ARCHER and ARCHER2
SP Quarterly Report

April – June 2020

EPCC
The University of Edinburgh

Document Information and Version History

Version:	1.0
Status	Release
Author(s):	Alan Simpson, Anne Whiting, Paul Clark, Andy Turner, Linda Dewar, Stephen Booth, Jo Beech-Brandt
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	03/07/20	Initial Draft	Anne Whiting, Jo Beech-Brandt, Linda Dewar, Lorna Smith, Stephen Booth
0.2	14/07/20	Initial Edits	Jo Beech-Brandt
0.3	14/07/20	Reviewed	Alan Simpson
0.4	15/07/20	Added more ARCHER2 Information	Jo Beech-Brandt
1.0	15/07/20	Version for UKRI	Alan Simpson and Jo Beech-Brandt

1. The ARCHER Service

1.1 Service Highlights

This is the report for the ARCHER SP Service for the Reporting Periods:
April 2020, May 2020 and June 2020.

- Utilisation over the quarter was 82% which is somewhat lower than the previous quarter where the utilisation was 88%.

- The results of the 2019 annual ARCHER User Survey have been analysed. 157 responses were received 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)	2019 Mean Score (out of 5)
Overall Satisfaction	4.4	4.3	4.3	4.4	4.5	4.3
Hardware	4.1	4.1	4.2	4.3	3.9	3.8
Software	4.0	4.0	4.2	4.1	3.8	3.8
Helpdesk	4.5	4.5	4.5	4.6	4.5	4.4
Documentation	4.1	4.1	4.2	4.2	4.0	3.9
Website	4.1	4.2	4.2	4.2	4.0	3.9
Training	4.1	4.1	4.2	4.1	4.3	4.2
Webinars	3.6	3.9	3.9	4.2	3.9	4.0
Online training	-	4.0	4.1	4.2	3.9	3.9

As can be seen, users are still generally very positive at this late stage in the service with a score of 4.3 for the overall service. The scores for the hardware, software and documentation are down on the previous year, but this could be due to the age of the system and its software, and the expectation of a move to ARCHER2. As with previous years the helpdesk has received the highest score.

- Analysis is being carried out following the recent security incident which the ARCHER service to ensure that the root causes of the incident are understood, that service improvements are identified and that the implementation of improvements is tracked. The effectiveness of the measures taken will be measured and monitored to minimise the risk of such an incident occurring again in the future. Similar analysis and improvement measures are being carried out for the power outage incident.

- High priority queues have been created on standard and high memory nodes to enable the urgent computing request for COVID-19 projects.

- In preparation for the ARCHER2 Service we have created a new installation of the SAFE https://safe.epcc.ed.ac.uk/. This is installed within a new hosting environment and provides a single unified view of all systems managed by the SAFE including ARCHER, ARCHER2 and the Tier-2 services. This is fully synchronized with the current SAFE instances which will continue to operate while ARCHER remains in service. The new SAFE uses an updated look and feel (though users will be able to revert to the old style if they prefer). In addition we have introduced a number of new features including enhancements to the Service Desk, improved support for the management of ssh-keys and faster report generation.

- SAFE users will now be asked to revalidate their email address and other details on an annual basis. This is to ensure that the information we hold is up to date and to help identify inactive accounts that should be retired.

1.2 Forward Look

- The ARCHER Service has a further extension to allow the service to continue while ARCHER2 is installed and prepared for users. Users are kept informed of developments via the weekly emails from UKRI and news items on the ARCHER/ARCHER2 websites, and a minimum of 2 weeks notice will be provided for the end of the ARCHER Service.

- Work continues to prepare the user community for the end of the ARCHER service. The Service Exit Plans are in place and will be activated as we approach the end of service. EPCC will continue to work with EPSRC and NERC to provide assistance to them in planning the transition of user data and projects from ARCHER to ARCHER2.

2. Contractual Performance Report

This is the contractual performance report for the ARCHER SP Service.

2.1 Service Points and Service Credits

The Service Levels and Service Points for the SP service are defined as below in Schedule 2.2.

- **2.6.2 - Phone Response (PR):** 90% of incoming telephone calls answered personally within 2 minutes for any Service Period. *Service Threshold: 85.0%; Operating Service Level: 90.0%.*
- **2.6.3 - Query Closure (QC):** 97% of all administrative queries, problem reports and non in-depth queries shall be successfully resolved within 2 working days. *Service Threshold: 94.0%; Operating Service Level: 97.0%.*
- **2.6.4 - New User Registration (UR):** Process New User Registrations within 1 working day.

Definitions:

Operating Service Level: *The minimum level of performance for a Service Level which is required by the Authority if the Contractor is to avoid the need to account to the Authority for Service Credits.*

Service Threshold: *This term is not defined in the contract. Our interpretation is that it refers to the minimum allowed service level. Below this threshold, the Contractor is in breach of contract.*

Non In-Depth: *This term is not defined in the contract. Our interpretation is that it refers to Basic queries which are handled by the SP Service. This includes all Admin queries (e.g. requests for Disk Quota, Adjustments to Allocations, Creation of Projects) and Technical Queries (Batch script questions, high level technical 'How do I?' requests). Queries requiring detailed technical and/or scientific analysis (debugging, software package installations, code porting) are referred to the CSE Team as In-Depth queries.*

Change Request*: This term is not defined in the contract. There are times when SP receives requests that may require changes to be deployed on ARCHER. These requests may come from the users, the CSE team or Cray. Examples may include the deployment of new OS patches, the deployment Cray bug fixes, or the addition of new systems software. Such changes are subject to Change Control and may have to wait for a Maintenance Session. The nature of such requests means that they cannot be completed in 2 working days.*

2.1.1 Service Points

In the previous Service Quarter, the Service Points can be summarised as follows:

Period	**Apr 20**		**May 20**		**Jun 20**		**20Q2**
Metric	**Service Level**	**Service Points**	**Service Level**	**Service Points**	**Service Level**	**Service Points**	**Service Points**
2.6.2 – PR	100.0%	-5	100.0%	-5	100.0%	-5	-15
2.6.3 – QC	98.18%	-1	94.79%	3	98.98%	-1	1
2.6.4 – UR	1 WD	0	1 WD	0	1 WD	0	0
Total							**-14**

The details of the above can be found in Section 2.2 of this report.
It is worth noting that the low Service Level for May was caused by the unavailability of the ARCHER Service and so some queries could not be completed within the required 48 hours.

