HECToR to ARCHER

The transition of the new Service went very smoothly. This was due largely to the continuity in people, processes and software allowing a quick transition of the users across to the new Service. We also saw the benefit of having both systems overlapping to enable users to transition as and when they were ready. As to be expected there were a large number of queries during the initial transition phase (3322 in 1st Quarter) but they were all dealt with promptly and within the agreed service level.

Analysis of Job Efficiency

Using SAFE records, a job efficiency analysis was made to determine the Scheduling Coefficient; this is defined as the ratio of runtime to runtime plus wait time. The result showed that generally queuing times are quite short. The only cases where longer wait times than runtimes were encountered were either for very short jobs (as there is always a scheduling overhead) or for very large jobs (where the system has to drain compute nodes to make space for the jobs).

Analysis of Project Use versus Allocations

At the request of the EPSRC, reporting in SAFE was enhanced to enable analysis of project usage profiles versus their allocations for the period. This was further enhanced by the addition of automatic mailings from the SAFE to Consortia Leaders and PIs being implemented. This includes a graph of used resources plotted against allocated resources to assist PIs and Consortia Leaders in managing their allocations. This is especially important for large users with the introduction of period allocations.

Production of YouTube videos

User documentation was enhanced with the addition of YouTube videos for common SAFE tasks. These are being added to as new functionality becomes available and are available on the ARCHER website and also on the ARCHER YouTube channel.

System Reservations

All users are now able to make system reservations requests via the SAFE. Users have been making use of this new functionality to guarantee nodes for specific purposes such as industrial users requiring longer job runtimes and also large consortia holding training courses. A premium rate is charged for reservations.

Public-Facing Webpages

The ARCHER website was further developed to include public-facing web pages to help demonstrate the impact of ARCHER and encourage new users to apply for access to ARCHER. There is also a new dedicated Industry section that features case studies and described methods of access for Industrial users.

Phase 2 Upgrade and Photo Blog

The System was upgraded during this period and the upgrade was completed on schedule. A photo blog was updated regularly to allow users to follow the progress of the upgrade. .

ARCHER Annual User Survey

The ARCHER Annual User Survey took place and 152 responses were received (more than 10% of the active user base). Users were asked to rate and comment on the following aspects of the service: overall satisfaction, hardware, software, helpdesk, documentation, website, training, webinars. The response was generally very positive and the analysis report is included as an appendix to this annual report (see Appendix 1).

5. Operations and Systems Group (OSG)

5.1 Service Failures

The Service Provider has been responsible for 3 Service Failures in 2014 (details of Technology Failures are reported in Section 7.3).

Incident	Date	Description
1	03-Mar-14	Maintenance Overrun – CLE Upgrade work
2	09-Apr-14	Maintenance Overrun - CLE Upgrade work
3	05-Nov-14	Phase2 Maintenance Start Time Incorrect

Three failures were classified as Service Provider failures during 2014

- Two service failures occurred when agreed maintenance sessions overran their approved time.
- One service failure occurred when the maintenance session for the Phase 2 Preparation work started early. It was approved by the Authority to start at 0900 but work started at 0830. The work was completed well within the approved 8 hours.

5.2 OSG Service Activities

The activities of OSG (Operations and Systems Group) are divided into four main activity areas:

Maintaining Day-to-Day Operational Support

Day-to-day operational support includes regular tasks (system healthchecks, monitoring and responding to queries) being undertaken by two shift teams (each with a shift leader and shift administrator) to support the operational day. Very close liaison is maintained with both the on-site Cray support staff and the Helpdesk. Development activities have included additional security monitoring and logging, and enhancements to the backup provision

Planning Service Enhancements in a Near-to Medium Timeframe

OSG is responsible for installing systems level software on the XC30 and applying security fixes and patch sets to OS components. Software changes, even those that may not be directly invasive to the running service, are generally undertaken within defined maintenance windows; these may be "at-risk" (and hence concurrent with the live user service) or "dedicated" (whereby the interactions with the system require that the service be taken down). Sessions are undertaken generally twice per month, and, wherever possible, these are "at-risk." Maintenance sessions require a significant degree of planning and coordination to ensure that all software dependencies are maintained. Typically, two future sessions are planned at any time, with the very detailed planning for more major upgrades (such as for CLE major releases) being undertaken across a few months. Principal user-visible service enhancements included:

- CLE 5.1 release
- Monthly releases of CDT
- Changes to PBS to support serial jobs
- Changes to PBS to support short debugging jobs
- Streamlining of the way that course reservations are handled

