



ARCHER Service 2018 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	2019-01-10	Inputting initial information	Anne Whiting
0.2	2019-01-11	Adding USL information	Anne Whiting, Jo Beech-Brandt
0.3	2019-01-15	Added CSE information	Andy Turner, Alan Simpson, Lorna Smith
0.4	2019-01-16	Added CSE executive summary items	Andy Turner
0.5	2019-01-17	Added HPCSG information	Linda Dewar
0.6	2019-01-18	Added Cray CoE information	Harvey Richardson
0.7	2019-01-21	Added Cray service information	Felipe Popovics
0.8	2019-01-22	Added Maintenance session detail	Linda Dewar
1.0	2019-01-30	Version for EPSRC	Alan Simpson

Table of Contents

Document Information and Version History.....	2
Table of Contents.....	3
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. HPC Systems Group (HPCSG).....	10
5.1 Service failures.....	10
5.2 HPC Systems Group (HPCSG) Service activities.....	10
5.3 Principal activities undertaken (in addition to day-to-day operational activities) included:.....	10
6. Computational Science and Engineering (CSE).....	11
6.1 Benchmarking UK national HPC facilities.....	11
6.2 Understanding parallel I/O use on ARCHER in collaboration with CoE.....	12
6.3 Delivering a high quality, fair and transparent eCSE programme to enhance codes on ARCHER.....	13
6.4 Reaching new research communities through innovative training.....	14
6.5 Enhanced outreach materials and activities.....	14
6.6 Celebrating Diversity: Women in HPC Chapter.....	15
7. Cray Service Group.....	16
7.1 Summary of Performance and Service Enhancements.....	16
7.2 Reliability and Performance.....	16
7.3 Service Failures.....	16
7.3.1 Full Service failures.....	16
7.3.2 Partial Service failures.....	16
7.3.3 Scheduled maintenance activities.....	16
8. Cray Centre of Excellence (CoE).....	17
8.1 The LASSi framework and related work.....	17
8.2 Training and workshops.....	18
8.3 ARCHER queries and software.....	18
8.4 Support of the eCSE programme.....	19

1. Introduction

This annual report covers the period from 1 Jan 2018 to 31 Dec 2018.

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 HPC Systems Group;
- 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 HPC Systems Group report in Section 5 describes activities and progress in their main areas of responsibility.

In Section 6 the CSE team describe a number of highlights of the work carried out by the team in 2018.

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/2018.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 for the overall ARCHER service. More details are provided in the appropriate section of the document.

- EPCC were delighted to be able to announce that they had been awarded ISO 27001:2013 Information Security certification after the external audit in September 2018. The scope covered all the HPC and Data Services delivered by EPCC, including ARCHER, the RDF, the KNL and Cirrus. The certification demonstrates the importance EPCC places on keep user data secure and in following best practice in the era of data analytics and GDPR.
- The last user survey received very positive scores and feedback from the user community with an increase in mean score from 4.3 in 2016 to 4.4 in 2017 (out of a maximum of 5.0). As with previous years the highest mean score was achieved by the Helpdesk (4.6); the greatest positive increase in score was for webinars, from 3.9 to 4.2.
- Two new key senior staff joined the HPC Systems team, bringing a wealth of experience with them. Paul Clark joined EPCC as Director of High Performance Computing Systems, leading both the Tier-1 and Tier-2 systems teams to provide an integrated cross-service team. Calum Muir joined EPCC as Data Centre Manager bringing expertise in the electrics, cooling and plant side of things so vital to keeping ARCHER and Cirrus running optimally.
- The CSE team has been leading an open source benchmarking effort to run the benchmarks across the national HPC services, e.g., ARCHER and Tier 2. The aims are to provide information for users to help them choose the correct HPC resource for their research and to provide service providers with useful data to support future procurements. This work involves collaboration with the Tier 2 RSE teams and the wider UK HPC community.
- In 2018, one of the main achievements of the Cray CoE was to bring the LASSi I/O analysis framework to maturity and to use it both to investigate issues and understand performance of the Lustre filesystem.
- The Cray LASSi development has enabled the CSE team to collaborate with the Cray CoE and the IO-500 developer (Julian Kunkel, University of Reading) to integrate parallel I/O performance data from the Cray LASSi tool into SAFE. This will allow users to understand how their applications use parallel I/O on the system without the need for any alterations to their application or intervention on their part. This will also allow us to investigate how different research communities and applications on ARCHER use parallel I/O to improve the service and to provide advice for future procurements.
- The CSE Training team has trialed innovative, online training modeled on the approach used by the US XSEDE programme but modified to suit the ARCHER community to engage different communities of HPC, and potential HPC, users. These courses allow users to access the technical expertise of the CSE trainers without having to travel or take a large single block of time out of their normal working schedule. This allows the training to be available to attendees who would find it difficult to attend our standard, face-to-face courses.
- 2018 saw the final eCSE call and panel take place. Over its lifetime, the eCSE programme has funded 100 projects to improve software on ARCHER and to develop the technical expertise of UK-based researchers. Our impact analysis shows that that the eCSE programme has had wide ranging impacts: 13% of the projects were for new communities to allow them to exploit ARCHER for their research; 40% of the top 40 codes have benefitted from eCSE development effort; and as of January 2019 eCSE work has resulted in an estimated saving of £24.5M worth of compute time on ARCHER.
- Cray Service Group noted that 2018 has been another excellent year for the ARCHER service with a very stable technology environment facilitating high resource utilisation for the user community.

