



ARCHER SP Service Quarterly Report

Quarter 3 2016



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	04/10/16	Initial Draft	Anne Whiting
0.2	06/10/16	Second draft	Anne Whiting
0.3	07/10/16	Added Graphs	Jo Beech-Brandt
0.4	07/10/16	Reviewed	Alan Simpson, Andy Turner
1.0	07/10/16	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 2016, August 2016 and September 2016

- Utilisation on the system during 16Q3 was 95% as compared to 93% in 16Q2. The continued high utilisation of the service supports the need for ongoing investment in HPC whilst presenting challenges to the user community experiencing slow turnaround of their jobs.
- After discussions with CSE and EPSRC, the SP Service has adjusted the job priority formula in the ARCHER scheduling system to make queue times more equitable across different job sizes. Analysis following this change has shown a dramatic reduction in the number of jobs that are queuing for very long times and a balancing of queue times across different job sizes. The volume of users registering their concerns with queuing times has greatly reduced since this change has been made. Feedback and analysis was provided to the Research Councils.
- EPSRC agreed with Cray a 12-node XC40 KNL system that is due to be available to users in October. Preparatory work is ongoing to incorporate the KNL system into ARCHER. Technical verification is being prepared. The SAFE is being modified to facilitate access and reporting. Training material and a KNL Driving Test are being developed. Access routes have been agreed with EPSRC and expressions of interest from users are being collected.
- Patches and configuration changes were made to resolve the long-running login node failure problem. These changes have been successful and no further login node failures have occurred since.
- The maintenance schedule has been modified with the approval of the Research Councils, reducing the number of full maintenance days to one a month to minimise user impact. As ever, the maintenance schedule is published well in advance and timely reminders sent out as a maintenance period approaches to allow the user community to plan their work around it.
- The second ARCHER Champions Workshop was successfully held in Oxford in September with 30 attendees contributing to many interesting discussions.
- Feedback on the new version of SAFE was positive and user suggestions have been tracked and incorporated. Access to the older version was removed on 1st October. Prior to its removal the two remaining users were contacted individually and their requirements discussed and logged for action. The new version of SAFE has also been rolled out to the users of the DIRAC and EPCC systems both to deliver the same benefits as for the ARCHER community and to avoid confusion for users who have accounts on more than one system.
- Package account management functionality has been rolled out in SAFE to enable users to easily and consistently request access to packages such as VASP and CASTEP. Standardising this process will also provide easier access to reporting data around package usage.
- We have produced a risk analysis on implementing a high-level, fair-share scheduling system on ARCHER and circulated this to the Research Councils for discussion. The analysis identified substantial risks associated with trialing such a scheme due to the limited scope for testing outwith a production environment.

1.2 Forward Look

- Work will continue to prepare, test and release the KNL system for use by the user community. The system is due to be available in the second half of October. This will be provided with similar network, compiler and tools as the main system. During the first month it will be freely available to all users with no CPU limit although with the following constraints to ensure all users can get reasonable access:
 - Max running jobs = 1
 - Max job length = 12 hours
 - Max number of queued jobs per user = 2
 - Max job size = 8 nodes

After the first month:

- A guaranteed allocation for all users who used it in the first month of 1% of monthly cycles
 - A similar allocation for all ARCHER users who want to use the KNL system
 - All other users can obtain access through a KNL driving test
 - Allocation for successful eCSEs, with amount to be decided
- An upgrade to CLE 5.2UP04 is to be delivered in 3 approved maintenance sessions between 26 October 2016 and 23 November 16. Testing is commencing on the TDS.
 - Investigations will be carried out to prepare ARCHER to upgrade to CLE 6, delivering the new functionality that this brings and ensuring the continued full support of the system's operating system by Cray.
 - SAFE functionality to allow users to register publications will be released in the SAFE. This information will contribute to the service benefits realization data collected.
 - The next ARCHER Champions Workshop will take place in Leeds on 9-10 February, co-located with HPC-SIG and following the decision for the meetings to be bi-annual and distributed around the UK.
 - Work will continue to prepare for the ISO 9001:2015 certification audit. The external stage 1 document audit for the ISO 9000 certification will take place in December 2016. The full external audit should take place in 17Q1.