Planning Major Service Enhancements

The principal major service enhancement was the Phase 2 upgrade that was undertaken in the period October to November. Substantial electrical and mechanical infrastructure works needed to be undertaken in order to accommodate the new hardware, and these had to be performed concurrently with the 24x7 provision and support of the live operational service . While the bulk of the Phase 2 planning and provision was undertaken by Cray, there were some system-level changes that were required to be undertaken to support the enhanced configuration, such as

those to the PBS batch configuration. The bulk of this work was undertaken concurrently with the physical hardware upgrade

Developing Associated Services that Underpin the Main Operational Service

In order to support the main external service on the XC30, there is a considerable number of associated systems that provide the external-facing services for the helpdesk, SAFE, and the web site along with internal-facing systems that provide authentication and backup services, software repositories, security logging and so forth, along with a complex set of internal management networks. OSG provides and supports these systems (across two physical sites), with multiple-replicated servers giving redundancy, resiliency and Disaster Recovery capabilities. Development activities have centred on the methods used internally for the processing of user-generated tickets (e.g., for password changes, quota modifications, …) to minimise turnaround of such requests.

6. Computational Science and Engineering (CSE)

Smooth Startup in Very Tight Timescales

The timescales for award of CSE contract to start of service were very compressed in Q4 2013. Despite this, the ARCHER CSE service started extremely smoothly at full capacity to ensure that users could exploit the new ARCHER service successfully from day 1. This success is illustrated in a number of ways during the early period of the ARCHER Service (November 2013 – March 2014):

- the number of queries handled by the CSE Service (236 excluding course registrations);
- the amount of training provided: 18.5 days across 6 different locations.
- the fact that the first eCSE call was opened in November 2013 and handled a record number of proposals (19 proposals with 14 funded);
- that the Consortium Contact programme was in place from November 2013;
- ARCHER Best Practice Guide was in place in November 2013 to assist users;
- that the Technical Forum webinar programme started in November 2013 to share technical expertise and experience throughout the ARCHER community and beyond

The fact that the CSE service was running at full capacity so quickly contributed to the high levels of ARCHER utilisation in the early period of the service, smoothed the transition process for users migrating from HECToR, and helped to provide high levels of user satisfaction with the new Service. All of these are key to the success of any new HPC facility.

Consortium Contact Programme

To get the maximum impact out of any HPC facility it is important to ensure that the users can use the system efficiently; a major part of this is making sure that the service engages with the user community effectively. If done properly, the users will be able to use the facility with the minimum overhead and focus on progressing their research. The Consortium Contact programme has proved to be an positive, novel way of engaging with the EPSRC and NERC consortia on ARCHER, allowing them to use the Service effectively. Each of the EPSRC and NERC scientific consortia has a named contact from the ARCHER CSE team. The contact works with the consortia members to develop an in-depth understanding of their requirements from the ARCHER Service so that the CSE service can provide the most effective support. So far, contacts have engaged with the consortia in a number of ways:

- providing technical advice on the use of ARCHER;
- attending and presenting at consortia meetings;
- assisting in the production of eCSE, RAP and grant proposals;
- co-authoring research papers and conference submissions;
- acting as an advocate for the consortium within the Service;
- providing a conduit for feedback on the Service to assist with continual service improvement.

We aim to build on this success by ensuring that the consortium contact is known to and has contacted all the major users in each of the consortia and by providing consortia-specific sections on the ARCHER website where they can host information that is useful to their members.

ARCHER Image Competition: Showcasing Research on ARCHER

The ARCHER Image Competition was run at the end of 2014 and had two major goals:

- Showcase the research being performed on ARCHER in an accessible way;
- Increase the engagement of the ARCHER user community with the service.

Thirty-two, high-quality entries were received. Each entry consisted of both an image and a short description of the research that the image illustrated aimed at the interested non-specialist. Judging took both the image and the description into account and was undertaken by members of the ARCHER user community, the ARCHER CSE team and other experts. The winning image along with a number of other entries were selected for inclusion in an ARCHER 2015 calendar that was produced and sent to a wide range of people from the UK scientific, HPC and funding communities. Throughout 2015 we will continue to use the image competition entries to raise the profile of ARCHER and encourage user engagement by adding the entries to the ARCHER website and highlighting each one throughout the year on the ARCHER homepage.

Competitive eCSE Programme

In addition to the 6 FTEs in the centralised CSE team, the embedded CSE (eCSE) programme provides funding for 14 FTEs embedded directly into the scientific community through a series of competitive, peer-reviewed calls. The programme aims to build up scientific software development expertise within the scientific community itself whilst gaining maximum benefit from ARCHER via improvements to key codes running on the system.

The demand for eCSE funding from the community has been extremely high throughout the first year of operation, reflecting the unique opportunity the calls provide for both developing research software in a sustainable way and for increasing the software development skills of the staff involved. In total, 68 proposals have been received from four, regularly spaced calls with 33 projects funded (310 person months of staff effort) from the first three calls (the fourth call panel is still to meet). Technical staff are drawn from 15 different institutions and the project PIs are from 19 different institutions so far. The regular, predictable nature of the call programme has allowed applicants to plan their proposals and time them to meet the requirements of the project.

