



ARCHER SP Service Quarterly Report

Quarter 1 2017



Document Information and Version History

Version:	1.0
Status	Final
Author(s):	Alan Simpson, Anne Whiting, Andy Turner, Mike Brown, Stephen Booth, Jo Beech-Brandt
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	03/04/17	Initial Draft	Anne Whiting
0.2	05/04/17	Added graphs	Jo Beech-Brandt
0.3	06/04/17	Added information from OSG	Anne Whiting
0.4	06/04/17	QA review	Alan Simpson
0.5	06/04/17	Update post review	Anne Whiting
1.0	06/04/17	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 2017, February 2017 and March 2017.

- Utilisation on the system during 17Q1 was 90% as compared to 94% in 16Q4. Utilisation in January was 82% but by February it was back up to 94%. The continued high utilisation of the service supports the need for ongoing investment in HPC.
- We are delighted to announce that we have passed ISO 9001:2015 certification for the provision of National Tier 1 and Tier 2 HPC services, making use of established high quality ITIL service delivery processes. This certification covers: helpdesk, systems administration, training, in-depth support, eCSE and outreach. The ISO 9001:2015 quality management system provides a framework to continually improve services for our users.
- The number of responses to the ARCHER user survey was 161 compared to 153 in 2014 and 230 in 2015. The feedback was very positive, with the mean score for overall satisfaction being 4.3 (on a scale of 1 (unsatisfactory) to 5(excellent)). The helpdesk feedback was particularly favourable with a mean score of 4.5. Positive comments were received on the improvements to the queuing times after the priority changes were made. Five lucky winners have been selected randomly from the users who replied to the survey and have been offered a prize of either 3000 KAUs on ARCHER, or a £10 donation to Save the Children.
- The KNL system is being well used by the wider user community both to run appropriate jobs and to compare performance of codes between ARCHER and the KNL:
 - 49% utilization over the period
 - 267 users
 - 2672 jobs
 - 4574 kAUs used
 - KNL user survey run
 - Consultation was carried out with users on memory configuration options
 - Benchmarking is ongoing of code performance between the two architectures
- The ARCHER Champions Workshop took place in Leeds on 10 February, co-located with HPC-SIG. The Workshop went very well and started the process of integrating Tier 2 into the Champions community.

1.2 Forward Look

- The CLE upgrade to 5.2 UP04 will take place on 12th and 26th April 2017.
- The next ARCHER Champions Workshop will take place on Monday June 26th and Tuesday June 27th (lunchtime to lunchtime) and will be at Hartree, Daresbury. It will be a joint Tier 1 / Tier 2 Champions event.
- We will repeat the scheduler analysis that we undertook throughout summer 2016 to ensure that the scheduler configuration remains well suited to the ARCHER workload.

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 17		Feb 17		Mar 17		17Q1
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	97.2%	-2	98.8%	-2	97.9%	-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 were no unplanned outages where responsibility lies within the terms of the SP Contract.

However, there was an unplanned outage from around 1700 on 23rd January until around 1000 on 25th January due to disruption of the high voltage distribution network in the south Edinburgh area. The source of the problem was beyond the reasonable control of the University. The Distribution Network Operator have not commented on the reasons for the disruption, but have only confirmed that it occurred. The University will report to EPSRC on the consequences of the event under the terms of the Lease Agreement.

Details of planned maintenance sessions 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 17Q1.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Jan 17	Feb 17	Mar 17	17Q1
Phone Calls Received	36 (17)	21 (4)	34 (6)	91 (27)
Answered 2 Minutes	36	21	34	91
Service Level	100.0%	100.0%	100.0%	100.0%

The volume of telephone calls remained low in 17Q1. Of the total of 91 calls received above, only 27 were actual ARCHER user calls that either resulted in queries or answered user questions directly.