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 16		Aug 16		Sep 16		16Q3
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	99.0%	-2	99.9%	-2	99.8%	-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 Service Failures in the period as defined in the metric. 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 16Q3.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Jul 16	Aug 16	Sep 16	16Q3
Phone Calls Received	27(12)	12(2)	17(3)	56
Answered 2 Minutes	27	12	17	56
Service Level	100.0%	100.0%	100.0%	100.0%

The volume of telephone calls remained low in 16Q3. Of the total of 56 calls received above, only 17 were genuine ARCHER user calls that either resulted in queries or answered user questions directly.

2.2.2 Query Closure (QC)

	Jul 16	Aug 16	Sep 16	16Q3
Self-Service Admin	258	222	345	825
Admin	138	142	112	392
Technical	17	15	15	47
<i>Total Queries</i>	413	379	472	1264
<i>Total Closed in 2 Days</i>	411	376	468	1255
Service Level	99.5%	99.2%	99.1%	99.3%

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 16Q3:

	Jul 16	Aug 16	Sep 16	16Q3
Change Requests	1	1	2	4

2.2.3 User Registration (UR)

	Jul 16	Aug 16	Sep 16	16Q3
No of Requests	65	45	79	264
Closed in One Working Day	65	45	79	264
Average Closure Time (Hrs)	0.5	0.6	0.5	0.8
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:

	Jul 16	Aug 16	Sep 16	16Q3
CRAY	5	11	10	26
ARCHER_CSE	62	62	100	224
ARCHER_SP	727	636	855	2218
Total Queries Assigned	794	709	965	2468
Total Assigned in 1 Hour	794	709	965	2468
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

A change in the maintenance arrangements was agreed with the Authority during this quarter. There is now a single day each month (fourth Wednesday of each month) that 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 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.

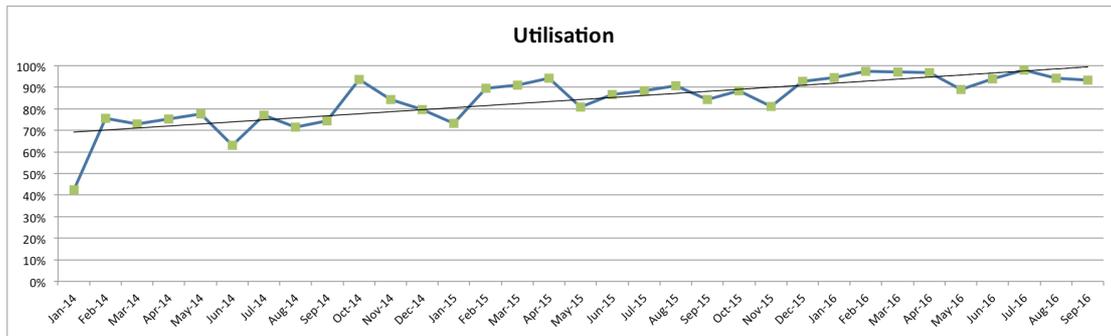
Date	Start	End	Duration	Type	Notes	Reason
13/07/16	0900	1700	8 hrs	At-Risk	EPSRC Approved 0900 – 1700	Reverting to earlier version of GPFS
27/07/16	0900	1605	7 hrs 5 mins	Full Outage	EPSRC Approved 0900 – 1700	Work on eslogin and espp nodes
10/08/16	0900	1700	8 hrs	At-Risk	EPSRC Approved 0900 – 1700	Reboot of login nodes
24/08/16	0900	1700	8 hrs	RDF Outage	EPSRC Approved 0900 – 1700	RDF hardware maintenance, CDT installation, deletion of old modules
14/09/16	0900	1700	8 hrs	At-Risk	EPSRC Approved 0900 – 1700	Implement budget on DAC Configuration changes on RDF
22/06/16	0800	1717	9hrs 17mins	Full Outage	EPSRC Approved 0800 – 1800 (CRAY SESSION)	Hardware and firmware updates

