ARCHER SP Service Quarterly Report

Quarter 4 2017

Document Information and Version History

Version:	1.1
Status	Release
Author(s):	Alan Simpson, Anne Whiting, Andy Turner, Linda Dewar, Stephen Booth, Jo Beech-Brandt
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	19/12/17	Initial Draft	Anne Whiting
0.2	03/01/18	Added quarterly metrics	Anne Whiting
0.3	03/01/18	Added cumulative usage graphs	Anne Whiting
0.4	04/01/18	Added further graphs and phone data	Jo Beech-Brandt
0.5	09/01/18	Added OSG highlights and maintenance details	Anne Whiting
0.6	17/01/18	Reviewed	Alan Simpson
0.7	18/01/18	Post review changes made	Anne Whiting
1.0	19/01/18	Version for release to EPSRC	Alan Simpson
1.1	13/03/18	Added comment to utilization graph re consortia renewal period	Anne Whiting

1. The Service

1.1 Service Highlights

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

October 2017, November 2017 and December 2017.

- Utilisation on the system during 17Q4 was 86%, as compared to 85% in 17Q3. The continued high utilisation of the service supports the need for ongoing investment in HPC.

- A new dashboard is being trialed by the Systems team providing enhanced statistics about ARCHER. Data is collected from servers, switches, storage systems, rack power bars, from the cooling systems and, potentially, from the electrical systems as well. The dashboard will provide additional information to assist with service optimization and troubleshooting, and potentially some of it could be of interest to the user community.

- A reminder was sent out to the user community that jobs which fall outside the limits for normal jobs on ARCHER because, for example, they need to run over more than 48 hours, may be accommodated via a reservation. Information about how to request a reservation and the charges, lead time requirements and restrictions are available on the ARCHER website at http://archer-www.epcc.ed.ac.uk/documentation/user-guide/batch.php#sec-5.14

- Work has been carried out in collaboration with the CSE team to incorporate a new feed of data into the SAFE from the ARCHER system capturing (amongst other metrics) peak memory usage and total energy usage for all user jobs. The SAFE reporting provided allows the user community to calculate the energy costs and memory usage of HPC applications, memory use by projects/users, and memory use by difference science areas.

- The scheduler analysis, originally undertaken throughout summer 2016, has been re-run and this indicates that the scheduler configuration remains well suited to the ARCHER workload.

- With the support of EPSRC, EPCC is working with Zenotech in a trial to see if the Zenotech Elastic Private interactive Cloud (EPIC) portal could be used to provide simplified access to Archer for users primarily from SME companies. Representatives from Zenotech now have access to the TDS for testing of the EPIC job submission front end.

- Improved external connectivity has been implemented in the ACF, which is now is directly connected to JANET and therefore no longer reliant on any other networks to carry traffic to/from the JANET network. This change will enable future upgrades of link speed and thus faster connectivity to the services housed at the ACF.

- Project and user accounting for the Peta-5 facility (led by University of Cambridge) has been added to the Tier 2 SAFE, joining HPC Midlands Plus (led by Loughborough University) and Cirrus (led by University of Edinburgh).

- A GAP analysis has been completed to identify the areas needing work in preparation for the ISO 27001 information security audit. Following ISO 27001 will allow us to demonstrate that best practice is embedded whilst handling user data.

1.2 Forward Look

- The new version of PBS, 13.408, is due to be tested to ensure it does not adversely affect the service before upgrading.

- Investigations are planned to look at the options for increasing the ACF to JANET bandwidth.

- The 2017 User Survey will be launched in February 2018 and details will be sent out to the user community soon to obtain user feedback on all aspects of the service and any suggestions for improvement.

- Work is underway to prepare for the next ISO 9001:2015 external audit due in February 2018. The audit is to ensure the focus on service delivery to our users and continual improvement has been maintained since the last external audit.

- Preparation for ISO27001 information security certification has just started with the aim of certification during 2018 for EPCC managed services hosted at the ACF including ARCHER, Cirrus and the RDF. This will demonstrate that best practice is followed when handling user data.

- Testing is being carried out 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. Subject to successful testing the repurposed node should be available for use soon.

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	Oct 17		Nov 17		Dec 17		17Q4
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	98.9%	-2	97.6%	-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.

