ARCHER SP Service Quarterly Report

Quarter 1 2018

Document Information and Version History

Version:	1.0
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	03/04/18	Initial Draft	Anne Whiting
0.2	03/04/18	Phone info and graphs	Jo Beech-Brandt
0.3	04/04/18	Added heatmaps	Clair Barrass
0.4	04/04/18	Reviewed	Alan Simpson
0.5	04/04/18	Updated post review	Linda Dewar, Anne Whiting
1.0	04/04/18	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:

January 2018, February 2018 and March 2018.

- Utilisation on the system during 18Q1 was 91%, as compared to 85% in 17Q4. The dip in utilisation experienced at the end of 2017 was thought to be related to consortium members being involved in the 3-year consortium renewal process.

- The results of the 2017 annual ARCHER User Survey have been analysed. 164 responses were received with the mean results shown below (scores 1 representing "Very Unsatisfied" and 5 representing "Very Satisfied"):

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)
Overall Satisfaction	4.4	4.3	4.3	4.4
Hardware	4.1	4.1	4.2	4.3
Software	4.0	4.0	4.2	4.1
Helpdesk	4.5	4.5	4.5	4.6
Documentation	4.1	4.1	4.2	4.2
Website	4.1	4.2	4.2	4.2
Training	4.1	4.1	4.2	4.1
Webinars	3.6	3.9	3.9	4.2
Online training	-	4.0	4.1	4.2

As can be seen users have provided very positive feedback for the service and the most significant improvement is in the score for webinars. The full report can be found at http://www.archer.ac.uk/about-archer/reports/.

- The ARCHER service staff worked through the Red Alert weather problems at the end of February, keeping the service running uninterrupted despite weather warnings, snow drifts, and buildings and transport being shut down. The ARCHER service declared a Major Incident and used the experience to test and improve our Major Incident and Disaster Recovery processes. We were very pleased how well the service ran, and very much appreciated the support and good wishes from the user community.

- The ISO 9001:2015 annual external audit took place in February 2018 and we are delighted to announce we passed with only 1 minor finding. The audit is to ensure the focus on service delivery to our users and continual improvement has been maintained since the last external audit.

- Two new key senior staff have joined our Systems team, bringing a wealth of experience with them. Paul Clark has joined us as Director of High Performance Computing Systems, leading both our Tier-1 and Tier-2 systems teams to provide an integrated cross-service team. Calum Muir has joined us as Data Centre Manager bringing expertise in the electrics, cooling and plant side of things so vital to keeping ARCHER and Cirrus running optimally.

- Work has successfully been completed 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. The repurposed node should be available for use soon.

- All patching and system configuration changes implemented this quarter have been carried out during at-risk sessions rather than taking ARCHER down for a full maintenance session. Whilst there will undoubtedly be instances where a full outage is required in the future, we are aiming to continue to reduce the number of planned outages to minimise the impact on users.

1.2 Forward Look

- The new version of PBS, 13.408, is being tested to ensure it does not adversely affect the service before upgrading. As well as providing new functionality, it should resolve issues that we have experienced where jobs that cannot run prevent other jobs being scheduled. The new version will be tested on the TDS before being put on the KNL and ARCHER systems.

- The next ARCHER Champions meeting will be held in Manchester at the Museum of Science and Industry on Wednesday 25th April 2018. It will be a full-day meeting that will follow the UK/US SSI Impact Event taking place on Tuesday 24th April. Topics will include the eCSE programme, training and MOOCs, Tier 2, and a discussion session on the future of Champions.

- ARCHER service policies are being reviewed to ensure they comply with the requirements of the General Data Protection Regulation (GDPR), with the deadline of 25 May 2018 when the new legislation comes into use. We are aiming to send out information on the updated policies by the end of April.

- Preparation for ISO27001 information security certification is well underway 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 data.

