



ARCHER SP Service Quarterly Report

Quarter 1 2019



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	01/04/19	Initial Draft	Anne Whiting
0.2	01/04/19	Added cumulative graphs	Anne Whiting
0.3	02/04/19	Added cumulative usage graphs	Anne Whiting
0.4	03/04/19	Added scheduling coefficient data	Clair Barrass
0.5	05/04/19	Added utilisation and usage graphs. Added phone data	Jo Beech-Brandt
0.6	08/04/19	Reviewed	Alan Simpson
1.0	11/04/19	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:

January 2019, February 2019 and March 2019.

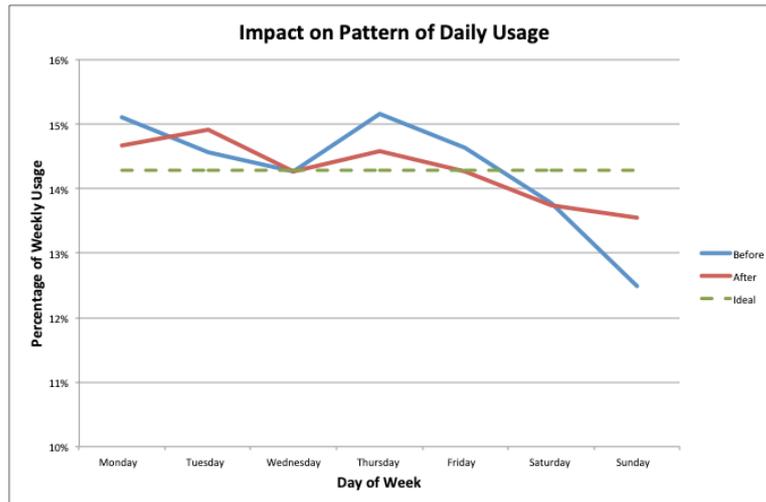
- Utilisation over the quarter was 86%, maintaining the same level as the previous quarter.
- The results of the 2018 annual ARCHER User Survey have been analysed. 188 responses were received with the mean results shown below (scores 1 representing “Very Unsatisfied” and 5 representing “Very Satisfied”). Overall satisfaction with the service was at the highest level achieved during its lifetime. A few of the ratings have gone down slightly, for example for hardware and software, which perhaps reflects the age of the service and that we are nearing the end of its lifetime; user comments support this. The helpdesk continues to perform exceptionally well, with an overall rating of 4.5 and a range of positive comments from users.

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

Users who have supplied contact details will be contacted to discuss any need for support or to discuss feedback provided further.

- We are keen to minimise the impact of service outages for full maintenance sessions on the user community and aim to carry out maintenance tasks during at-risk sessions wherever possible. We are pleased to be able to say that we have taken no maintenance outages this quarter.
- We are pleased to announce that we have passed our annual ISO 9001:2015 external quality management audit. This provides us with a framework to measure the effectiveness of service improvements made, and reflects the importance placed on delivering the best possible service to our users.
- The programming environment has been upgraded to the final version of the Cray Programming Environment for CLE5.2UP04. This resolved a number of issues reported by users including multiple Fortran bugs.

- The weekend queue, which was enabled the previous quarter and runs between 12:00 on Saturday to 12:00 on Monday every weekend, continues to be popular. 29 different projects have used the queue so far including all 8 of the EPSRC HEC, all 3 of the NERC consortia and 110 different users. The graph shows the impact of the queue on the pattern of daily usage, with weekend usage increased and the levelling out of daily usage.



- A new feature has been added to the SAFE to support credential and quota management for the S3 compatible object store that we run as part of the Cirrus Tier-2 service. This enables users of any system including ARCHER to access the object store depending on the access policies set.

1.2 Forward Look

- With the continued use of the current Cray programming environment in this its last version, more frequent interactions are ongoing with the other UK Cray sites that are on the same programming environment to share experiences and work-arounds to issues.
- To improve support for the service, more frequent interactions with third party hardware vendors, Mellanox (for site network), DDN, OCF (both RDF) are ongoing to discuss open support issues and to improve communications between the organisations involved.
- Continuous enhancements are being made to the system monitoring tools to help catch issues before they impact users.
- Plans are underway for increasing the ACF external and internal network links to 100GB improving communication speeds for the user community.
- A review of the software and firmware levels of the RDF components will be carried out in order to plan for attachment of the RDF to ARCHER 2.
- Work is underway to prepare for a combined ISO 9001 Quality Management and ISO 27001 Information Security external audit in the autumn. Moving to a combined system and audit leverages the strengths of the process-based quality management approach with the controls provided by the information security management system to deliver the best and most secure service to our users.

