



# ARCHER SP Service Quarterly Report

Quarter 3 2018



## Document Information and Version History

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

<b>Version</b>	<b>Date</b>	<b>Comments, Changes, Status</b>	<b>Authors, contributors, reviewers</b>
0.1	01/10/18	Initial Draft	Anne Whiting
0.2	07/10/18	Added graphs	Jo Beech-Brandt
0.3	08/10/18	Added phone information	Jo Beech-Brandt
0.4	10/10/18	Added final details and heatmaps	Anne Whiting
0.5	11/10/18	Review	Alan Simpson
0.6	14/10/18	Updates post-review	Anne Whiting
1.0	15/10/18	Version for EPSRC	Alan Simpson

# 1. The Service

## 1.1 Service Highlights

This is the report for the ARCHER SP Service for the Reporting Periods:

July 2018, August 2018 and September 2018.

- Utilisation over the quarter was 84%, up from 82% the previous quarter. Moreover the trend across the quarter was of increasing utilisation with a July figure of 79%, August 85% and for September utilisation was 89%, reflecting the changes made by EPSRC to allocations.
- Demonstrating the importance EPCC places on keeping users' data secure and following best practice, we are very pleased to be able to announce that EPCC has passed our ISO 27001 Information Security certification audit. ISO 27001 is an internationally recognised standard ensuring that an organisation has the controls and processes in place to handle and keep data secure as regards confidentiality, integrity, accessibility and availability. The scope of certification covers all HPC and Data Services run by EPCC, including ARCHER, Cirrus and the RDF. The feedback from the auditor was very positive, with very few minor issues to address. Subsequent audits happen on an annual basis.
- At the request of an ARCHER user, the short development queue availability times have been increased from the current 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.
- 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.
- EPCC have also been continuing our review and update of the SAFE user interface as well as adding features to improve its robustness and scalability.
- Work has successfully been completed with Cray and NCAS on the repurposed compute node configured as a serial node, known as a "mamu" node. This node allows multiple user jobs to be run simultaneously in a manner similar to the serial nodes, and is expected to support the NERC community in compiling their UM code. The repurposed node is now operation on ARCHER and is being tested by the NCAS community.
- Work to replace the Uninterrupted Power Supply (UPS) batteries in the RDF has now been completed with all batteries replaced. The UPS is used to ensure the RDF stays up long enough to close down safely in the case of a power cut, minimising the risk of user data corruption.

## 1.2 Forward Look

- After analysis of the pattern of usage of ARCHER, it was found that queues tend to be quieter at the weekend, primarily between Sunday lunchtime and Monday lunchtime. With the approval of EPSRC, a new weekend queue is therefore going to be configured which will be in operation between 12:00 on Saturday to 12:00 on Monday. Jobs submitted to the weekend queue will receive a 50% discount in charging. Details of how the new queue will work will be circulated once the queue is live.
- At the request of a user, plans are in place to increase the maximum job length on the KNL from 8 hours to 24 hours.
- A survey has been commissioned to review power distribution in the computer room housing ARCHER. Based on the results of the survey, a plan will be created to improve power redundancy for the ARCHER control rack, thus reducing the likelihood of outages that impact the ARCHER service, such as the one that happened on 22 June 2018.
- The new version of PBS, 13.408, is being tested to ensure it does not adversely affect the service before upgrading. As well as providing new functionality, it should resolve issues that we have experienced where jobs that cannot run then prevent other jobs being scheduled.
- An upgrade to the Sonnexion software which supports the work file system on ARCHER is being planned in order to ensure that the Sonnexions are at a level which will be fully supported until the end of the Archer service.
- Plans are being finalised to upgrade the KNL to CLE 6 UP04 which will enable the Meltdown patches to be applied to the KNL system.

## 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	Jul 18		Aug 18		Sep 18		18Q3
Metric	Service Level	Service Points	Service Level	Service Points	Service Level	Service Points	Service Points
2.6.2 – PR	100%	-5	100%	-5	100%	-5	-15
2.6.3 – QC	99.2%	-2	99.5%	-2	99.1%	-2	-6
2.6.4 – UR	1 WD	0	1 WD	0	1 WD	0	0
<b>Total</b>		<b>-7</b>		<b>-7</b>		<b>-7</b>	<b>-21</b>

The details of the above can be found in Section 2.2 of this report.

### 2.1.2 Service Failures

There were no unplanned outages where responsibility lies within the terms of the SP Contract.

