ARCHER Service
2019 Annual Report

Document Information and Version History

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

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2020-02-18	Inputting initial information	Anne Whiting
0.2	2020-03-02	Additional information	Lorna Smith
0.3	2020-03-31	Addition of Cray CoE text	Harvey Richardson
0.4	2020-04-07	Additional text and review	Lorna Smith
0.5	2020-04-08	Minor corrections	Anne Whiting
0.6	2020-04-10	Text from Cray Systems	Felipe Popovics, Steve Jordan
0.7	2020-05-01	Reviewed	Alan Simpson
1.0	2020-05-05	Version for the Authority	Alan Simpson

Table of Contents

1. Introduction

This annual report covers the period from 1 Jan 2019 to 31 Dec 2019.

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

- The ARCHER system continued to be busy, with utilisation over the year at 88%. A total of 7,246 queries were answered by the Service Provider with 99% resolved within 2 days.
- EPCC passed a four-day combined external audit of ISO 9001 Quality Management and ISO 27001 Information Security Management. This ensures EPCC delivers the best and most secure service to our users.
- A Business Continuity and Disaster Recovery (BCDR) scenario test was carried out in October. There were no interruptions to service and the user community was unaware that it had taken place. This test helps prepare our staff and to improve our processes to ensure we keep the services running as best we can in case of major incidents in the future.
- In order to facilitate faster data movement to the RDF GPFS filesystems from the ARCHER login nodes, the existing bonded pair of 10gbit links from the ARCHER core switches to the RDF sitewide network has been upgraded to a pair of 40gbit links. The increase in transfer speed will assist users in transferring data from ARCHER to the RDF.
- The ARCHER eCSE programme has provided funding to the ARCHER user community to develop software in a sustainable manner to run on ARCHER and on future Tier-1 services. To date cumulative benefits of almost £35M have been shown using our benefits realisation techniques, representing a 5-fold return on investment.
- With the wealth of national HPC facilities available for researchers to choose from for their computational work (ARCHER, DiRAC, Tier-2 HPC), we have performed a significant amount of work comparing the performance of HPC applications across different platforms. This work has resulted in a public repository of benchmarks.
- Recently, the ARCHER CSE service has led an initiative that coordinates with HPC Champions and the wider HPC RSE community to consider how we could create a community UK HPC technical knowledge base that would allow us to share and reuse useful technical information and experience.
- In 2019, the CoE took further advantage of our LASSi I/O analysis framework both in analysing user problems and to support studies into I/O usage of applications and communities on ARCHER.

3. Service Utilisation

3.1 Overall Utilisation

Utilisation over the year was 88%, up slightly from 86% in 2018.