3. Service Utilisation


3.1 Overall Utilisation

Utilisation over the year was 86%, slightly down from 88% in 2017.



3.2 Utilisation by Funding Body

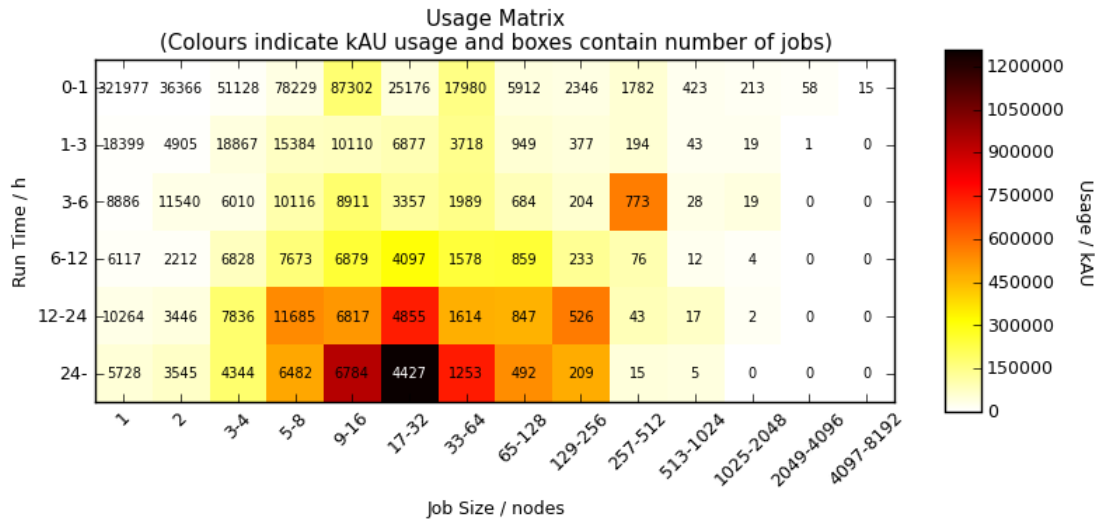
The utilisation by funding body relative to their allocation can be seen below.



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 2018.



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 levels met. A total of 6,551 queries were answered by the Service Provider, down from 10,082 queries during 2017, reflecting the maturity of the service. Over 99% were resolved within 2 days. In addition to this, the Service Provider passed on 149 in-depth queries to CSE and Cray.

	18Q1	18Q2	18Q3	18Q4	TOTAL
Self-Service Admin	1269	1411	969	1372	5021
Admin	333	318	298	344	1293
Technical	63	59	59	56	237
Total Queries	1665	1788	1326	1772	6551

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 Registrations.