2.2.2 Query Closure (QC)

	Jan 17	Feb 17	Mar 17	17Q1
Self-Service Admin	648	546	1103	2297
Admin	151	188	212	551
Technical	32	25	21	78
<i>Total Queries</i>	831	759	1336	2926
<i>Total Closed in 2 Days</i>	808	750	1308	2866
Service Level	97.2%	98.8%	97.9%	97.9%

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 17Q1:

	Jan 17	Feb 17	Mar 17	17Q1
Change Requests	0	0	3	3

2.2.3 User Registration (UR)

	Jan 17	Feb 17	Mar 17	17Q1
No of Requests	78	84	129	291
Closed in One Working Day	78	84	129	291
Average Closure Time (Hrs)	0.9	0.7	0.6	0.7
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 Additional Metrics

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 17	Feb 17	Mar 17	17Q1
CRAY	16	10	11	37
ARCHER_CSE	106	73	188	367
ARCHER_SP	1239	1158	1879	4276
Total Queries Assigned	1361	1241	2078	4680
Total Assigned in 1 Hour	1361	1241	2078	4680
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.

2.3.2 Maintenance

Maintenance now takes place on a single day each month (fourth Wednesday of each month). This is marked as a full maintenance session for a maximum of 8 hours taken. There is an additional “at-risk” session that is scheduled for the second Wednesday of each month. This reduces the number of sessions taken, which then reduces user impact since the jobs running on the service have to be drained down once per month and not twice. It also eases the planning for training courses running on ARCHER.

Such Maintenance Periods are defined as ‘Permitted Maintenance ’ and recorded in the Maintenance Schedule. A 6-month forward plan of maintenance has been agreed with the Authority.

Where possible, SP will perform maintenance on an ‘At-risk’ basis, thus maximising the Availability of the Service. The following planned maintenance took place in the Service Quarter.

No ‘permitted maintenance ’ sessions were taken this quarter.

2.3.3 Quality Tokens

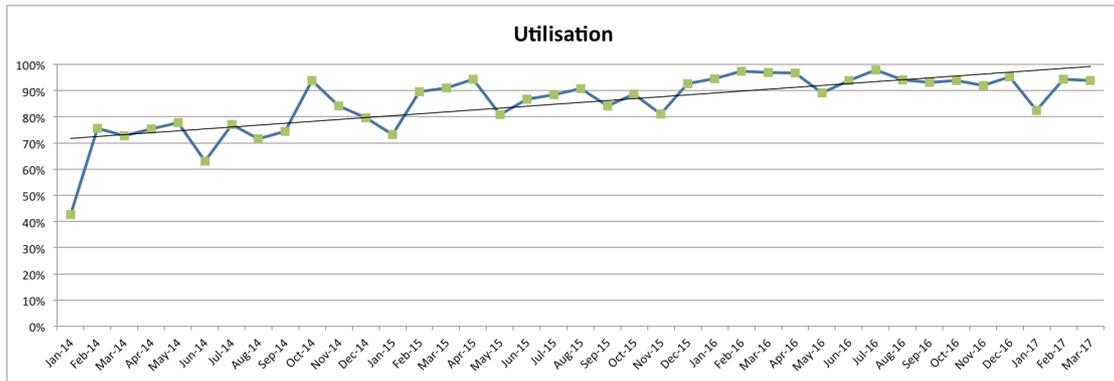
One quality token has been received this quarter. This was positive (3 stars) with the comment included ‘The queue times have been much better in the past several weeks! Also, my reservations have been great for my long simulations.’ The user has been responded to.

