ARCHER SP Service Quarterly Report

Quarter 4 2018

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	20/12/18	Initial Draft	Anne Whiting
0.2	03/01/19	Added December and total q4 data	Anne Whiting
0.3	03/01/19	Added graphs and phone data	Jo Beech-Brandt
0.4	10/01/19	Updated Highlights, Forward Look and Maintenance sections.	Linda Dewar
0.5	10/01/19	Reviewed	Alan Simpson
0.6	11/01/19	Updates post review	Anne Whiting
1.0	14/01/18	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:

October 2018, November 2018 and December 2018.

- Utilisation over the quarter was 86%, up from 84% the previous quarter. The remedial actions taken to increase utilisation would seem to be having a positive impact, particularly considering the typical dip over the festive period.

- Weekend queue was made live and is now available between 12:00 on Saturday to 12:00 on Monday every weekend. Jobs submitted to the weekend queue are receiving a 50% discount in charging. After analysis of jobs submitted to the weekend queue since its launch, 1159 jobs were submitted by 21 projects or 49 users, using 101, 482 kAUs and thus saving the projects concerned 101, 482 kAUs to spend on further science.

- At the request of a user, the maximum job length on the KNL was increased from 8 hours to 24 hours, enabling users to run longer jobs without exceeding the permitted job length.

- Reviews of user requirements for module versions on the system have been carried out with the users concerned to ensure that the planned module deletions minimise any negative impact on users.

- The SAFE is now importing statistics of operations on the work file-systems, with data being stored for each job run on the system. This allows reports to be written to analyse IO-use on the system and the IO requirements of different application codes. A number of enhancements have been made to the report generator to help analyse this data.

- We have made a number of changes to the SAFE to reduce the impact on other SAFE users when large and complicated reports are being processed. These include various improvements to reduce the resources needed by report generation as well as checks to detect reports that are taking too long to run.

- The configuration of the power supplying the ARCHER Management and Storage racks in the computer room housing ARCHER has been modified to improve the resiliency of this equipment in the event of a power distribution unit failure.

- Training has been provided for several new members of the OSG team based at Daresbury, thus adding to the overall level of expertise and experience with the team.

- EPSRC and EPCC reached agreement to extend the ARCHER SP service to end on 18 November 2019 to allow additional time for the procurement process for ARCHER 2 to be completed. Work started to prepare the ARCHER SP service to be ready to transition to an ARCHER 2 SP service whoever the selected SP provider is.

1.2 Forward Look

- Plans are being finalised to delete old module versions on the system to allow for the final version of the Cray Programming Environment for CLE5.2UP04 to be installed.

- 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.

- The 2018 User Survey will be launched in February 2019 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 2019. The audit is to ensure the focus on service delivery to our users and continual improvement has been maintained since the last external audit. Improvements made include adopting the more rigorous approach taken to risk management used within our ISO 27001:2013 Information Security Management System and applying this to our ISO 9001:2015 Quality Management 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	Oct 18		Nov 18		Dec 18		18Q4
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 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)

	Oct 18	Nov 18	Dec 18	18Q4
Phone Calls Received	30 (5)	31 (3)	19 (4)	**80 (12)**
Answered in 2 Minutes	30	31	19	**80**
Service Level	**100.0%**	**100.0%**	**100.0%**	**100.0%**

The volume of telephone calls remained low in 18Q4. Of the total of 80 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)

	Oct 18	Nov 18	Dec 18	18Q4
Self-Service Admin	500	600	272	**1372**
Admin	130	117	97	**344**
Technical	17	29	10	**56**
Total Queries	647	746	379	**1772**
Total Closed in 2 Days	642	738	377	**1757**
Service Level	**99.2%**	**98.9%**	**99.5%**	**99.2%**

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:

	Oct 18	Nov 18	Dec 18	18Q4
Change Requests	1	0	0	**1**

2.2.3 User Registration (UR)

	Oct 18	Nov 18	Dec 18	18Q4
No of Requests	78	61	30	**169**
Closed in One Working Day	78	61	30	**169**
Average Closure Time (Hrs)	0.9	0.5	1.0	**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.

