



ARCHER Service 2017 Annual Report



Document Information and Version History

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

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2017-12-19	Inputting initial information	Anne Whiting
0.2	2017-12-20	Added highlights form OSG	Linda Dewar
0.3	2018-01-04	Added SP metrics and graphs, USL and CSE sections	Anne Whiting, Jo Beech-Brandt, Andy Turner
0.4	2018-01-09	Added further details to OSG section	Linda Dewar
0.5	2018-01-09	Added Cray Service Group section	Felipe Popovics
0.6	2018-01-11	Highlights added	Anne Whiting
0.7	2018-01-12	Added Cray CoE section	Harvey Richardson
0.8	2018-01-19	Reviewed	Alan Simpson
0.9	2018-01-24	Updates after review	Anne Whiting
1.0	2018-01-25	Version for EPSRC	Alan Simpson, Anne Whiting

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	15
7.1 Summary of Performance and Service Enhancements.....	15
7.2 Reliability and Performance	15
7.3 Service Failures	15
7.4 Scheduled maintenance activities.....	16
8. Cray Centre of Excellence (CoE)	17
8.1 Longer-term projects.....	17
8.2 Training and workshops	18
8.3 Tier-2 HPC integration.....	18
8.4 ARCHER queries and software	18
8.5 Support of the eCSE programme	19

1. Introduction

This annual report covers the period from 1 Jan 2017 to 31 Dec 2017.

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 carried out by the team in 2017.

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/2017.php>

2. Executive Summary

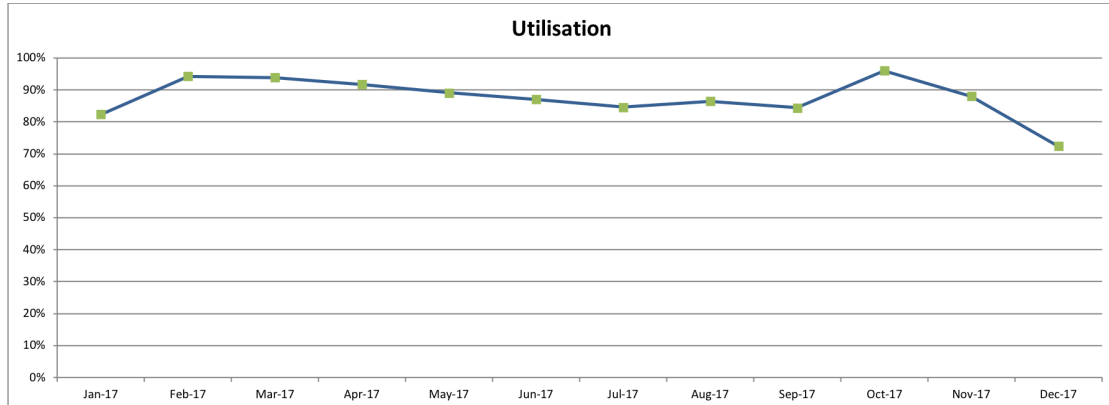
The sections from the various teams describe highlights of their activities. This section gives a brief summary of highlights from the last year of the overall ARCHER service. More details are provided in the appropriate section of the document.