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 17Q4.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Oct 17	Nov 17	Dec 17	17Q4
Phone Calls Received	28 (2)	20 (5)	16 (4)	**64 (11)**
Answered in 2 Minutes	28 (2)	20 (5)	16 (4)	**64 (11)**
Service Level	**100.0%**	**100.0%**	**100.0%**	**100.0%**

The volume of telephone calls remained low in 17Q4. Of the total of 64 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)

	Oct 17	Nov 17	Dec 17	17Q4
Self-Service Admin	334	549	330	**1213**
Admin	142	165	107	**414**
Technical	11	28	22	**61**
Total Queries	487	742	459	**1688**
Total Closed in 2 Days	483	734	448	**1664**
Service Level	**99.0%**	**98.9%**	**97.6%**	**98.6%**

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 17Q4:

	Oct 17	Nov 17	Dec 17	17Q4
Change Requests	0	2	2	**4**

2.2.3 User Registration (UR)

	Oct 17	Nov 17	Dec 17	17Q4
No of Requests	63	52	54	**169**
Closed in One Working Day	63	52	54	**169**
Average Closure Time (Hrs)	0.43	0.94	0.38	**0.56**
Average Closure Time (Working Days)	0.05	0.10	0.04	**0.06**
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:

	Oct 17	Nov 17	Dec 17	17Q4
CRAY	8	4	3	**15**
ARCHER_CSE	159	60	34	**253**
ARCHER_SP	787	1026	666	**2479**
Total Queries Assigned	954	1090	703	**2747**
Total Assigned in 1 Hour	954	1090	703	**2747**
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 only 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 EPSRC.

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
22/11/17	09:00	14:41	5 hrs 41 mins	Full outage	EPSRC Approved 0900 – 1700	Field notice and patch set update

2.3.3 Quality Tokens

No quality tokens have been 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 85%. The plot below shows a slight decrease in utilisation in December but an overall increase over the lifetime of the service. It was suggested by several different consortium members that the decrease in utilisation could be related to consortium members being involved in the 3 year consortium renewal process.

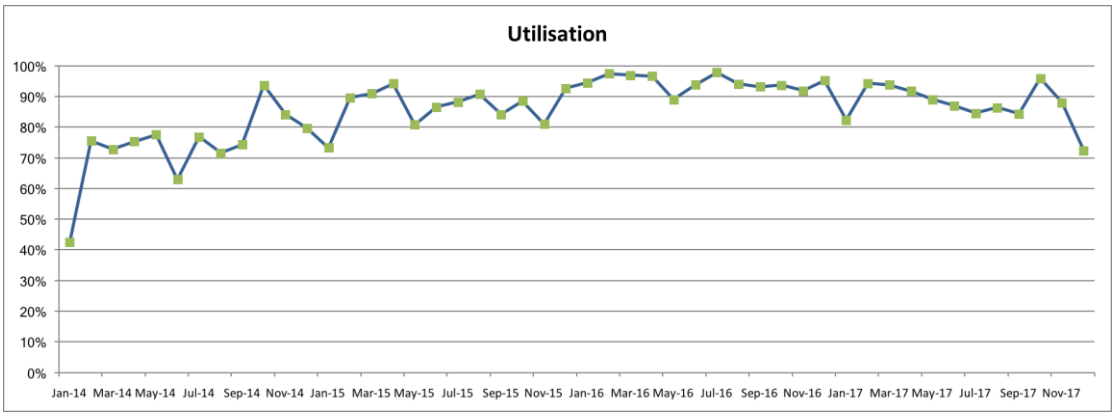

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that NERC exceeded their target of 23% with their usage at 26.75%, whereas EPSRC did not meet their 77% target as their usage was 59.47%.