	18Q1	18Q2	18Q3	18Q4	TOTAL
Phone Calls Received	79	60	95	80	314
Change Requests	4	1	2	1	8
User Registration Requests	250	186	179	169	784

It is worth noting that the volume of telephone calls was low throughout the year. Of the 314 calls received in total, only 38 (12%) 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 2017

The results of the 2017 annual ARCHER User Survey was run in February 2018. 164 responses were received compared to 161 in 2017, 230 in 2015 and 153 in 2014, with the mean results shown below (scores 1 representing "Very Unsatisfied" and 5 representing "Very Satisfied"):

Service Aspect	2014 Mean Score (out of 5)	2015 Mean Score (out of 5)	2016 Mean Score (out of 5)	2017 Mean Score (out of 5)
Overall Satisfaction	4.4	4.3	4.3	4.4
Hardware	4.1	4.1	4.2	4.3
Software	4.0	4.0	4.2	4.1
Helpdesk	4.5	4.5	4.5	4.6
Documentation	4.1	4.1	4.2	4.2
Website	4.1	4.2	4.2	4.2
Training	4.1	4.1	4.2	4.1
Webinars	3.6	3.9	3.9	4.2
Online training	-	4.0	4.1	4.2

As with previous years the highest mean score was achieved by the Helpdesk (4.6); the greatest positive increase in score was for webinars, from 3.9 to 4.2. Mean satisfaction had risen from 4.3 in 2016 to 4.4 in 2017. The full report can be found at <http://www.archer.ac.uk/about-archer/reports/>.

ISO 27001:2013 Information Security Certification

EPCC were delighted to be awarded ISO 27001:2013 Information Security certification after the external audit in September 2018. The scope covered all the HPC and Data Services delivered by EPCC, including ARCHER, the RDF, the KNL and Cirrus. The certification demonstrates the importance EPCC places on keep user data secure and in following best practice in the era of data analytics and GDPR.

Red Weather Major Incident

In February 2018 the ARCHER service staff worked through the snow of the red weather warning, keeping the service running uninterrupted despite blizzards, and buildings and transport being shut down. The ARCHER service declared a Major Incident and used the experience to test and improve our Major Incident and Disaster Recovery processes. EPCC were very pleased how well the service ran, and very much appreciated the support and good wishes received from the user community.

GDPR and ARCHER and SAFE data policies

The ARCHER service policies were reviewed to ensure they comply with the requirements of the General Data Protection Regulation (GDPR), to meet the deadline of 25 May 2018 when the new legislation came into use.

SAFE improvements

Changes have been made to the SAFE to enable it to support two-factor authentication. The functionality uses the same standard mechanism to implement additional time based one-time passwords as Google and Azure, so users can make use of the smart-phone apps with the SAFE. Two-factor authentication is optional with the standard password method still supported.

The SAFE is now importing statistics of operations on the work file-systems; data is stored for each job run on the system. This allows reports to be written to analyse IO-use on the system and the IO requirements of different application codes. A number of enhancements have been made to the report generator to help analyse this data.

Improvements made in response to user feedback

The ARCHER team prides itself on listening to its users and making requested improvements where practical. At the request of ARCHER users:

- the short development queue availability times have been increased from the previous times of 08:00 – 17:00, five days a week to run 24 hours a day, 7 days a week. This move supports the development work carried out by ARCHER users out of standard office hours, and reflects the importance that EPCC places on ensuring the ARCHER service meets user requirements.
- the maximum job length on the KNL was increased from 8 hours to 24 hours.

ISO 9001:2015 Quality Management Certification

The ISO 9001:2015 quality management certification for the Tier 1 and Tier 2 services was retained in February 2018, passing the external audit with very few audit findings. The feedback from the external auditor was very positive about the improvements made since the inaugural audit, with the improvements in risk management particularly highlighted. The use of the quality management system as a framework for service improvement has delivered service benefits to the user community. An example of this was the scenario test carried out of the Major Incident and Disaster Recovery processes which led to process improvements which helped us to deliver an uninterrupted service during the red weather shutdown in February 2018.

5. HPC Systems Group (HPCSG)

5.1 Service failures

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

5.2 HPC Systems Group (HPCSG) Service activities

New senior staff and team integration at the ACF

Two new key senior staff joined the HPC Systems team, bringing a wealth of experience with them. Paul Clark joined EPCC as Director of High Performance Computing Systems, leading both the Tier-1 and Tier-2 systems teams to provide an integrated cross-service team. Calum Muir joined EPCC as Data Centre Manager bringing expertise in the electrics, cooling and plant side of things so vital to keeping ARCHER and Cirrus running optimally.

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

One combined Systems team

Under the leadership of Paul Clark, the two systems teams, OSG which managed the Tier 1 ARCHER service and SDT which managed the Tier 2 Cirrus service, have been combined into one HPC Systems Group (HPCSG). This has facilitated enhanced knowledge sharing, and provided additional resilience in staff numbers and skills.