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	Jan 19		Feb 19		Mar 19		19Q1
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
Total		-7		-7		-7	-21

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

2.1.2 Service Failures

There was one unplanned outage this quarter:

15/02/19 07:05 – 13:20 duration 6 hours, 15 minutes

This problem was caused by a user running a job which exceeded the file limit for their project. Their job failed, causing more open connections than the current operating system supports. The system therefore became overloaded and unable to support running work. Parameters in the Systems Database Configuration are being tuned to reduce the likelihood of the issue reoccurring but it is a system limitation of the current operating system, CLE5.2UP04.

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 19Q1.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Jan 19	Feb 19	Mar 19	19Q1
Phone Calls Received	14 (3)	23 (3)	26 (6)	63 (12)
Answered in 2 Minutes	14	23	26	63
Service Level	100.0%	100.0%	100.0%	100.0%

The volume of telephone calls remained low in 19Q1. Of the total of 63 calls received above, only 12 were actual ARCHER user calls that either resulted in queries or answered user questions directly.

2.2.2 Query Closure (QC)

	Jan 19	Feb 19	Mar 19	19Q1
Self-Service Admin	956	522	318	1796
Admin	195	108	132	435
Technical	27	26	21	74
<i>Total Queries</i>	1178	656	471	2305
<i>Total Closed in 2 Days</i>	1148	652	461	2261
Service Level	97.5%	99.4%	97.9%	98.1%

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 18Q4:

	Jan 19	Feb 19	Mar 19	19Q1
Change Requests	1	3	0	4

2.2.3 User Registration (UR)

	Jan 19	Feb 19	Mar 19	19Q1
No of Requests	144	83	66	293
Closed in One Working Day	144	82	66	292
Average Closure Time (Hrs)	0.7	0.6	0.6	0.6
Average Closure Time (Working Days)	0.1	0.1	0.1	0.1
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:

	Jan 19	Feb 19	Mar 19	19Q1
CRAY	4	2	3	9
ARCHER_CSE	89	135	117	341
ARCHER_SP	1748	948	764	3460
Total Queries Assigned	1841	1085	884	3810
Total Assigned in 1 Hour	1841	1085	884	3810
Service Level	100%	100%	100%	100%

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.

There was a spike in the number of queries in January, which we have also seen in previous years. This is caused by an increase in new user accounts and project administration tasks, which saw a decline during December.

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.

No planned maintenance outages were taken this quarter.

2.3.3 Quality Tokens

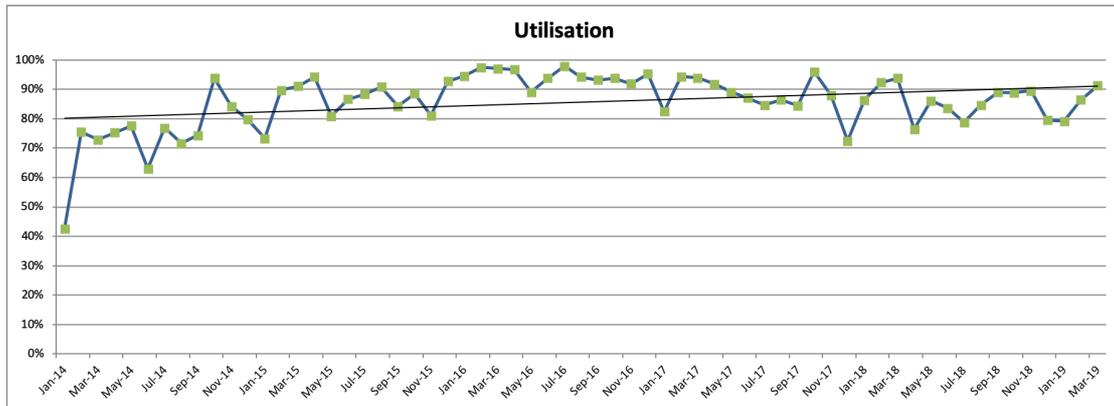
No quality tokens were received this quarter.