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 (-21), no Service Credits apply in 18Q3.

## 2.2 Detailed Service Level Breakdown

### 2.2.1 Phone Response (PR)

	<b>Jul 18</b>	<b>Aug 18</b>	<b>Sep 18</b>	<b>18Q3</b>
Phone Calls Received	32 (4)	33 (5)	30 (2)	<b>95 (11)</b>
Answered in 2 Minutes	32	33	30	<b>95</b>
<b>Service Level</b>	<b>100.0%</b>	<b>100.0%</b>	<b>100.0%</b>	<b>100.0%</b>

The volume of telephone calls remained low in 18Q3. Of the total of 95 calls received above, only 11 were actual ARCHER user calls that either resulted in queries or answered user questions directly.

### 2.2.2 Query Closure (QC)

	<b>Jul 18</b>	<b>Aug 18</b>	<b>Sep 18</b>	<b>18Q3</b>
Self-Service Admin	249	264	456	<b>969</b>
Admin	107	102	89	<b>298</b>
Technical	29	18	12	<b>59</b>
<i>Total Queries</i>	385	384	557	<b>1326</b>
<i>Total Closed in 2 Days</i>	382	382	552	<b>1316</b>
<b>Service Level</b>	<b>99.2%</b>	<b>99.5%</b>	<b>99.1%</b>	<b>99.2%</b>

The above table shows the queries closed by SP during the period.

In addition to the Admin and Technical queries, the following Change Requests were resolved in 18Q3:

	<b>Jul 18</b>	<b>Aug 18</b>	<b>Sep 18</b>	<b>18Q3</b>
Change Requests	1	0	1	<b>2</b>

### 2.2.3 User Registration (UR)

	Jul 18	Aug 18	Sep 18	18Q3
No of Requests	61	48	70	179
Closed in One Working Day	61	48	70	179
Average Closure Time (Hrs)	0.57	0.24	0.50	0.46
Average Closure Time (Working Days)	0.06	0.03	0.05	0.05
<b>Service Level</b>	<b>1 WD</b>	<b>1 WD</b>	<b>1 WD</b>	<b>1 WD</b>

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](mailto:helpdesk@archer.ac.uk). Users may choose not to receive such emails by mailing [support@archer.ac.uk](mailto:support@archer.ac.uk).

#### 2 – Tracking Identifier

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

	Jul 18	Aug 18	Sep 18	18Q3
CRAY	3	2	8	13
ARCHER_CSE	86	100	103	289
ARCHER_SP	627	583	846	2056
Total Queries Assigned	716	685	957	2358
Total Assigned in 1 Hour	716	685	957	2358
<b>Service Level</b>	<b>100%</b>	<b>100%</b>	<b>100%</b>	<b>100%</b>

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.

The following planned maintenance took place this quarter:

Date	Start	End	Duration	Type	Notes	Reason
25/07/18	09:00	14:48	5 hours 48 mins	Full outage	Approved by EPSRC	Implementation of MAMU nodes on ARCHER

### 2.3.3 Quality Tokens

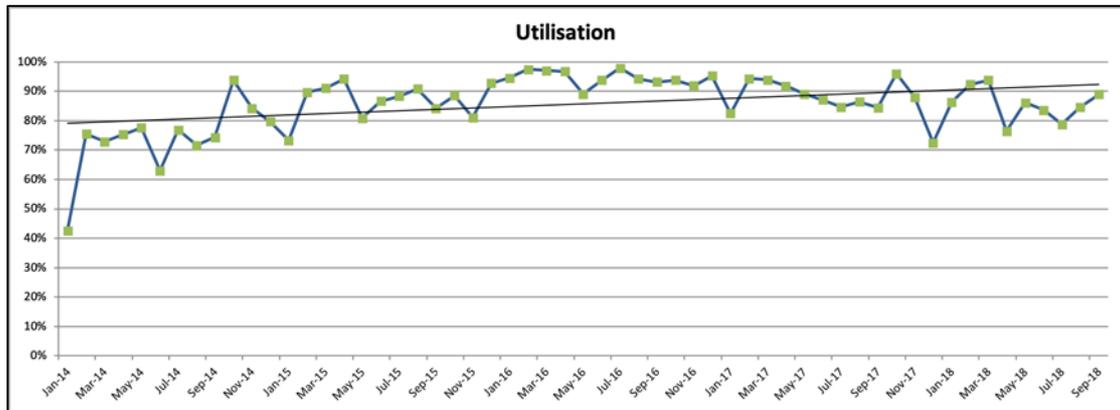
Two sets of positive quality tokens were received this quarter from the same user. A five-star token was to thank us for planning to make the short queue 24x7, which was a request the same user had made via a previous quality token. The second four-star token was to thank us for providing the report generator tools in SAFE.