- Led by the CSE team together with SP and Cray, a new feed of data has been incorporated into the SAFE from the ARCHER system capturing (amongst other metrics) peak memory usage and total energy usage for all user jobs. The SAFE reporting provided allows the user community to calculate the energy costs and memory usage of HPC applications, memory use by projects/users, and memory use by difference science areas.
- A focus of the ARCHER training programme for 2017 has been exploiting collaborations with UK, EU and overseas partners to enhance the range and quality of training provision to the UK HPC community. Highlights included working with the new Tier 2 HPC facilities to deliver training appropriate to their local facility and exploiting links with EU projects to bring trainers from outside the UK to deliver advanced courses to UK users.
- The ARCHER service has long recognised the importance of fostering early career researchers, thereby supporting scientific talent to sustain and enhance the world-class research enabled by ARCHER. In 2017 the ARCHER Service ran a competition, with funding from EPSRC, to promote and encourage interactions, networking and collaboration between early-career ARCHER users and groups in the USA.
- Following 11 successful calls for applications, the targets for the eCSE programme were met and resulted in significant investment in staff and software development expertise all across the UK. EPCC also made an additional commitment to invest all the money awarded in the contract into the programme. We were therefore able to run an additional call in 2017 for proposals and fund an extra 6 software development projects.
- Administration of the HPC Tier-2 projects run by Peta-5 (led by University of Cambridge), HPC Midlands Plus (led by Loughborough University) and Cirrus (led by University of Edinburgh) has been incorporated into the ARCHER SAFE.
- On-site Cray personnel developed diagnostic software for the Lustre filesystem which records a high level of detail about application usage, to help in diagnosing and investigating potential performance problems. OSG integrated this into their monitoring and metrics systems and produced a dashboard which enables all service partners to see performance bottlenecks in real-time, so that issues take less time to detect and less time to resolve.
- In total, the Service handled 10082 queries in 2017, up from 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 2016 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. Positive comments were received on the improvements to the queuing times after the priority changes were made. The annual user survey covering 2017 will be carried out in February 2018.

3. Service Utilisation

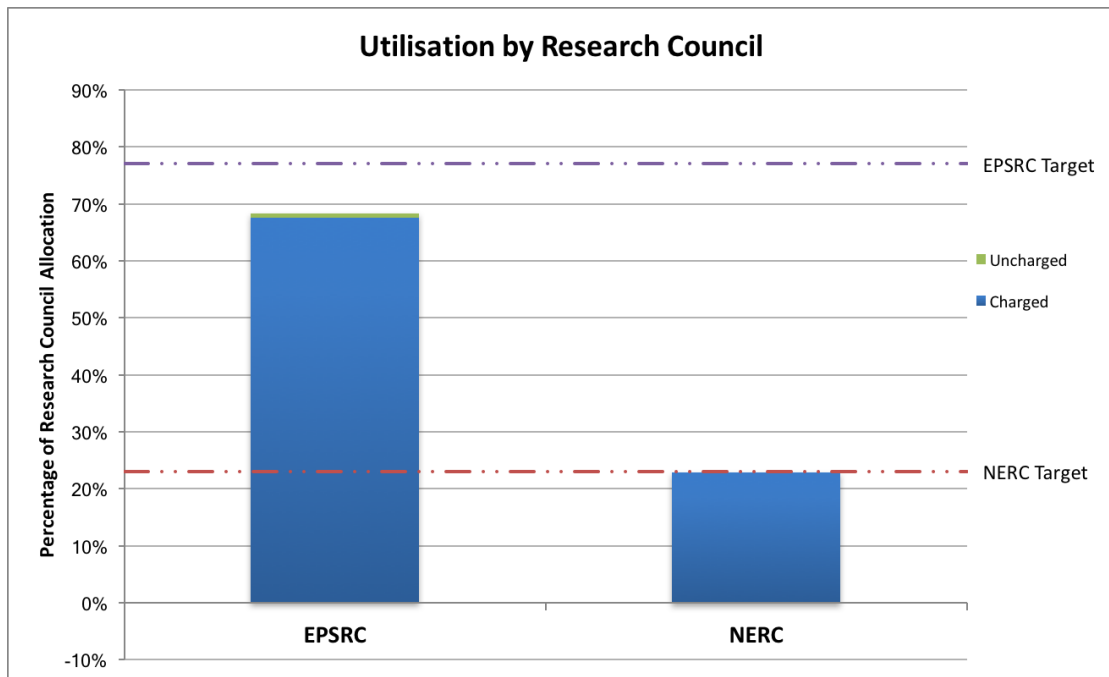
3.1 Overall Utilisation

Utilisation over the year was 88%, down from 94% in 2016.