Information Security

HPCSG had a key role in the work to achieve the ISO 90027:2013 Information Security Certification. As a part of their Information Security monitoring activities, they continuing to evaluate and respond to CERT advisory notices and apply mitigation measures. Work was carried out in 2018 to test mitigations for the well-publicised Spectre/Meltdown vulnerability and to plan the service approach to these.

Minimising user disruption through reduced numbers of maintenance outages

HPCSG has continued its efforts to reduce the number of planned maintenance sessions and thus the disruption to the service for users. Wherever possible tasks are now carried out *at risk* rather than requiring a systems outage, and where planned outages are required these are taken jointly with Cray to minimise user disruption. Proposed changes are also tested on the TDS before implementing these in production. In 2018, 4 planned maintenance sessions were taken as compared to 7 in 2017.

Working with Cray staff to maintain and improve ARCHER and reduce risk of service interruption

HPCSG worked closely with the Cray team to keep the system patched, applying field notices and patch sets according to Cray recommendations. Improvements have been made to the system monitoring tools used in order to proactively resolve issues before they become user or system affecting. Working has been carried out with Cray onsite staff to provide the systems support for the infrastructure required for the development of the Lassi tool.

Work on power at the ACF

Work has taken place to evaluating and modifying the power configuration in the computer room housing ARCHER to allow redundant connections to the ARCHER management and storage racks for mitigation against power events that may affect this equipment. Work was also carried out to mitigate the effects of problems with power supply units within login nodes as was experienced on 22 June.

Red weather major incident

The Red Weather incident in March 2018 meant that for the first time systems staff had to close the ACF and work remotely for a significant period of time. Due to planning and testing for such an eventuality, the service was kept running uninterrupted with all plant and supporting equipment regularly checked, the system monitored and with support provided remotely.

6. Computational Science and Engineering (CSE)

6.1 Benchmarking UK national HPC facilities

Following the start of the UK national Tier-2 HPC services and the publication of the updated ARCHER benchmarks¹ in 2017, the ARCHER CSE team has pioneered an “Open Source” benchmarking approach to compare the different UK national HPC services using the ARCHER benchmarks. This work has the aims of:

- providing UK researchers with the information required to make the best choice of HPC system for their research;
- understanding the differences in performance between different HPC systems for different benchmarks to help inform code development and future procurements;
- providing a useful library of performance data for other parties to compare to;
- sharing benchmarking good practice and experience across the UK and worldwide HPC community;
- demonstrating the leading role that the UK HPC community is taking in HPC benchmarking and prototyping novel HPC architectures.

The ARCHER CSE service has created a public repository on GitHub at: <https://github.com/hpc-uk/archer-benchmarks> with public access to all the details of the benchmarking effort. One of the major issues faced when using benchmarking results in the public domain is that a quantitative comparison is required but is often not possible as the published results do not include the required details for a valid comparison. To ensure that our results do not fall foul of these issues we have adopted an “Open Source” approach to our benchmarking work. This means that the public repository contains all the information and data to compare to our results, and to reuse our data. The items available in the repository include:

- Raw log file output from the benchmark runs (to allow for reuse of the data by interested parties)
- Full details of how the benchmark applications were compiled on different platforms
- Full sets of input data for the benchmarks
- Job submission scripts used to submit the benchmark runs on the different HPC systems
- Jupyter Python notebooks with the full detail of how data was extracted from the raw output logs, and how it was analysed to produce the performance data used for comparisons
- Freely-available reports comparing the performance across different HPC systems and providing advice to researchers on choosing the correct platform for their work

This benchmarking effort has also helped to build a collaborative community of RSEs involved in delivering UK national HPC services, as bringing together results across the Tier-1 and Tier-2 national HPC services has required collaboration, discussion and review of results across the different sites. The benchmarking effort to date has been led by the ARCHER CSE team but has also benefitted from substantial contributions from the University of Cambridge RSE team, the GW4 Isambard RSE team and the RSEs supporting the High-End Consortium for Biomolecular Simulation (HEC BioSim). Comparative benchmarking was also one of the main areas of discussion at the Tier-2 RSE/ARCHER Champions Workshop associated with the RSE’18 conference and this discussion was facilitated by the ARCHER CSE team.