2.3.1 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 18	Nov 18	Dec 18	18Q4
CRAY	12	0	3	**15**
ARCHER_CSE	160	114	59	**333**
ARCHER_SP	960	986	545	**2491**
Total Queries Assigned	1132	1100	607	**2839**
Total Assigned in 1 Hour	1120	1100	607	**2826**
Service Level	**99%**	**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 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
31/10/18	09:00	16:39	7 hrs 39 mins	Full outage	Approved by EPSRC 09:00 – 17:00	Improve power resiliency to management and storage racks
12/12/18	08:00	15:42	7 hrs 42 mins	Full outage	Approved by EPSRC 08:00 – 20:00	Update patches on Archer Compute/eslogins/es pps

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%, up from 84% the previous quarter. The usual festive dip in utilisation was less in 2018 than 2017, for December 2017 it was 72% but in 2018 79%.. Utilisation for October was 89%, for November 89% and for December 79%. 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.

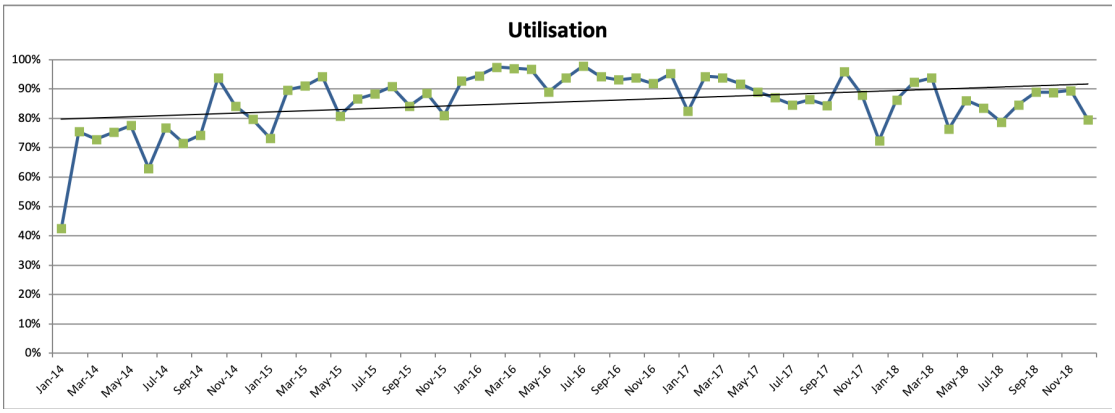

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that EPRSC did not meet their target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

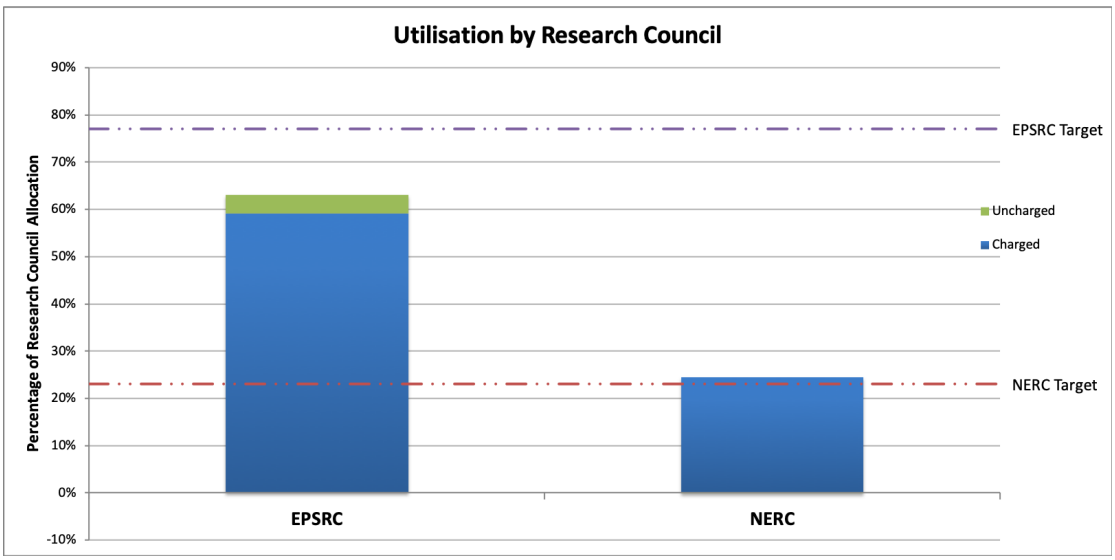

The cumulative allocation utilisation for the quarter by the Research Councils is shown below:

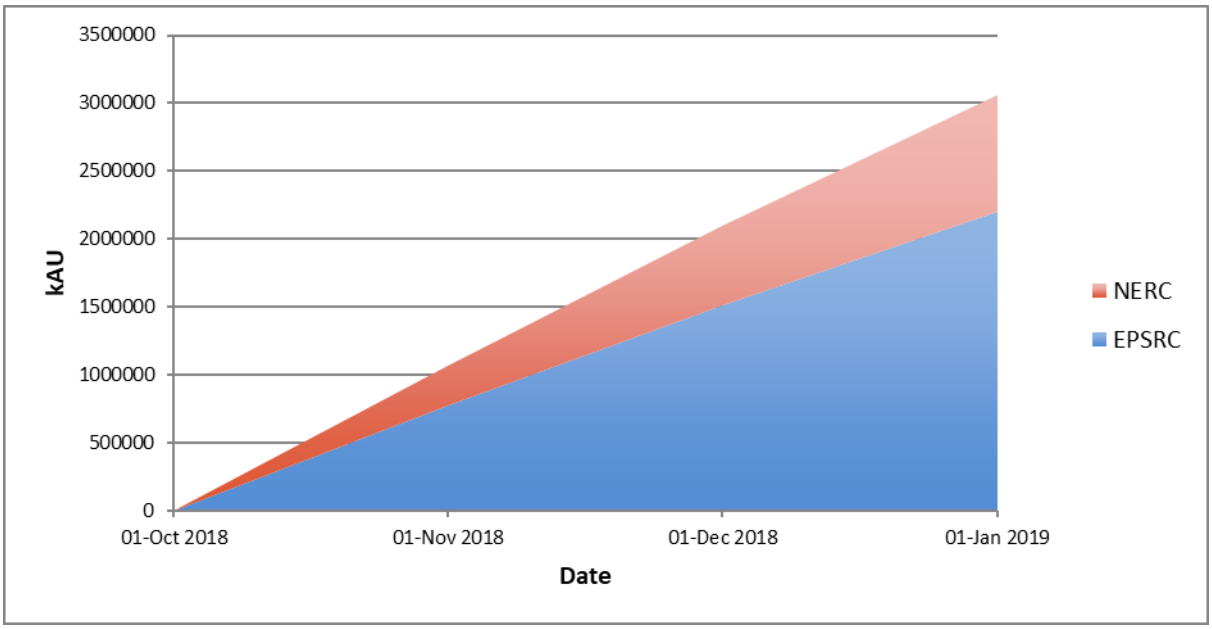

The cumulative allocation utilisation for the quarter by EPSRC broken down by different project types (see below) shows that the majority of usage comes from the scientific Consortia (as expected) with significant usage from research grants, ARCHER Leadership projects and ARCHER RAP projects. The times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