3. Service Statistics

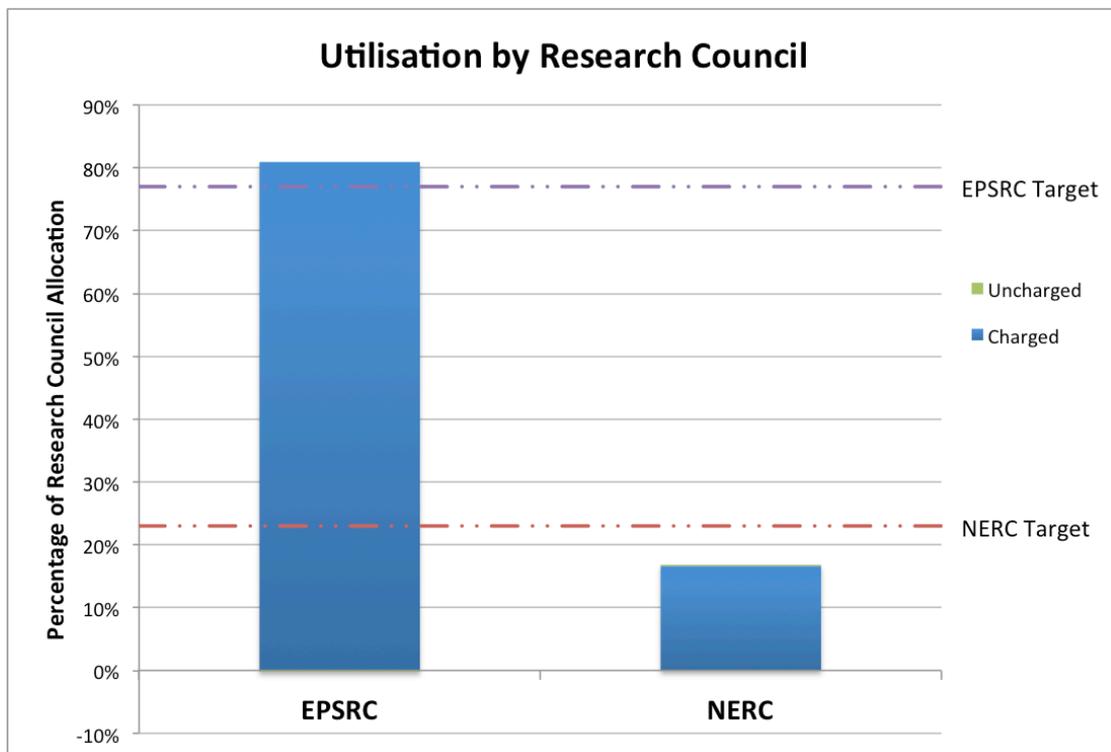
This section contains statistics on the ARCHER service as requested by EPSRC, SAC and SMB.

3.1 Utilisation

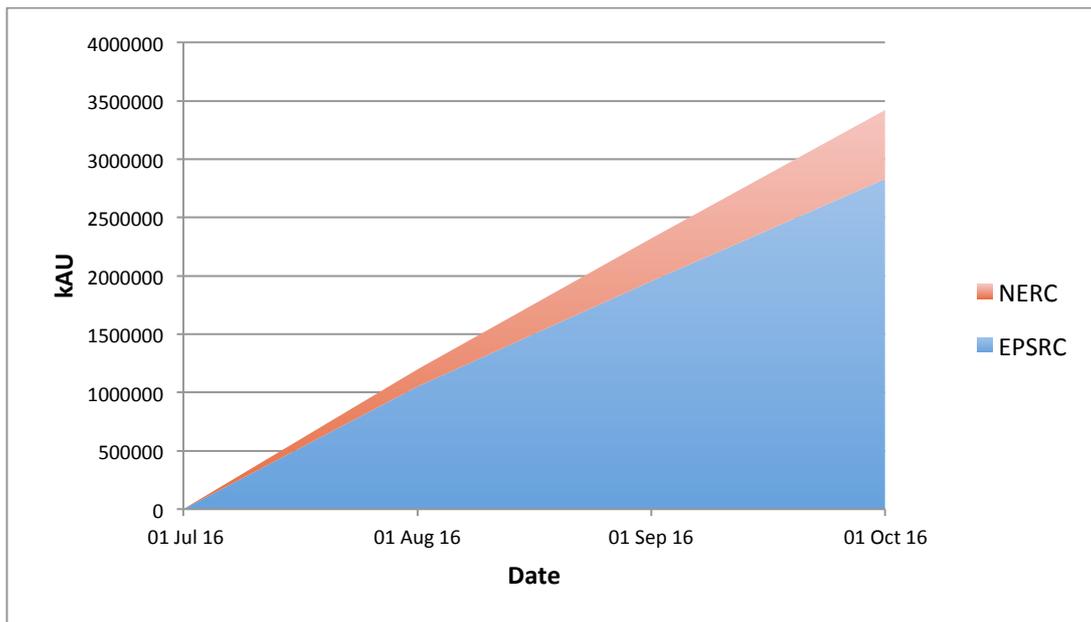
Utilisation over the quarter was 95%. The plot below shows a steady increase in utilisation over the lifetime of the service to Dec 2015 and since then the service has effectively been operating at maximum capacity as shown by the steady utilisation value:



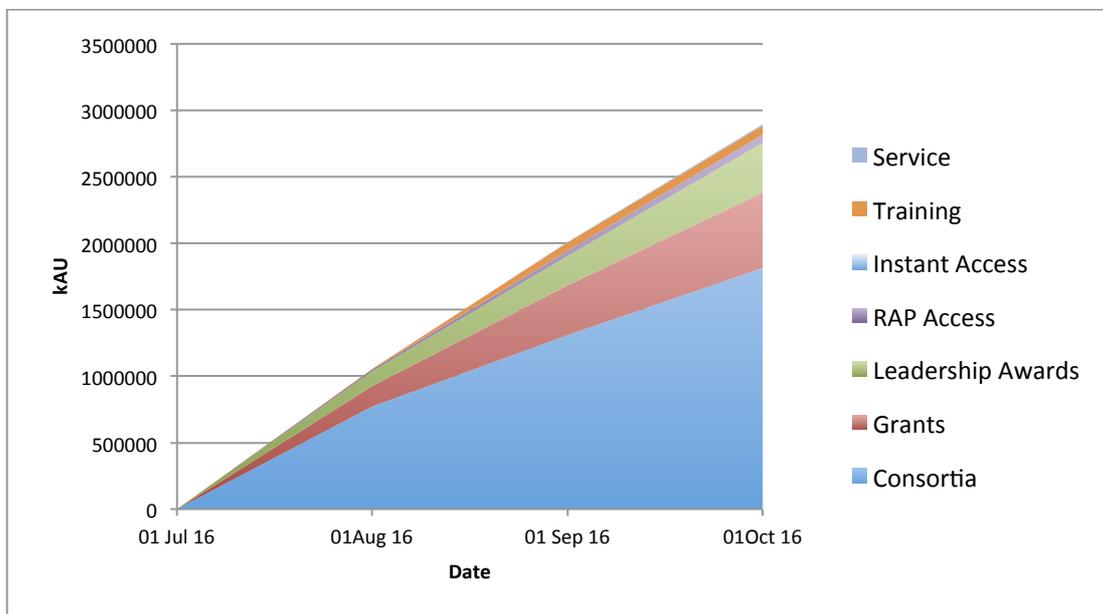
The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that the EPSRC utilisation exceeded their 77% target this quarter and was 81%, whereas NERC utilisation was 17% with their target being 23%.