### 3. Service Statistics

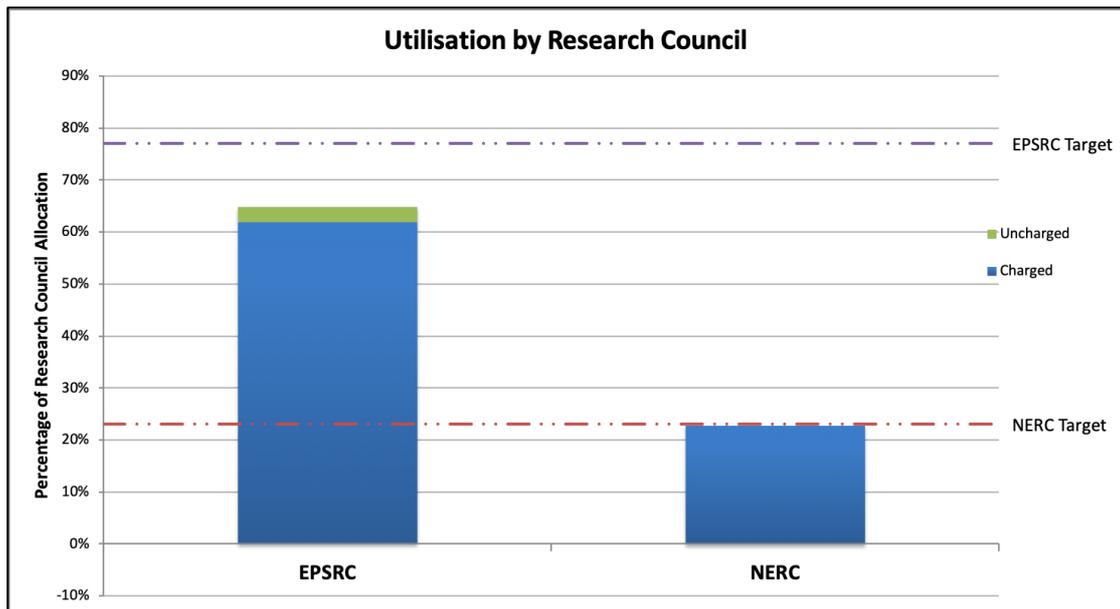
This section contains statistics on the ARCHER service as requested by EPSRC, SAC and SMB.

#### 3.1 Utilisation

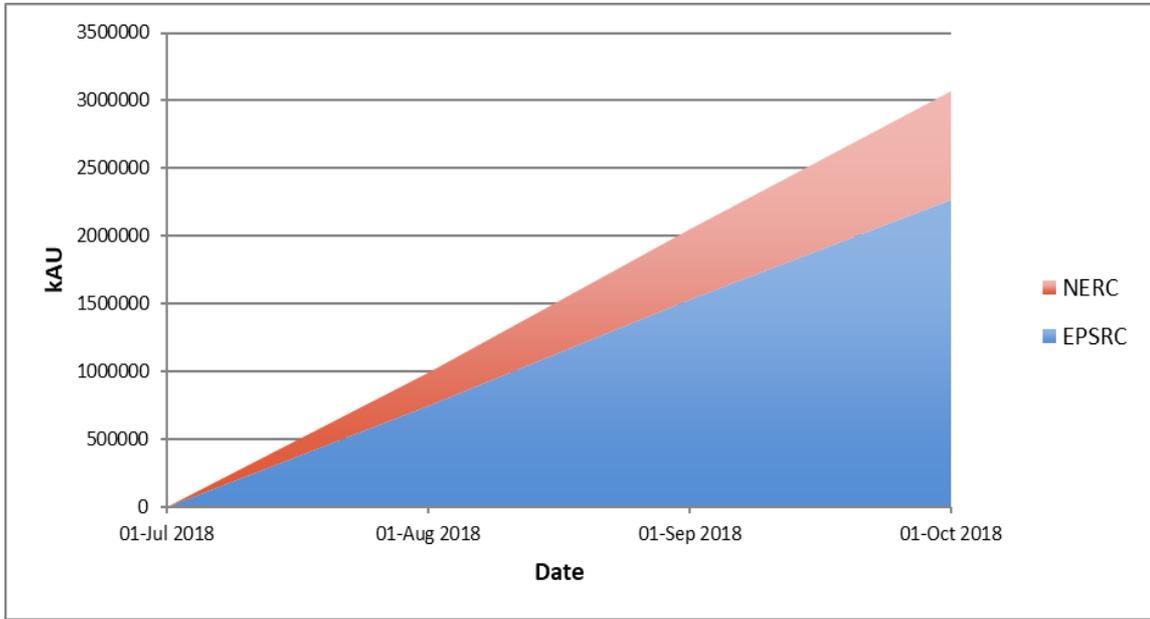
Utilisation over the quarter was 84%, up from 82% the previous quarter. The trend across the quarter was of increasing utilisation. Utilisation for July was 79%, for August 85% and for September 89%. The plot below shows a steady increase in utilisation over the lifetime of the service to Dec 2015 and since then the service has effectively been operating at maximum capacity as shown by the generally steady utilisation value.