3. Service Statistics

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

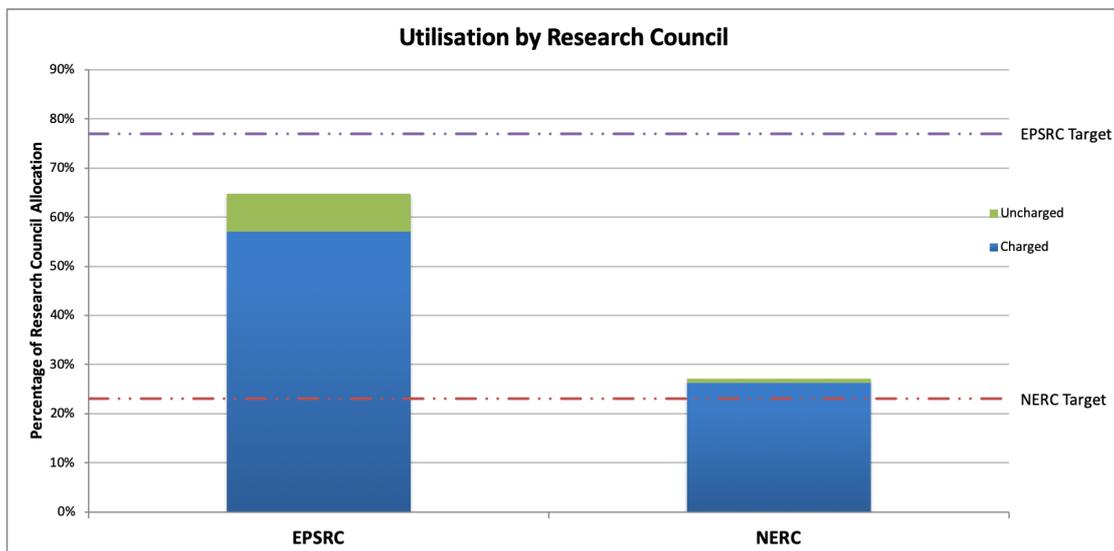
3.1 Utilisation

Utilisation over the quarter was 86%, the same as for the previous quarter. 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.

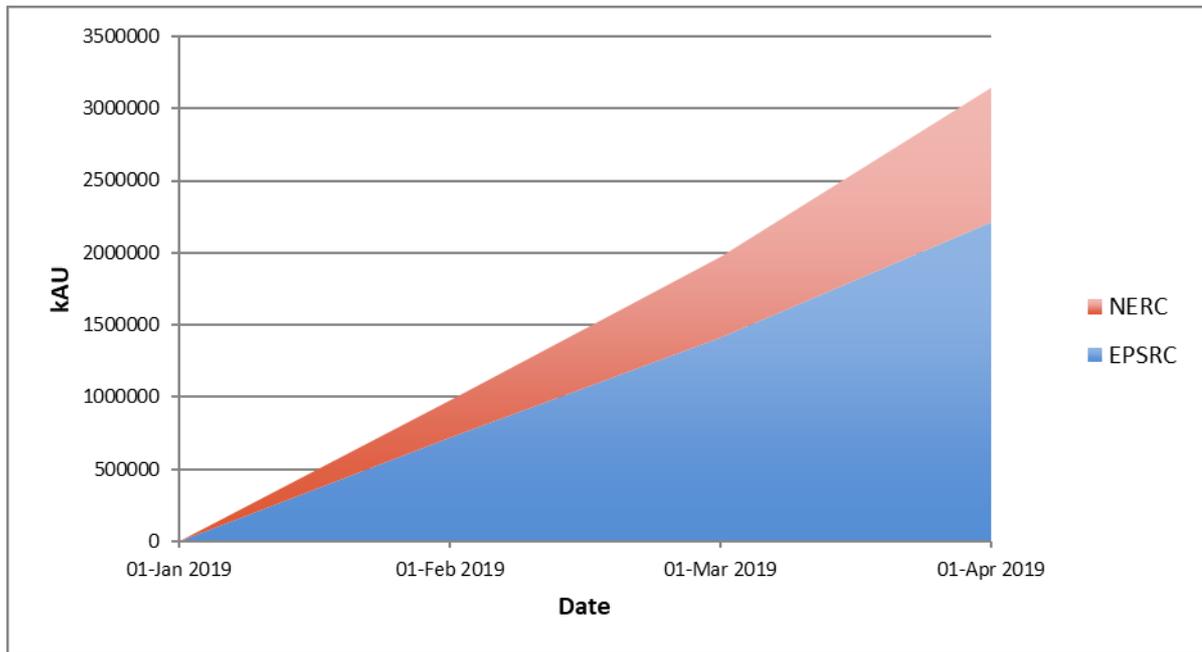


The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that EPSRC did not meet their target this quarter with their usage being at 65% (against their target of 77%) whereas NERC exceeded their target with utilisation being 27% (against their target of 23%). This compares with 63% for EPSRC and 25% for NERC for the previous quarter.