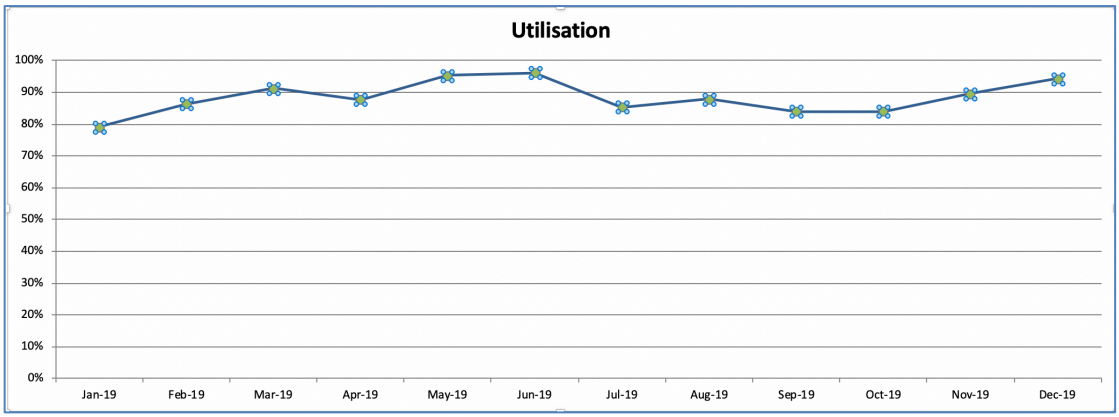

3.2 Utilisation by Funding Body

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

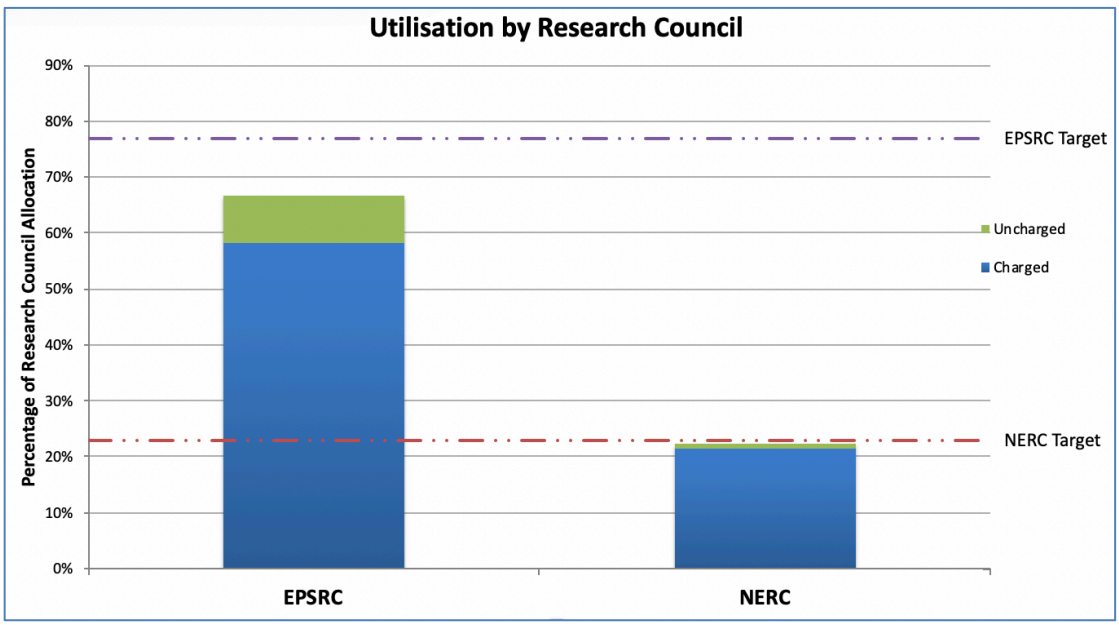

This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER.

3.3 Additional Usage Graph

The following heatmap provides a view of the distribution of job sizes on ARCHER throughout 2019.