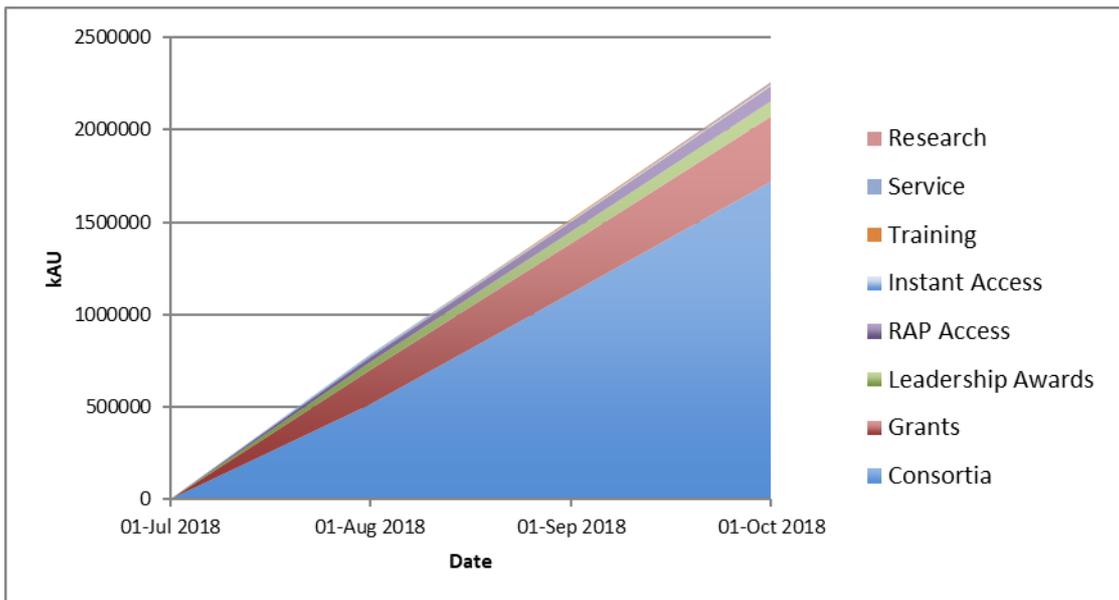
The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that EPSRC did not meet their target this quarter with their usage being at 65% (against their target of 77%) whereas NERC met their target with utilisation being 23% (against their target of 23%). This compares with 60% for EPSRC and 23% for NERC for the previous quarter.