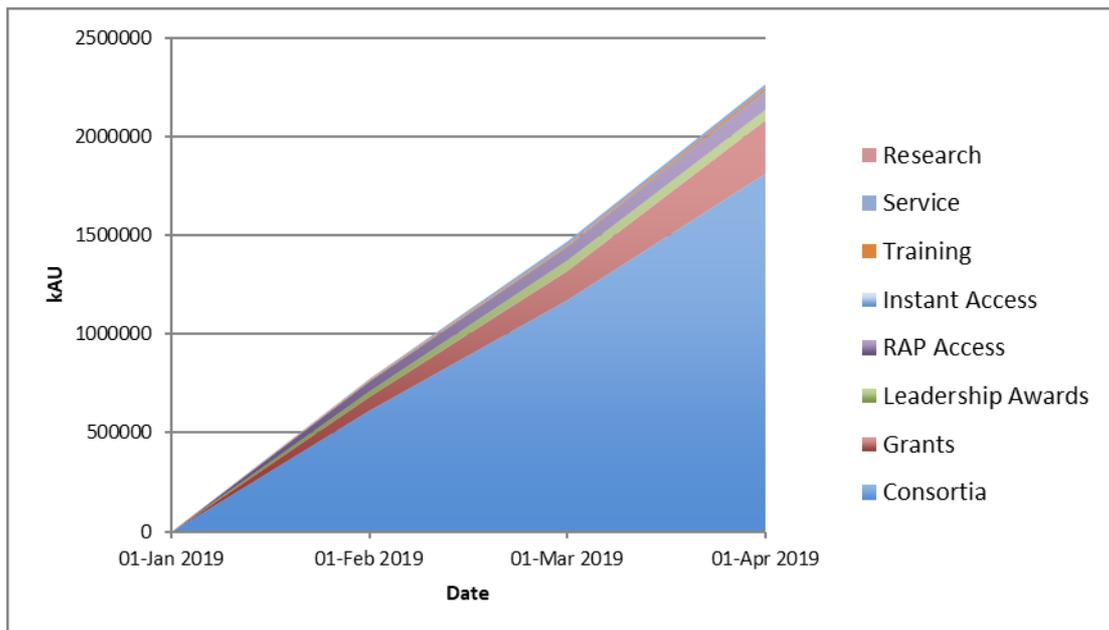
The increases for the uncharged portions are due to the introduction of the weekend queue on ARCHER that has a 50% reduction.