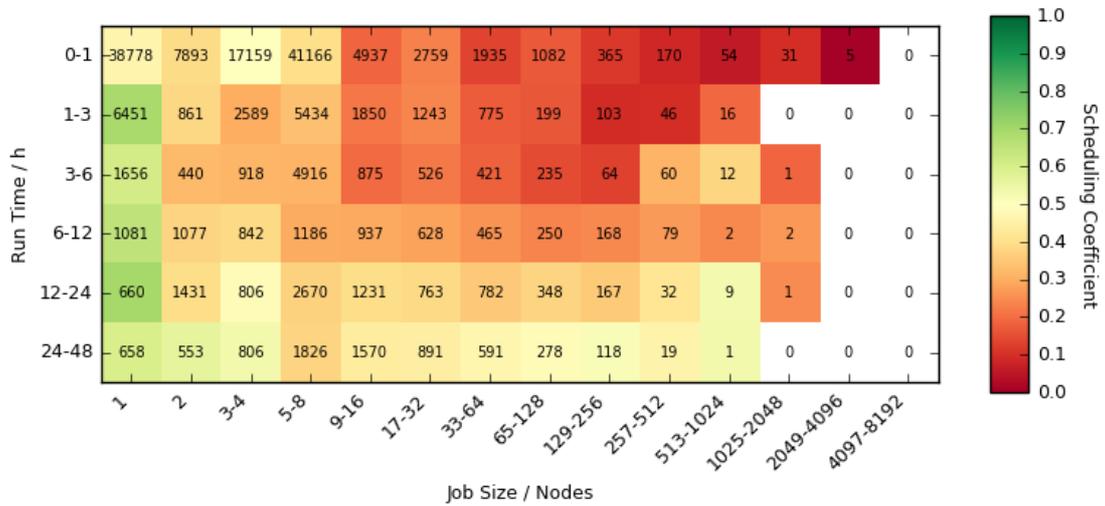
The cumulative allocation utilisation for the quarter by the Research Councils is shown below:



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

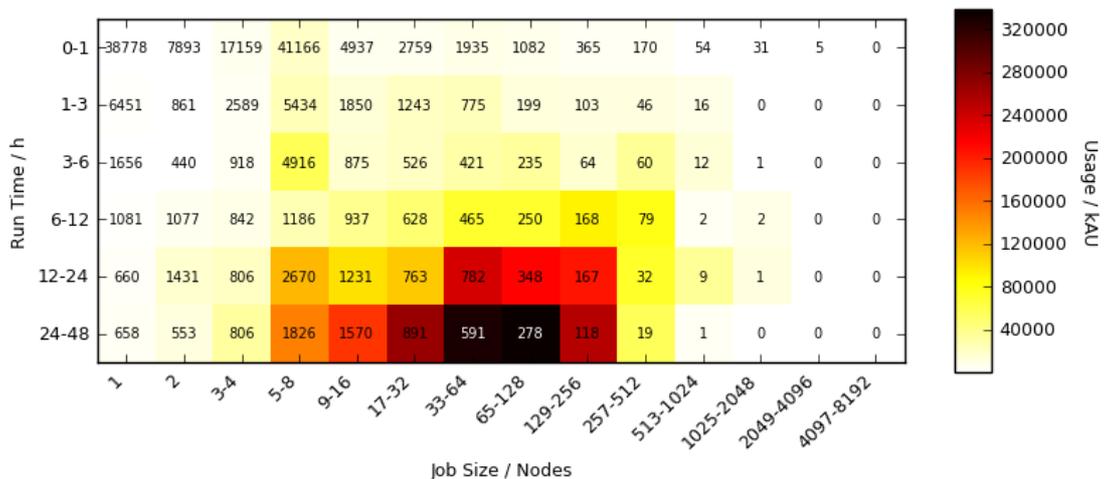


The colour in the matrix indicates the value of the Scheduling Coefficient. This is defined as the ratio of runtime to runtime plus wait time. Hence, a value of 1 (green) indicates that a job ran with no time waiting in the queue, a value of 0.5 (pale yellow) indicates a job queued for the same amount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