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

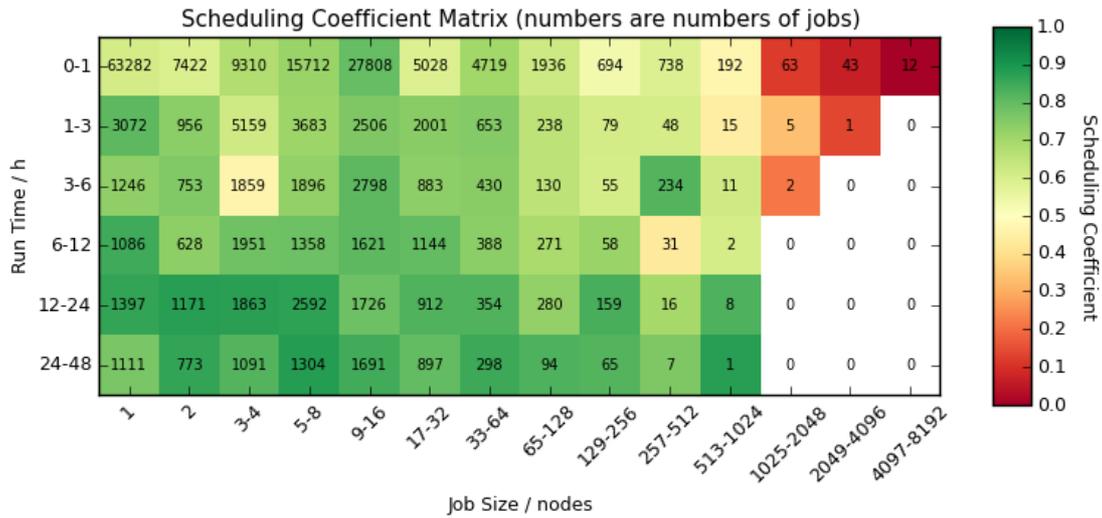


The cumulative allocation utilisation for the quarter by EPSRC broken down by different project types (see below) shows that the majority of usage comes from the scientific Consortia (as expected) with significant usage from research grants, ARCHER Leadership projects and ARCHER RAP projects. The times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



### 3.2 Scheduling Coefficient Matrix

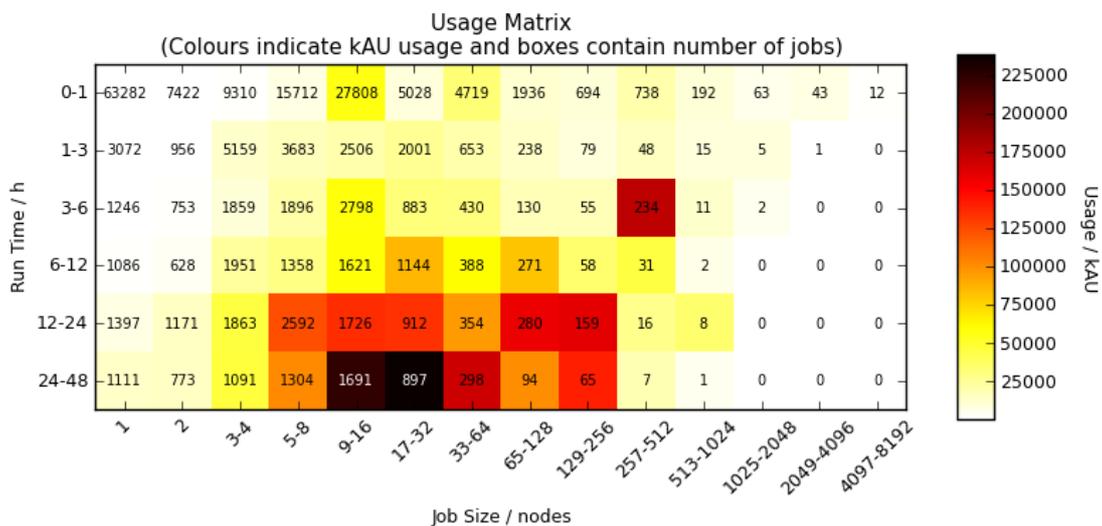
The colour in the matrix indicates the value of the Scheduling Coefficient. This is defined as the ratio of runtime to runtime plus wait time. Hence, a value of 1 (green) indicates that a job ran with no time waiting in the queue, a value of 0.5 (pale yellow) indicates a job queued for the same amount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