- Plans are being finalised to upgrade the KNL to CLE 6 UP04 which will enable the Meltdown patches to be applied to the KNL 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	Jan 18		Feb 18		Mar 18		18Q1
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, 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 18Q1.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Jan 18	Feb 18	Mar 18	18Q1
Phone Calls Received	32 (1)	25 (5)	22 (4)	**79 (10)**
Answered in 2 Minutes	32	25	23	**79**
Service Level	**100.0%**	**100.0%**	**100.0%**	**100.0%**

The volume of telephone calls remained low in 18Q1. Of the total of 79 calls received above, only 10 were actual ARCHER user calls that either resulted in queries or answered user questions directly.

2.2.2 Query Closure (QC)

	Jan 18	Feb 18	Mar 18	18Q1
Self-Service Admin	684	282	303	**1269**
Admin	107	107	119	**333**
Technical	14	24	25	**63**
Total Queries	805	413	447	**1665**
Total Closed in 2 Days	799	408	442	**1649**
Service Level	**99.3%**	**98.8%**	**98.9%**	**99.0%**

The above table shows the queries closed by SP during the period. It is worth noting that there was a significant drop in the number of self-service admin queries in February and March, mirroring a similar drop in the number of new users registered. Each user registration and new project creation creates multiple self-service admin queries.

In addition to the Admin and Technical queries, the following Change Requests were resolved in 18Q1:

	Jan 18	Feb 18	Mar 18	18Q1
Change Requests	2	1	1	**4**

2.2.3 User Registration (UR)

	Jan 18	Feb 18	Mar 18	18Q1
No of Requests	121	60	69	**250**
Closed in One Working Day	121	60	69	**250**
Average Closure Time (Hrs)	0.58	0.46	0.51	**0.53**
Average Closure Time (Working Days)	0.06	0.05	0.05	**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:

	Jan 18	Feb 18	Mar 18	18Q1
CRAY	2	4	8	**14**
ARCHER_CSE	166	183	98	**447**
ARCHER_SP	1280	718	764	**2762**
Total Queries Assigned	1448	905	870	**3223**
Total Assigned in 1 Hour	1448	905	870	**3223**
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 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 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. 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.

This quarter the changes that have been required have all been low risk, and so have been made during at-risk sessions. We have therefore not needed to take any full outages.

2.3.3 Quality Tokens

One quality token was received during the quarter; this gave a 5 star positive rating without any comment.

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 91%, up from 85% 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 at maximum capacity as shown by the generally steady utilisation value:. It was suggested by several different consortium members that the dip in utilisation during December 2017 could be related to consortium members being involved in the 3-year consortium renewal process.

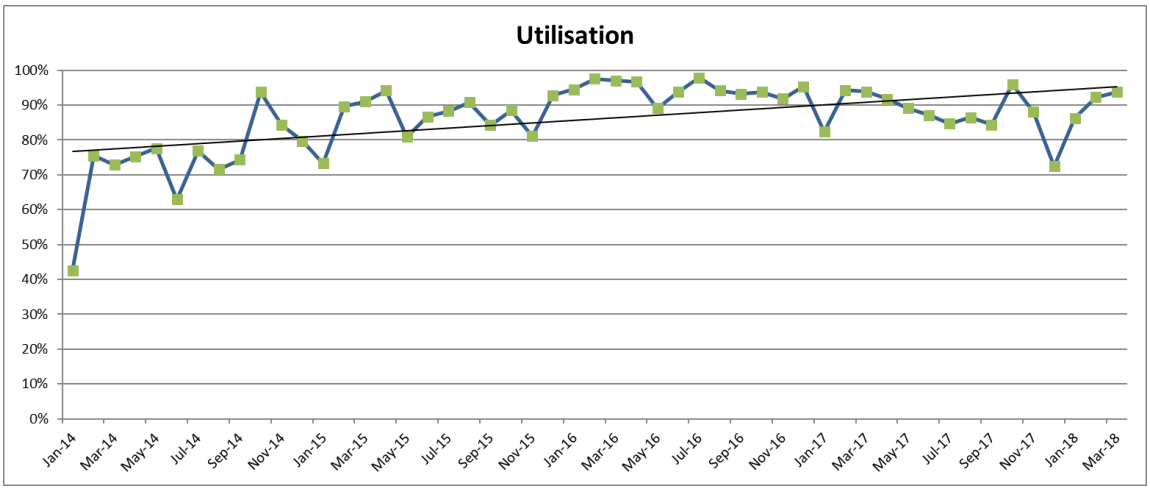

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that EPRSC did not meet their target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