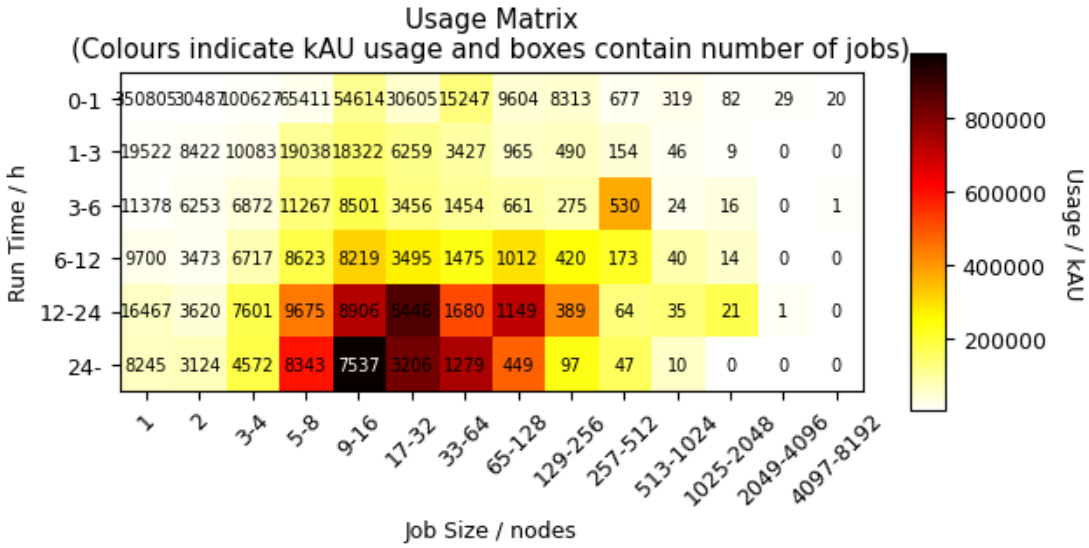

The heatmap shows that most of the kAUs are spent on jobs between 192 cores and 3,072 cores (8 to 128 nodes). The number of kAUs used is closely related to money and shows how the investment in the system is utilised.

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 7,246 queries were answered by the Service Provider, up from 6,551 queries during 2018. 99% were resolved within 2 days. In addition to this, the Service Provider passed on 141 in-depth queries to CSE and Cray.

	19Q1	19Q2	19Q3	19Q4	TOTAL
Self-Service Admin	1796	1361	1089	1365	**5611**
Admin	435	303	307	337	**1382**
Technical	74	64	69	46	**253**
Total Queries	***2305***	***1728***	***1465***	***1748***	***7246***

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.

	19Q1	19Q2	19Q3	19Q4	TOTAL
Phone Calls Received	63	63	47	48	**221**
Change Requests	4	2	1	3	**10**
User Registration Requests	293	212	246	225	**976**

It is worth noting that the volume of telephone calls was low throughout the year (221 calls). All phone calls were answered within 2 minutes, as required.

4.2 USL Service Highlights

User Survey 2018

The results of the 2018 annual ARCHER User Survey were run in February 2019. 188 responses were received compared to 164 in 2017, 161 in 2016, 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)	2018 mean Score (out of 5)
Overall Satisfaction	4.4	4.3	4.3	4.4	4.5
Hardware	4.1	4.1	4.2	4.3	3.9
Software	4.0	4.0	4.2	4.1	3.8
Helpdesk	4.5	4.5	4.5	4.6	4.5
Documentation	4.1	4.1	4.2	4.2	4.0
Website	4.1	4.2	4.2	4.2	4.0
Training	4.1	4.2	4.2	4.1	4.3
Webinars	3.6	3.9	3.9	4.2	3.9
Online Training	-	4.0	4.1	4.2	3.9

As with previous years the highest mean score was achieved by the Helpdesk (4.5). Mean satisfaction had risen from 4.4 in 2017 to 4.5 in 2018, the highest overall satisfaction score for the service to date. The full report can be found at http://www.archer.ac.uk/about-archer/reports/.

Combined ISO 27001:2013 and ISO 9001:2015 Certification Success

EPCC was delighted to be able to announce that they have passed a four-day combined external audit of ISO 9001 Quality Management and ISO 27001 Information Security Management. ARCHER and Cirrus, our Tier-2 service, are both in scope for these certifications. The success in achieving these certifications reflects the importance we place on delivering the best and most secure service to our users and to taking action on feedback received to improve our service.

BCDR Scenario Test

A Business Continuity and Disaster Recovery (BCDR) scenario test was carried out in October. The aims of the test are to verify the processes in place, identify improvements and to ensure that staff have had the necessary training to ensure that should such an event occur we can maintain an uninterrupted service to our users. More details of this successful scenario test are given in the CSE section.

Running weekend queue during the Christmas holidays

To contribute to the utilisation of ARCHER over the holidays, the weekend queue was in operation for the whole of the festive period. The use of the weekend queue contributed to ensuring a higher utilisation of ARCHER during the holiday period, with utilisation in December 2019 at 94% compared to 80% and 72% in December 2018 and December 2017 respectively.

Preparation for the end of the ARCHER2 service

A data migration webinar was run several times to provide assistance for the user community in planning the data migration required for ARCHER2. The recording was made available on the ARCHER website and the data migration guide has been updated with the same guidance information.