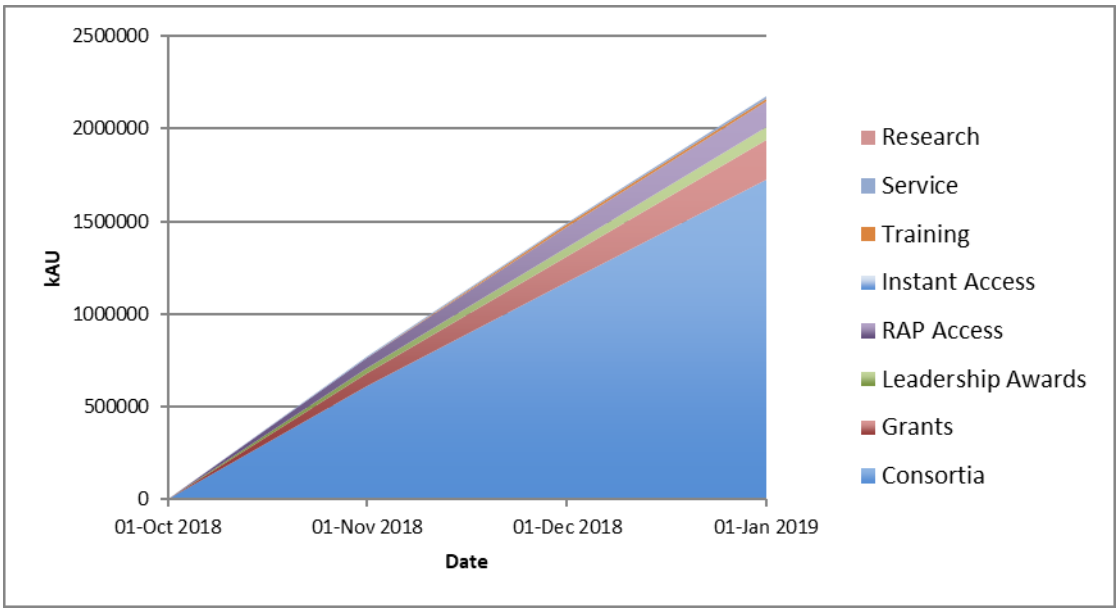

3.2 Scheduling Coefficient Matrix

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.

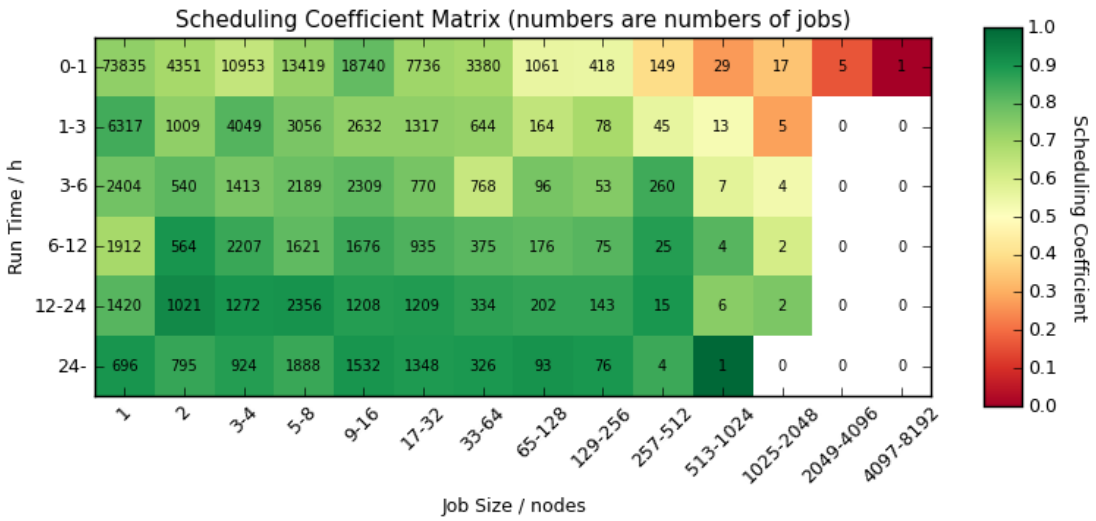

3.3 Additional Usage Graphs

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.