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

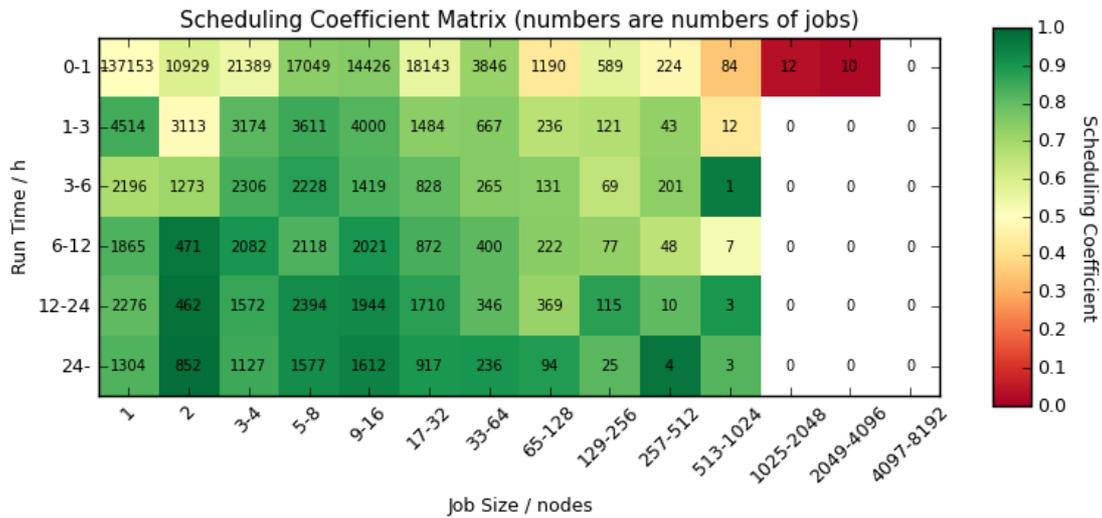


The cumulative allocation utilisation for the quarter by EPSRC broken down by different project types (see below) shows that the majority of usage comes from the scientific Consortia (as expected) with significant usage from research grants, ARCHER Leadership projects and ARCHER RAP projects. The times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



3.2 Scheduling Coefficient Matrix

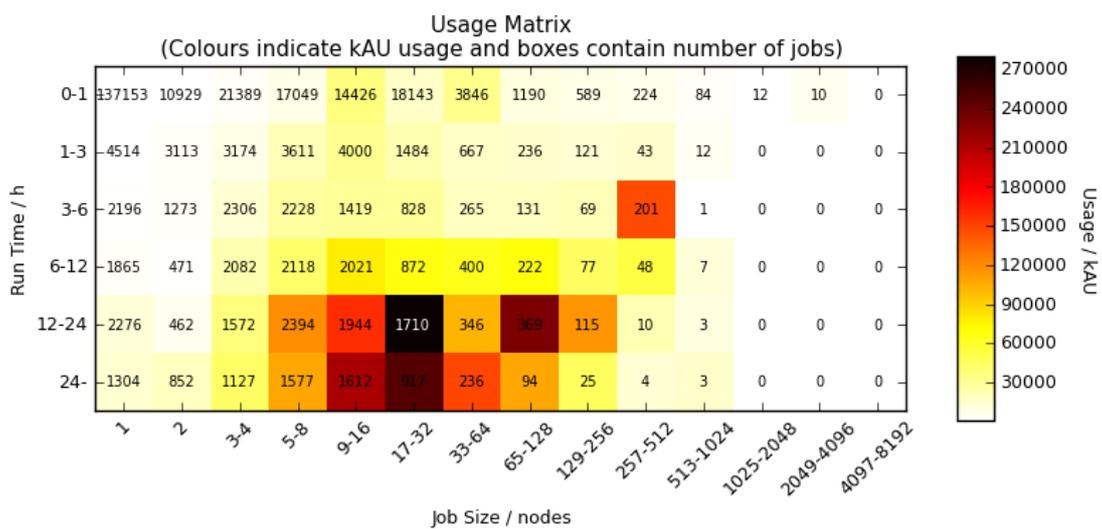
The colour in the matrix indicates the value of the Scheduling Coefficient. This is defined as the ratio of runtime to runtime plus wait time. Hence, a value of 1 (green) indicates that a job ran with no time waiting in the queue, a value of 0.5 (pale yellow) indicates a job queued for the same amount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