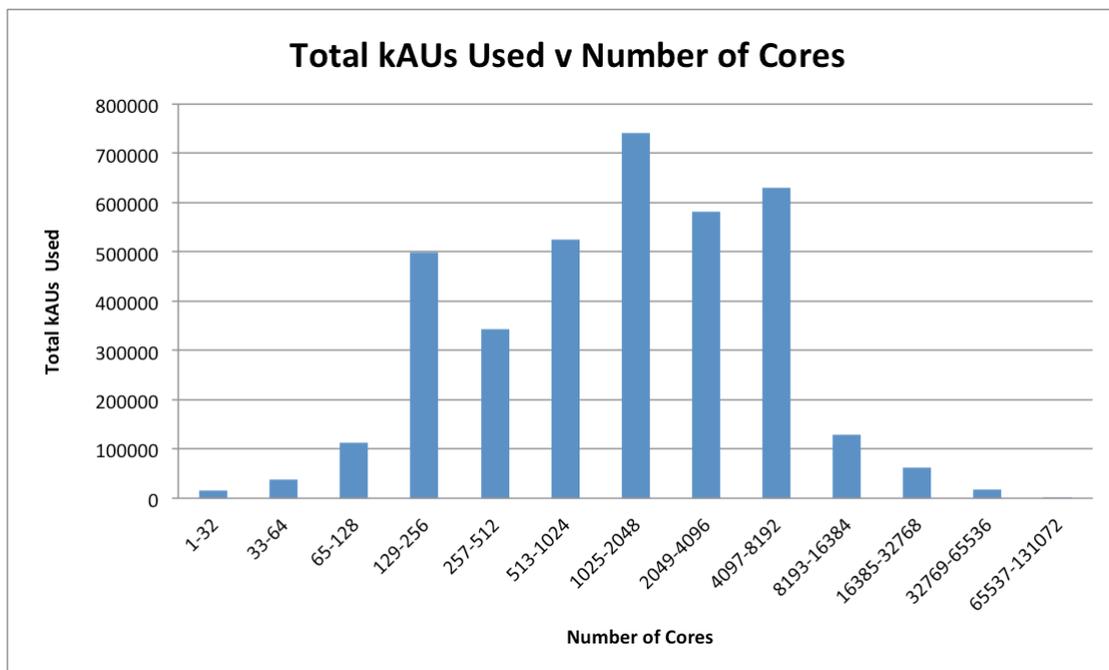
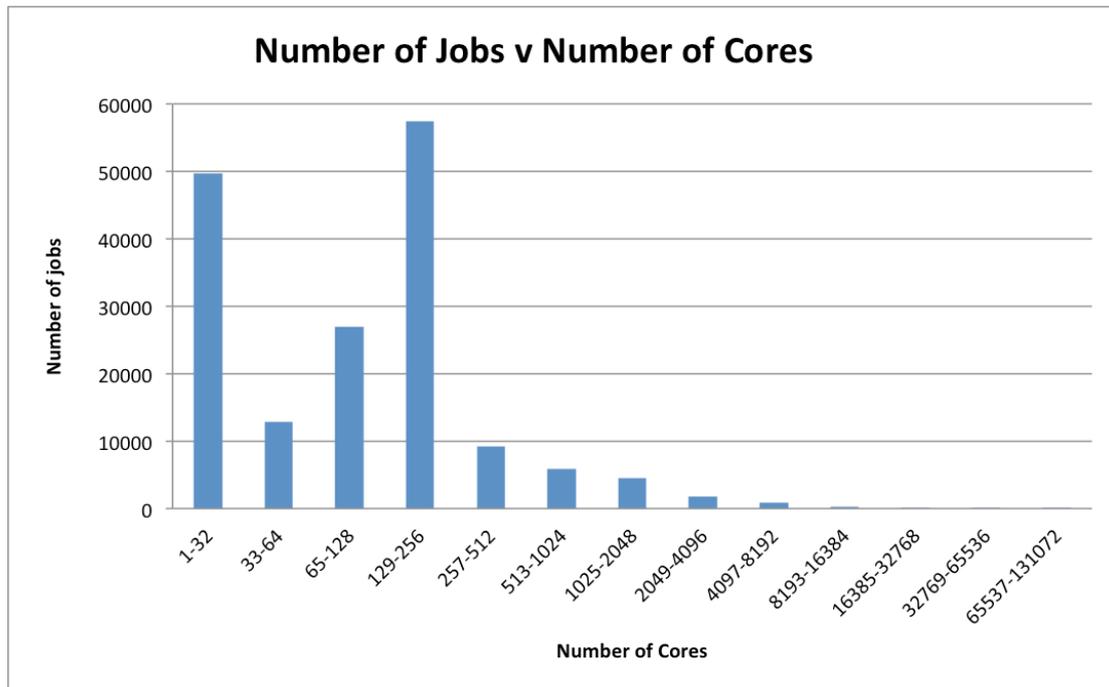
3.3 Additional Usage Graphs

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.