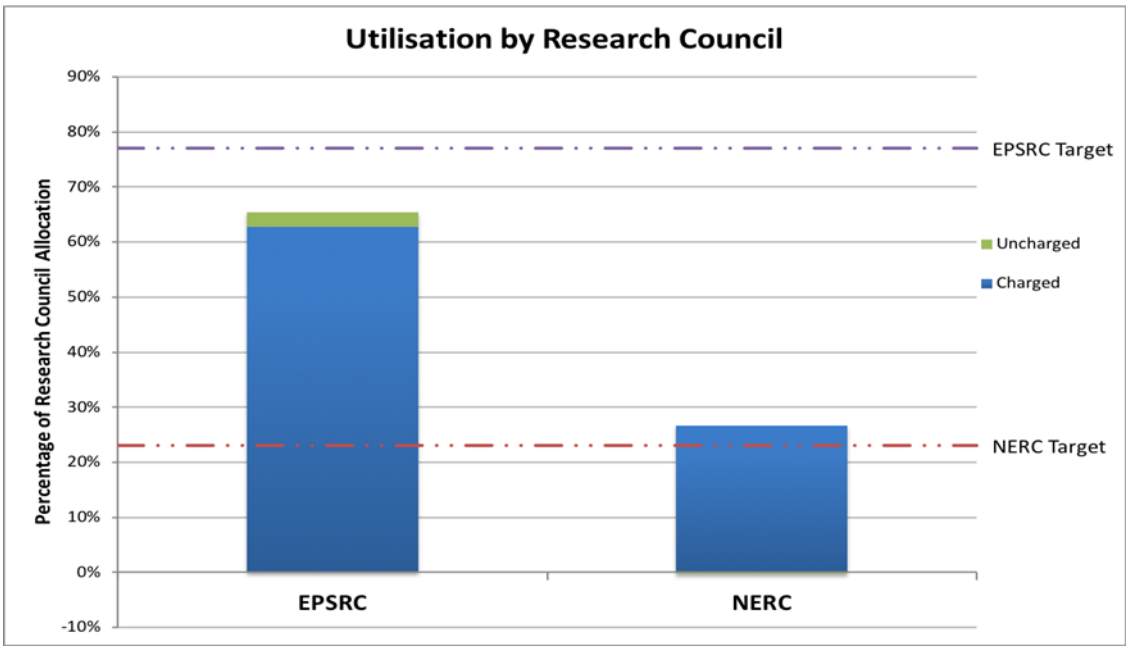

The cumulative allocation utilisation for the quarter by the Research Councils is shown below:

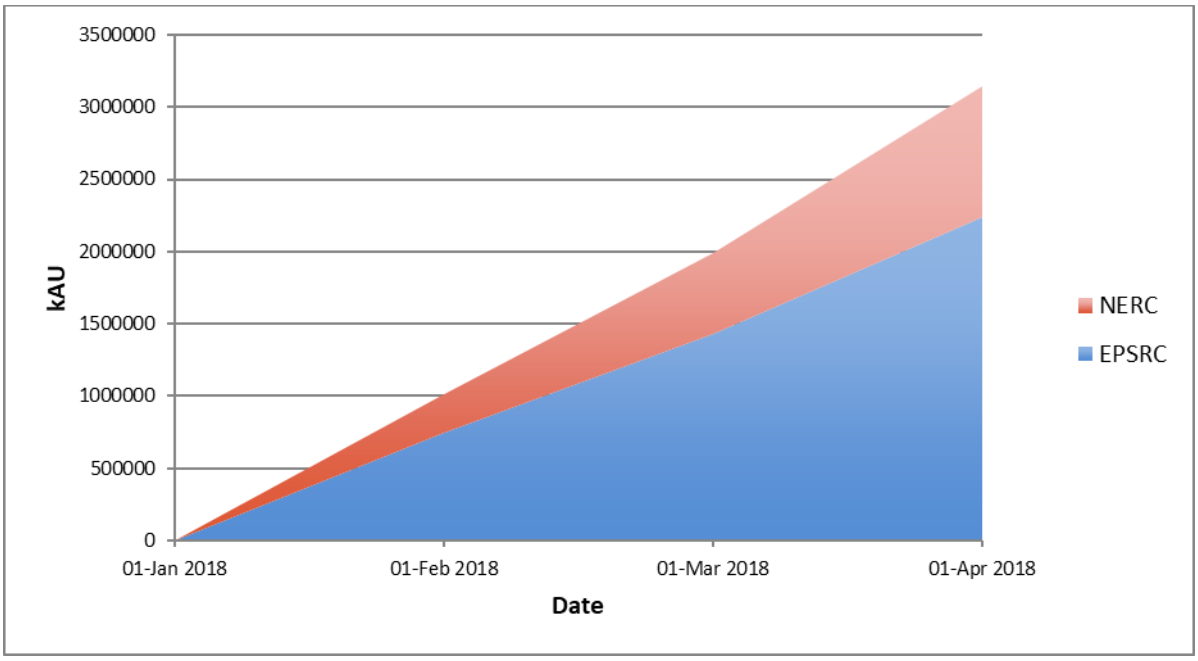

The cumulative allocation utilisation for the quarter by EPSRC broken down by different project types (see below) shows that the majority of usage comes from the scientific Consortia (as expected) with significant usage from research grants, ARCHER Leadership projects and ARCHER RAP projects. The times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

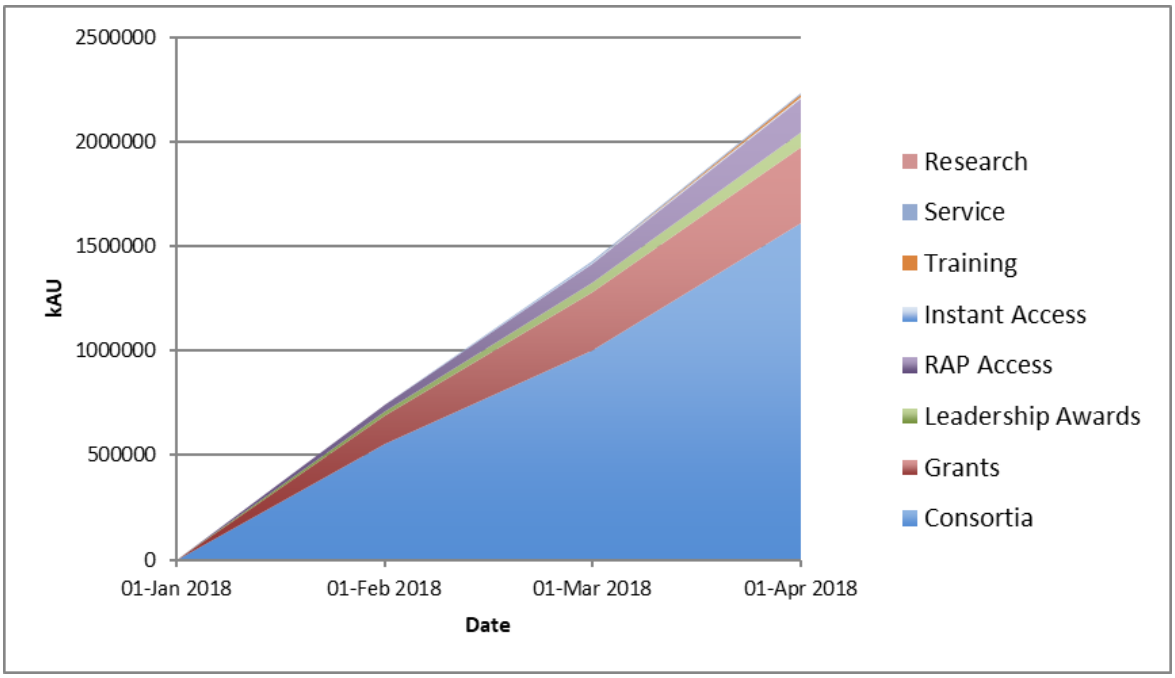

3.2 Scheduling Coefficient Matrix

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.