2.1.2 Service Failures

Details of planned maintenance sessions, if any, can be found in Section 2.3.2.

2.1.3 Service Credits

As the Total Service Points are negative (-14), no Service Credits apply in 20Q2.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Apr 20	May 20	Jun 20	20Q2
Phone Calls Received	4 (1)	4 (1)	6 (3)	**14 (5)**
Answered in 2 Minutes	4 (1)	4 (1)	6 (3)	**14 (5)**
Service Level	**100.0%**	**100.0%**	**100.0%**	**100.0%**

The volume of telephone calls remained low in 20Q1. Of the total of 14 calls received above, only 5 were actual ARCHER user calls that either resulted in queries or answered user questions directly.

2.2.2 Query Closure (QC)

	Apr 20	May 20	Jun 20	20Q2
Self-Service Admin	226	2476	1633	**4335**
Admin	96	203	311	**610**
Technical	8	7	24	**39**
Total Queries	330	2686	1968	**4984**
Total Closed in 2 Days	324	2546	1948	**4818**
Service Level	**98.18%**	**94.79%**	**98.98%**	**96.67%**

The above table shows the queries closed by SP during the period.
The sharp increase in Self-Service Admin queries during May and June was caused by the requirement that all users reset their password and upload a new SSH key for their accounts.

In addition to the Admin and Technical queries, the following Change Requests were resolved in 20Q2:

	Apr 20	May 20	Jun 20	20Q2
Change Requests	2	0	0	**2**

2.2.3 User Registration (UR)

	Apr 20	May 20	Jun 20	20Q2
No of Requests	46	43	96	**185**
Closed in One Working Day	46	39	96	**181**
Average Closure Time (Hrs)	0.3	5.5	0.4	**1.6**
Average Closure Time (Working Days)	0.03	0.6	0.05	**0.2**
Service Level	**1 WD**	**1 WD**	**1 WD**	**1 WD**

To avoid double counting, these requests are not included in the above metrics for “Admin and Technical” Query Closure.

2.3.1 Target Response Times

The following metrics are also defined in Schedule 2.2, but have no Service Points associated.

Target Response Times	
1	During core time, an initial response to the user acknowledging receipt of the query
2	A Tracking Identifier within 5 minutes of receiving the query
3	During Core Time, 90% of incoming telephone calls should be answered personally (not by computer) within 2 minutes
4	During UK office hours, all non telephone communications shall be acknowledged within 1 Hour

1 – Initial Response

This is sent automatically when the user raises a query to the address helpdesk@archer.ac.uk. Users may choose not to receive such emails by mailing support@archer.ac.uk.

2 – Tracking Identifier

This is sent automatically when the user raises a query to the address helpdesk@archer.ac.uk. Users may choose not to receive such emails by mailing support@archer.ac.uk. The tracking identifier is set in the SAFE regardless which option the user selects.

3 – Incoming Calls

These are covered in the previous section of the report. Service Points apply.

4 - Query Acknowledgement

Acknowledgment of the query is defined as when the Helpdesk assigns the new incoming query to the relevant Service Provider. This should happen within 1 working hour of the query arriving at the Helpdesk. The Helpdesk processed the following number of incoming queries during the Service Quarter:

	Apr 20	May 20	Jun 20	20Q2
CRAY	0	4	2	**6**
ARCHER_CSE	14	14	20	**48**
ARCHER_SP	519	3285	2364	**6168**
Total Queries Assigned	533	3303	2386	**6168**
Total Assigned in 1 Hour	533	3303	2386	**6168**
Service Level	**100.0%**	**100.0%**	**100.0%**	**100.0%**

The Service Desk assigns queries to all groups supporting the service i.e. SP, CSE and Cray. The above table includes queries handled by the other groups supporting the service as well as internally generated queries used to manage the operation of the service.

2.3.2 Maintenance

Maintenance now takes place on at most a single day each month (fourth Wednesday of each month). This is marked as a full outage maintenance session for a maximum of 8 hours taken. There are also additional “at-risk” sessions that may be scheduled for other Wednesdays. This reduces the number of sessions taken, which then reduces user impact since the jobs running on the service have to be drained down only once per month and not twice. It also eases the planning for training courses running on ARCHER. A 6-month forward plan of maintenance has been agreed with EPSRC.

Feedback has shown that the users would be happier if there were even fewer full outage maintenance sessions, and so we have been working to reduce these as much as possible. Some maintenance activities can only be done during a full outage (e.g., applying firmware updates), but for others the requirement to take a full outage can be evaluated on an individual basis based on potential risk.

There were no full maintenance sessions during Q2, all maintenance activities were carried out during at risk sessions.

2.3.2 Quality Tokens and query feedback emails

No quality tokens were received this quarter.

3. Service Statistics

This section contains statistics on the ARCHER service as requested by EPSRC, SAC and SMB.

3.1 Utilisation

Utilisation over the quarter was 82%, which is down from 88% the previous quarter. Utilisation for April was 83%, for May 80% and for June 82%. The plot below shows a steady increase in utilisation over the lifetime of the service to Dec 2015 and since then the service has effectively been operating around maximum capacity as shown by the generally steady utilisation value.

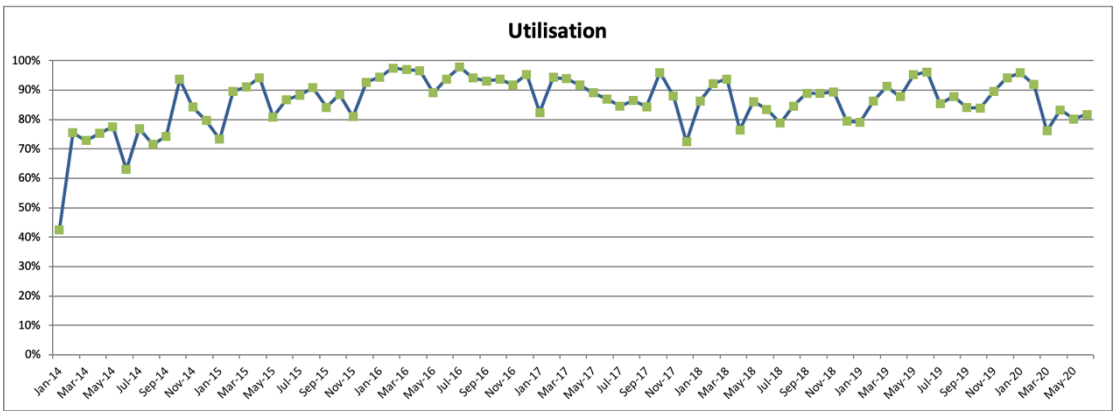

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that EPRSC did not meet their target this quarter with their usage being at 60% (against their target of 77%) and NERC also missed their target with utilisation being 16% (against their target of 23%). This compares with 71% for EPSRC and 19% for NERC for the previous quarter.