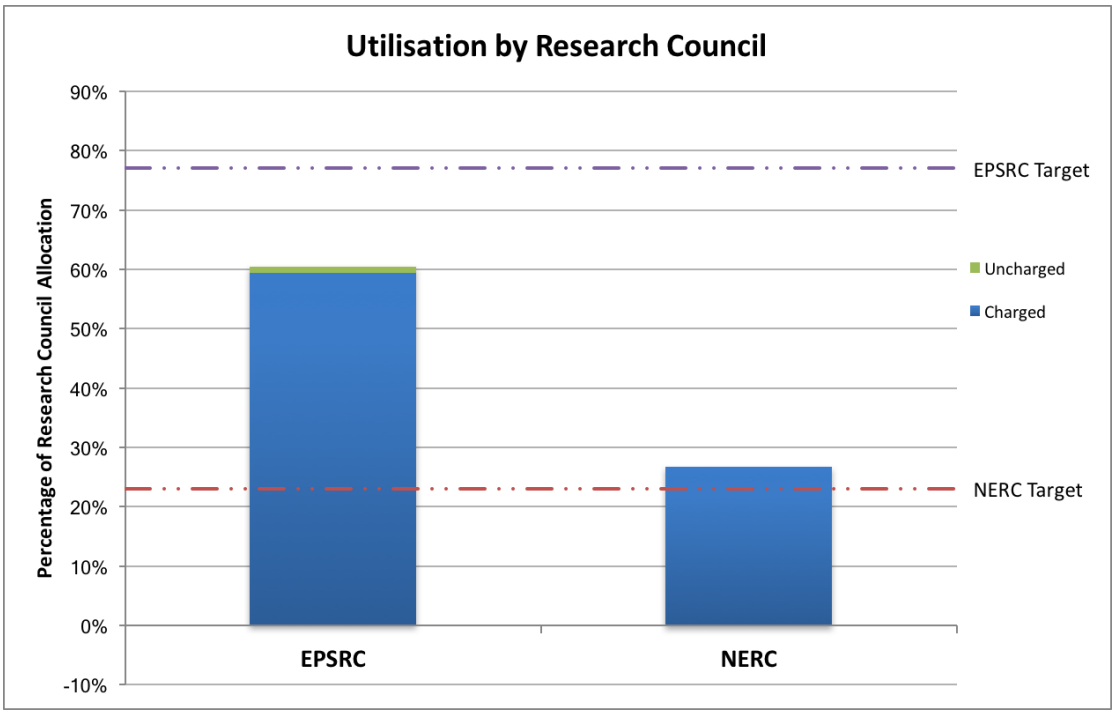

The cumulative allocation utilisation for the quarter by the Research Councils is shown below:

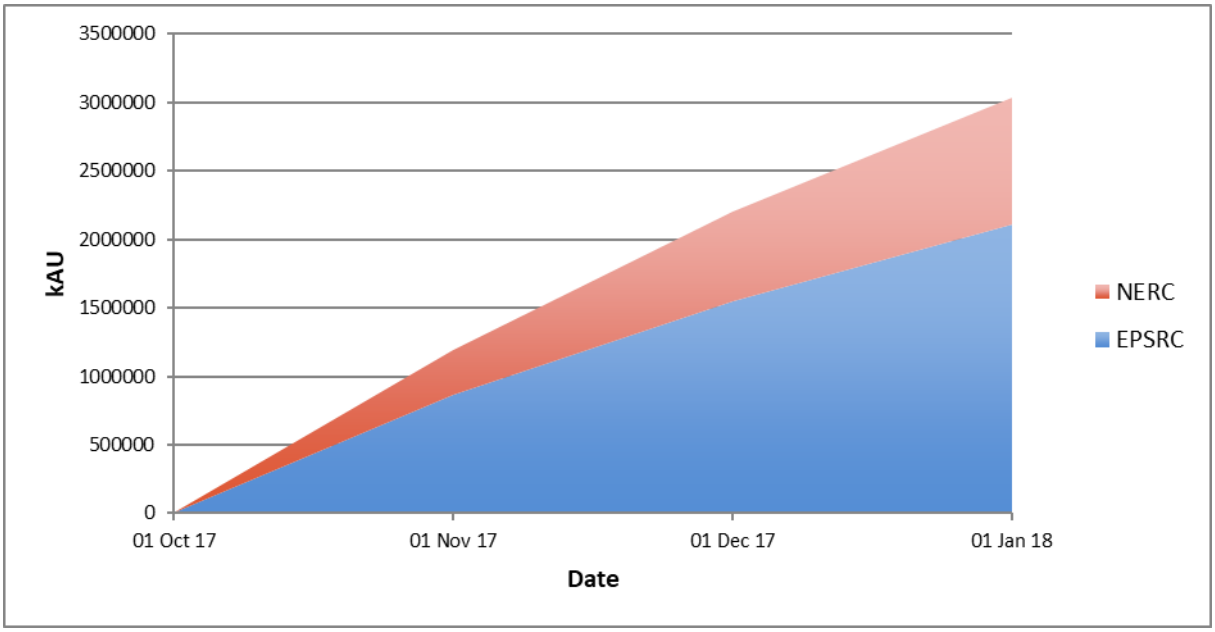

The cumulative allocation utilisation for the quarter by EPSRC broken down by different project types (see below) shows that the majority of usage comes from the scientific Consortia (as expected) with significant usage from research grants, ARCHER Leadership projects and ARCHER RAP projects. The times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