3.2 Utilisation by Funding Body

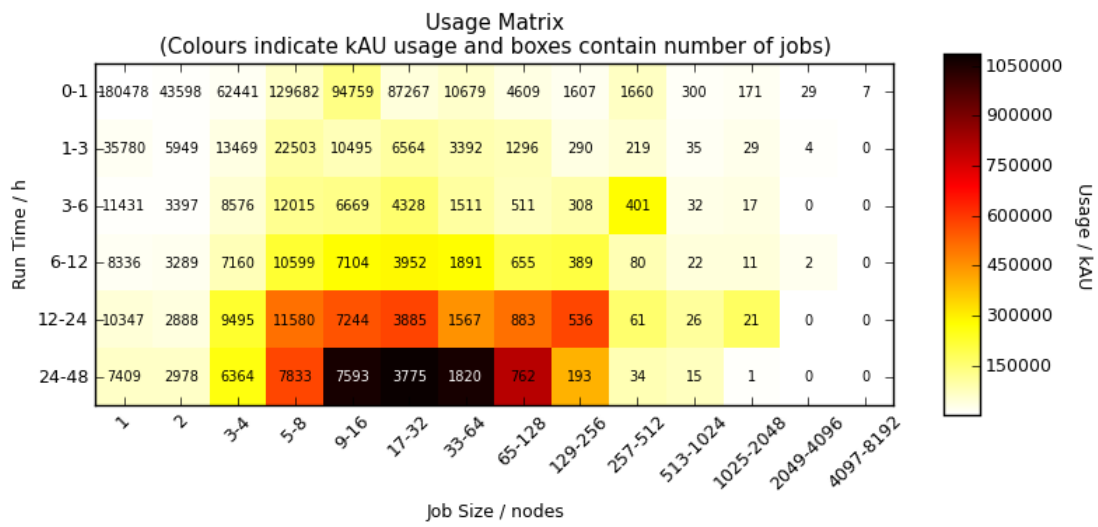
The utilisation by funding body relative to their allocation can be seen below.



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 throughout 2017.



The heatmap shows that most of the kAUs are spent on jobs between 192 cores and 3,072 cores (8 to 128 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 with all Service level met. A total of 10,082 queries were answered by the Service Provider, up from 7426 queries during 2016, and over 98% were resolved within 2 days. In addition to this, the Service Provider passed on 136 in-depth queries to CSE and Cray.

	17Q1	17Q2	17Q3	17Q4	TOTAL
Self-Service Admin	2297	2021	2517	1213	8048
Admin	551	411	381	414	1757
Technical	78	51	87	61	277
Total Queries	2926	2483	2985	1688	10082

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.

	17Q1	17Q2	17Q3	17Q4	TOTAL
Phone Calls Received	91	115	97	64	367
Change Requests	3	3	2	4	12
User Registration Requests	291	305	278	169	1043

It is worth noting that the volume of telephone calls was low throughout the year. Of the 367 calls received in total, only 66 (18%) were actual ARCHER user calls that resulted in queries. All phone calls were answered within 2 minutes, as required.

4.2 USL Service Highlights

User Survey 2016

The annual user survey was run in February 2017. The number of responses to the ARCHER user survey was 161 compared to 153 in 2014 and 230 in 2015. The feedback was very positive, with the mean score for overall satisfaction being 4.3 (on a scale of 1 (unsatisfactory) to 5 (excellent)). The helpdesk feedback was particularly favourable with a mean score of 4.5. Positive comments on the service included *"If we did not have access to ARCHER, we would not be able to do our science! It is an essential service for the UK"*. Positive feedback was also received from multiple users on the improvements to the queuing times after the priority changes were made.

Tier 2 integration and first RAP call

Tier 2 administration and helpdesk functionality have been integrated into the ARCHER SAFE for Peta-5 facility (led by University of Cambridge, HPC Midlands Plus (led by Loughborough University) and Cirrus (led by University of Edinburgh).

Early access projects were set up and run on Cirrus. The first Instant Access and RAP calls for Cirrus and other Tier 2 machines have been run and projects set up. The Tier 2 Technical Working group has been established to guide the Tier 2 infrastructure across the UK.

ARCHER Champions

Two very successful ARCHER Champions Workshops have been run in Leeds and Daresbury respectively in 2017 with the aim of promoting an integrated infrastructure and sharing best practice between Tier 1 and Tier 2 communities and support teams. Planning is underway for further events in 2018.