3. Service Statistics

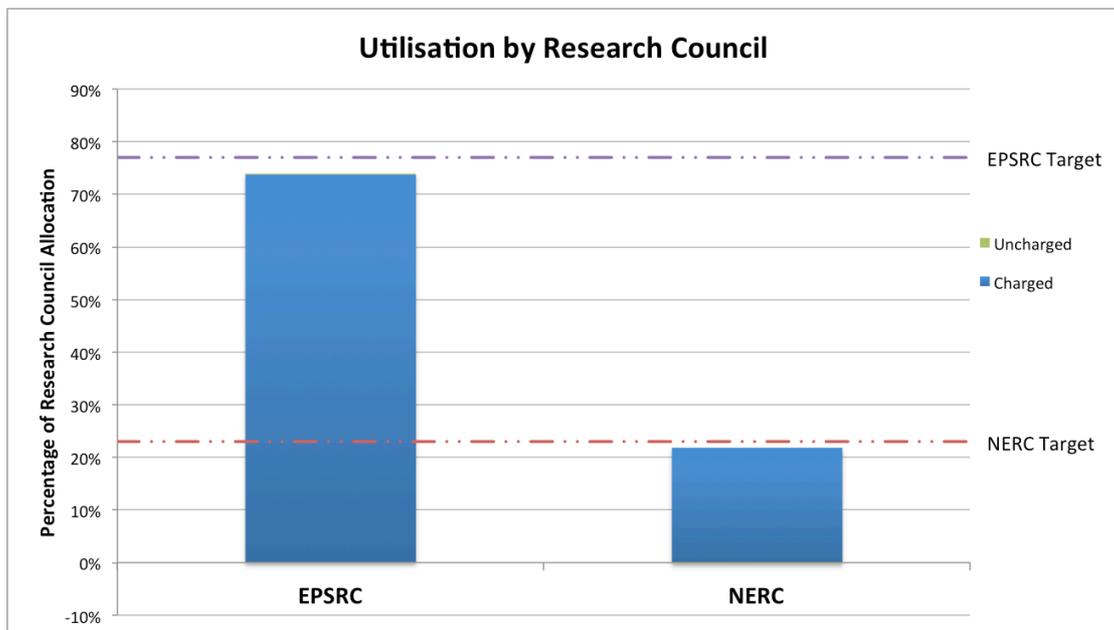
This section contains statistics on the ARCHER service as requested by EPSRC, SAC and SMB.

3.1 Utilisation

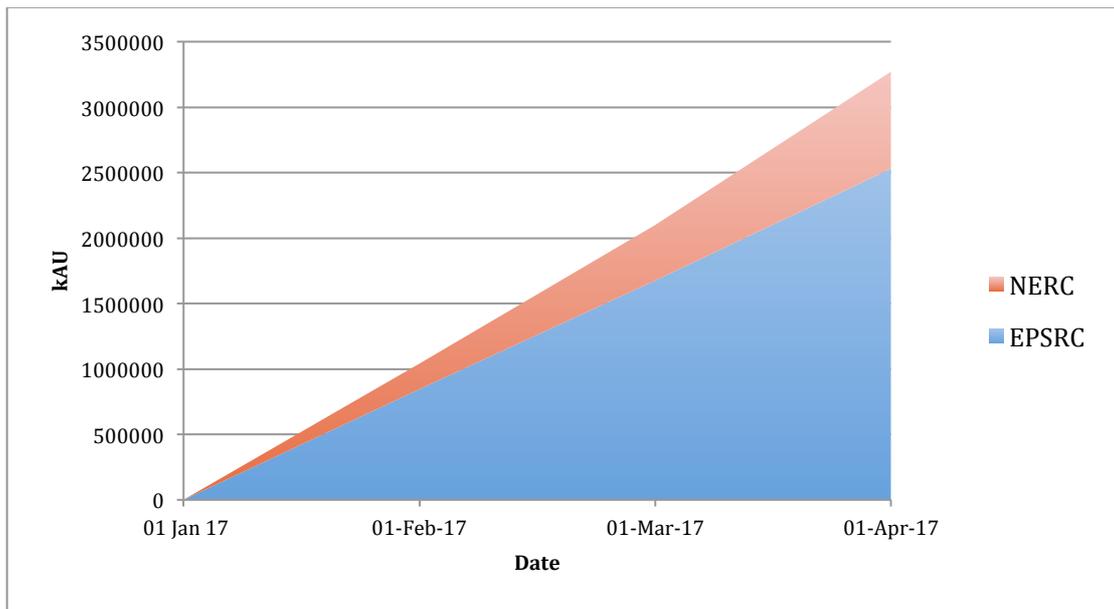
Utilisation over the quarter was 94%. The plot below shows a steady increase in utilisation over the lifetime of the service to Dec 2015 and since then the service has effectively been operating at maximum capacity as shown by the steady utilisation value:



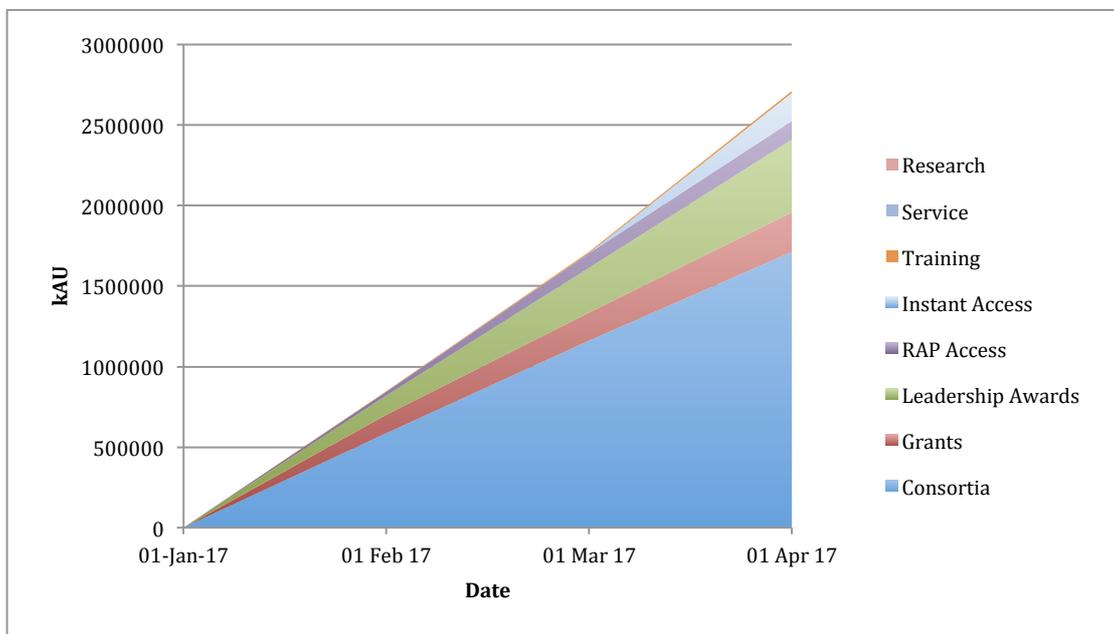
The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that both Research Councils did not meet their respective targets this quarter with EPSRC being at 73.8% (against their target of 77%) and NERC's utilisation being 21.82% (against their target of 23%).