We plan to continue to develop the benchmarking repository and community throughout 2019 with:

¹ http://www.archer.ac.uk/documentation/white-papers/benchmarks/UK_National_HPC_Benchmarks.pdf

- the addition of machine learning benchmarks to the set of benchmark applications;
- the incorporation of the IO-500 benchmark to allow for better comparison of the I/O performance of different systems;
- ensuring that benchmarking continues to feature in future UK HPC Champions workshops to continue to build the community;
- incorporation of results from the new multi-node ARM-based HPC systems (Isambard, Catalyst UK);
- continuing discussions with the DiRAC RSE group on incorporating their benchmarks into the performance comparison exercise.

6.2 Understanding parallel I/O use on ARCHER in collaboration with CoE

Although the performance of I/O in HPC applications has been of importance and interest to particular application classes and research areas for a long period of time, more recently it has become of wider importance and interest for a number of reasons, including:

- the increased scale of simulations and models run by researchers across many different research areas;
- the increased use of statistical methods that involve high numbers of model runs (e.g. parameter sweeps or meta-dynamics);
- the increasing use of programming languages such as Python where the I/O model works well on a local resource (such as a laptop or workstation) but runs into issues when scaled up onto a shared HPC resource with a shared parallel file system.

Key to addressing the challenges posed by the wider impact of I/O performance on HPC users is access to good quality, quantitative data on I/O use by users and service providers. Users require I/O performance data at the scheduler job level to assess the different I/O requirements of different run types and the performance variation they can expect to see across multiple instances of similar runs. Service providers require statistical I/O performance data aggregated across different classifiers; for example, different file systems, research areas, software applications, projects and combinations of the above. Ideally, this data should be automatically collected for all jobs run on the HPC system with little/no performance impact and made available automatically to the different audiences.

Throughout 2018, the CSE team and Cray CoE have been collaborating to integrate outputs of the Cray LASSi I/O monitoring tool into the EPCC SAFE service management and reporting framework to meet these requirements. We have also been working together to analyse the I/O usage on ARCHER and explore how outputs from LASSi can be best used to support the service and its users.

The Cray LASSi tool samples metrics from Lustre on how the file system is being used every 3 minutes, and associates metrics with particular application runs through the Cray ALPS scheduler logs. This effectively means that LASSi automatically produces Lustre I/O profiles for every job run on the system. In addition, LASSi produces derived metrics (at both a job and file system level) that quantify how the file system is being used (in terms of bandwidth and metadata use) and also quantify the risk level that users will see degraded I/O performance due to the conditions on the file system at a particular time. Furthermore, LASSi produces per-job risk level metrics, so that jobs that are potentially causing issues for the file system can be identified and investigated, often in collaboration with the user themselves.

Although the LASSi tool produces a lot of useful data, it does not have the ability to group metrics according to different classifiers, such as research area, application used, or project ID. LASSi also lacks a general user interface that would allow individual users (or project leaders) to query the data and analyse their own or their project's I/O use. Both of these features are required to realise the full potential of the metrics produced by the LASSi tool. The EPCC SAFE tool contains all the required classifiers, has the ability to link data from multiple sources, and provides a flexible reporting interface

that is already used by the users to analyse their use of ARCHER. We have configured SAFE to take a feed of per-job I/O performance metrics from ARCHER and link these to the other per-job records in the database. This has allowed us to start to analyse I/O use on ARCHER grouped by arbitrary classifiers over any period for which we have LASSi data. It is also allowing us to develop self-service reports for ARCHER users that allow them to explore the I/O performance that their applications are seeing on the system without the need for any instrumentation or intervention on their part.

In late 2018, the collaboration was widened to include Dr Julian Kunkel, a parallel I/O expert working at the University of Reading. His focus has been on using his IO-500 benchmark to provide a “control” application to understand better how different I/O patterns (in terms of bandwidth and metadata use) impact the metrics produced by LASSi, and how well synthetic benchmarks such as IO-500 model the I/O patterns seen for real HPC applications on a national service such as ARCHER.

We plan to submit a joint paper to the Cray User Group (CUG 2019) conference, which is taking place in Montreal in early May, describing the work and presenting initial analyses of the LASSi I/O data produced from ARCHER. We are also in the process of producing SAFE reports that users can use to query their own I/O use of the system. These should be available in Q1 2019 and we will publicise them to all users and seek feedback to further improve the reporting tools available. Finally, we will be producing a white paper with an analysis of I/O use on ARCHER using the SAFE/LASSi integrated reporting, which will be published on the ARCHER website in early 2019.