KNL usage

The Cray 12-node XC40 KNL system continues to be well used by the wider user community both to run appropriate jobs and to compare performance of codes between standard ARCHER nodes and the KNL nodes:

- 23.4% utilization over the period
- 272 users
- 20,867 jobs
- 8,962 kAUs used
- KNL user survey run
- Consultation was carried out with users on memory configuration options
- Benchmarking is ongoing of code performance between the KNL system and the main ARCHER system

SAFE improvements

Improvements have been made to the SAFE project overview reports, including the ability for PIs to track file usage over time for their consortia or project to ensure they stay within file quota limits. Excessive numbers of files can impact machine performance and are difficult to manage at the end of projects when files are moved off to other locations or deleted.

Weekly user digest mailings

Weekly digest user mailings communicating relevant service highlights and upcoming training courses are now being sent out every Tuesday with the mailing content available in the ARCHER website News section. The mailings help to ensure the user community has a consistent method of keeping up to date with upcoming events and news, and positive user feedback has been received.

ISO 9001:2015 Quality Management Certification

ISO 9001:2015 certification was achieved in February 2017, passing with flying colours and very few audit findings. The process of implementing ISO 9001 has been very positive, helping to embed a culture of continual service improvement and process improvement with the goal of improved service to the user community. Whilst previously service improvement was carried out on a more ad hoc basis, the quality management system provides a framework to ensure that the results of improvement initiatives are tracked, measured and documented to ensure they are effective. Work is underway to prepare for the second external audit, taking place in February 2018, which will assess how well processes are performing in delivering services to the user community.

ISO 27001:2013 Information Security Certification

EPCC has embarked on the process of implementing ISO 27001:2013 with an aim for certification in 2018. This standard specifies the requirements for establishing, implementing, maintaining and continually improving an information security management system. It also includes requirements for the assessment and treatment of information security risks tailored to the needs of the data owners and the organisation. Implementing ISO 27001 will allow us to demonstrate that EPCC follows best practice when handling user data, which is increasingly important in this era of big data and GDPR (data protection requirements).

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 activities) included:

- Operating system and applications software support:
 - Installing regular compiler and programming development upgrades
 - Supporting OS enhancements to external login nodes.
 - Successful upgrade to CLE5.2UP04
- Systems Administration:
 - Improved the throughput of the Serial queues by changing the default number of CPUs allocated to jobs.
 - Increased the amount of memory available to users on the eslogin nodes in response to a more memory-hungry version of the Cray compiler as a result of the May Programming Environment upgrade.
 - Implemented enhancements to the system monitoring by including the internal nodes and automatic paging of on-call Cray staff.
- Supporting Cray operations on Site:
 - Worked closely with Cray on-site staff on the root cause analysis of the eslogin/espp failures in Q1.
 - Developed graphical representations of Lustre data to identify some of the causes of poor filesystem performance.
 - Shared maintenance sessions with Cray to reduce the overall number of outages.
- Security:
 - Continued respond to CERT advisory and apply mitigation measures – in particular the well-publicised “Stack-clash” vulnerability.
 - Installed Cray-supplied security field notices.
 - Successfully tested the use of 2FA on the TDS.
- Outreach:
 - Attended two UKCSF meetings: - ACF in March and Met Office in September.
 - Attended Supercomputing 17 and Computing Insight UK in December
 - Attended ARCHER Champions in Leeds
 - Working with Zenotech to test the EPIC job submission portal on the TDS.
- Looking ahead:
 - Investigating options to increase the ACF-Janet bandwidth.
 - Planning to upgrade to PBS 13.408 in Q1.

6. Computational Science and Engineering (CSE)

Memory use analysis to enhance future national services

One of the key decisions to be made when designing a new HPC service is to understand the requirements for memory per node needed by user applications. It is generally expensive to supply higher amounts of memory so knowing what the maximum memory requirements for any service is key information. In the past, the amount of memory required has been specified using anecdotal evidence from the user community or measured using small numbers of runs on benchmarks. On ARCHER, there is the potential to capture memory use data for **all** running applications automatically to build up a true statistical picture of how memory capacity is used on the UK national service. The ARCHER centralised CSE team led a project to enable this functionality on ARCHER and to provide a useful quantitative analysis that could be used in the design of future UK national supercomputing services. This project required the CSE team to coordinate input from Service Provision (for enabling data collection and data access), from Cray (to understand what data could be collected and its format), from the SAFE development team (to incorporate the data into the ARCHER SAFE database), and from the CSE team itself (to provide data analysis expertise and interpretation of the results).