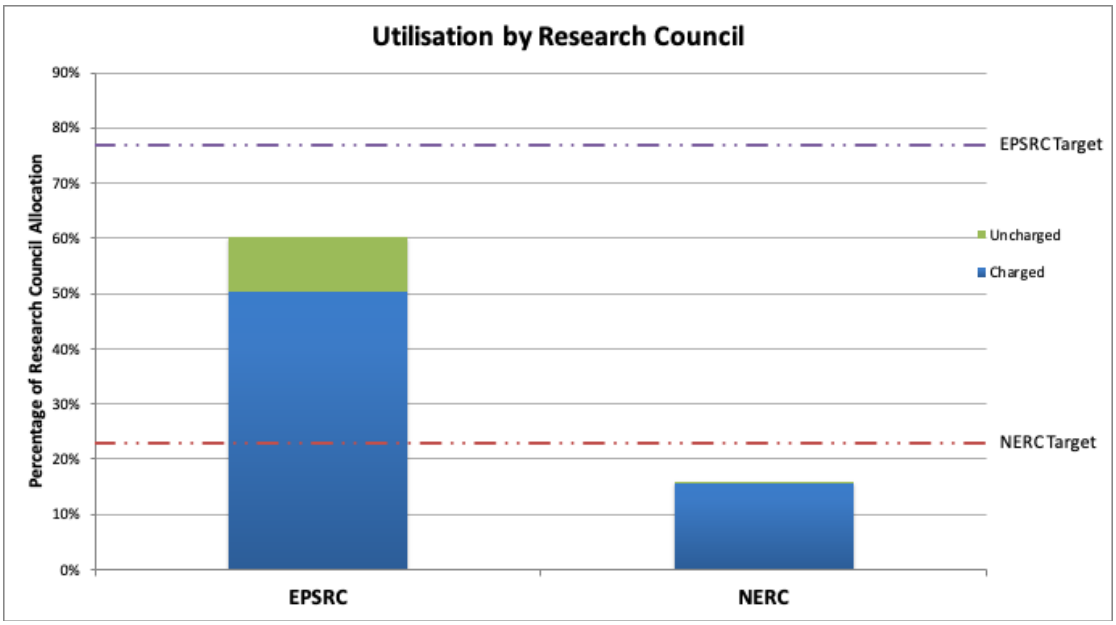

The cumulative allocation utilisation for the quarter by the Research Councils is shown below:

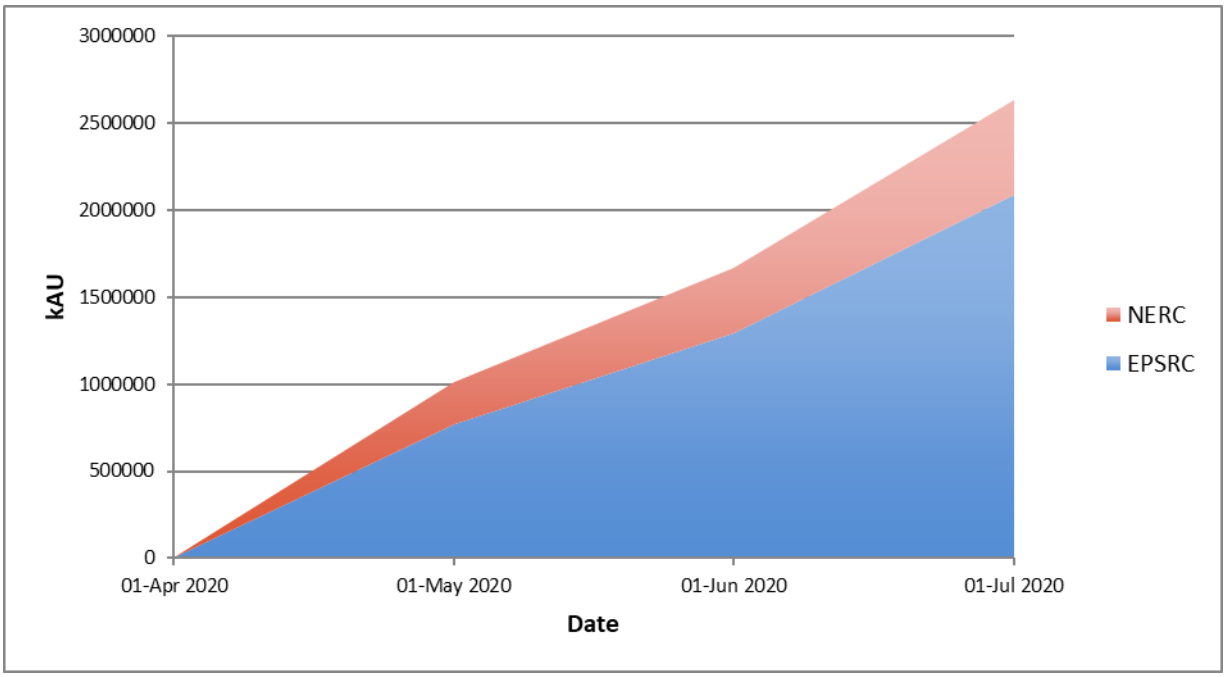

The cumulative allocation utilisation for the quarter by EPSRC broken down by different project types (see below) shows that the majority of usage comes from the scientific Consortia (as expected) with significant usage from research grants and ARCHER RAP projects. The total time used by Instant Access projects is very small.