The eCSE panel is made up of scientific and technical experts drawn from science and industry within the UK and chaired by Walter Lioen from SURFsara in the Netherlands. The panel provides a broad range of HPC and scientific expertise to prioritise and rank proposals. Continual improvements are made to the submission process: for example, allowing applicants to submit proposals online via the SAFE to streamline the application and review process. We are also continually improving the submission guidance to provide applicants with the best information possible.

One the aims of the eCSE programme is to attract new communities to ARCHER and so for the fourth call we updated the proposal form to help applicants from areas of science not presently exploiting ARCHER submit a proposal. It is important to recognise that such applicants may not always have existing ARCHER performance data and that benefits to ARCHER may not appear immediately and are likely to be part of a multi-stage process. In addition, for the fourth call, a certain amount of staff effort is ring-fenced for proposals from new communities.

Distributed Training: Courses Located where Users Need Them

To ensure that high quality training is available to ARCHER users regardless of their location in the UK, we have carefully chosen the mechanisms and locations of course delivery. The entire training programme is subject to approval from an external Training Panel representing various HPC communities, ensuring that training is delivered when and where it is required based on user requirements.

Specific examples up to the end of 2014 include:

- training at 15 different locations across the UK (where, for example, we count London as a single location despite having used several different venues in the city);
- placing all lecture material on the web for future reference under a creative commons license granting free re-use for educational purposes, thus increasing the impact of the ARCHER investment;
- running monthly interactive virtual tutorials with recordings made available via the ARCHER YouTube channel to allow users to benefit from the training even if they are unable to attend in person;
- developing an online ARCHER Driving Test for users to test their knowledge and to increase the skills base of the UK HPC community;
- packaging up lectures, exercises and videos into an online introductory course which leads onto the ARCHER Driving Test so that this can be sat from anywhere at the user's own pace;
- setting up training hubs in Southampton, London and the north of England (via the N8 partnership) to enable coherent sets of courses to be delivered at key locations to cater for researchers who want to follow a programme to develop their skills beyond a basic level;
- development of two new courses in data management and analysis, to be delivered in 2015, in response to the growing interest in Big Data and the challenges faced by many

ARCHER users in handing the datasets produced by simulations that are continually increasing in size and number.

Women in HPC

The official launch of the Women in HPC network took place in Edinburgh in April 2014, with a conference bringing together over 30 women working in HPC from around the UK, showcasing work by leading women in the field of HPC from across the UK, and providing the early career women with the opportunity to network with other women in their field, meet role models and potential mentors. The launch finished with a lively panel discussion on the reasons behind the gender inequality in the field of HPC, whether this is different from other areas, and what can be done about it.

As part of the launch we also ran the first ever 'Women in HPC' training event: A Hands-On introduction to HPC. This course, while not excluding male attendees, was led entirely by female lecturers and demonstrators and comprised 80% female attendees. Verbal feedback suggested that many of the attendees had been nervous of attending such training before, even though some also admitted that there wasn't a lot in the session that they didn't already know, but they had never had the confidence to attend such a course previously. This again enabled the establishment of collaborations between the attendees and the development of mentoring connections.

In November we ran the first ever international Women in HPC workshop at Supercomputing 2014 in New Orleans, USA. With 46 attendees and 11 early career researchers presenting their work to the first female dominated audience they had probably ever worked in and for every single early career presenter their first opportunity to attend such a high-profile conference, the workshop was a resounding success. The poster session and panel discussion in particularly proved very popular and we hope to repeat this at ISC15 in Frankfurt and SC15 in Austin.

Women in HPC was also an invited member of the Intel Community Hub discussion on Women in STEM at SC14 and we also co-hosted the Women in HPC BoF in its 5th year.

The panel discussions at the launch in April and at Supercomputing in November have directed the path for the Women in HPC network in 2015. The discussions highlighted the importance of bringing women together at such events, so they can meet female peers, where otherwise they may not know any, and also to start a series of best practice guides for individuals and employers which we will develop throughout 2015. As the panel discussions were so popular and productive we will continue these in 2015 but with directed themes to encourage the ability to draw firm conclusions from the sessions.

Throughout 2014, Women in HPC distributed three electronic newsletters, advertising events and funding opportunities to women working in HPC. By the end of 2014 the network had 118 people registered on the Women in HPC website.

7. Cray Service Group

7.1 Summary of Performance and Service Enhancements

2014 has been an excellent year for the ARCHER service with a very stable and reliable technology environment enabling high resource utilisation for the user community. Improvements in new software releases have helped resolve issues that have impacted users and provided new features to enrich the user experience. Feedback to Cray from key stakeholders has been positive.