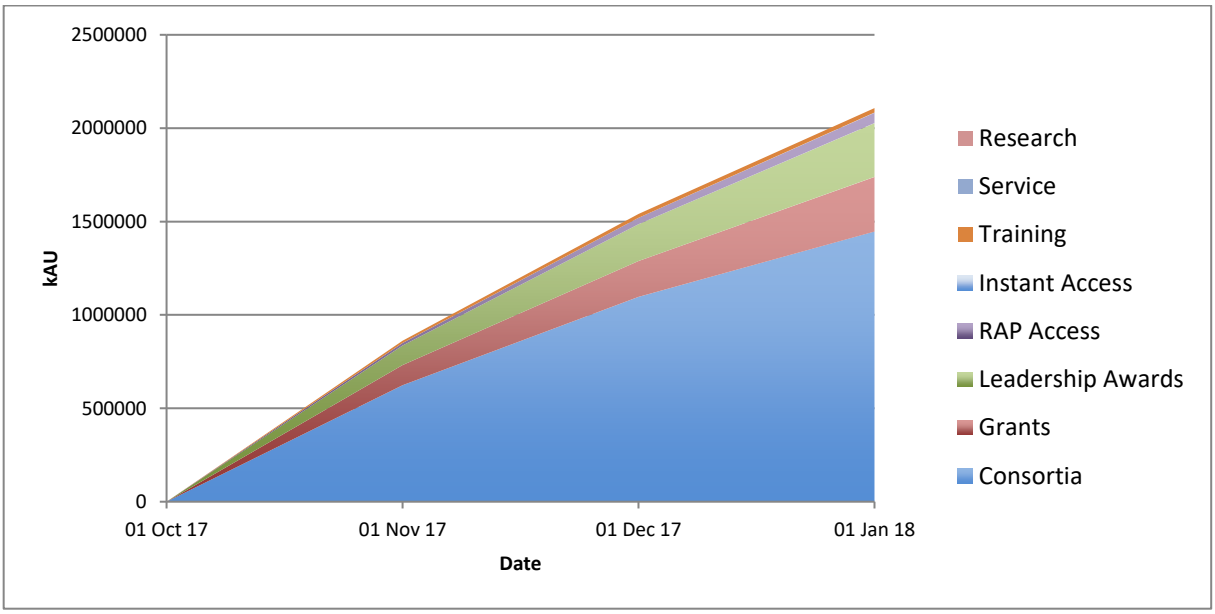

3.2 Scheduling Coefficient Matrix

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.

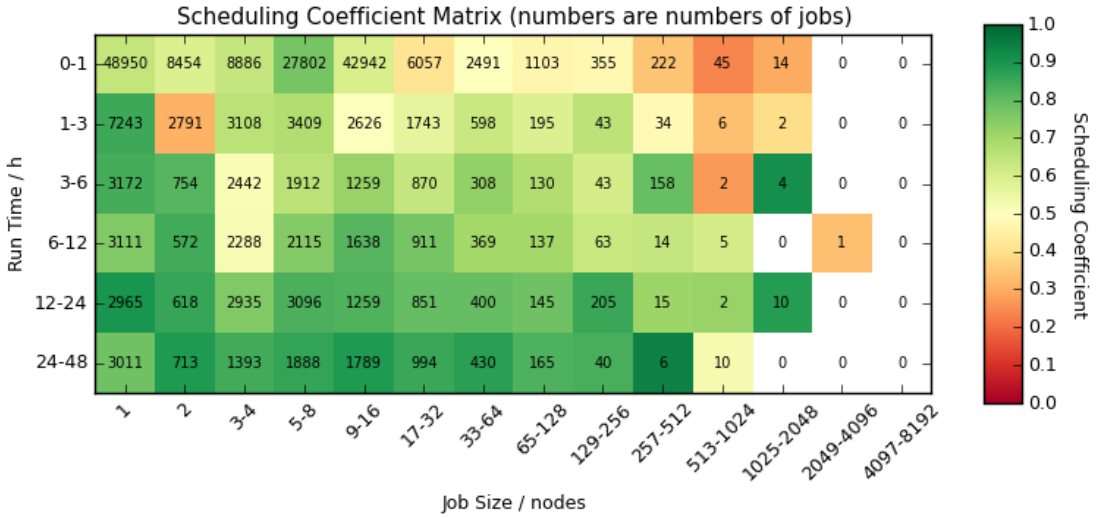

3.3 Additional Usage Graphs

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.