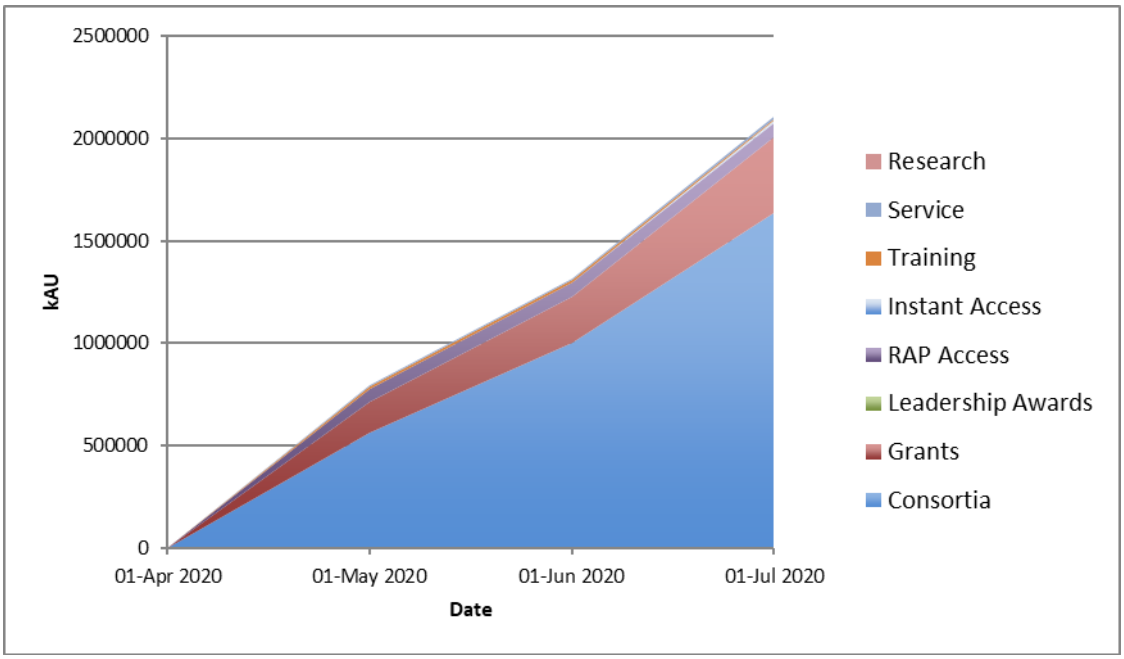

3.2 Scheduling Coefficient Matrix

The colour in the matrix indicates the value of the Scheduling Coefficient. This is defined as the ratio of runtime to runtime plus wait time. Hence, a value of 1 (green) indicates that a job ran with no time waiting in the queue, a value of 0.5 (pale yellow) indicates a job queued for the same amount of time that it ran, and anything below 0.5 (orange to red) indicates that a job queued for longer than it ran.

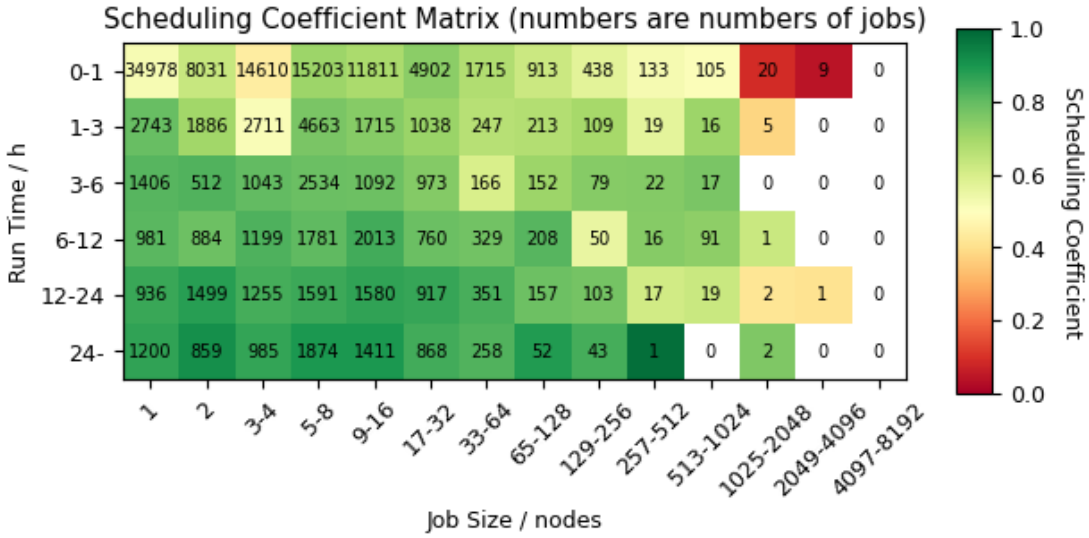

The usage heatmap below provides an overview of the usage on ARCHER over the quarter for different job sizes/lengths. The colour in the heatmap indicates the number of kAUs expended for each class, and the number in the box is the number of jobs of that class.

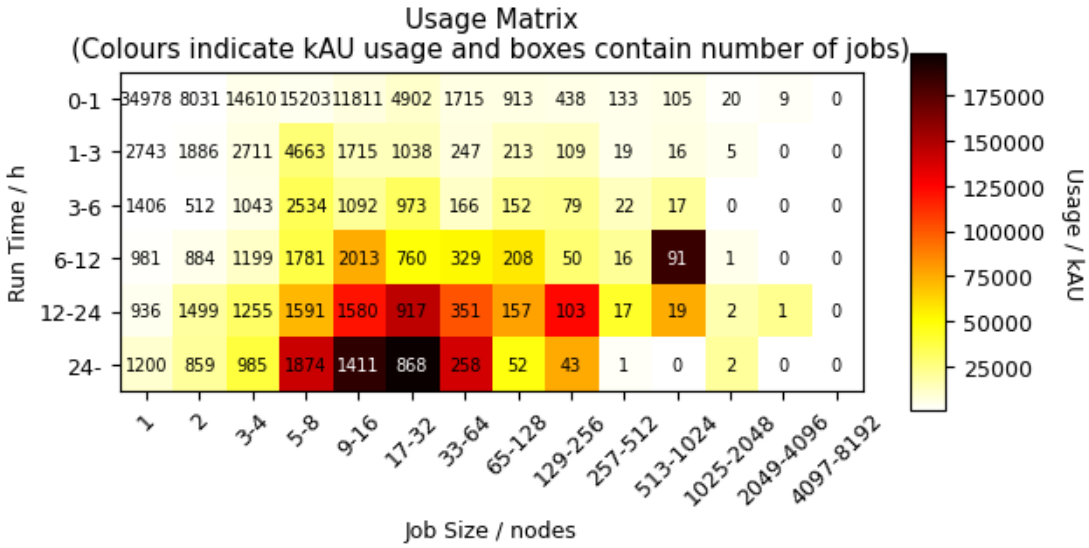

3.3 Additional Usage Graphs

The following charts provide different views of the distribution of job sizes on ARCHER.