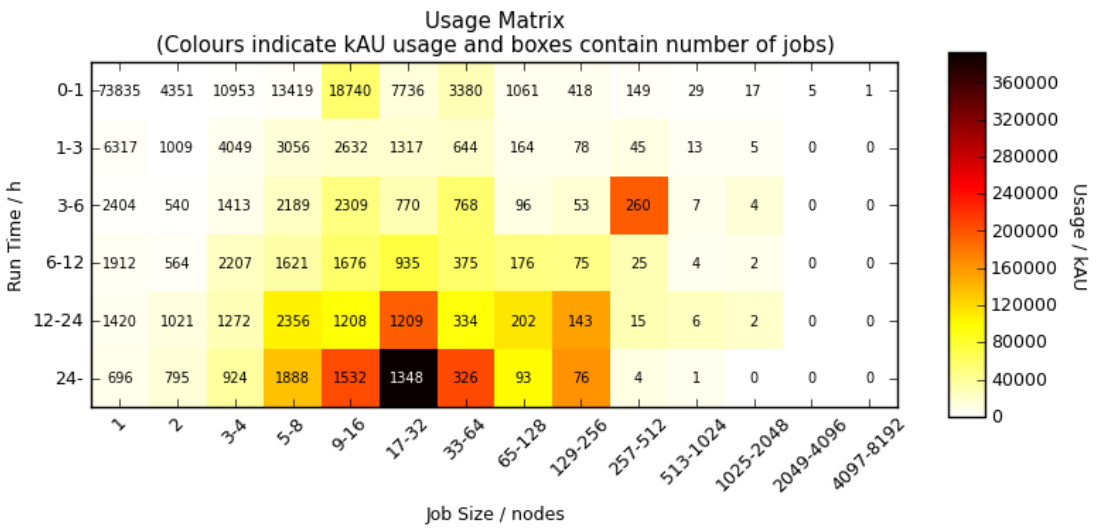

Analysis of Job Sizes

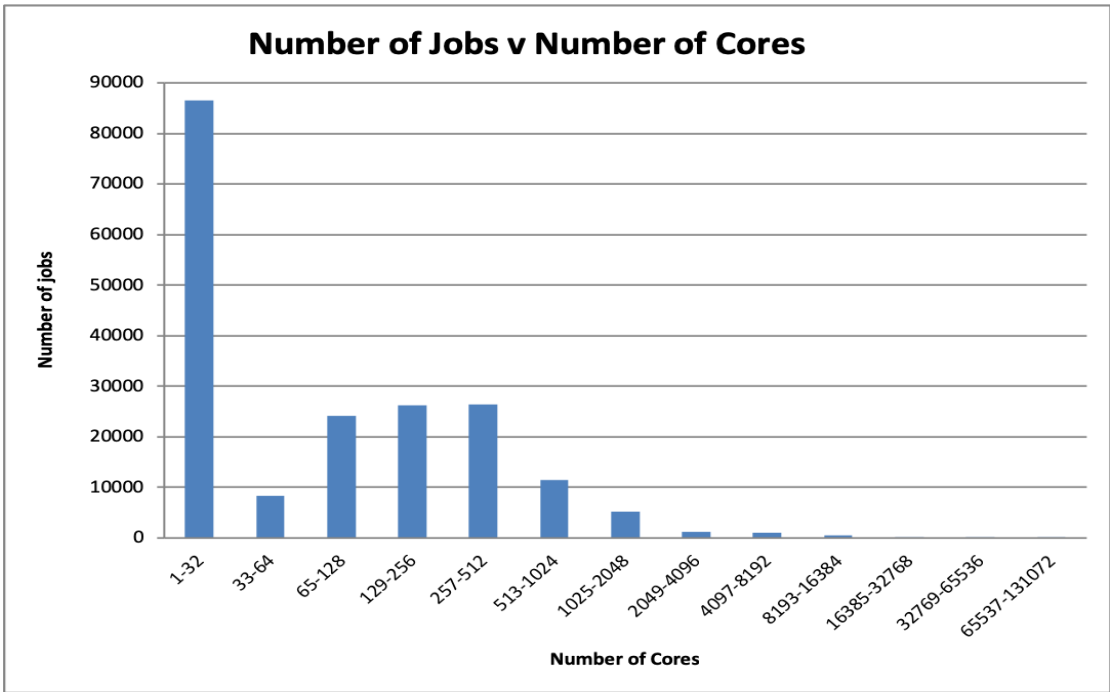

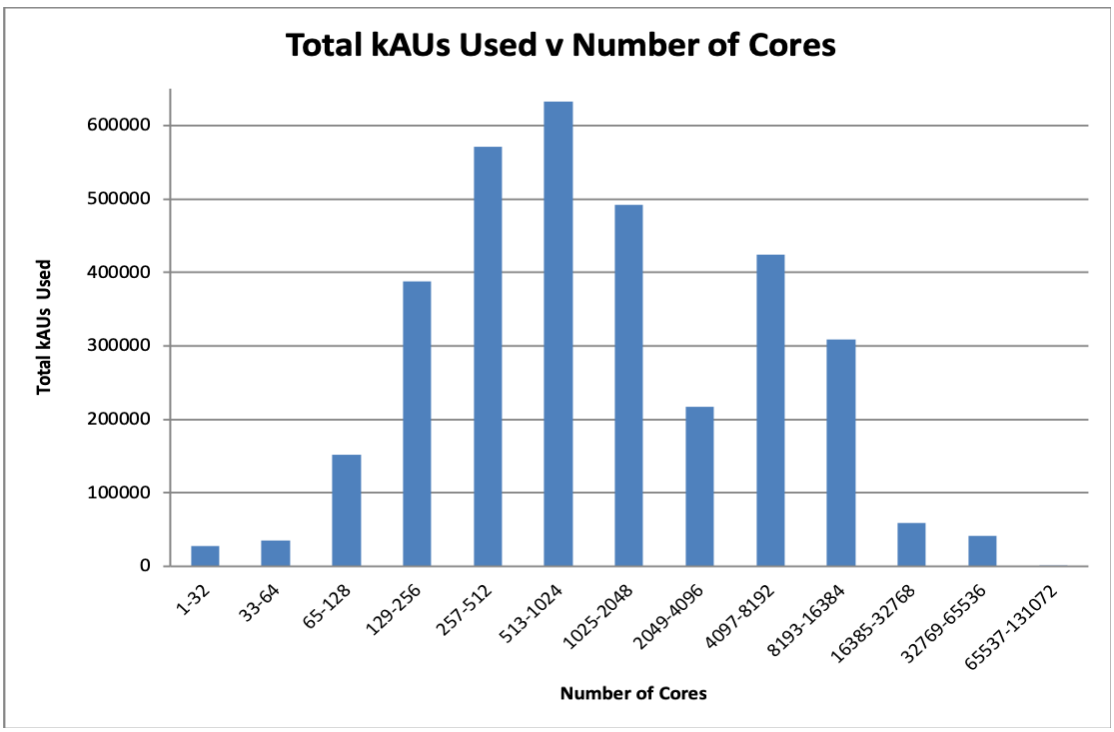

The first graph shows that, in terms of numbers, there are a significant number of jobs using no more than 1024 cores. However, the second graph reveals that most of the kAUs were spent on jobs between 65 cores and 16384 cores. The number of kAUs used is closely related to money and shows better how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