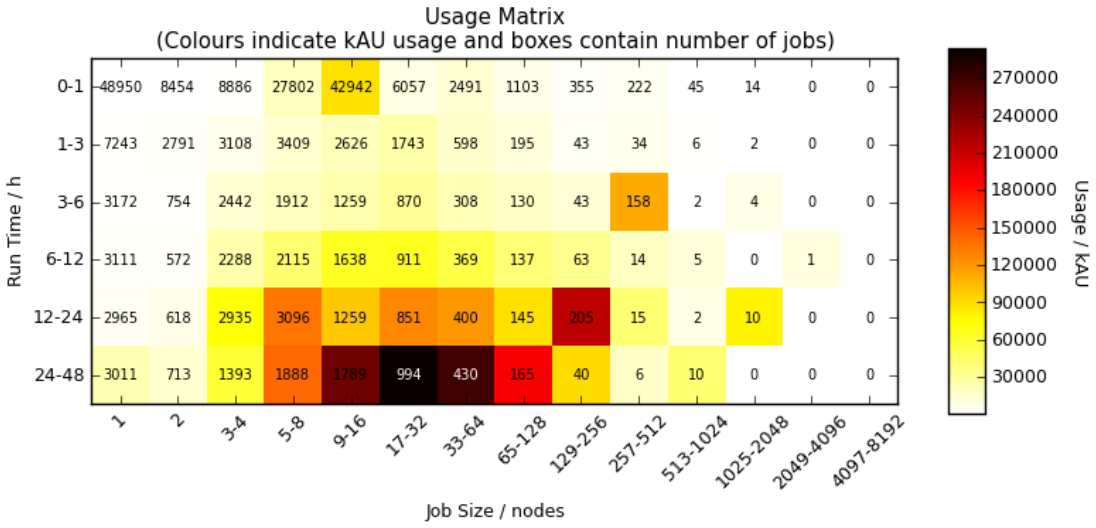

Analysis of Job Sizes

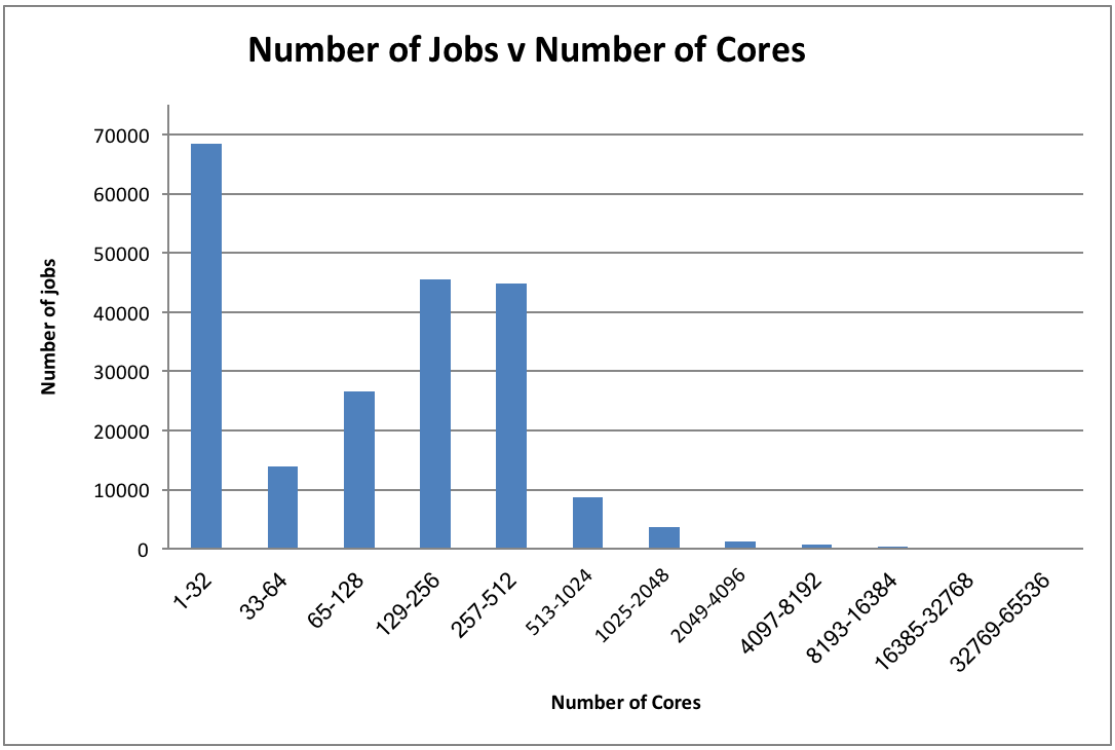

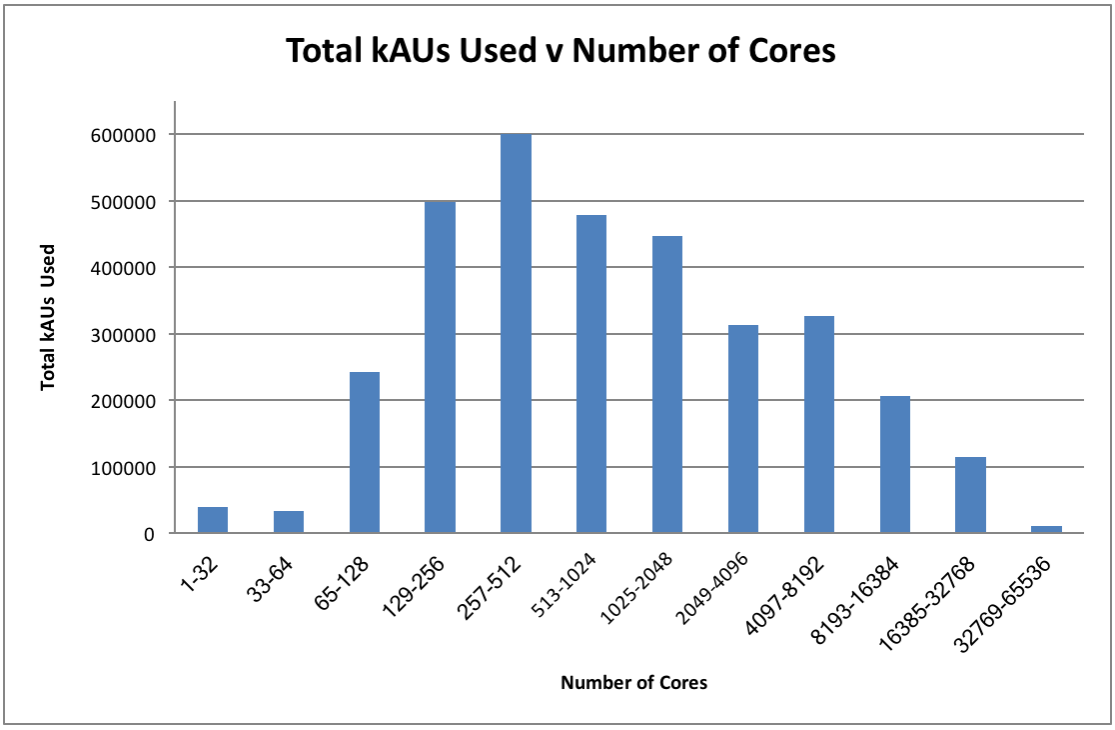

The first graph shows that, in terms of numbers, there are a significant number of jobs using no more than 1024 cores. However, the second graph reveals that most of the kAUs were spent on jobs between 129 cores and 8192 cores. The number of kAUs used is closely related to money and shows better how the investment in the system is utilised.