Analysis of Job Sizes

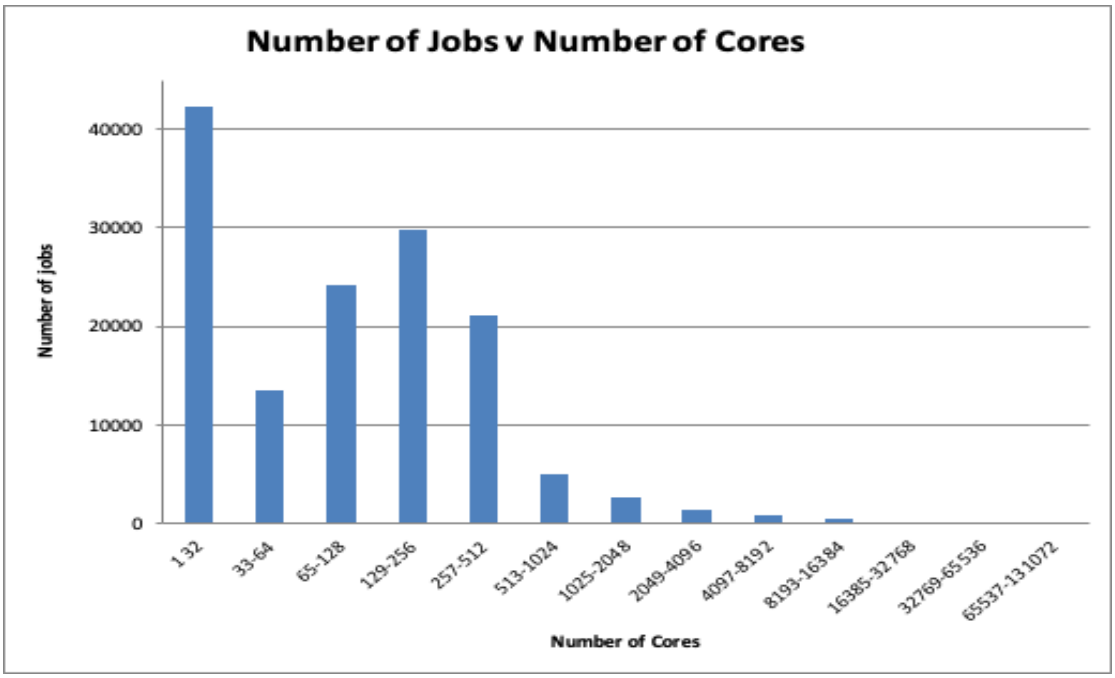

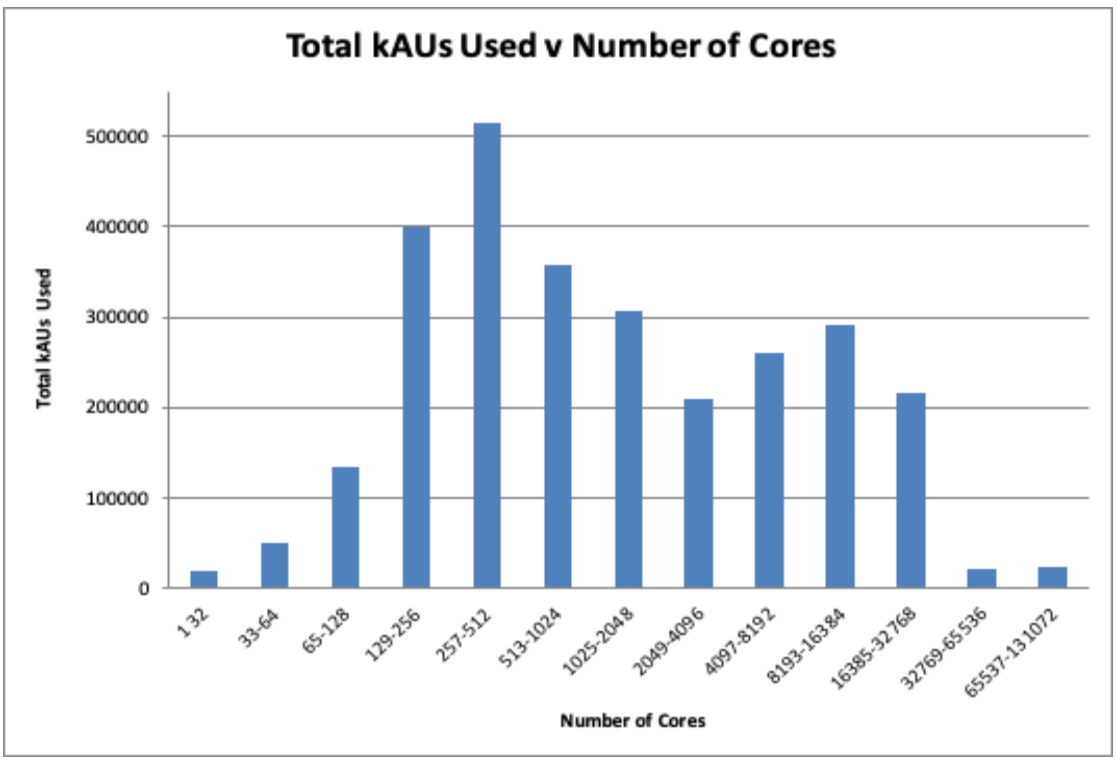

The first graph shows that, in terms of numbers, there are a significant number of jobs using no more than 512 cores. However, the second graph reveals that most of the kAUs were spent on jobs between 129 cores and 16384 cores. The number of kAUs used is closely related to money and shows better how the investment in the system is utilised.