This work resulted in a new feed of data into the SAFE from the ARCHER system that captures (amongst other metrics) peak memory usage and total energy usage for all user jobs. Incorporating this data into the SAFE database allows us to analyse this data in several interesting ways; for example, memory use by HPC applications, memory use by projects/users, and memory use by difference science areas. This analysis is only possible in SAFE where we are able to link multiple data streams. Using this new data, the CSE team and Prof. Simon McIntosh-Smith (a HPC hardware expert from the University of Bristol) analysed the memory use by all applications over a full year of ARCHER use to provide insight into memory use patterns by job size and software application type.

The analysis, published as a paper at Supercomputing 2017, showed that the majority (65%) of use on ARCHER in the year used less than 12 GiB/node of memory (0.5 GiB/core) and 85% of use required less than 24 GiB/node (1 GiB/core). These numbers are significant as they suggest that a large majority of users on ARCHER use would still be able to run if a future service had new memory technology that offered high performance gains at the expense of reduced capacity.

Community engagement and leadership in training

A major feature of this year's ARCHER training programme has been the way in which we have leveraged collaborations with UK, EU and overseas partners to enhance the range and quality of training provision to the UK HPC community.

Examples from the 2017 programme include:

- Collaborating with the CDTs in Next Generation Computational Modelling at the University of Southampton and Condensed Matter at Heriot-Watt University to deliver training for their PhD students in parallel programming, software development and data carpentry skills.
- Running an Introduction to HPC course for the Alan Turing Institute's Doctoral Training scheme.
- Leveraging support from our PRACE Advanced Training Centre to offer travel bursaries to courses, including the week-long ARCHER Summer School in Edinburgh.
- Working with Women in HPC to promote diversity at specific events such as an MPI course in London, and general awareness-raising through a virtual tutorial.
- Using our extensive contacts in the SSI project to develop a pool of helpers to provide support at software and data carpentry courses.
- Working with the new Tier 2 HPC facilities to deliver training appropriate to their local facility, for example, a GPU course in Daresbury, and a KNL course in Cambridge.
- Exploiting links with EU projects to bring trainers from outside the UK to deliver advanced courses to UK users, such as PGAS training from INTERTWinE and HPC tools training from VI-HPS members in Germany.

- Using material from our MSc in HPC to develop a new course in Data Analytics for HPC which runs on the RDF Data Analytics Cluster.
- Promoting the free Supercomputing MOOC, developed under PRACE, to the UK community. 25% of the 6,000 joiners for the two runs in 2017 were from the UK.

We continue to exploit new collaborations to develop innovative material and delivery methods. We are active in the development of the new HPC Carpentry course which we plan to offer to ARCHER users in 2018. We have also worked with trainers from the US XSEDE project to understand the way they use interactive web technology to live stream technical courses to a remote audience while still offering face-to-face support via a distributed network of training hubs. We will use this delivery method to run an online MPI course at the start of 2018, with technical help supplied from the ARCHER Champions and UK RSE communities.

ARCHER eCSE programme: new science above and beyond requirements

The Embedded CSE (eCSE) programme provides funding to the ARCHER user community to develop and enhance the software underpinning research output on ARCHER. Investment in the software directly enhances the research output and resulting economic and societal impact. The programme was required to award and fund an average of 14 full time software engineers across the 5 years of the service. Following a series of regular and successful calls for applications (11 in total), this target was met and resulted in significant investment in staff and software development expertise all across the UK.

In addition, EPCC made a commitment to invest any remaining money awarded in the contract into the programme. This not-for-profit approach was designed to ensure the programme provided maximum benefit to the UK research community. Having fulfilled our contractual obligations, with the remaining funds we were able to run an additional call for proposals and fund an extra 6 software development projects.

These 6 projects provide over 3 years of staff effort and are dispersed across the UK: from Plymouth, through Loughborough, Leeds, Daresbury and Durham to Glasgow, St Andrews and Edinburgh. They involve a range of computational developments including enhanced parallelisation strategies, new libraries, implementation of new algorithms and enhanced interfaces.