This first full year of the ARCHER service has also seen a significant upgrade of computing resources with the Phase 2 Upgrade in November 2014. The execution of the upgrade went entirely according to the detailed plans agreed with EPSRC and the other service providers.

7.2 Reliability and Performance

The reliability of ARCHER hardware and software has proven to be excellent during 2014. Such large and complex HPC systems inevitably encounter occasional technology failures but the importance of minimising the effects of such failures on the user community has been key to the early success of the ARCHER service. Good system design and configuration that utilises high availability components and failover technologies has resulted in very high availability and generally good user experience. Concurrent maintenance tools to enable hardware component replacement to take place whilst the system remains in operation have helped to reduce the amount of dedicated maintenance time required for both preventive maintenance and hardware problem resolution when compared to older systems, such as HECToR.

Software and firmware enhancements to resolve problems encountered and provide additional features are constantly in development. These enhancements are made available to the service provider following extensive in-house testing by Cray on dedicated systems. There have been improvements across all areas of the service during 2014.

7.3 Service Failures

Four incidents classified as full technology service failures were encountered during 2014 (for details of Service Provider failures see Section 3.3):

Incident	Date	Description
1	21-May-14	All running user jobs lost due to procedural error
2	20-Jun-14	All running user jobs lost due to conflict with a software patch installation
3	26-Nov-14	System reboot required to clear high speed network routing error
4	31-Dec-14	System reboot required to clear high speed network routing error

The details of these four technology service failures were:

- One service failure occurred whilst investigating a PBS batch subsystem issue. A procedural error resulted in the loss of running user work. Site procedures have been reviewed and improved in light of this incident.
- One service failure occurred when all running user work was lost following complications resulting from the installation of a software patch.
- Two failures occurred when unrecoverable high speed network errors resulted from dynamic reconfiguration of routing information. Improvements in new versions of Cray supplied software will be installed by OSG early in 2015 to help address these issues.

7.4 Phase 2 Upgrade

The composition of the Phase 2 upgrade was decided following detailed consultation with the user community. The overwhelming response from users was a desire for addition of the same technology as already deployed in the successful Phase 1 implementation. The building block for the system is the compute node comprising of dual socket Intel Ivybridge processors and 64GB of memory. Table 1 shows the increase in user compute node and core counts for the expanded system.

	Cray XC30 cabinets	Number of compute nodes	Number of compute cores
Phase 1	16	3008	72,192
Phase 2	26	4920	118,080

Table 1. Summary of user compute resources for ARCHER Phase I and Phase 2

The Phase 2 upgrade was completed in three stages:

- Stage 1: The physical installation and testing of the new 10 cabinet Cray XC30 system whilst the Phase 1 production system remained in operation.
- Stage 2: The amalgamation of the Phase 1 16 cabinet system with the new 10 cabinet system.
- Stage 3: Acceptance and early user testing.

The Phase 2 equipment arrived at the site on Tuesday 4th November and the physical installation work began immediately. By the end of the day, all 10 Cray XC30 and 7 cooling cabinets had been positioned in the computer room. Electrical and chilled water connections were then completed by site contractors and the 10 new cabinets were configured as a standalone Cray XC30 system for QA testing.

On 11th November the user service was closed to enable the combining of the original 16 cabinet system with the new 10 cabinets. This mainly involved the addition of optical high speed network cables to link the two systems together along with necessary software changes to create a single system. Cabinet cooling enhancements were also made to the original 16 Cray XC30 cabinets at this time. Once this work was completed, further QA testing of the combined system took place.

The final stage of the upgrade began with the successful completion of all QA testing. A period of early user stress testing then took place to put the system through its paces prior to the resumption of the full production service. The system returned to full production service on 24th November as planned.

8. Cray Centre of Excellence (CoE)

8.1 CoE Project Highlights

Python in HPC

Python is increasingly being used within HPC and following user consultations two major issues were highlighted:

- People are using additional Python packages beyond what is shipped with the standard Python distributions which come with Linux
- The startup time associated with Python applications can be high, especially for applications which are run at scale. This is not a Cray-specific issue, but rather a generic Python issue when Python is used in large scaling applications.

To help address these issues, the CoE has instigated a project to investigate a variety of approaches for resolving these issues, and assess them on a number of criteria (portability, cost, ease of use, performance). The CoE is working closely with EPCC as the Service Provider and CSE Provider in this project. A new project on ARCHER has been created (d61) and within this project evaluations have already started on a number of technologies. As an example the DLFM tool, which is developed as an internal Cray project and which aims to help with the scalability issue associated with Python applications was evaluated on ARCHER. The DLFM developers ported and tested the DLFM package on the ARCHER system with a workload mimicking that used by some key ARCHER applications that utilise Python. Very successful results have been obtained with Python benchmarks and user applications including GPAW. Based on this information, DLFM has been configured and built on ARCHER by Cray, and is now maintained as a user module by the CSE team.