It is worth noting that this work would not have been possible without the close day-to-day collaboration between the ARCHER CSE team at EPCC and the ARCHER Cray Centre of Excellence. The co-location of ARCHER CSE staff and Cray CoE staff has been invaluable to this initiative.

6.3 Delivering a high quality, fair and transparent eCSE programme to enhance codes on ARCHER

The eCSE programme has allocated funding to the UK computational science community through a series of funding calls over a period of 6 years. The eCSE programme has been a significant source of funding for RSEs across the UK providing an important mechanism for new RSE groups to retain expert staff. The goal throughout has been to deliver a funding programme that is fair, transparent, objective and consistent. The programme has enhanced the quality, quantity and range of science produced on the ARCHER service through improved software, while developing the computational science skills base, and providing expert assistance embedded within research communities, across the UK.

During 2018, the last of the funding was successfully allocated. In total there were 13 calls, with 100 projects funded. The programme was not-for-profit, with all funding awarded to projects. This resulted in an additional 32 project months being funded, over and above the contractual 14 FTEs per year. The range of institutions and people involved in the eCSE programme has been significant, with 45 different organisations involved, distributed all across the UK, and 71 technical people funded to work on eCSE projects. In addition to projects in all the major traditional scientific areas, 13% of the projects developed codes for new communities, helping to grow and enhance the ARCHER user base.

The enhanced quality of the software can be seen on the ARCHER system. In the last 6 months over 40% of the 40 most heavily used codes on the system have benefited from eCSE funding investment. Many eCSE projects also generate significant financial savings, which in turn are reinvested to allow scientists to achieve enhanced science from the same amount of resource. While not all projects contribute to this particular benefit, to date the financial saving is estimated at £24.5M. This is balanced against an overall cost of the eCSE programme of around £6M, showing a 4x return on investment.

6.4 Reaching new research communities through innovative training

A focus of the ARCHER training programme in 2018 was to try and reach new HPC users and communities. We addressed this by developing a series of programming courses delivered entirely online, so that all potential users could attend without having to travel.

The US XSEDE programme has developed an HPC Workshop Series, where courses are delivered as interactive webinars. The MPI course illustrates the format: a presenter at Pittsburgh Supercomputer Centre delivers lectures via videoconference, with users able to ask questions via an associated WebEx session. Attendance is restricted to around a dozen satellite locations across the US, so attendees can choose a nearby location to minimise travel.

Having attended an XSEDE MPI Workshop as a special guest satellite site, we decided to modify this format for the ARCHER training programme. The changes we made were:

- Use a single technology - Blackboard Collaborate - for both slide delivery and interaction with users. This approach has already proved successful for the ARCHER Virtual Tutorials and also enables lectures to be recorded and uploaded to YouTube.
- Enable users to attend individually from their own laptop rather than having to attend specific locations (although particular sites can still choose to use a single room for multiple attendees).
- Run a two-day course across four successive Wednesday afternoons rather than as a solid block, to give time for attendees to get to grips with the practical examples and to review the material.
- Provide opportunities for peer support via a shared chat page.

This format has proved very successful and in 2018 we ran versions of both the MPI and OpenMP courses. The complete courses, including PDF slides, practical exercise material and videos of the lectures, are available at:

<http://www.archer.ac.uk/training/courses/2018/01/OnlineMPI/>
http://www.archer.ac.uk/training/courses/2018/10/openmp_online/

As well as delivering lectures, we can also run interactive quiz sessions or perform live programming examples to illustrate key points.

We plan repeat runs of both these courses in 2019. To improve them further, we will encourage attendees to make more use of the chat pages to provide peer support for the practical exercises between the webinars. We will also explore identifying additional contacts from the ARCHER community, HPC Champions and wider RSE community who can provide one-to-one practical advice to reduce the support load on the single trainer.

6.5 Enhanced outreach materials and activities

Through the lifetime of the ARCHER service, EPCC has been very active in developing and delivering outreach activities to school children and the general public. Our aim has been to explain the real-world benefits of the ARCHER service and to encourage young people to consider careers in computational science. We have participated in various large-scale outreach events across the UK and developed our own suitcase-sized Supercomputer, Wee Archie. Wee Archie has been hugely popular and we have used it at various events across the UK. Having established Wee Archie and a range of demos, in 2018 we focused on enhancing the activities we offer to complement this, with a particular focus on programming and algorithms as well as providing a more visual experience of ARCHER.