Analysis of Jobs Length

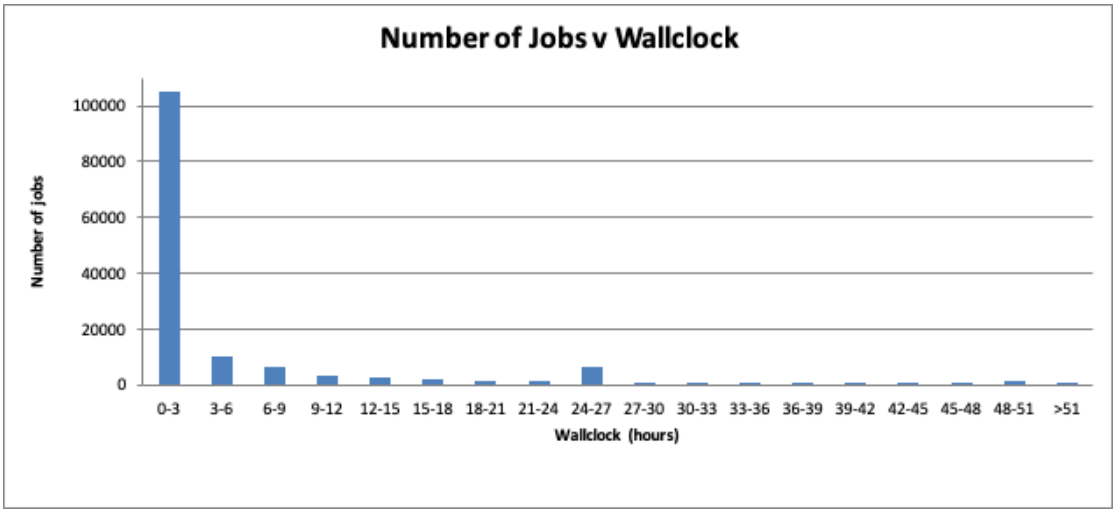

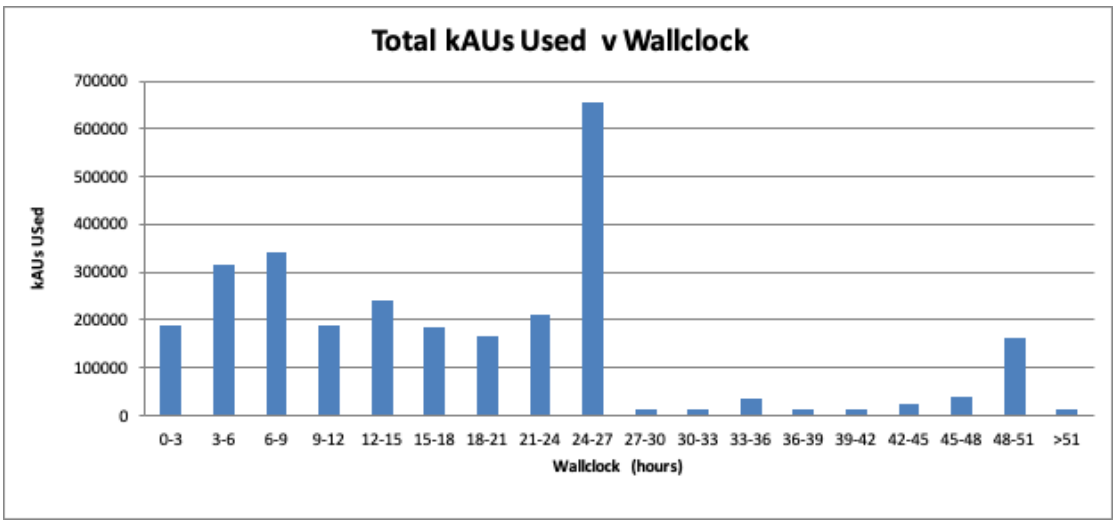

From the first graph, it would appear that the system is dominated by short jobs. However, the second graph shows that actual usage of the system is more spread and dominated by jobs of around 24 hours with a second peak for jobs around 48 hours.