In addition the Cray CoE for ARCHER has made contact with the developers at Continuum Analytics (http://continuum.io/) who develop the Anaconda package which aims to simplify the packaging and distribution of Python and associated modules. The Anaconda developers have been provided an account on ARCHER, under d61, and did some initial porting and testing work. The CoE took over this porting effort and created an initial port of the Anaconda package, which was then passed onto the CSE team for proper installation in module format on the system. It is felt that the addition of these two packages will provide a significant benefit to the ARCHER Python user community.

Power Monitoring

As part of the CRESTA project a paper was prepared by the ARCHER CoE in conjunction with Cray R&D and users from TUD on power monitoring on Cray XC30 systems, and presented at CUG2014 in Lugano. The paper was very well received and was a runner-up for best paper award at the conference (https://cug.org/CUG2014). The lessons learned from preparing this material are now being used in CoE discussions with the CSE and SP teams about how best to undertake power monitoring on the ARCHER system, what could be shown to users, and how to present the results.

Following on from the very successful power monitoring research a user-level library, called pmlib, has been produced. The pmlib library allows for easier user-level access to power monitoring counters that can be called directly from an application. This library will soon be available to users on an open source, limited support basis as a public download.

UM

The ARCHER CoE has worked with the UM users in group n02 to assist with proper job setups and decompositions for running on the ARCHER system. Many of their users were running jobs in exactly the same way as they were run on HECToR; however, with some simple changes to the way they launch their jobs, they will achieve much better performance. The CoE team visited the NCAS users in February to work with them on this issue, which was resolved.

Work continues with the n02 users to address slowdown seen in some UM jobs which appears to be IO related. The CoE has worked closely with the n02 users and the SP and CSE teams to

identify where this slowdown is coming from and provide the users with workarounds for this issue. This work is ongoing in two directions. Firstly the CoE has been providing advice to the n02 users about use of the UM IO server logging, which would provide more information about the IO performance when they hit a slow period. Secondly, the Cray systems staff are working on implementing some server side IO logging which will help identify how the filesystem is being used by the various applications on the system. It is thought these two approaches will provide good information to help diagnose any further slowdowns.

Use of Hadoop in HPC

Some users and EPCC have expressed an interest in running Hadoop/Big Data workloads on ARCHER. The CoE worked together with people at EPCC to understand more about the nature of these Hadoop workloads, and possible avenues for running this on ARCHER.

Access to an in-house Cray system with the Cray Hadoop framework installed was also provided to EPCC and was used as part of a University of Edinburgh MSc project exploring big data workloads on traditional HPC architectures.

CASTEP

The ARCHER CoE has worked with the CASTEP developers to ensure that the latest versions of CASTEP work well with CCE. When the project started earlier this year, about 20% of the CASTEP regression suite passed with CCE. In the latest release (cce/8.3.0) nearly all tests now pass, and use of the Cray Compiler has become a real option for this user community, which gives them added flexibility in how they build and run their workload.

In addition the CoE has worked with the CASTEP developers to obtain the most recent versions (pre-release) of CASTEP to ensure the continued success of using CASTEP with the Cray tools.

8.2 Training and User Visits

The CoE have been involved in many training activities over the course of the year, both in terms of running CoE workshops for users and engaging in workshops ran by the CSE team. Initial Cray XC30 training material was also provided to the CSE team as a basis for producing some of the CSE training material. Presentations were also given to the ARCHER Technical Forum. International engagement came through presentations and tutorials at SC14 and CUG2014.

A number of visits to user groups around the country were undertaken during the year including to Reading, London and Southampton. This allowed for engagement with user groups to work on particular issues which required a closer, sustained period of working.

8.3 eCSE meetings

Jason Beech-Brandt attends the ARCHER eCSE meetings and acts as an outside technical advisor for these meetings reviewing proposals that cannot be reviewed by the CSE team due to conflict of interest.