Firstly, we looked to provide a better sense of what ARCHER really is, how it works and what it looks like. To do this we have produced a set of to-scale display panels, showing the inside of a cabinet front

and back, and have also obtained a real Cray XC30 board for participants to handle. This has been particularly popular at science events and we are grateful to Cray for providing both the board and photographs.

A particular highlight has been the development of a series of practical demonstrations based around programming, problem solving and algorithms. Problem solving and giving clear instructions are at the heart of good programming and we have developed a set of logic puzzles to demonstrate this in a fun and accessible way, these puzzles range from chasing zombies to crossing rivers with chickens! Coupled with other new algorithm demonstrations, such as a practical example of the towers of Hanoi, we now have a suite of demonstrations that can be tailored to different audiences and events.

These new activities have been tested at events across the UK, for example at the Big Bang Fair in Birmingham, New Scientist Live in London, and at the Edinburgh International Science Festival.

Throughout 2018 we have also placed a focus on engaging students and secondary school pupils, while increasing and training more ARCHER staff members in outreach skills.

A key highlight has been our work experience opportunities, which have been offered to a series of secondary school pupils. Working with the ARCHER team for a week, the pupils helped develop and test outreach activities, while spending time visiting and learning about the ARCHER system. The aim was to provide opportunities for students to experience the world of work, looking to raise aspirations and help pupils make the right career choices. The input of these young people around our outreach activities has also helped enhance our outreach delivery. Coupled with this we had summer students working with us to programme the LED lights attached to each node of Wee ARCHIE, and to develop our Faces of HPC material. We plan to continue to develop our work experience programme in 2019, with an emphasis on students from programmes targeted at widening participation.

6.6 Celebrating Diversity: Women in HPC Chapter

EPCC has a long-established commitment to promoting diversity in computational science and STEM more broadly. We feel passionate about inclusivity and about supporting young people to consider careers in STEM, no matter their gender identity, sexual orientation, social background or ethnicity. As one of the founders of the Women in HPC organisation, a highlight of 2018 has been our success in becoming a Women in HPC Chapter. The UK-wide chapter aims to celebrate and encourage diversity in computational science and allows us to act as an ambassador for the UK. Initial efforts have focused on enhancing our case studies (See Faces of HPC: <https://www.hpc-diversity.ac.uk/faces-of-hpc>), on developing a workshop focused on secondary school children, and on organising a launch event in 2019.

7. Cray Service Group

7.1 Summary of Performance and Service Enhancements

2018 has been another excellent year for the ARCHER service with a very stable technology environment facilitating high resource utilisation for the user community.

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 most 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.

7.3 Service Failures

7.3.1 Full Service failures

Only one incident classified as a full-service failure was encountered during 2018.

ARCHER full service failures 2018		
Incident	Date	Description
1	21/08/18	Expired SSL certificates on the ARCHER work filesystems caused a Sonexion storage node to crash and would not reboot after its High Availability partner failed due to a hardware failure. Users lost the ability to use one of the three work filesystems and 466 XC-30 nodes were temporarily unavailable during the incident.

7.3.2 Partial Service failures

Although most technology failures do not impact on the user community, there were exceptions during 2018, with three incidents of Elogin node rebooting due to PCI faults on slots containing Myricom 10G-PCIe cards. Users running on the nodes at the time of the failures lost their connections. The latest firmware for the card has been downloaded and installation is planned for Q1 2019.

7.3.3 Scheduled maintenance activities

None of the permitted four dedicated maintenance sessions was taken by Cray in 2018, ensuring that the ARCHER system remained in the hands of the user community as much as possible.

8. Cray Centre of Excellence (CoE)

In 2018, one of the main achievements of the CoE was to bring the LASSi I/O analysis framework to maturity and to use it both to investigate issues and understand performance of the Lustre filesystem. This is described below along with other activities during the year.

8.1 The LASSi framework and related work

We had a range of activities relating to I/O this year. The helpdesk periodically receives 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. Last year we developed tools to assist with this: the LASSi analysis tool and a profiling report tool.