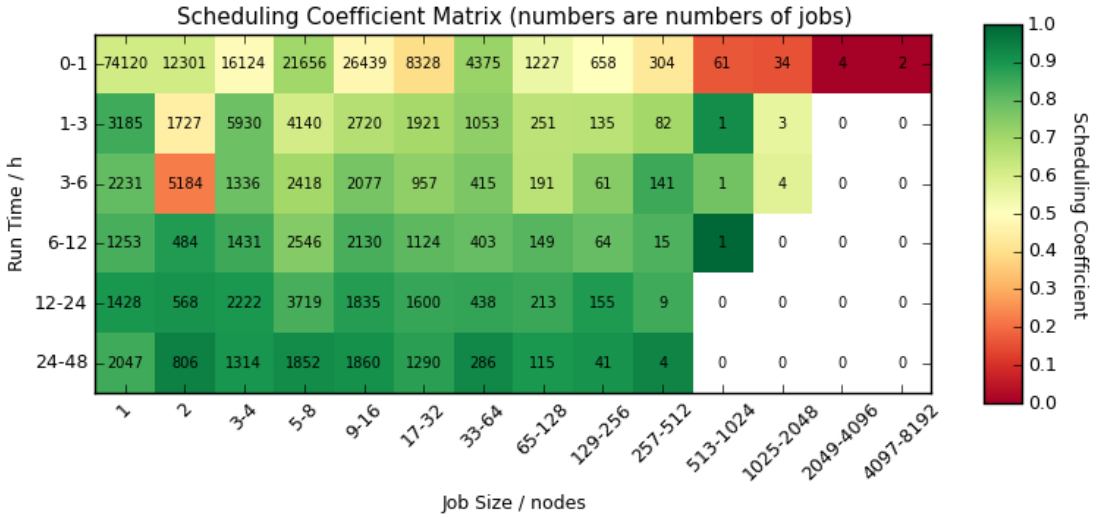

3.3 Additional Usage Graphs

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

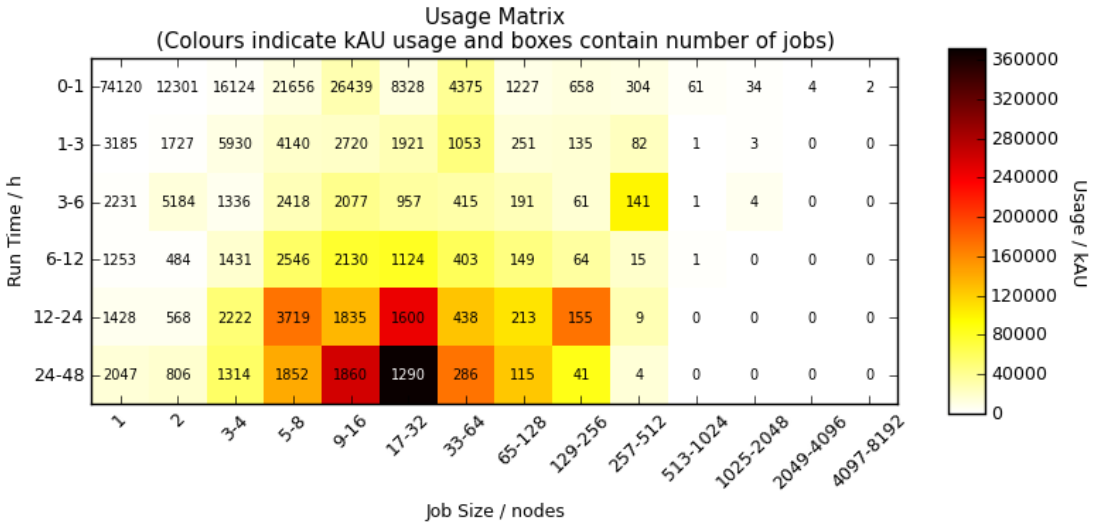

Analysis of Job Sizes

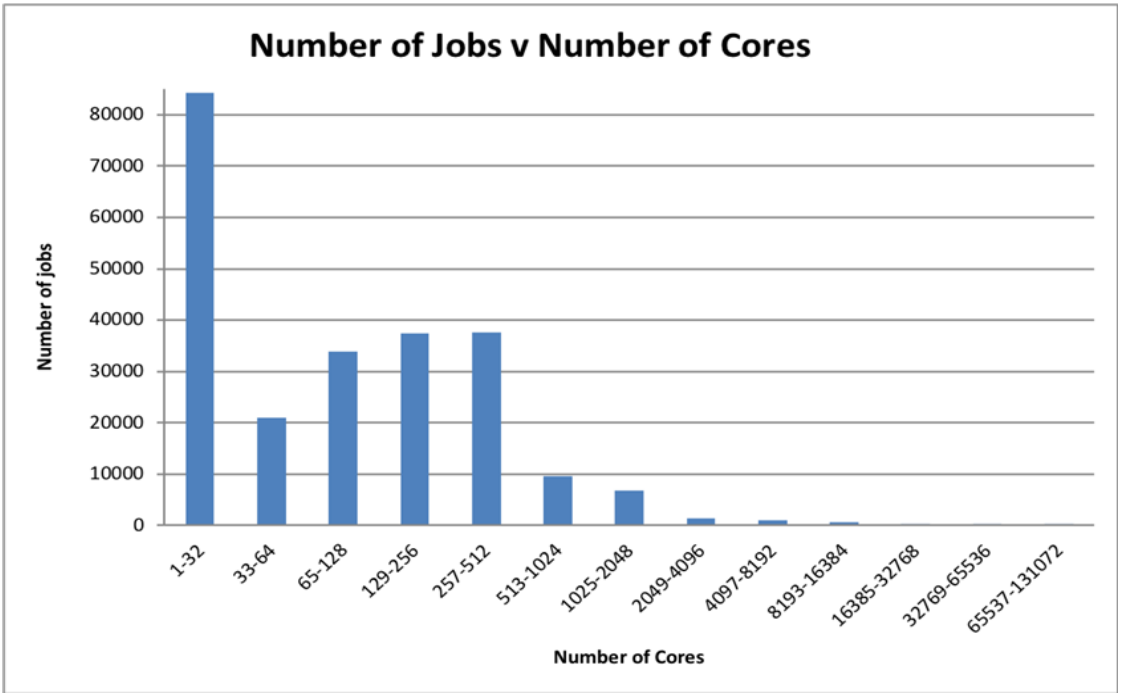

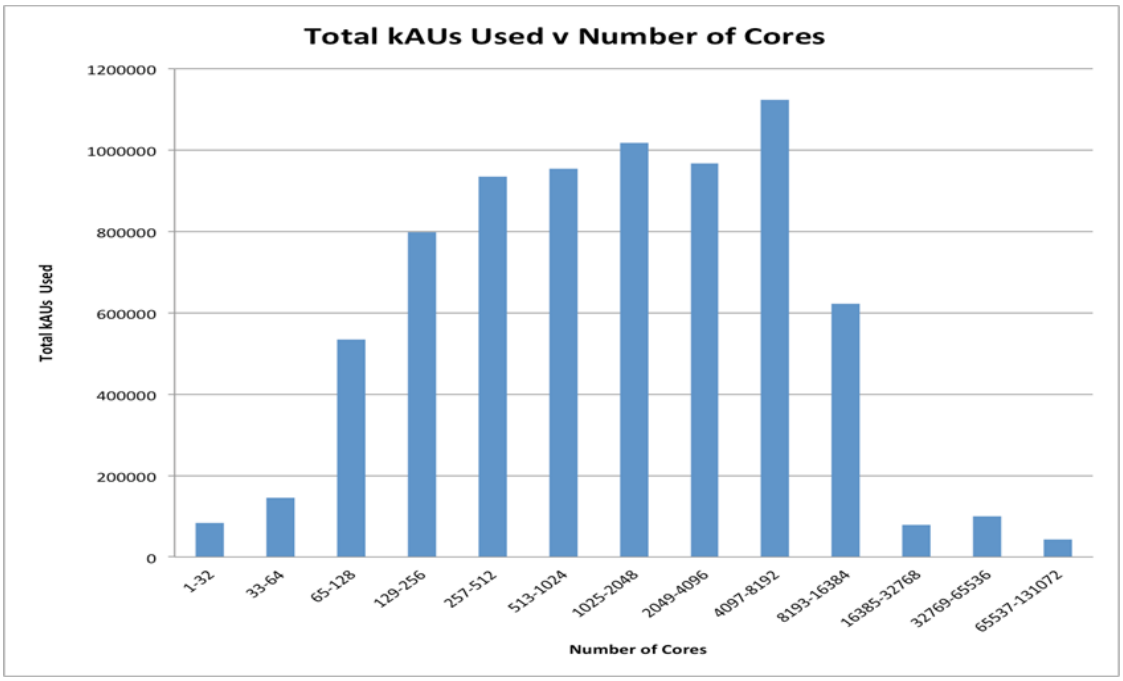

The first graph shows that, in terms of numbers, there are a significant number of jobs using no more than 1024 cores. However, the second graph reveals that most of the kAUs were spent on jobs between 65 cores and 8192 cores. The number of kAUs used is closely related to money and shows better how the investment in the system is utilised.