Appendix A: ARCHER User Survey

A.1 Description of the Survey

The ARCHER User Survey was opened on 8 December 2014 and closed on 7 January 2015. 153 responses were received from ARCHER users. The survey asked for ratings (on a scale of 1 to 5) with the following questions:

1. Please rate your overall experience of the ARCHER Service (required) [Very Unsatisfied (1) – Very Satisfied (5)]
2. Has the ARCHER hardware configuration met the requirements of your research? (required) [Not met any requirements (1) – Exceeded requirements (5)]
3. Has the software on ARCHER met the requirements of your research? (required) [Not met any requirements (1) – Exceeded requirements (5)]
4. If you have used the ARCHER helpdesk, please rate your experience [Very Unsatisfied (1) – Very Satisfied (5)]
5. If you have used the ARCHER documentation, did it provide the information you required? [Did not provide the information I required (1) – Provided all the information I required and more (5)]
6. If you have used the ARCHER website, please rate the quality of the content and ease of navigation [Very poor (1) – Excellent (5)]
7. Please rate your experience of any ARCHER Training you have used (either online or face-to-face)? [Very Unsatisfied (1) – Very Satisfied (5)]
8. If you have attended any Technical Forum Webinars, did you find the session worthwhile? [A complete waste of time (1) – Extremely interesting and useful (5)]

Only the first three questions were compulsory for all survey responders but over 92% of responders also provided feedback for at least some of the optional questions. Users were also provided with the opportunity to offer comments or suggestions under all of the above headings and provided with space for any other comments or suggestions at the end of the survey.

The survey was constructed using Google Forms and embedded directly into the ARCHER website.

A.2 Selected Quotes

The following quotes reflect the tone of the majority of responders to the survey with regard to the ARCHER service:

> *"best machine around. I have access also to Titan and Mira, and ARCHER is by far the best"*
> *"The staff I have interacted with via email have been top-notch. Very knowledgeable, with great communication."*
> *"I am very happy with the service they have provided, and my research would have been severely slowed down if ARCHER had not been available."*
> *"ARCHER is an excellent service, thank you all very much!"*
> *"Our model now runs so fast on ARCHER (as compared to HECToR) that our main problem is moving the output off so as not to exceed the memory allowance (I mean this as a positive!)"*

The old argument between having accelerators or not as part of the national service was also evident:

> *"Would be great if there are GPU nodes"*
> *"Please don't go down the accelerator route until someone (anyone!) has code that actually runs on X thousand GPUs..."*
> *"[Please provide] GPU, Xeon Phi or other acceleration."*

"Please stick to homogeneous hardware."

Quotes on the helpdesk (which also reflect on the centralised CSE team) echo the extremely high ratings for this aspect in particular that are shown below:

"Extremely good: my questions are usually very technical and the information I get back is invariably very strong."
"I would like to emphasize the satisfactory experience of helpdesk team. Their response is always fast, efficient and trying to help as much as they can. "
"Not only very helpful but also friendly, it is a pleasure to contact you!"
"I have always had prompt, courteous responses, with quick resolutions to my reported problems/questions. They have kept me informed in cases where the solution would take some time."

A.3 Ratings

All questions asked responders to rate their satisfaction with each particular aspect of the survey on a scale of 1 to 5 with 1 representing "Very Unsatisfied" and 5 representing "Very Satisfied". Table 1 summarises the ratings for each aspect and reveals the all aspects of the ARCHER Service are rated highly by users.

Service Aspect	Total responses	Mean Score (out of 5)	Median Score (out of 5)
Overall Satisfaction	153	4.4	4
Hardware	153	4.1	4
Software	153	4.0	4
Helpdesk	129	4.5	5
Documentation	142	4.1	4
Website	144	4.1	4
Training	81	4.1	4
Webinars	41	3.6	4

Table 1: Summary of scores for different aspects of the ARCHER Service

As can be seen from Figure 1, the overall satisfaction with the ARCHER service is extremely high with no responders rating the service below 3 on a 1-5 scale from "Very Unsatisfied" to "Very Satisfied". The mean rating is 4.4 and the median rating is 4.

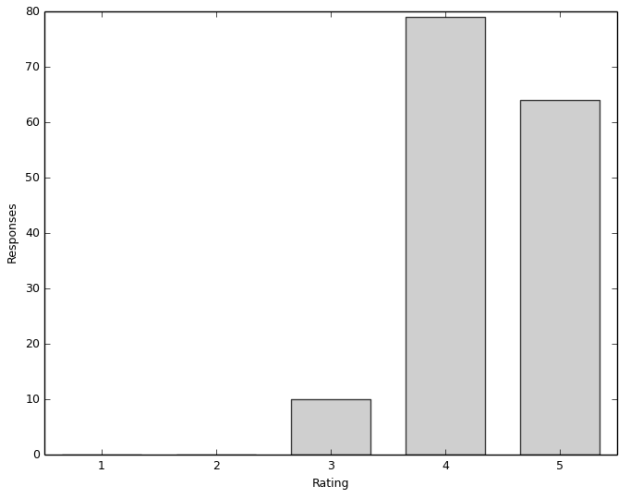

Figure 1: Distribution of scores for overall satisfaction with the ARCHER service (153 responses in total).