SAFE improvements

After user feedback received, two new reports were added to SAFE to enable users and group leaders to run the Cray Lassi reports to show parallel IO use by user and by group. The data includes a summary of data written and read, and statistics from individual jobs giving data written and read and write operations.

5. HPC Systems Group (HPCSG)

5.1 Service failures

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

5.2 Principal activities undertaken

In addition to day-to-day operational activities, principal activities undertaken included the following:

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 with a particular aim to provide stability for ARCHER towards the end of service.

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. In 2019, only one full maintenance session was taken, compared with 4 in 2018 and 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.

Scheduler Upgrade

The PBS scheduler was upgraded from version 12.2.401 to 13.0.412 enabling users to benefit from the new features offered.

Faster data movement to the RDF

In order to facilitate faster data movement to the RDF GPFS filesystems from the ARCHER login nodes, the existing bonded pair of 10gbit links from the ARCHER core switches to the RDF sitewide network has been upgraded to a pair of 40gbit links. The increase in transfer speed will assist users in transferring data from ARCHER to the RDF.

6. Computational Science and Engineering (CSE)

6.1 Business Continuity and Disaster Recovery Scenario

The CSE and SP teams ran a joint test of ISO-9001 Business Continuity and Disaster Recovery processes on Tuesday 8th October. This is in line with our commitment to run a full and realistic scenario, as part of continual service improvement, every two years.

This particular test was based around a scenario in which a significant portion of the team was affected by food poisoning and unable to work at short notice, following on from an EPCC work function. The test was very successful:

- Staff involved are confirmed to be better equipped to deal with any real major incident that may occur.
- There was no impact on the actual ARCHER service.
- Lessons learned from the previous test were confirmed to have been implemented effectively.
- A set of recommended improvements have been agreed as an outcome of the test.

A detailed report on the preparation, execution, and outcomes of the test has been produced and shared with UKRI. A set of potential improvements were identified in the course of the test, including a review of interfaces to other University of Edinburgh agencies and an update to induction material for new staff to give sufficient emphasis to business continuity planning and testing.

This test and the previous one demonstrate the benefits of regular BCDR testing, and encourage more frequent testing. Consideration is being given to running different types of BCDR testing that could be run more frequently. For example, a table-top exercise to review the response to various major-incident scenarios.

6.2 Demonstrating benefit to the community from the eCSE programme

The ARCHER eCSE programme has provided funding to the ARCHER user community to develop software in a sustainable manner to run on ARCHER and on future Tier-1 services. The programme ran throughout the ARCHER service with 13 calls receiving 222 proposals leading to 100 awarded projects. It was set up to provide at least 14 FTEs (8.4 FTEs in the final year) of effort *embedded* across the UK ARCHER community and ran as a not-for-profit service with any remaining money used to fund extra person months. In the end the programme funded 973 person months of effort which included an extra 32 person months from the remaining funds.

Projects are not funded to do scientific research itself, but instead focus on software developments which lead to a number of tangible improvements; these in turn aid researchers' abilities to carry out their scientific research. Many of these improvements can be quantified to show the return on investment using benefits realisation techniques developed during the programme. For example, using data on how much faster a code runs it is possible to compare how much code usage would have cost before and after the code improvements. From this comparison it is possible to see how much cost saving a project has provided in financial terms.

To date cumulative benefits of almost £35M have been shown using our benefits realisation techniques. The programme as a whole cost almost £7M to run therefore showing around a 5-fold return on investment, a figure which will continue to grow as further usage of the improved codes brings further cost savings.

6.3 Continual Service Improvement (CSI) on ARCHER

In collaboration with user groups and the other Service partners, the CSE service identified several priority service improvement areas to invest technical effort from the centralised CSE team. This section provides highlights of the CSI projects implemented in 2019.