Analysis of Jobs Length

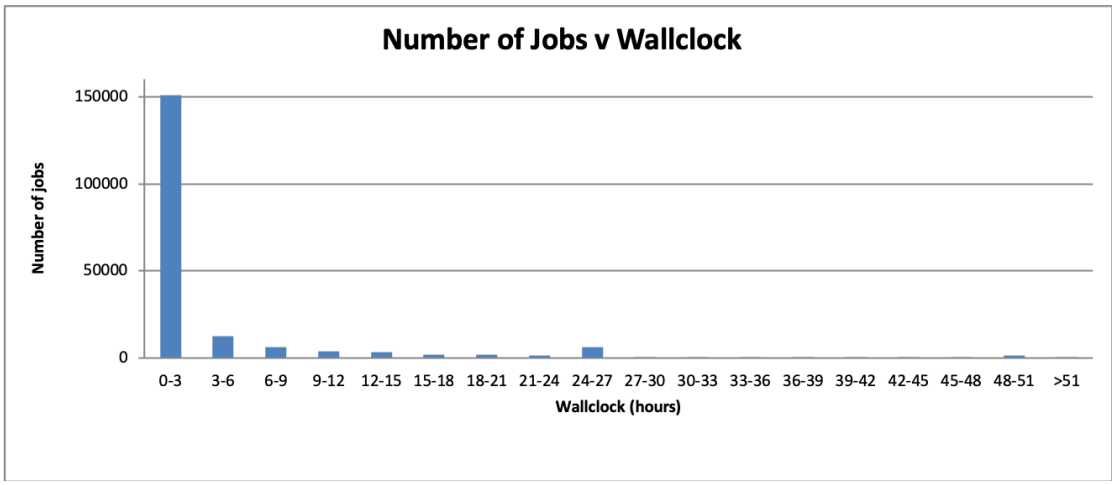

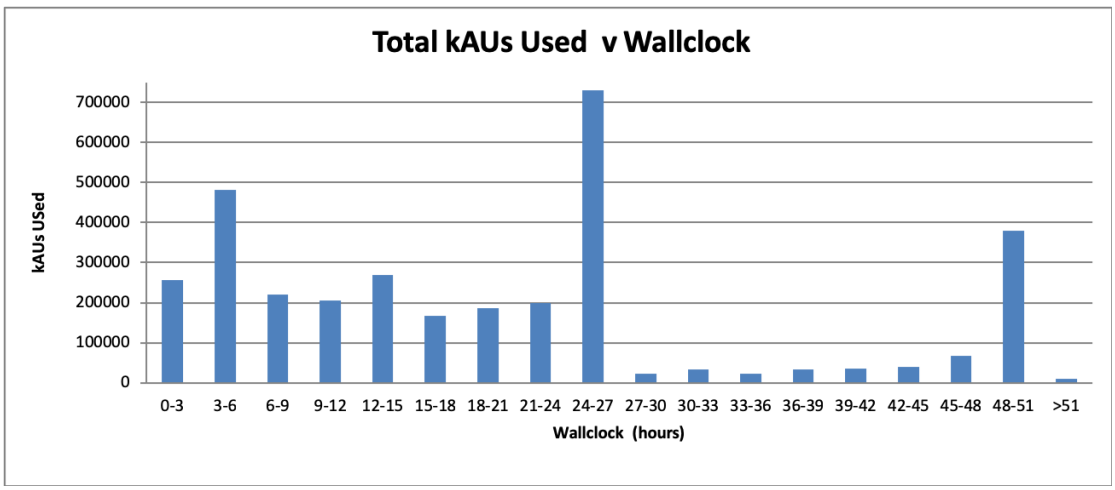

From the first graph, it would appear that the system is dominated by short jobs. However, the second graph shows that actual usage of the system is more spread and dominated by jobs of up to 27 hours with a second peak for jobs at 48-51 hours.