Similarly, for the hardware and software (Figure 2 and Figure 3 respectively), the overall satisfaction with the service is high with only one user rating the hardware below 3 and three users rating the software below 3. The single rating of 1 ("Very Unsatisfactory") for the software on ARCHER was not accompanied by any additional comments and the responder indicated that they did not want to be contacted about their response to the survey. The mean rating for hardware is 4.1 (median is 4) and the mean rating for the software is 4.0 (median is 4).

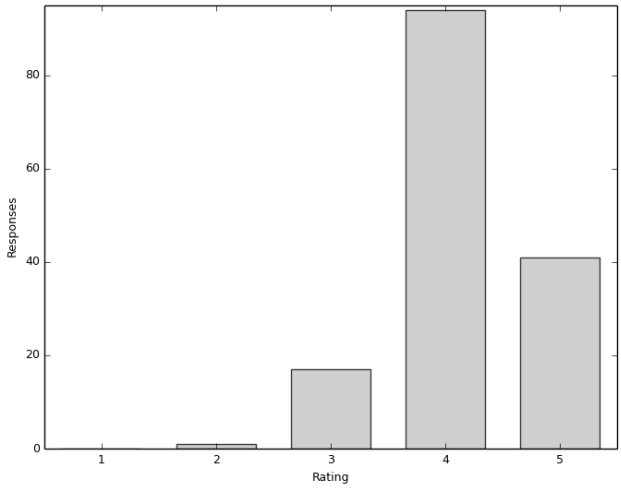

Figure 2: Distribution of scores for satisfaction with the ARCHER hardware (153 responses in total).

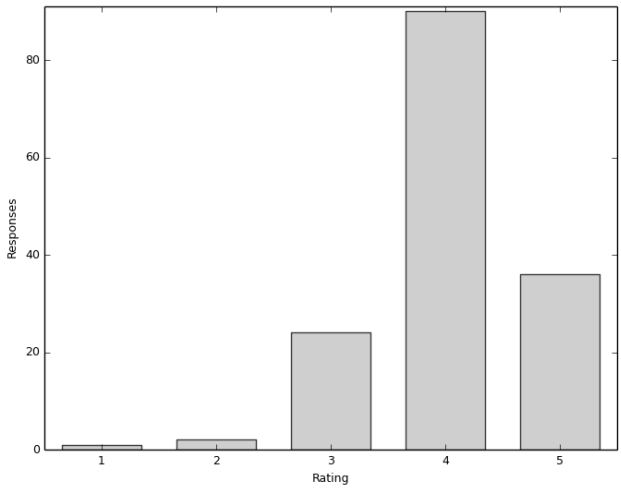

Figure 3: Distribution of scores for satisfaction with the ARCHER software (153 responses in total).

The distributions of ratings for the ARCHER Helpdesk (Figure 4) show that this is the highest rated aspect of the ARCHER service out of those surveyed, with a mean rating of 4.5 (median is 5). Both of the users who left the two ratings of 2 on the helpdesk indicated that they did not wish to be contacted regarding their responses to the survey. These high ratings, in particular, are testament to the staff that work on both the frontline helpdesk and to those who answer the user queries from the Service Provision, CSE and Cray teams.

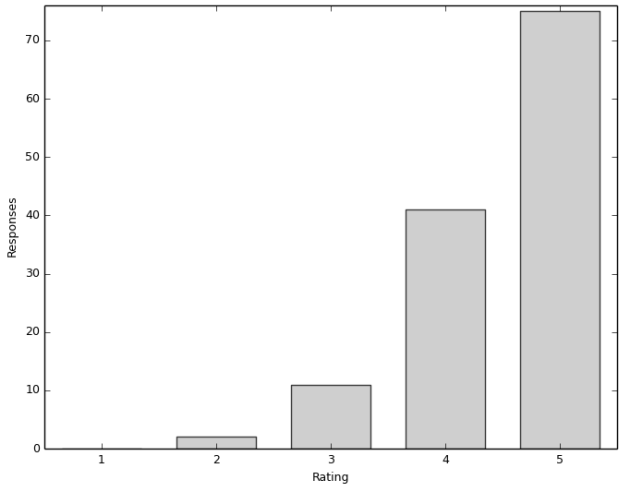

Figure 4: Distribution of scores for satisfaction with the ARCHER helpdesk (129 responses in total).

ARCHER documentation (Figure 5, mean = 4.1, median 4), website (Figure 6, mean = 4.1, median 4), and training (Figure 7, mean = 4.1, median = 4) all show the same high level of satisfaction as that shown for the overall service and have high respondent rates. The results for ARCHER training are consistent with the course survey results presented in the CSE Service quarterly reports.

The Technical Forum webinars (Figure 8) have a much lower responder rate (possibly due to the fact that the technical nature of the webinars is of interest to only a subset of ARCHER users) and also shows a slightly lower satisfaction rating (mean = 3.6, median = 4). Although this is still above 3 and only 1 user rated the webinars below a value of 3. From the comments we can see that the lower rating is due to the webinar software used (Blackboard Collaborate). The CSE Service is currently evaluating other webinar solutions to see if there is a better option available.