Analysis of Jobs Length

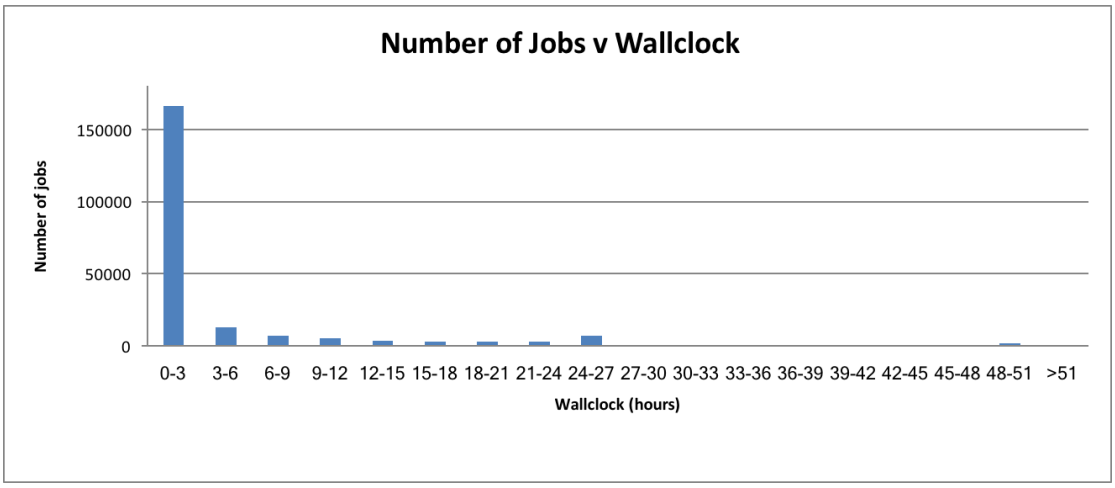

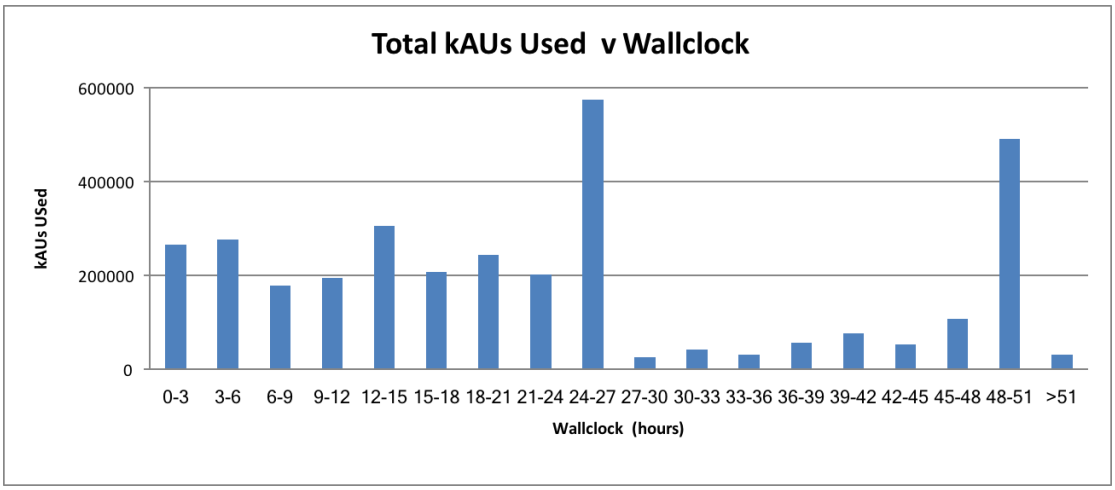

From the first graph, it would appear that the system is dominated by short jobs. However, the second graph shows that actual usage of the system is more spread and dominated by jobs of up to 27 hours with a second peak for jobs at 48-51 hours.