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

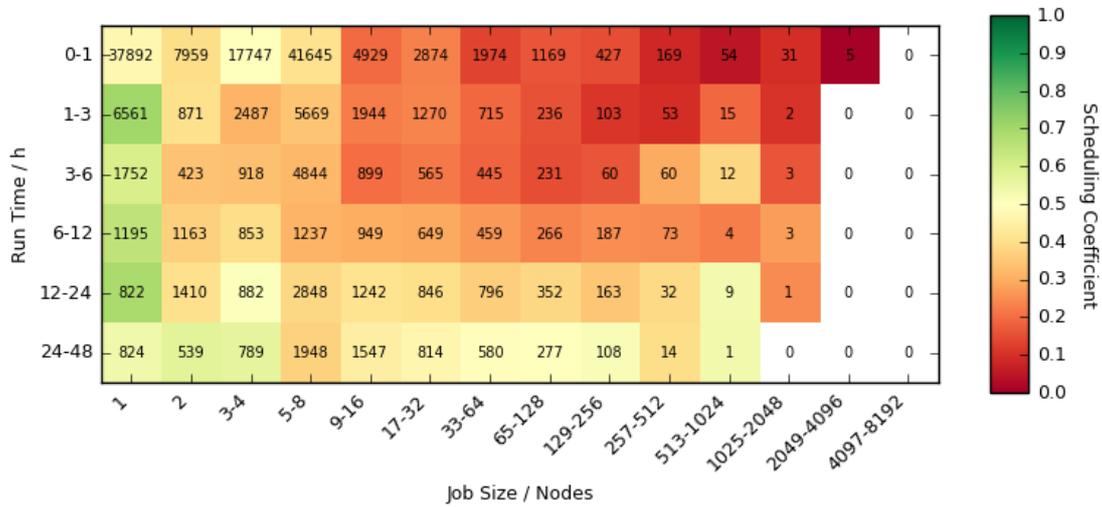


The cumulative allocation utilisation for the quarter by EPSRC broken down by different project types (see below) shows that the majority of usage comes from the scientific Consortia (as expected) with significant usage from research grants, ARCHER Leadership projects and ARCHER RAP projects. The times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



3.2 Scheduling Coefficient Matrix

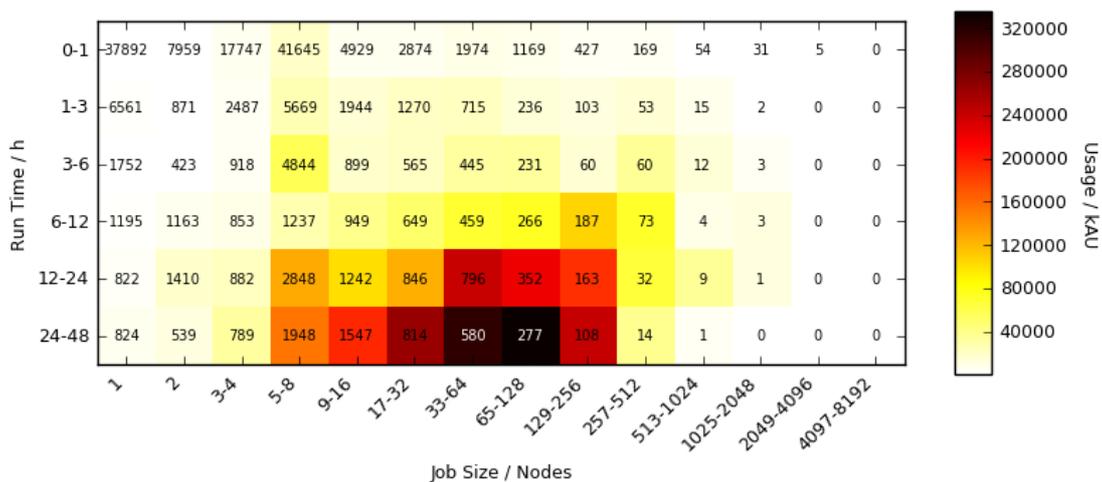
The colour in the matrix indicates the value of the Scheduling Coefficient. This is defined as the ratio of runtime to runtime plus wait time. Hence, a value of 1 (green) indicates that a job ran with no time waiting in the queue, a value of 0.5 (pale yellow) indicates a job queued for the same amount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