Core Hours per Job Analysis

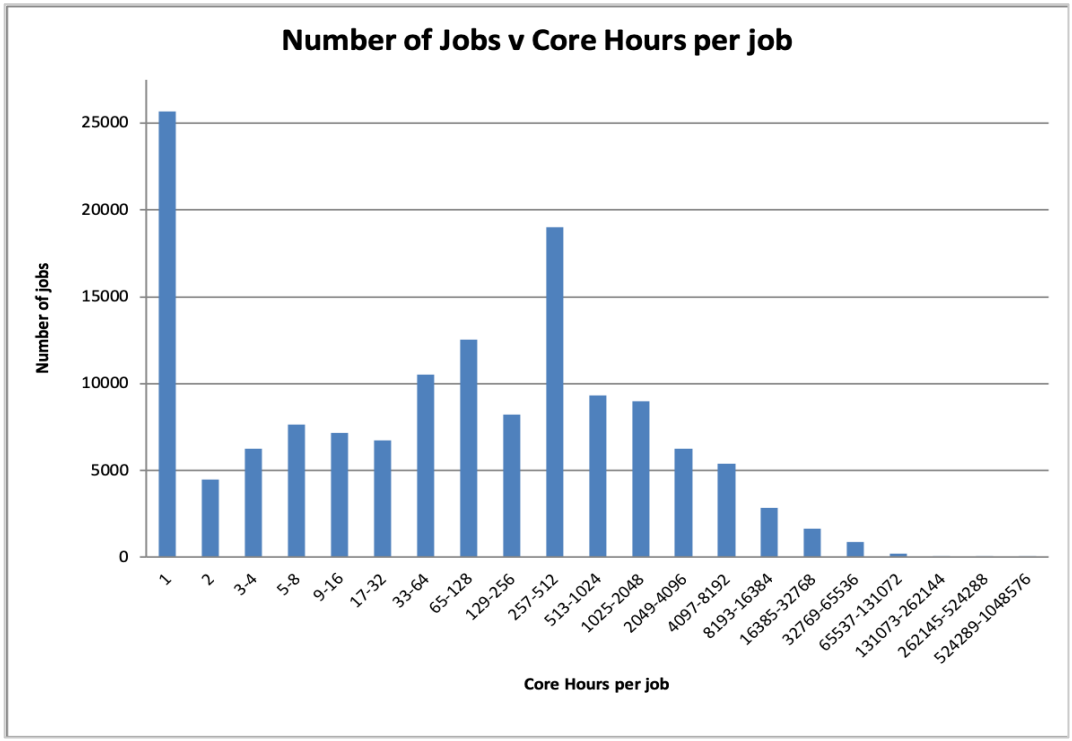

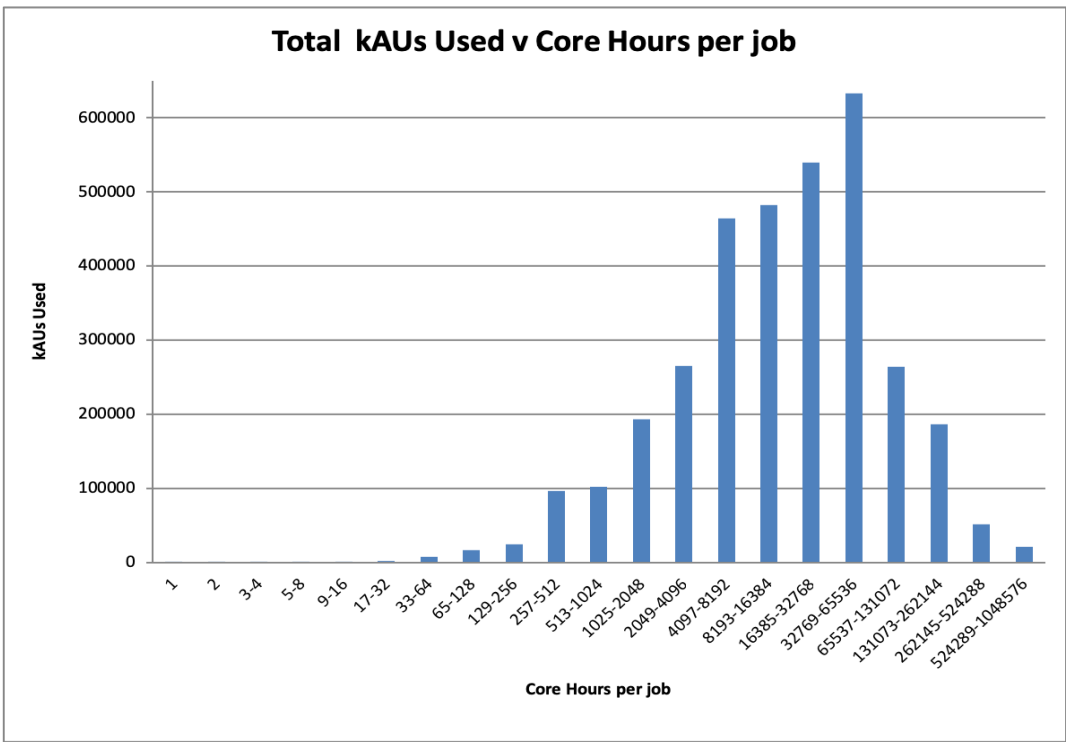

The above graphs show that, while there are quite a few jobs that use only a small number of core hours per job, most of the resource is consumed by jobs that use tens of thousands of core hours per job.