Analysis of Jobs Length

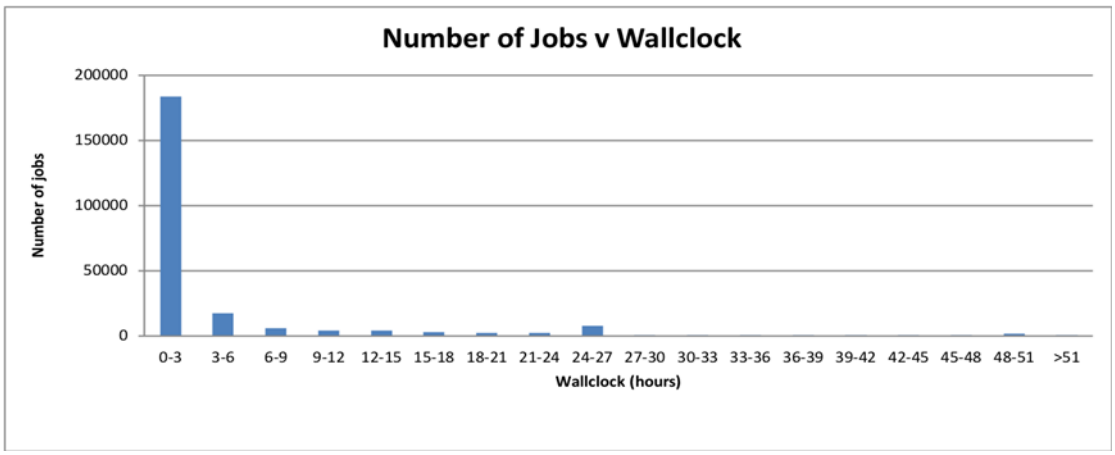

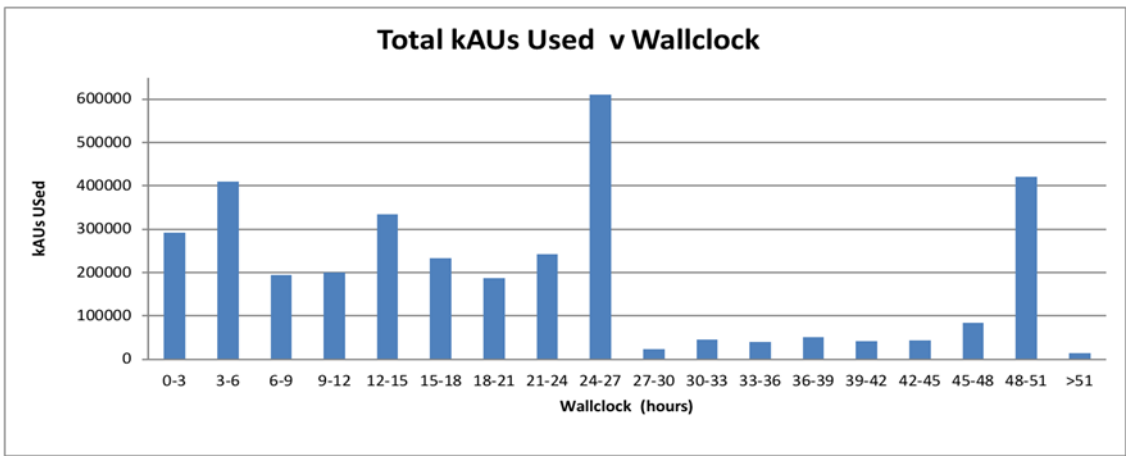

From the first graph, it would appear that the system is dominated by short jobs. However, the second