Core Hours per Job Analysis

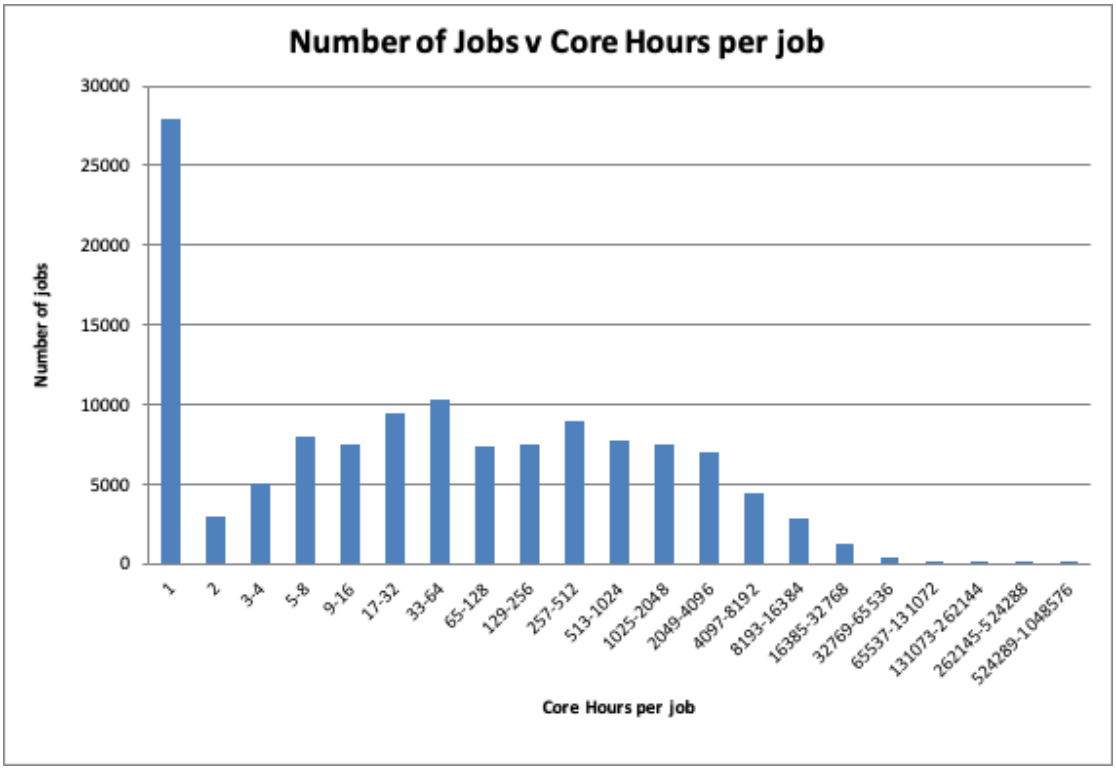

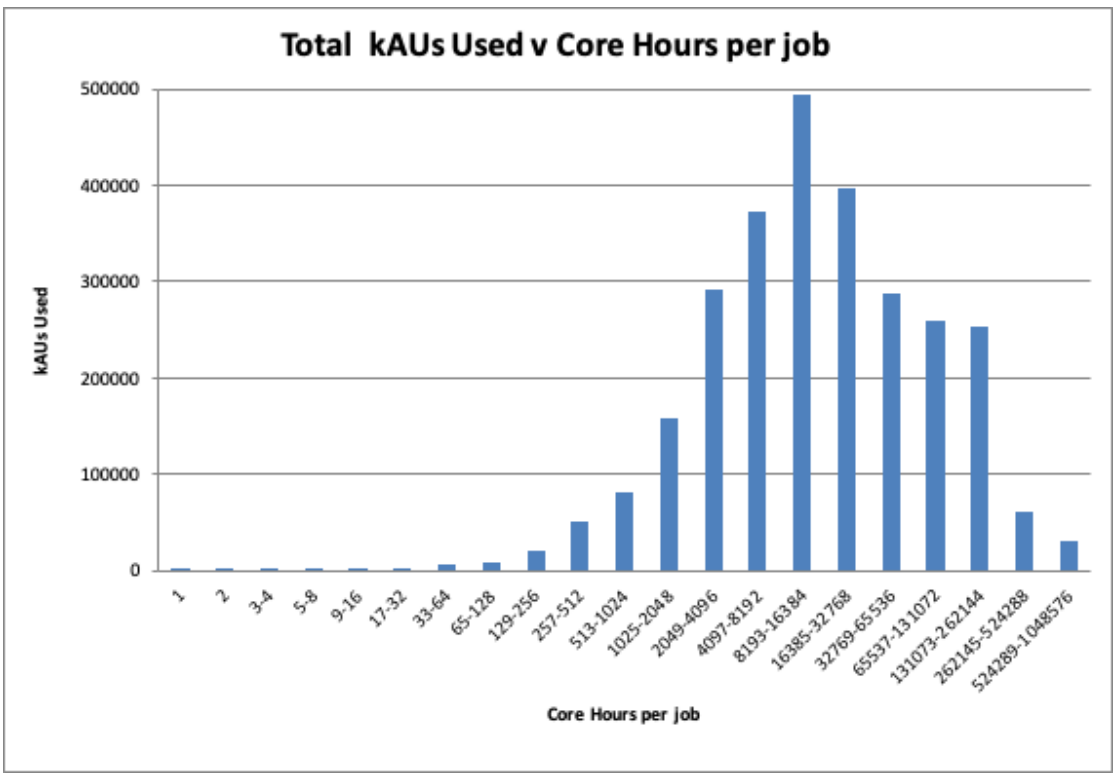

The above graphs show that, while there are quite a few jobs that use only a small number of core hours per job, most of the resource is consumed by jobs that use tens of thousands of core hours per job.

1 The ARCHER2 Service

1.1 Service Highlights

- The ARCHER2 SP Service Desk team has been in operation since the 1st April 2020 ahead of the operational commencement date of 6th May 202. Support hours have been increased to cover 08:00 – 18:00 on working days, with the ability to log support calls in the SAFE provided 24x7. The address for the ARCHER2 Service Desk is support@archer2.ac.uk.

- There were only 2 SP queries during this quarter and they met all of the ARCHER2 Service Levels. Two users provided feedback and scored the Service Desk Excellent and Very Good.

- The HPC Systems have worked with the Cray/HPE team to provide the Test and Development System (TDS). This has allowed testing and preparation work to get underway with user documentation, training courses and user application code support.

- The ARCHER2 website is operational, with training events, virtual tutorials and the eCSE call advertised on the website. These have seen significant traffic.

- Policies and procedures have been put in place for updates to the website and to user documentation, to ensure material is appropriately reviewed before publication.

- An Accessibility policy has been developed for the ARCHER2 website.

1.2 Forward Look

- The 4-Cabinet ARCHER2 System is due to arrive in Edinburgh on July 13th and the installation team will prepare the system for the ARCHER2 users. A daily blog and tweets will be used to publicise the installation of system.

- We will work with EPSRC and NERC to ensure the seamless transition of users from ARCHER to ARCHER2.

- With the importance placed in ARCHER2 on having robust business continuity and disaster recovery plans and processes in place, EPCC is planning to start working towards obtaining ISO 22301 business continuity certification.

- Plans are underway for increasing the ACF external and internal network links to 100GB improving communication speeds for the user community.

2 ARCHER2 Performance Report

This is the contractual performance report for the ARCHER2 SP Service for the Reporting Periods from 6th May 2020 (ARCHER2 OCD) until end of June 2020.

2.1 Service Points and Service Credits

The Service Levels and Service Points for the SP service are defined by EPSRC in Schedule 2.2 of ARCHER2 SP Service Contract.

The Working Day (WD) for the ARCHER2 Service is 10 Working Hours (WH) as the Service operates from 0800-1800. The Median Time to Resolution is measured in WD.