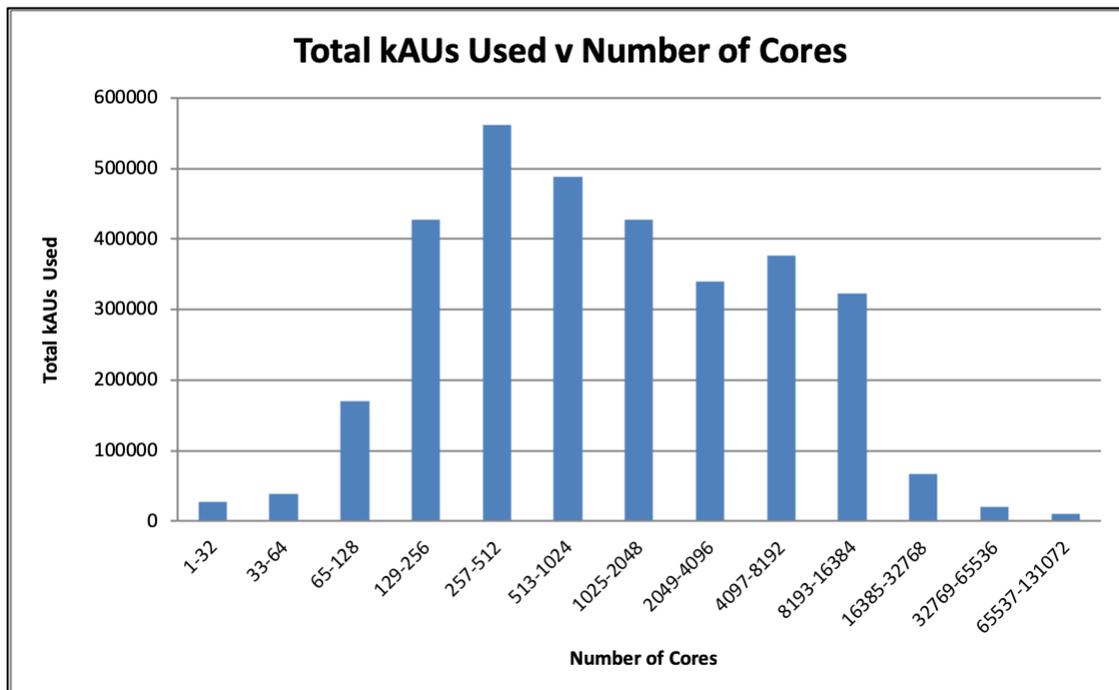
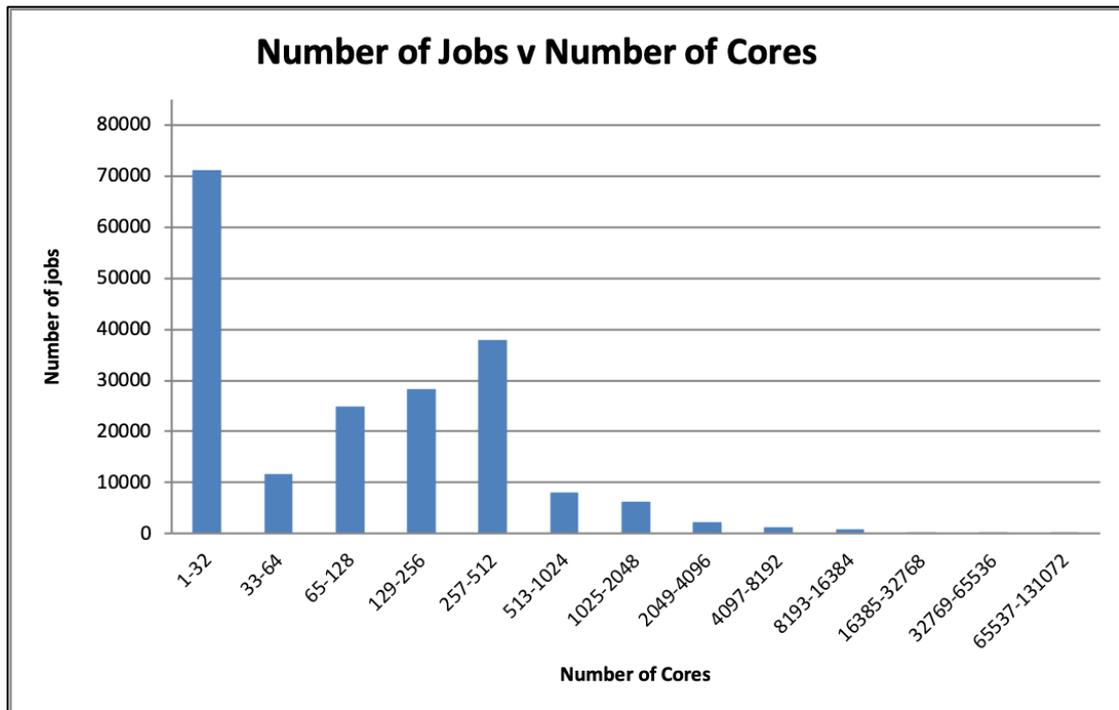
### 3.3 Additional Usage Graphs

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

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.