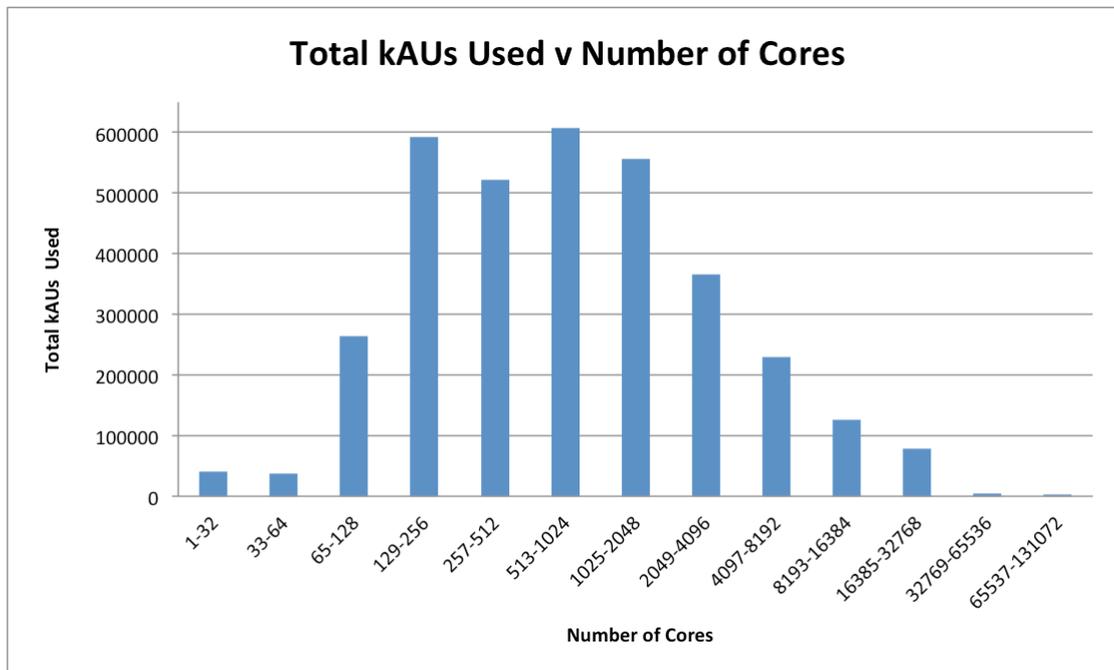
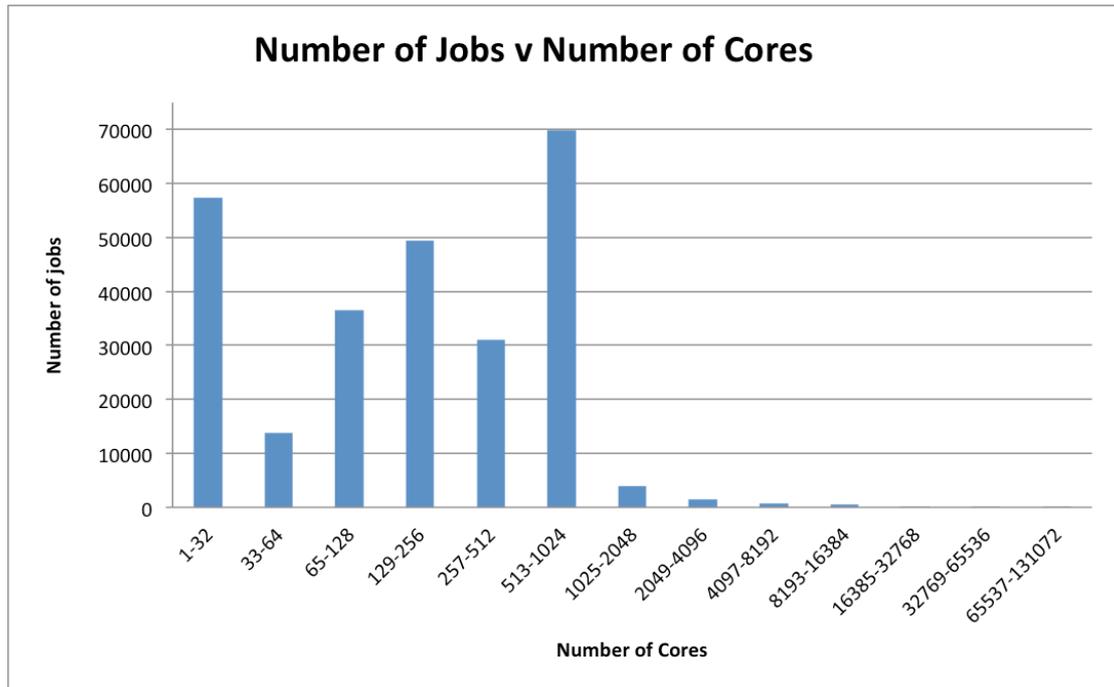
3.3 Additional Usage Graphs