The resulting research enabled on ARCHER is broad and significant:

- enhanced seismology simulations;
- new, larger and more complex computations in offshore renewable energy and offshore and coastal engineering;
- enhanced simulations of polydisperse granular materials – used in, for example, geomaterials, avalanches and landslides, crushing of mining ores and food processing;
- a step change in the ability to model the effects of clouds, currently the greatest uncertainty in weather and climate prediction;
- high-resolution flood risk modelling at regional to national scales;
- more accurate prediction of complex jet noise, aimed at reducing the impact of the expansion of aircraft fleets and airports – and thus aircraft noise - on communities.

Overall, the not-for-profit approach of the service has resulted in significant benefit, providing additional employment and skills development in the UK computational science community and facilitating enhanced societal and economic impact from the research conducted on ARCHER.

Early-career researchers: building links with the USA

The ARCHER service has long recognised the importance of fostering early career researchers, thereby supporting scientific talent to sustain and enhance the world-class research enabled by ARCHER.

In this vein, we recently ran a competition, with funding from EPSRC, to promote and encourage interactions, networking and collaboration between early-career ARCHER users and groups in the USA.

The aim was to capture the best science being carried out on ARCHER by early career researchers and to enable these researchers to generate new collaborations with the USA to further enhance their research.

Each winning entry received £3000 in travel money, and this is being used for a range of travel: to spend time embedded in a research group, to visit multiple researchers to facilitate new collaborations and to broaden horizons through attendance and presentation at relevant conferences.

These winners of the competition come from a range of disciplines across EPSRC's remit and showcase the exciting science being carried out by early career researchers on ARCHER. Highlights include:

- helping to prevent pipeline blockages in the oil and gas industry;
- improving the performance of solar [power] by studying photovoltaic panel materials;
- and simulating combustion engines to improve efficiency and to reduce environmental impact.

The Awards were presented at an evening reception at the Royal Geographical Society in London on 28th September 2017 in a 'Celebration of new and enhanced international research opportunities in computational science' providing an opportunity for the winners to network with those in the ARCHER community and with Research Software Engineers from across the UK. The winners will come together once again in April 2018 to share the impact of the awards they received with the wider research community.

An area of the ARCHER web site has been created to showcase the work of these and other early career researchers on ARCHER and highlight the importance of early career researchers to the service. See: <http://www.archer.ac.uk/community/earlycareer/earlycareerindex.php>

Women in HPC: Recognition, Growth and Sustainability

Women in HPC is now four years old and is an ongoing highlight of the ARCHER service. The organisation held its second full-day workshop at SC17 in Denver in November 2017. Women in HPC has run the greatest number of events to date this year, with events in the UK and internationally, finishing with a 'Women in HPC partner' event in South Africa in December.

WHPC was also once again recognised by the community in winning the HPCWire Readers' Choice Award for Diversity in 2017. This was the third year that the initiative has been recognised, despite growing competition, and this is recognition that the community welcomes the impact that the movement is making both in the UK and abroad.

Our full day workshop at SC17, our seventh international workshop, was the best received so far, with 105 participants and 16 early career women presenting their work. Alongside this event WHPC organised for all early career presenters to be offered a mentor. The mentoring began in late September, and included training for mentors as well as mentees. This is possibly the first example of a dedicated HPC mentoring program that has taken place almost entirely online. WHPC also organised a mentoring breakfast where mentors and mentees could meet in a relaxed environment at SC17, hopefully taking the opportunity to plan their week ahead.

2017 has also seen the increase in WHPC's activities specifically to promote careers for women, both by improving recruitment and by providing career networking opportunities. At both ISC17 and SC17 we held networking events with an emphasis on careers. At ISC17 we focused on asking our attendees what employers could do to improve their recruitment practices and this formed the basis for our SC17 BoF on improving recruitment and selection procedures to diversify the workforce. At

SC17 our evening networking event brought together our supporters who were seeking to employ more diverse candidates with women seeking their next opportunity.