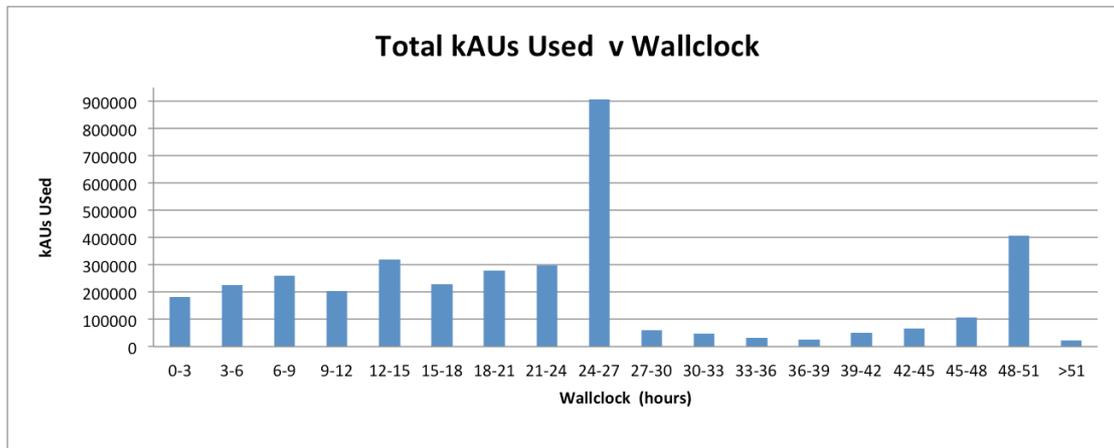
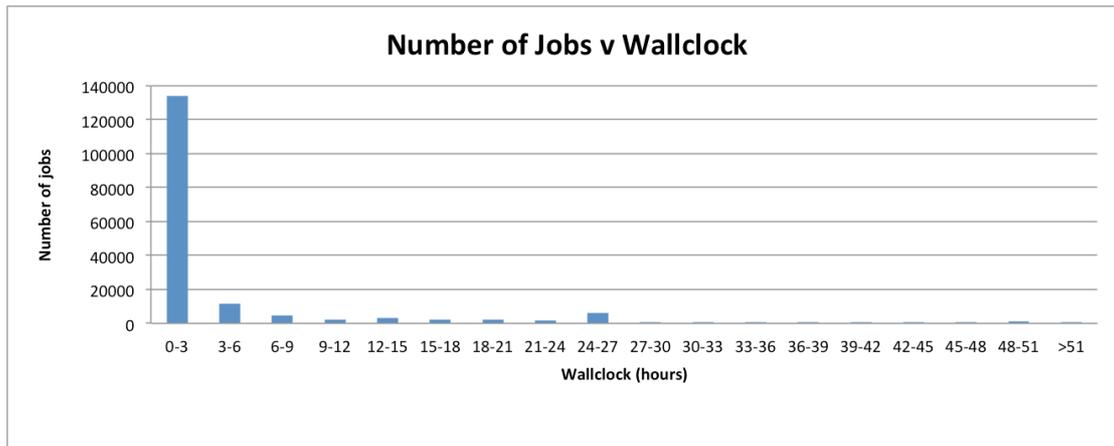