Core Hours per Job Analysis

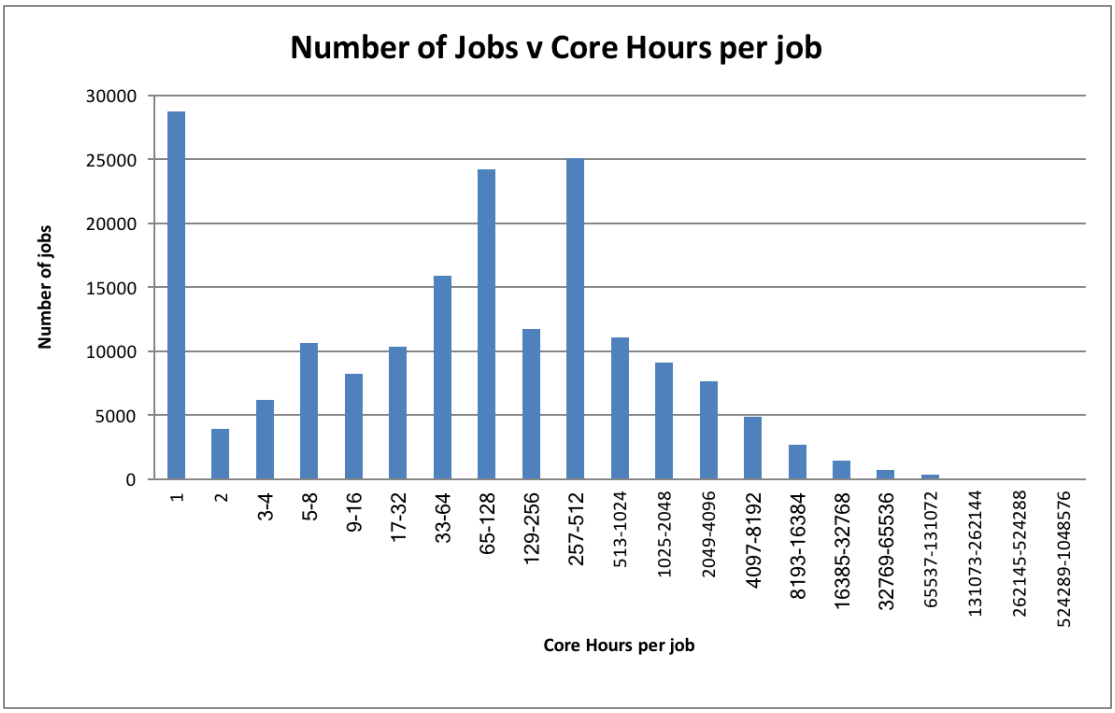

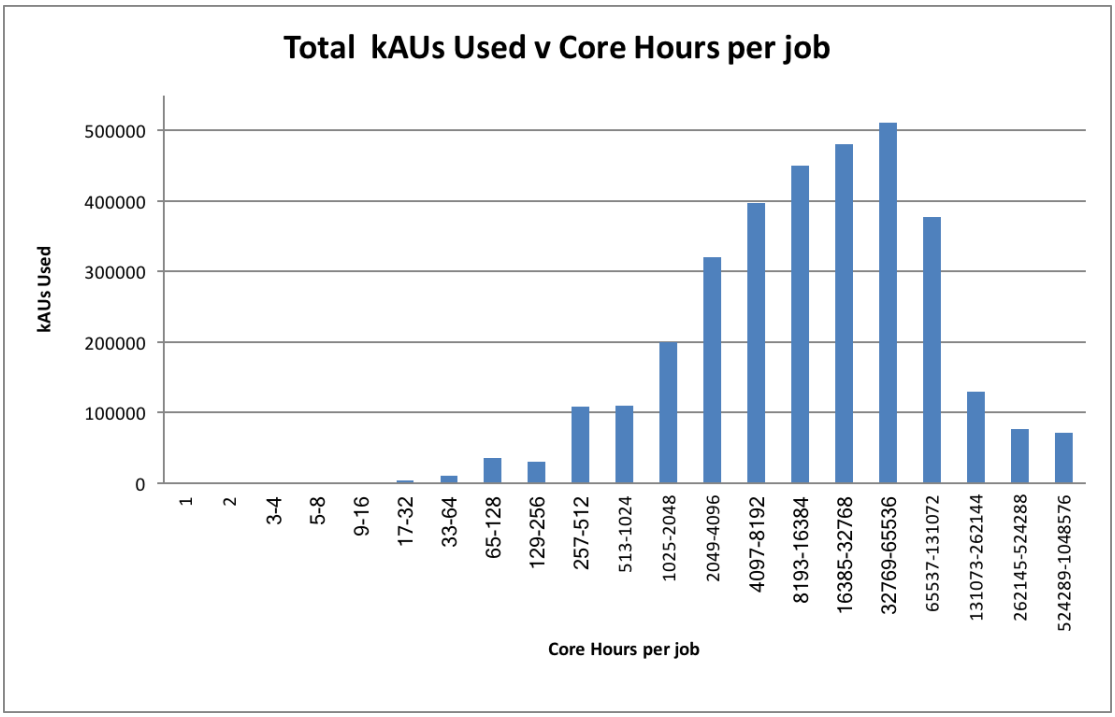

The above graphs show that, while there are quite a few jobs that use only a small number of core hours per job, most of the resource is consumed by jobs that use tens of thousands of core hours per job.