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 kAU expended for each class, and the number in the box is the number of jobs of that class.

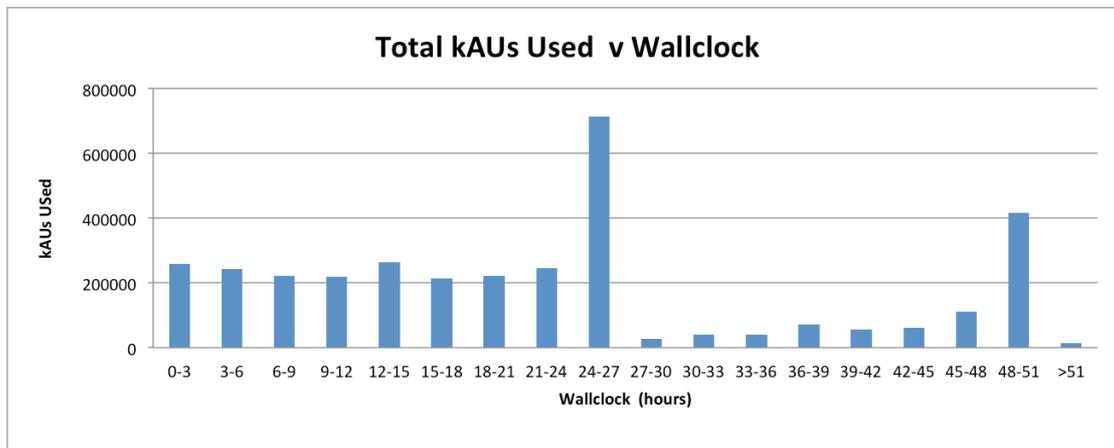
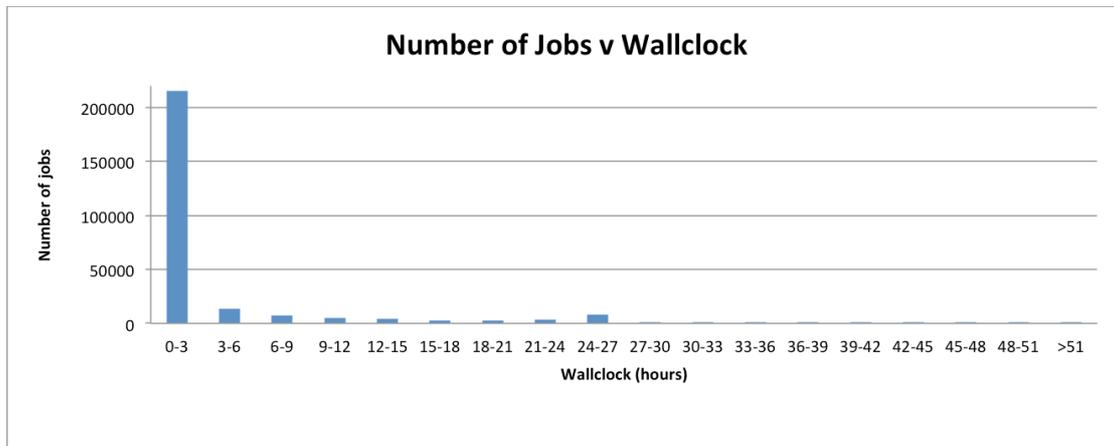


Analysis of Job Sizes



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 8192 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