Analysis of Job Sizes



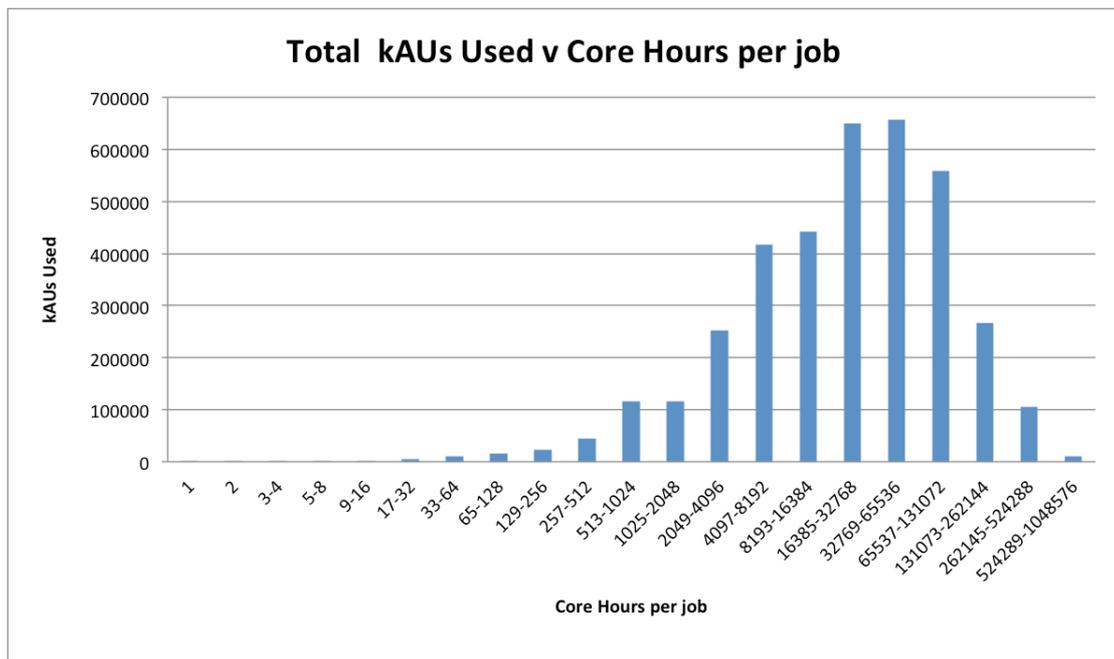
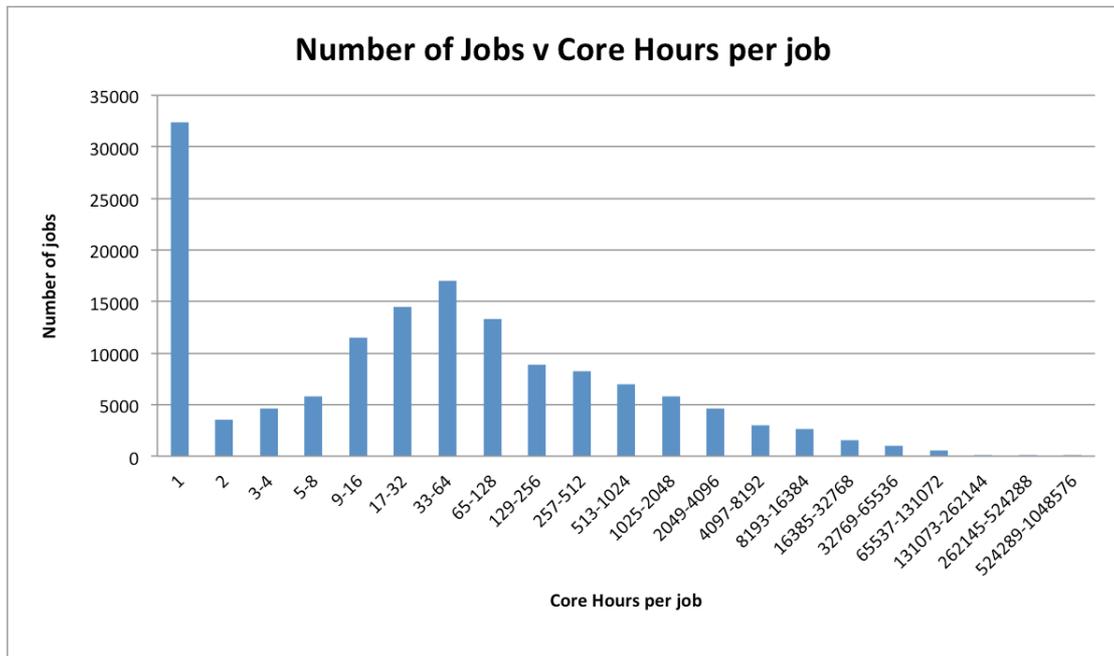
The first graph shows that, in terms of numbers, there are a significant number of jobs using no more than 256 cores. However, the second graph reveals that most of the kAUs were spent on jobs between 257 cores and 8192 cores. The number of kAUs used is closely related to money and shows better how the 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



Appendix – Infrastructure report

There is nothing to report regarding infrastructure work within this quarter.