- **Availability**: *Service Threshold: <=96.5%; Operating Service Level: >98.0%, ≤ 98.5%.*
- **ARCHER2_SP_Level1 (MTR)**: The Median Time to Resolution, of all SP queries falling within Level 1 resolved by the Contractor in the Reporting Period. *MTR Service Threshold: >1 WD; Operating Service Level: >0.3 WD, ≤ 0.45 WD.*
- **ARCHER2_SP_Level2 (MTR)**: The Median Time to Resolution, of all SP queries falling within Level 2 resolved by the Contractor in the Reporting Period. *MTR Service Threshold: >8 WD; Operating Service Level: >2 WD, ≤4 WD.*
- **ARCHER2_SP_Level3 (MTR)**: The Median Time to Resolution, of all SP queries falling within Level 3 resolved by the Contractor in the Reporting Period. *MTR Service Threshold: >25 WD; Operating Service Level: >12 WD, ≤16 WD.*
- **Initial Response to Queries (%):** The percentage of the total number of SP queries assigned to the Contractor in the Reporting Period responded to within 3 Working Hours. *Service Threshold: <96.00%; Operating Service Level: 98.00 – 98.99%.*
- **Query User Satisfaction (%):** The percentage of the total number of query satisfaction surveys completed in each Reporting Period, rating the quality of the resolution of Queries by the Contractor as "Good", "Very Good" or "Excellent". *Operating Service Level: 82.00 – 87.99%*

2.1.1 Service Points

Metric	May 2020 (starting from OCD)		June 2020		Q2 2020	
	Perf	Points	Perf	Points	Perf	Points
Availability	-	-	-	-	-	-
ARCHER2_SP_Level1 (MTR)	-	-	0.01WD	-2	0.01 WD	-2
ARCHER2_SP_Level2 (MTR)	0.01WD	-2	-	-	0.01 WD	-2
ARCHER2_SP_Level3 (MTR)	-	-	-	-	-	-
Initial Response to Queries (%)	100%	-1	100%	-1	100%	-2
Query User Satisfaction (%)	-	-	-	-	-	-
Total		-		-		-6

2.1.2 Service Credits

As the Total Service Points are negative (-6), no Service Credits apply in 20Q2.

2.2 SP Query Statistics

The metrics were specified by EPSRC in Schedule 2.2 of ARCHER2 SP Service Contract.

- **No. of Assigned:** The number of SP queries assigned to the Contractor within each query resolution category in the Reporting Period.
- **No. of Resolved:** The number of SP queries resolved by the Contractor within each query resolution category in the Reporting Period.
- **Backlog:** The number of SP queries assigned to the Contractor that remained unsolved within each query resolution category in the Reporting Period
- **Avg No. of Correspondence:** The average number of pieces of correspondence generated for SP queries in each query resolution category.
- **Avg Time of First Responses:** The average time taken for the Contractor to first respond to the Originator of the SP query.

May 2020 (from OCD)					
Service level	**Assigned**	**Resolved**	**Backlog**	**Average Correspondence**	**Average Time First Response**
ARCHER2_SP_Level1	0	0	0	0	0
ARCHER2_SP_Level2	1	1	0	6	0.01
ARCHER2_SP_Level3	-	-	-	-	-
June 2020					
Service level	**Assigned**	**Resolved**	**Backlog**	**Average Correspondence**	**Average Time First Response**
ARCHER2_SP_Level1	0	1	0	4	0.01
ARCHER2_SP_Level2	0	0	0	0	0
ARCHER2_SP_Level3	0	0	0	0	0
Q2 2020 (from OCD)					
Service level	**Assigned**	**Resolved**	**Backlog**	**Average Correspondence**	**Average Time First Response**
ARCHER2_SP_Level1	1	1	0	4	0.01
ARCHER2_SP_Level2	1	1	0	6	0.01
ARCHER2_SP_Level3	-	-	-	-	-

2.3 Query Resolution

Metric	**May 20 (starting from OCD)**		**Jun 20**		**Q2 2020**
Service Level	**MTR**	**Number Resolved**	**MTR**	**Number Resolved**	**Number Resolved**
ARCHER2_SP_Level1	0	0	0.01	1	1
ARCHER2_SP_Level2	0.01	1	0	0	1
ARCHER2_SP_Level3	0	0	0	0	0
Total	**0**	**1**	**0**	**1**	**2**

There were 0 queries that failed the maximum completion time during this period.

A total of two queries were resolved by the ARCHER2 SP Service in the Reporting Period. The percentage of queries responded to within 3 hours was 100%. The low number of SP queries is expected as the ARCHER2 Service has not yet launched and both of these queries were from EPCC staff.

2.4 Query Feedback

There was no query feedback received during this period.

2.4 Maintenance

As the ARCHER2 Service has not yet been launched, there were no maintenance sessions during this period.

Appendix: ARCHER Major Incidents

Security Incident

11th – 21st May

ARCHER was one of many computers that were compromised in the UK and elsewhere in Europe as part of a major issue across the academic community. EPCC worked with the National Cyber Security Centre (NCSC), JISC, other European HPC centres and Cray/HPE in order to better understand the position and plan effective remedies.

This ARCHER incident was initially raised by Cray/HPE who identified a footprint from an exploit which they had been warned about by another site.

Due to the severity of the situation, the decision was taken to take the ARCHER service offline until all evidence of compromise could be identified and removed; any further access to the system using the compromised accounts could be prevented; and controls put in place to mitigate the risk of similar attacks happening in the future.

The following steps were taken to return ARCHER to service:

1. User passwords and ssh keys reset
2. Administrator and system credentials changed
3. Additional hardening and protection
4. Local filesystems reset to their state prior to the compromise.
5. Mounted filesystems scanned for compromised files.
6. Additional monitoring for potentially malicious activity

In addition to the above, on return to service users were required to use both a password and ssh key to prevent similar modes of compromise in the future

Protecting the ARCHER system against the threat of security violation is an ongoing activity and investigating known vulnerabilities, applying security patches and improving detection techniques are a key responsibility of the HPC Systems team.

Power Incidents

There were three power incidents caused by the repeated trip/failure of ARCHER Power Distributions Units (PDUs). These occurred on:

- 11/06/2020 10:00 until 12/06/2020 18:55
- 22/06/2020 14:30 until 22/06/2020 20:15
- 23/06/2020 10:00 until 25/06/2020 10:30

In all cases the cause of the power failure is understood to be the erroneous tripping of the moulded case circuit breakers (MCCB) in the non-uninterruptible power supply (UPS) PDUs in Computer Room 3 (CR3). The possibility that this has been caused by under-voltage in the supply provided from the National Grid is being investigated by analysing the UPS logs. Next steps will be determined based on the analysis of the logs