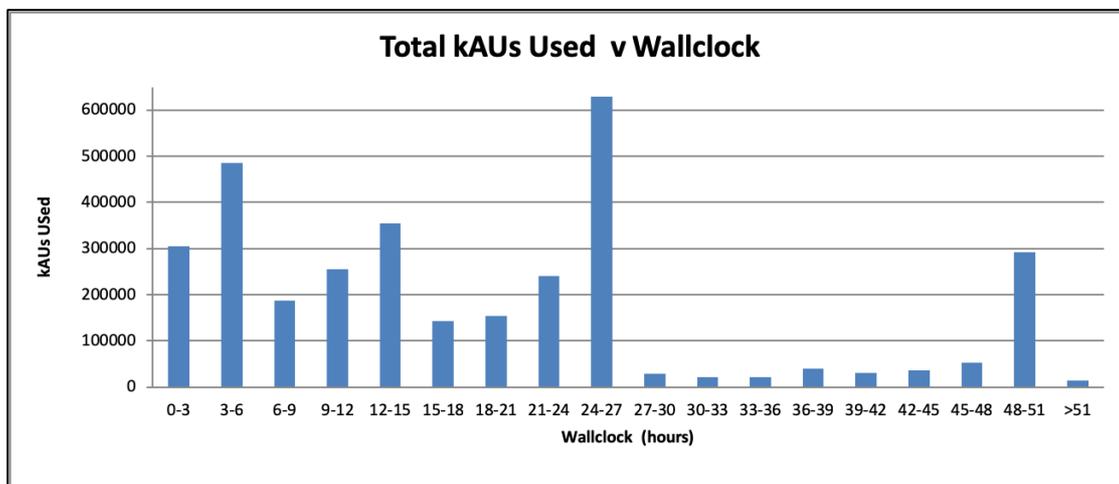
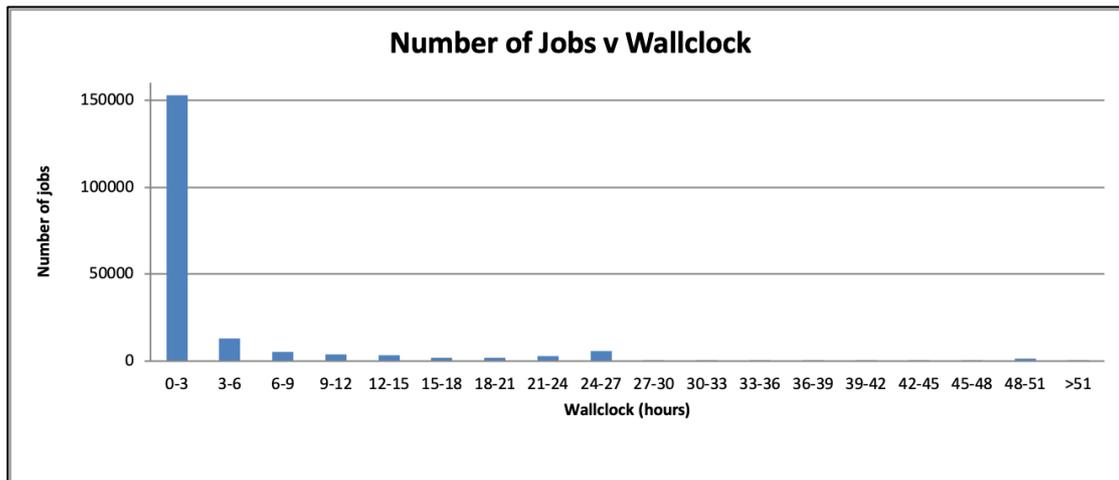