Benchmarking MPI performance on UK HPC facilities

With the wealth of national HPC facilities available for researchers to choose from for their computational work (ARCHER, DiRAC, Tier-2 HPC), we performed a significant of work comparing the performance of HPC applications across different platforms. This work has resulted in a public repository of benchmarks, results and performance analysis at:

https://www.github.com/hpc-uk/archer-benchmarks

and two reports on application benchmark performance to help users choose the correct facility for their research (https://doi.org/10.5281/zenodo.1288378, and https://doi.org/10.5281/zenodo.2616549). The next step is to provide more information on the performance differences seen on different platforms. As all of the application benchmarks (and the overwhelming majority of parallel HPC applications) use the MPI library to implement their distributed memory parallelism, understanding the performance of MPI libraries across the different HPC platforms is critical to understanding the performance of parallel HPC applications.

Initial results show significant differences in performance across platforms and MPI libraries, and analysis is ongoing. Up to date results and analysis are publicly available in the benchmarking Github repository (https://www.github.com/hpc-uk/archer-benchmarks).

Investigating the effects of lustre of striping on I/O performance

We have been investigating the effects of lustre striping on the I/O performance of VAMPIRE on ARCHER. VAMPIRE is an open source software package that simulates magnetic materials at an atomic level. The application was previously ported to ARCHER and optimised for its I/O performance in the eCSE project eCSE07-9 (https://www.archer.ac.uk/community/eCSE/eCSE07-09/eCSE07-09.php).

As part of the previous study on ARCHER, VAMPIRE's I/O performance was investigated and optimised for three different data output methods: file-per-process, file-per-node, and single-shared-file via MPI-IO. While great scaling was obtained for file-per-process and file-per-node (with file-per-node being better at high node counts), the MPI-IO method fared substantially worse than the other two.

Initial findings show that all previous results were based on using a constant lustre stripe count of 4, which may lead to the previous limit of the I/O performance improvement. We have been testing whether increasing the lustre stripe count can further improve the I/O performance of VAMPIRE using MPI-IO and analysis is ongoing. We plan to provide the analysis conclusion and recommendation on further optimisations to the VAMPIRE developers and ARCHER/future Tier-1 system users before the end of the ARCHER extension period, to help them better understand and consider the effects of lustre striping when tuning the I/O performance of their applications.

6.4 Towards a shared UK HPC knowledge base

Recently, the ARCHER CSE service has led an initiative that coordinates with HPC Champions and the wider HPC RSE community to consider how we could create a community UK HPC technical knowledge base that would allow us to share and reuse useful technical information and experience. Much of these thoughts came out of discussions at the HPC Champions that took place on 16 September 2019 alongside the UK RSE Conference 2019 in Birmingham, UK together with subsequent discussions at the monthly HPC RSE calls.

Although there has been a large amount of successful work to increase information and experience sharing between the different UK national HPC services and also to UK institutional HPC services, there are still areas where this coordination and collaboration could be improved. One area for improvement has been highlighted by a question that has been raised within the current community meetings:

> *How can we share technical solutions and information that we have locally across the community and potentially publicly to the wider HPC community (including users, RSEs and service providers)?*

Currently, this useful information is stored in a number of different locations with different access levels; *e.g.* internal knowledge bases, wikis, service desks, public repositories and websites. Accessibility and information sharing could be substantially improved by creating a shared UK HPC knowledge base that would be used by to expose and collect this type of useful information. Based on discussions within the community, the requirements of a public UK HPC knowledge base should be:

- **Publicly visible through a web browser:** available for anyone to view the information and make use of it.
- **Searchable and well-indexed:** the most powerful search tools are internet search engines. The knowledge base should be well indexed and available to be found easily through standard search engines.
- **Permissions to allow public addition of knowledge base entries:** so all can contribute. This, however, means that we need a way to ensure that publicly added answers are technically correct - which leads to the next point.
- **Rate answers based on correctness:** to ensure that entries are kept up to date and that publicly added entries are reviewed.

Based on the requirements above, two initial potential solutions for the knowledge base have been identified: *Stack Overflow and related Stack Exchange sites* and *ask.cyberinfrastructure (Ask.CI)*.