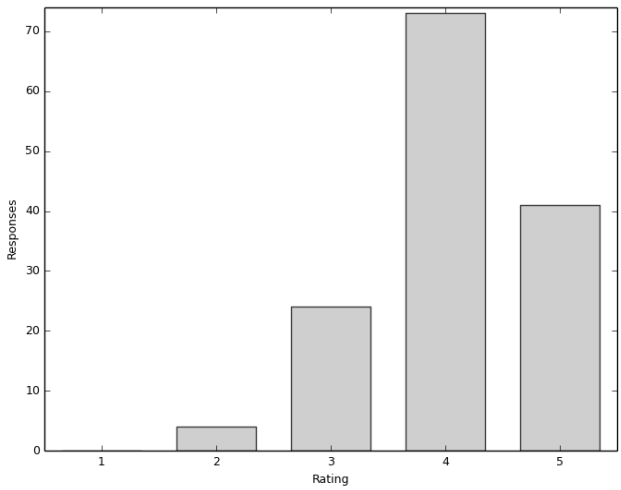

Figure 5: Distribution of scores for satisfaction with the ARCHER documentation (142 responses in total).

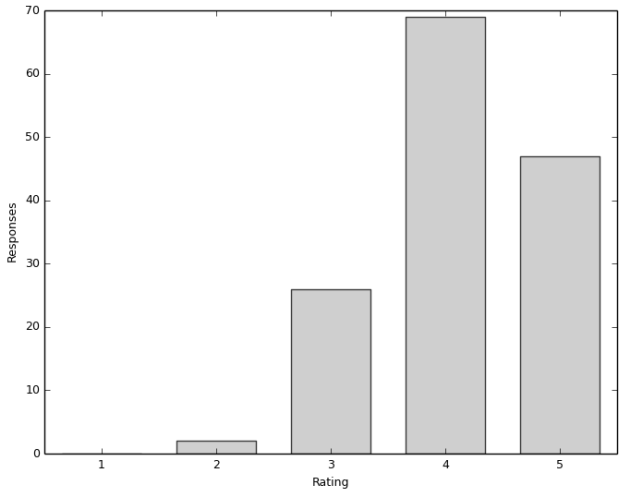

Figure 6: Distribution of scores for satisfaction with the ARCHER website (144 responses in total).

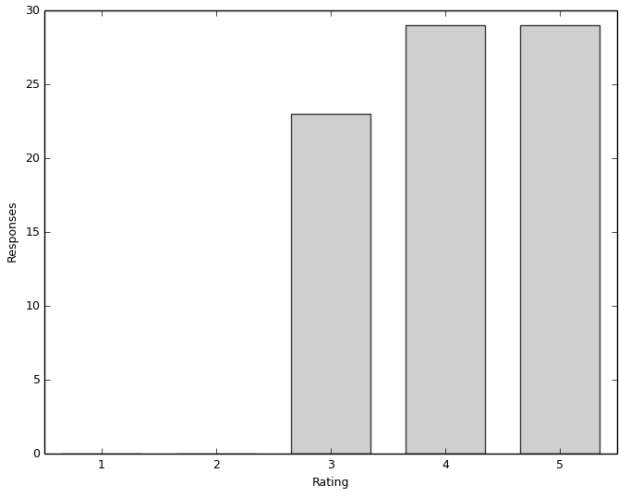

Figure 7: Distribution of scores for satisfaction with the ARCHER training (81 responses in total).

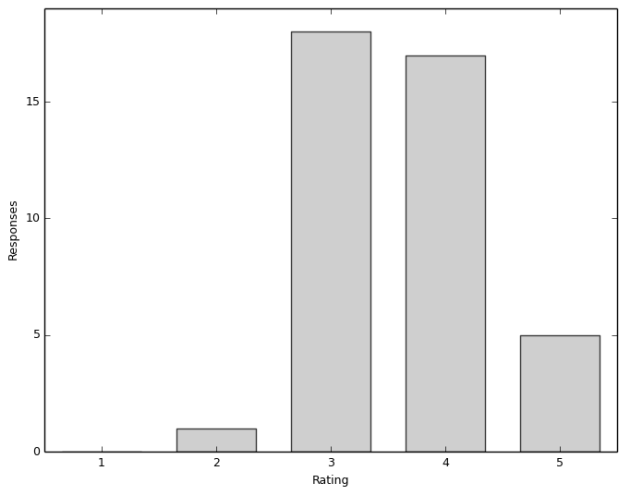

Figure 8: Distribution of scores for satisfaction with the ARCHER Technical Forum webinars (41 responses in total).