LASSi (Log Analytics for Shared System resources with instrumentation) analyses application I/O by combining Lustre filesystem I/O statistics with application information from the Cray ALPS (Application Level Placement Scheduler) logs. The Lustre data is ingested from a bespoke tool (previously developed by the onsite Cray system staff) that collects raw Lustre data at 3 minute intervals. This effectively means that LASSi automatically produces Lustre I/O profiles for *every* job run on the system. In addition, LASSi produces derived *risk* and *ops* metrics that relate to unusually high application IO behaviour. *risk* metrics measure the quantity of IO and *ops* metrics measure the quality. The metrics at both a job and file system level quantify how the file system is being used by applications and the *risk* that users will see degraded I/O performance due to the conditions on the file system at a particular time. This year LASSi was moved to a custom VM on the RDF cluster and a workflow created to ingest data, analyse data and generate daily reports. New automated daily reports (uploaded to a website) help system support staff not only analyse slowdown events reported by users but also identify and investigate jobs that are potentially causing issues for the file system.

LASSi was enhanced to work at a finer time resolution (6 minutes) and to provide raw I/O statistics in addition to the risk-based data. LASSi was updated to produce (manually) detailed risk and IO reports/plots of individual applications. In addition, a new summary data export was also implemented that will eventually allow users to obtain I/O summary information about jobs on ARCHER from SAFE.




Figure 1: Application metadata statistics

A new tool was developed to directly export and generate graphs of Lustre statistics for individual applications or multiple applications being run as a task farm (see example Figure 1).

LASSi is being used to assist with understanding user complaints of slow filesystem performance and it greatly speeds up the time taken to identify what is causing an issue. It has also enabled us to pinpoint applications that make unusual demands of the filesystem or use it inefficiently (for example by making small I/O requests). In collaboration with the ARCHER CSE support team we have started to

analyse application I/O behaviour more generally, looking for patterns in the I/O usage of applications, groups or science fields.

Following initial analysis of LASSi data, the CoE started to contact users during the year in order to understand more about particular risky jobs; we started with jobs showing high metadata rates. In some cases, the CoE also used another tool (previously developed) to gain insights from a system call profile of the jobs of interest. This activity is ongoing and we hope to use the knowledge gained as a basis for an I/O workshop in 2019.

8.2 Training and workshops

The CoE undertook or attended various events during the year:

- Harvey Richardson attended the EPSRC Centre for Doctoral Training in Pervasive Parallelism Industrial Engagement event, 10 April, to present to the Students about Cray's interests, possible internships and discuss their projects (some using ARCHER)
- ARCHER Champions Meeting, 25 April, Karthee Sivalingam gave a presentation outlining the latest status of the LASSi project, there was good interest and questions from the audience and follow up discussions.
- EPCC Seminar, 9 May, Martin Lafferty (Cray, on-site system team): ARCHER hardware talk and board hands-on
- EPCC seminar, 18 May, Karthee Sivalingam: LASSi – How CSE and Helpdesk can use it?
- ARCHER Webinar, Cray Programming Environment Update, May 16. Harvey Richardson outlined how and why updates are performed on ARCHER and gave a detailed explanation of which new features arrived with various updates installed and planned to be installed.
- NCAS, 9 July (discussion of I/O optimization/characterization and possible collaboration)
- UK RSE Conference, 3-4 September
- UKCSF, 11th September at ACF (presentation on LASSi and discussion of approach with service representatives from ARCHER, ECMWF and MetOffice)
- ECMWF HPC Workshop, 24th to 28th September

As agreed with EPSRC, a small number of ARCHER CoE days have been allocated to support the GW4 Tier-2 centre and for general Tier-2 integration. This largely focused on the preparations for the 2nd Hackathon activities.

8.3 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 investigated a complex issue with compilation scripts running on newly provisions MAMU (Multiple Application Multiple User) nodes. The problem was traced to a network configuration that was limiting port availability.
- We resolved a problem in building GROMACS by providing a workaround to a cmake-related issue.
- We provided and tested a recipe to support native R access to Cray MPICH via the pbdMPI package.
- Unusually we were asked to explain why Cray MPICH had improved performance for point-to-point operations as the change had made custom application optimizations irrelevant. We

pointed to specific improvements in the MPI software after consulting with the MPI developers.

8.4 Support of the eCSE programme

The CoE completed technical assessments and final reviews for the eCSE calls during the year, and staff attended the eCSE Panel meeting in September. The CoE worked with parts of the service to investigate the new MPI dynamic process management (DPM) feature support in Cray MPI as this was required by one of the proposals.