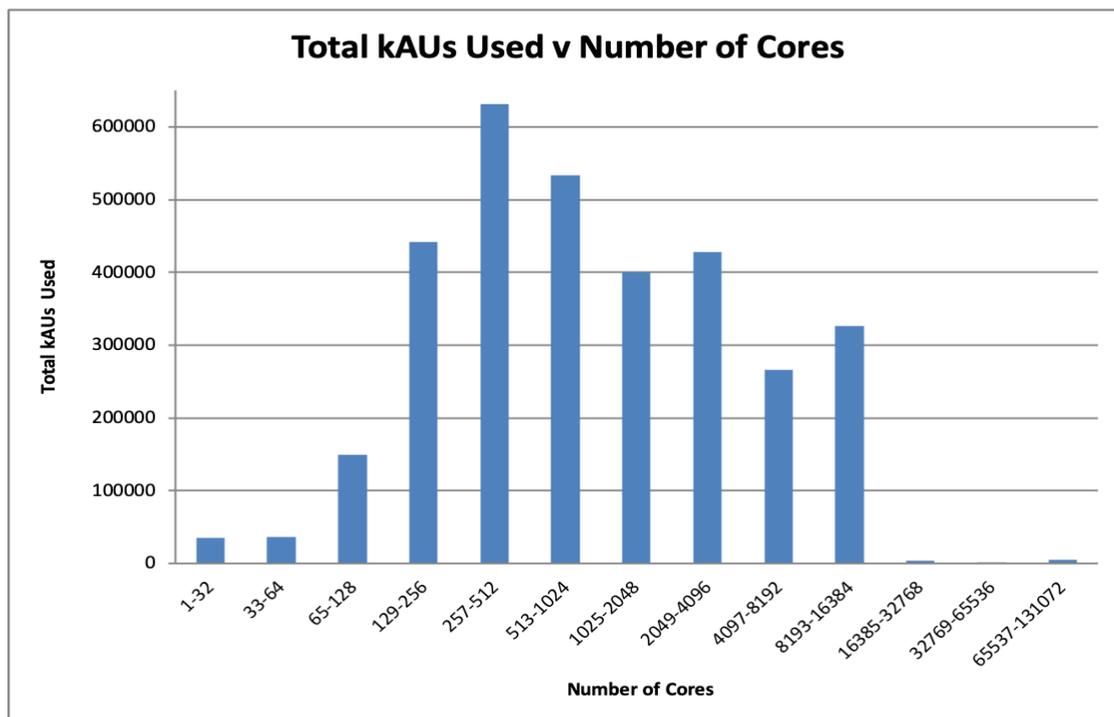
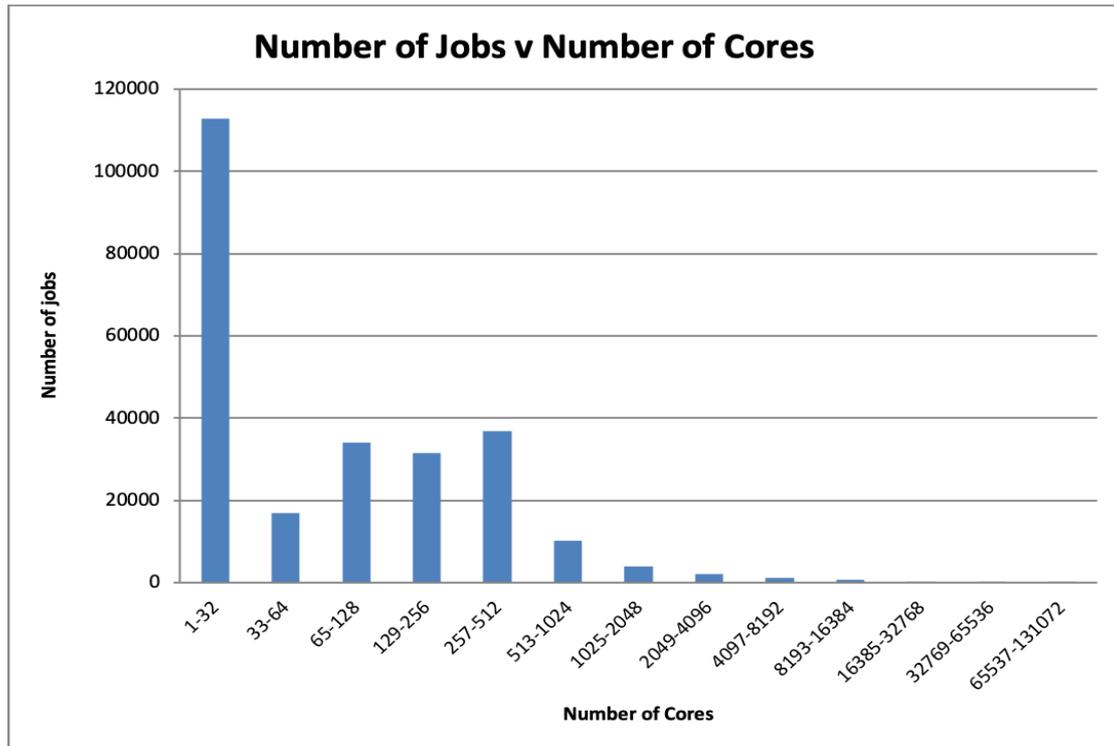
3.3 Additional Usage Graphs

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.