## Analysis of Job Sizes



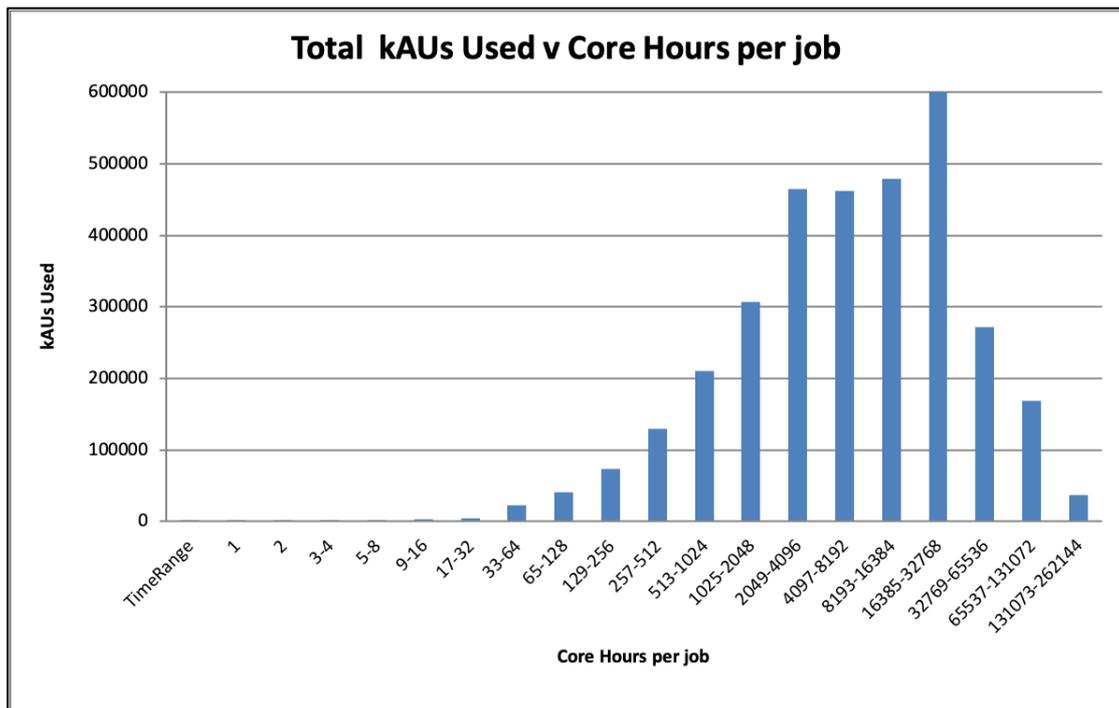
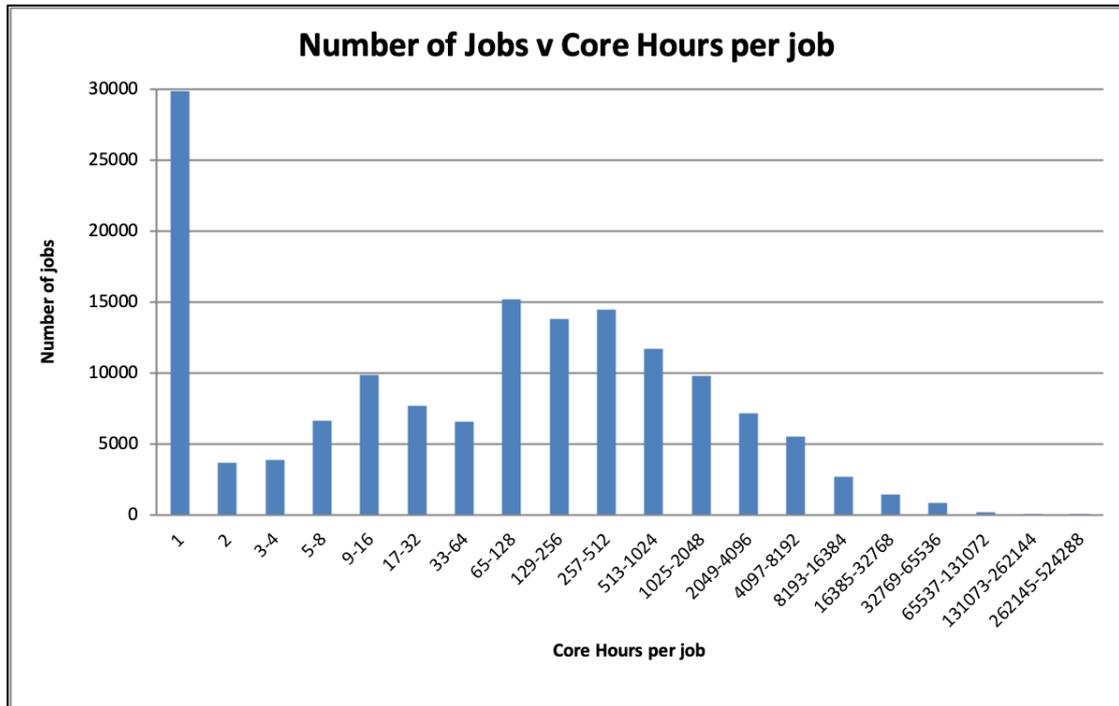
The first graph shows that, in terms of numbers, there are a significant number of jobs using no more than 1024 cores. However, the second graph reveals that most of the kAUs were spent on jobs between 65 cores and 16384 cores. The number of kAUs used is closely related to money and shows better how 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 up to 27 hours with a second peak for jobs at 48-51 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.