At the end of 2017 WHPC also reorganised its management. It is now bigger than the few individuals from the ARCHER team who initiated the project back in 2013. Funding now needs to incorporate a more diverse range of resources and, as such, a new Executive Committee has been appointed and will focus on introducing a sustainable model for WHPC in early 2018.

7. Cray Service Group

7.1 Summary of Performance and Service Enhancements

2017 has been another excellent year for the ARCHER service with a very stable technology environment facilitating high resource utilization for the user community. New features and fixes were introduced to the ARCHER service with the installation of Cray's CLE 5.2 UP04 operating system in April 2017. In addition, regular updates to the Cray Programming Environment are provided monthly to ensure the latest compilers and tools are available to users.

7.2 Reliability and Performance

The ARCHER system, based on Cray XC30 technology, continues to provide excellent levels of reliability and performance. Large and complex supercomputer systems such as ARCHER inevitably encounter component fallout on a regular basis but most failures have little or no impact on the vast majority of users.

Cray's dedicated support team for the ARCHER service utilise concurrent maintenance techniques to routinely replace failed components whilst the ARCHER system remains operational.

On-site Cray personnel developed diagnostic software for the Lustre filesystem which records a high level of detail about application usage, to help in diagnosing and investigating potential performance problems. OSG integrated this into their monitoring and metrics systems and produced a dashboard which enables all service partners to see performance bottlenecks in real-time, so issues take less time to detect and less time to resolve.

7.3 Service Failures

7.3.1 Full Service failures

Two incidents classified as full service failures were encountered during 2017.

ARCHER full service failures 2017		
Incident	Date	Description
1	23/02/17	PBS server crashed causing the loss of user batch jobs
2	30/11/17	A redundant power supply unit failed in a login server located in a critical IO cabinet. This resulted in an external circuit breaker trip and the loss of power to the cabinet, requiring a system reboot.

In addition to the full service failures above, there was a delayed return to service affecting all users following a scheduled maintenance session on 23 August

7.3.2 Partial Service failures

Although most technology failures do not impact on the user community, there were several exceptions during 2017. The most significant technology areas of the ARCHER service where issues were encountered in 2017 were:

- Four incidents of Cray XC30 cabinet crashes due a firmware bug affecting cabinet controller devices. A patch was provided and, since its installation, no additional incidents of this type have been seen.
- One incident of XC30 compute nodes being "admindown" due to a hung task on a Sonexion filesystem. This prevented new job placement on 9.4% of compute nodes until the hung threads were cleared.

7.4 Scheduled maintenance activities

Only one of the permitted four dedicated maintenance session was taken by Cray in 2017, ensuring that the ARCHER system remains in the hands of the user community as much as possible. The dedicated Cray maintenance session was used to install new firmware on power rectifier components and for High Speed Network cable maintenance.

8. Cray Centre of Excellence (CoE)

In 2017, the driving motivation for the CoE was to focus on projects that would be beneficial beyond individual applications and to impact communities rather than specific research groups. The areas we concentrated on, with a long-term outlook, were I/O optimization, monitoring and observation and auto-tuning. Other activities are outlined in the sections below

8.1 Longer-term projects

We had a range of activities relating to I/O this year. The helpdesk receives many queries relating to I/O performance and we want to be able to deal with these faster, understand the I/O requirements better, and offer advice on application design or usage that will benefit users.

I/O Projects

At the start of the year we took a hard look at the problem of how to optimize the I/O (or more broadly data movement) in applications of interest to the ARCHER and wider HPC communities. The use cases for this work were distilled from seven main application examples and what resulted is a set of requirements (for example staging, post-processing analysis, I/O avoidance, parallel I/O etc.). We took this project as far as an initial design for the components required to deliver the framework. We discussed this project with a limited number of potential partners. Although we had intended to use existing components as far as possible it soon became clear that this project would require more effort than possible under the ARCHER CoE and the work is continuing within the wider Cray EMEA Research Lab at Cray (which the CoE is part of). The CoE continued to work on I/O monitoring and analysis as described in the following sections.

I/O monitoring and analysis

Understanding how application I/O performs on a shared filesystem requires detailed information to be gathered (system-side) about the filesystem performance along with an application-centric view of I/O. This year we have developed two approaches to gain better insights into filesystem performance. Both are built on a Lustre monitoring framework developed by the Cray staff who support ARCHER on site.