Stack Overflow is a public, online knowledge base aimed at sharing technical knowledge relevant to programming. The site uses peer review of answers to determine the most useful and accurate answers and allows tagging of questions. Other public knowledge base sites of interest here use the same underlying technology, e.g., Computational Science Stack Exchange. On all Stack Exchange sites it is perfectly acceptable to ask a question and provide an answer. This can be useful when transferring knowledge from an internal resource (e.g., a service desk ticket) into the public knowledge base. The strengths of Stack Overflow (and related sites) include its high ranking in technical question searches using standard search engines, its well-designed interface, its strong definition of allowed questions and their format, and the fact that it is already the top resource for this type of technical information on the internet.

Ask.CI is a site to aggregate expert knowledge so that is widely distributed throughout the research computing community. The site provides a Q&A interface aimed at all people involved in research computing (e.g. researchers, users, support staff, RSEs, systems administrators). Like the Stack Exchange sites, Ask.CI is free to use and publicly available, allows for tagging of questions by topic and allows people to ask a question and answer it themselves. Unlike Stack Exchange sites, it allows discussion-style questions and allows topics from across the whole range of research computing. It does not have a peer review functionality to allow for rating of answers; people can simply like posts that they think have merit. The strengths of Ask.CI are that it is focused particularly on research computing (though this is broader than just HPC), it already has a dedicated, committed community asking and answering questions, and it allows all types of questions. Weaknesses are the low weight leading to answers being low in search engine rankings, lack of peer review, the limited community involved, the name (*cyberinfrastructure* has little meaning outside the USA) and the lack of specific

question guidelines that mean that many of the questions descend into discussion with no indication of what the accepted answer is.

Based on our analysis, it seems that using the Stack Exchange sites, primarily Stack Overflow, will provide the strongest solution. The key features that point towards this solution are the ability to search well using standard search engines and the strong rules around allowable questions, which drive the responses towards being a useful knowledge base rather than a question-and-answer resource where posts tend to result in discussions with no clear answer to the technical question being asked. These types of questions have a place, but they do not fit as well into the idea of a useful technical knowledge base.

The next step is for the community as a whole to take a decision on how to take this idea of a shared technical HPC knowledge base forward and boostrap its use. One initial option that has been mooted is to run a community coding day where the community would come together to seed the chosen option with initial questions and come up with documentation and guidelines to help sites adopt the chosen solution. The UK national HPC services (ARCHER2, DiRAC and Tier-2) will play a key role in taking this important initiative forward.

7. Cray Service Group

7.1 Summary

In 2019 the ARCHER service provided a very stable technology platform for the user community, facilitating high utilisation of the system resources.

7.2 Reliability and performance

The ARCHER Cray XC30 system provided excellent levels of reliability and performance throughout the period. Large and complex supercomputer systems such as ARCHER inevitably encounter component fallout on a regular basis but, due the redundancy of critical components, most failures have little or no impact on users. Cray's dedicated system engineers routinely utilise concurrent maintenance techniques to replace failed components without impacting the service.

Full Service failures

There was one incident classified as a full-service failure caused by a technology incident during 2019:

ARCHER full-service failures 2019		
Incident	**Date**	**Description**
1	29/09/2019	System High Speed Network failed to recover from a network quiesce operation following a link failure, resulting in the system becoming unresponsive and requiring a reboot to clear the issue.

Partial Service failures

Although most technology failures do not impact the user community, there were some exceptions during 2019:

- Two incidents of a PBS Pro MOM node losing contact to the Network File System (NFS) servers in the Boot and the SDB management nodes, requiring a reboot of the affected MOM nodes to re-establish the connections.
- One incident of Cray XC30 cabinet c7-2 controller failing due to uncorrectable memory errors, requiring the power down of the affected cabinet in order to replace the controller.

7.3 Scheduled maintenance activities

None of the permitted four dedicated maintenance sessions was taken by Cray in 2019.

8. Cray Centre of Excellence (CoE)

In 2019, the CoE took further advantage of our LASSi I/O analysis framework both in analysing user problems and to support studies into I/O usage of applications and communities on ARCHER. 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. We made the following enhancements to the LASSi I/O analysis tools which were originally developed to provide faster triage of reported filesystem performance problems and to help gain a better understanding of how applications stress the filesystem.