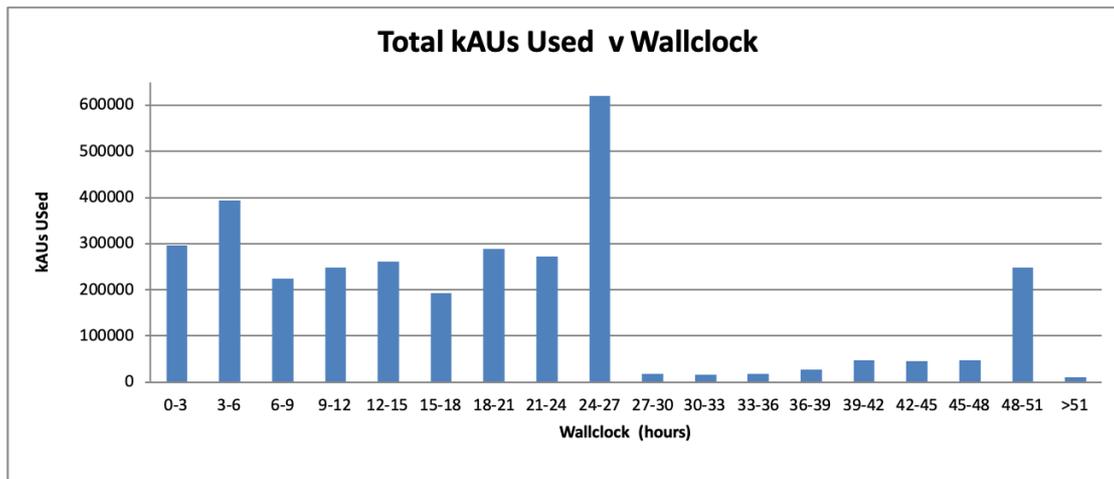
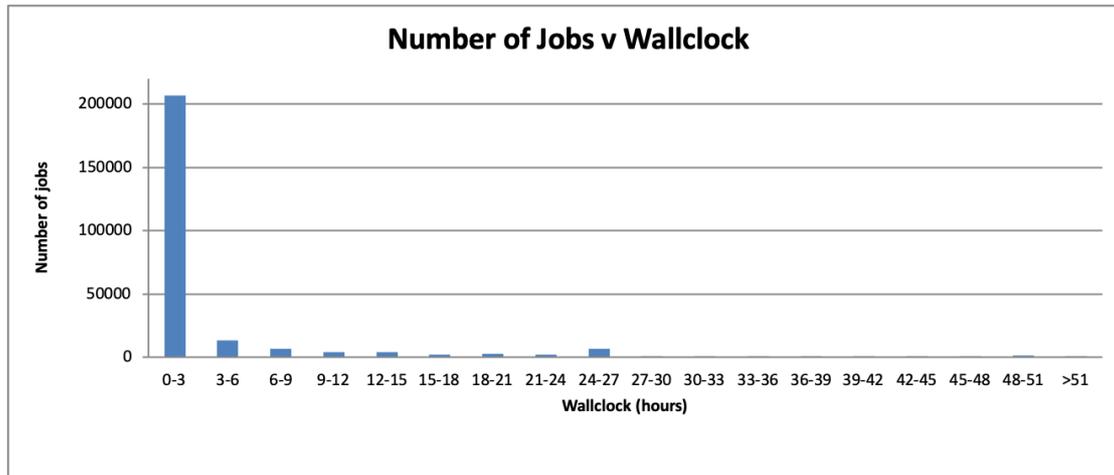