The first approach is to attempt to identify applications that behave abnormally (have a high runtime variation) on the basis that these may be perturbed if the filesystem is busy and of course they may be generating high I/O requirements as well. This approach was an extension of work done at HLRS (presented at CUG2017): we re-implemented a Spark code to ingest Lustre statistics and optimized it. This approach did highlight applications that we had already identified as high I/O usage based on helpdesk queries. This approach is limited as it requires detailed knowledge of expected application runtimes.

The second approach is to directly analyse the Lustre statistics for all applications and look for outliers that consume significant filesystem resource compared to the 'average'. We have developed software to do this and are working on new derived metrics that can identify sub-optimal application I/O. This approach is proving promising and is giving us a much faster way to check the I/O "profile" of an application at any given time when queries on filesystem performance are received. We have been feeding back our experiences and site requirements to the Cray R&D team.

I/O analysis for applications

As mentioned above, a filesystem view of application I/O is useful but we also need a better understanding of how applications interact with the file system in application terms i.e. (open(), close(), read(), write()). The CoE has developed a tool that can analyse a profile of an application's I/O calls (or any system call interaction) and display this information in useful ways. The tool proved useful first to understand the I/O performance of a python driven application by using the option to generate a timeline of I/O operation sizes. This showed that there was a large number of small reads but at the same time many large (1MB) lustre-friendly reads. This sort of behaviour is hidden by gross

averages. The tool has proved useful in identifying other issues with applications that were non-obvious. We continue to refine this tool and use it to directly understand applications.

Auto-Tuning

This year we have continued our interaction with NCAS and have successfully used the auto-tuning framework to tune input parameters for runs of the Unified Model. This is a limited use of the framework which can for example inject code and parameters into applications. Initial effort has enabled runtime tuning of parameters (set in namelists and run scripts) relating to parallel decomposition and segment size.

New features to support parameter injection in input namelists, perform scientific validation of runs, and support multiple line build and execution commands were added. The source code along with the documentation has been moved to a git repository. The results were very positive showing how the framework can automatically determine the best from the set of parameters.

8.2 Training and workshops

The CoE assisted with various workshops during the year.

- Harvey Richardson presented two talks (covering Lustre architecture, internals and tuning) at the ARCHER Parallel I/O course in Durham on 30th March. It was interesting to hear about the challenges of other system administrators in supporting diverse I/O workflows on Lustre and GPFS.
- Karthee Sivalingam assisted with the *Hands-on Porting and Optimisation Workshop: Making the most of ARCHER* course held in Birmingham in April
- Harvey Richardson, Karthee Sivalingam and Adrian Tate visited the NCAS group to discuss current project status.
- Harvey Richardson attended the CDT in Pervasive Parallelism Student Showcase event held in May where he was able to discuss projects with students and, in particular, offer advice on projects using ARCHER.
- Karthee Sivalingam attended the first day of the ARCHER Champions meeting at STFC in June.
- Adrian Tate and Harvey Richardson attended and prepared for the GW4 ARM Hackathon in Bristol Nov 1st and 2nd.

8.3 Tier-2 HPC integration

As agreed with EPSRC, a small number of ARCHER CoE days have been allocated to support the GW4 Tier-2 HPC centre and for general tier-2 integration. Around 8 days of support were allocated to this activity, largely focused on the preparations for the Hackathon activities.

8.4 ARCHER queries and software

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

Of particular note were the following interactions.

We were able to identify why the EPCC benchio benchmark was performing sub-optimally for shared HDF5 files and provide an application workaround to flush metadata and recover the expected performance.

One project was looking for a way to analyse network traffic for a CFD application in a mosaic view with traffic attributed to levels in the network hierarchy. We developed a python tool to postprocess

the XML output from perftools to graph and export this information at scale. As a result the project was better able to understand the performance of alternative preconditioners.

We spent significant time understanding a longstanding (and very infrequent) issue with the parallel build of XIOS, after detailed profiling of file accesses in the build (to fractions of a second) we determined it was a problem with the build system and suggested various workarounds

8.5 Support of the eCSE programme

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