- Optimised the workflow and added automated daily reporting
- Introduced new reports targeting I/O quality (to highlight small I/O operations)
- Collaborated with EPCC to provide LASSi data export to SAFE to give users a high level I/O summary.

We continued to use LASSi tools this year to identify the cause of reported filesystem performance problems and to support investigation into how particular applications were stressing the filesystem. One example was the Elk application which was writing the same data to a file from each MPI process. In addition, Elk uses a pattern of I/O where direct-access files are accessed by multiple processes. For each I/O operation: a process opens the file, writes one record and then closes the file. We were able to suggest an update to the developers which was accepted and integrated into the next released version.

The figures below show the improvement from one of the fixes and cover just over one iteration, the second figure shows far fewer multi-second I/O operations.

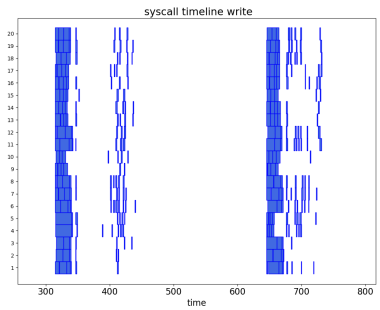

Figure 1: writes taking >1s

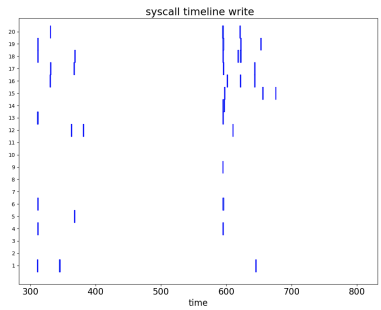

Figure 2: writes taking > 1s (after fix)

In addition, we attended the following conferences to promote this area of our work and where we presented results of new analysis of the data we had collected in various phases up to Nov 2019.

- *LASSi: metric based I/O analytics for HPC*, presented at the 27th High Performance Computing Symposium (HPC 2019) at SCS Spring Simulation Multi-Conference (SpringSim'19) in Tucson, AZ, US. This paper introduces LASSi and how it was used successfully to analyse application slowdown on ARCHER. (https://ieeexplore.ieee.org/document/8732903)
- *Analysis of parallel I/O use on the UK national supercomputing service, ARCHER using Cray's LASSi and EPCC SAFE* accepted to 2019 Cray User Group conference. This paper was written in collaboration with EPCC and the University of Reading. It brings together our experience with LASSi, EPCC application grouping analysis and some perturbation benchmarking from Reading University into one paper. (https://cug.org/proceedings/cug2019_proceedings/includes/files/pap118s2-file1.pdf)

- Details of the latest analysis of I/O on ARCHER were presented by Karthee Sivalingam at the Computing Insight UK conference in December. The talk covered our development of and use of tools along with the new analysis of the filesystem usage mentioned above. See the presentation for more details. During the conference we met many users of UK HPC facilities.

As an example of the analysis undertaken, the following figure show the I/O attributed to various science areas for Reads (total 59PB) and Writes (total 192PB) for the period Apr 2017-Nov 2019.