Analysis of Job Sizes



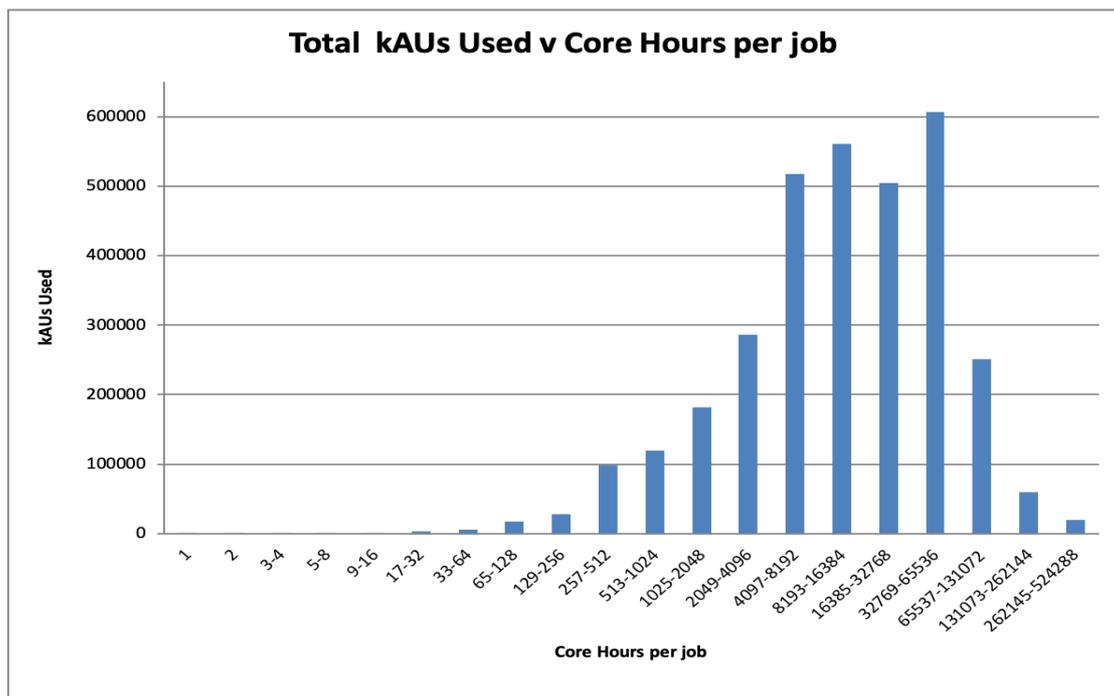
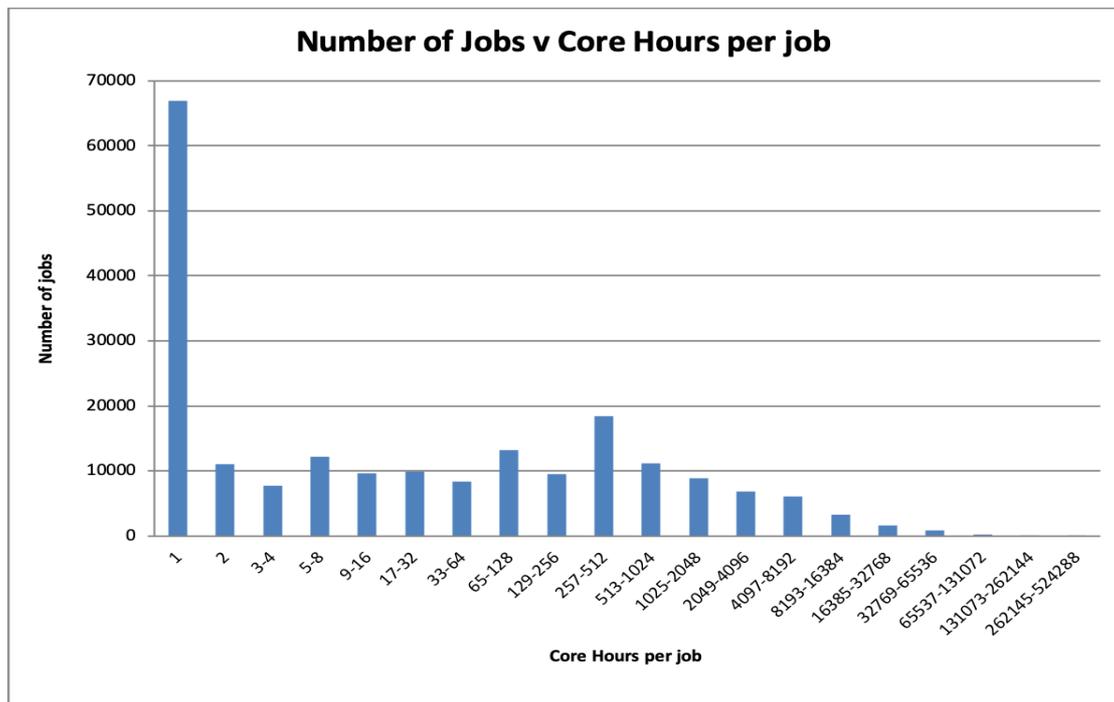
The first graph shows that, in terms of numbers, there are a significant number of jobs using no more than 1024 cores. However, the second graph reveals that most of the kAUs were spent on jobs between 65 cores and 16384 cores. The number of kAUs used is closely related to money and shows better how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