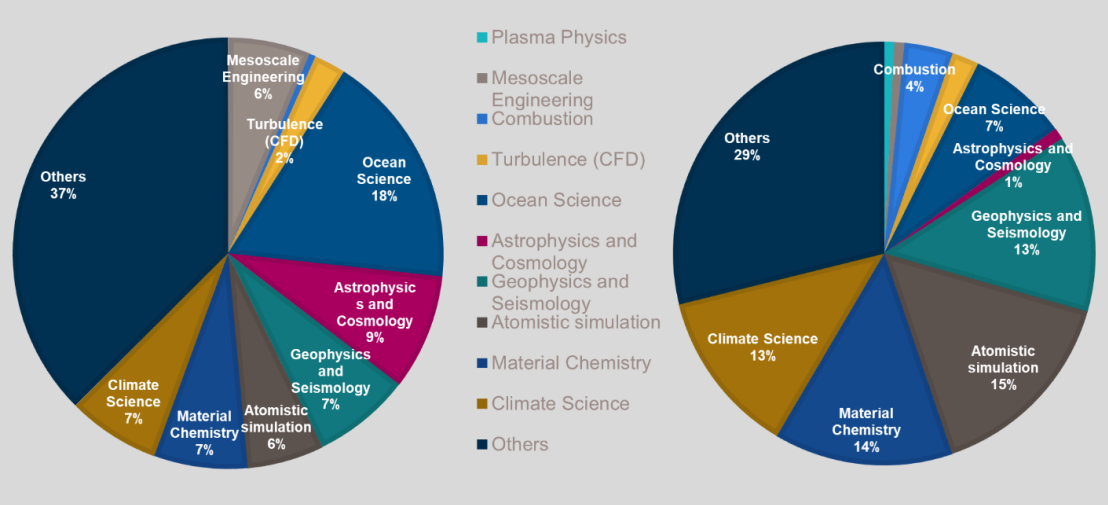

Figure 3: I/O usage of projects on ARCHER. Reads on left, Writes on right.

More detailed analysis and trend data can be found from the conference presentations.

8.2 Events attended

The CoE undertook or attended various events during the year:

- CDT in pervasive parallelism industrial engagement event, 19th February (not an ARCHER event but useful to interact with users of ARCHER)
- Storage Challenges in the UK workshop, Reading, 6th March, Karthee Sivalingam presented LASSi
- Harvey Richardson attended the ARCHER Champions Meeting on 9th May and gave a presentation entitled *I/O Observation on ARCHER: System call Profiling,* which presented a non-invasive approach for I/O profiling of applications showing examples of ARCHER applications.

8.3 Significant investigations

As agreed with EPRSC, 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 assisting GW4 with a performance investigation related to network scaling of collective communication patterns in MPI. We also investigated a problem where large performance variation was observed in FFTW transpose operations when nodes were spread across the network. Cray developers are planning to improve the communication scheduling in FFTW to improve this. We also performed an investigation into the use of new rootless container technologies on ARCHER. We were able to build a Tensorflow container using uDocker. We had hoped that containers could be used to optimize I/O in Lustre but unfortunately uDocker unpacks the container image and this destroys any potential advantage.

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/or need interaction with Cray R&D experts.

Of particular note were the following interactions:

1. Helped a user by providing a workaround to a known Lustre problem that affects getcwd() calls made by the Intel compiler runtime, the call will return an error code, not be retried and the application exits. Built library fix for user and gave instructions how this could be linked into the application.
2. User complained of an issue with MPI_Comm_split() but after engaging with the user we discovered the problem was due to a side-effect of a compilation with no arguments supplied. This was raised with the Cray developers although this is not classified as a bug.
3. Two reports of MPI programs failing when using Fortran 2008 bindings. Traced this to different errors in an include file that is distributed by ANL MPICH. Bug submitted.
4. Investigated a problem where very infrequently some job output files would be extremely large, having a 'hole' in the file. We understand the structure of the file but without a reliable reproducer this is hard to make progress on. We have given advice on what the user can do to obtain more information if this happens again and we have provided instructions on how to convert the 'huge' file to one of the expected size.
5. Two queries relating to use of the craypat profiling tools. Advice was given on alternative reports in one case and for the other we were able to provide a workaround suggested by the developers.
6. Investigated a problem where the Unified Model running on ~580 nodes would fail with an *Out of Memory* error when profiled with the Cray profiling tools. We offered various suggestions and contacted the Cray development team for advice. We determined that valid trace files were written and the performance report could be regenerated on the PP nodes which have a higher memory limit. The workaround was acceptable to the user; we note that newer versions of the performance tools use less memory when generating a report.

The Programming environment was updated during the year and the CoE assisted with checking this on the test system we have before installation.

8.5 Support of the eCSE programme

We continued to support this program by completing four final project reviews.