graph shows that actual usage of the system is more spread and dominated by jobs of up to 27 hours with a second peak for jobs at 48-51 hours.

Core Hours per Job Analysis

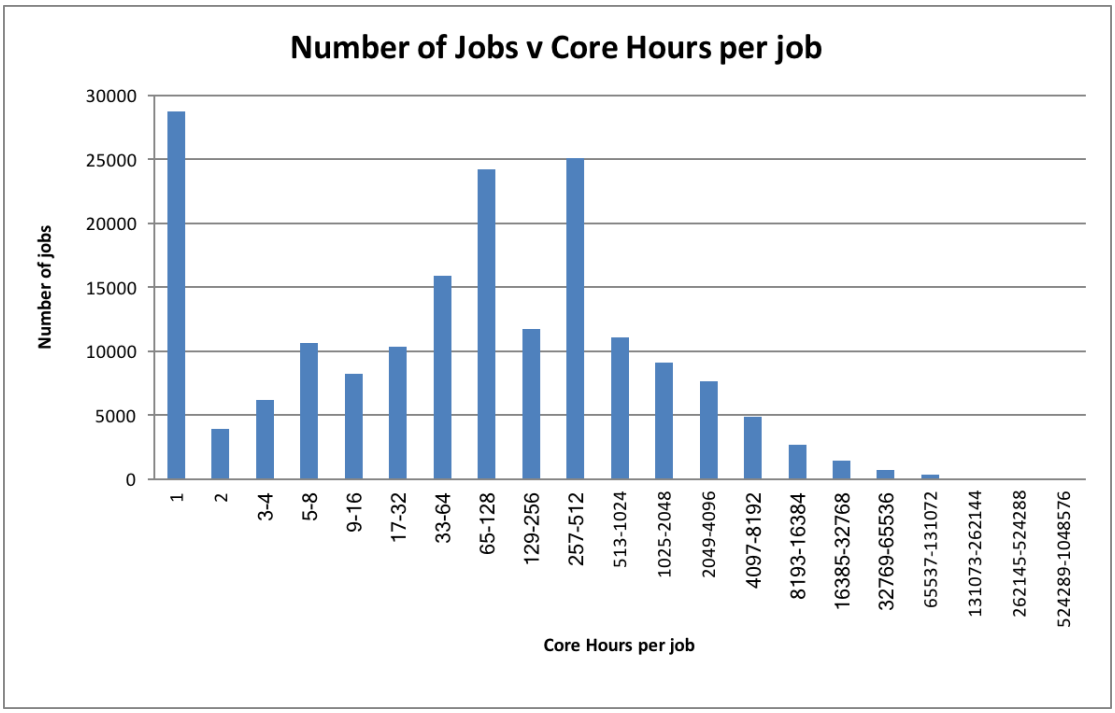

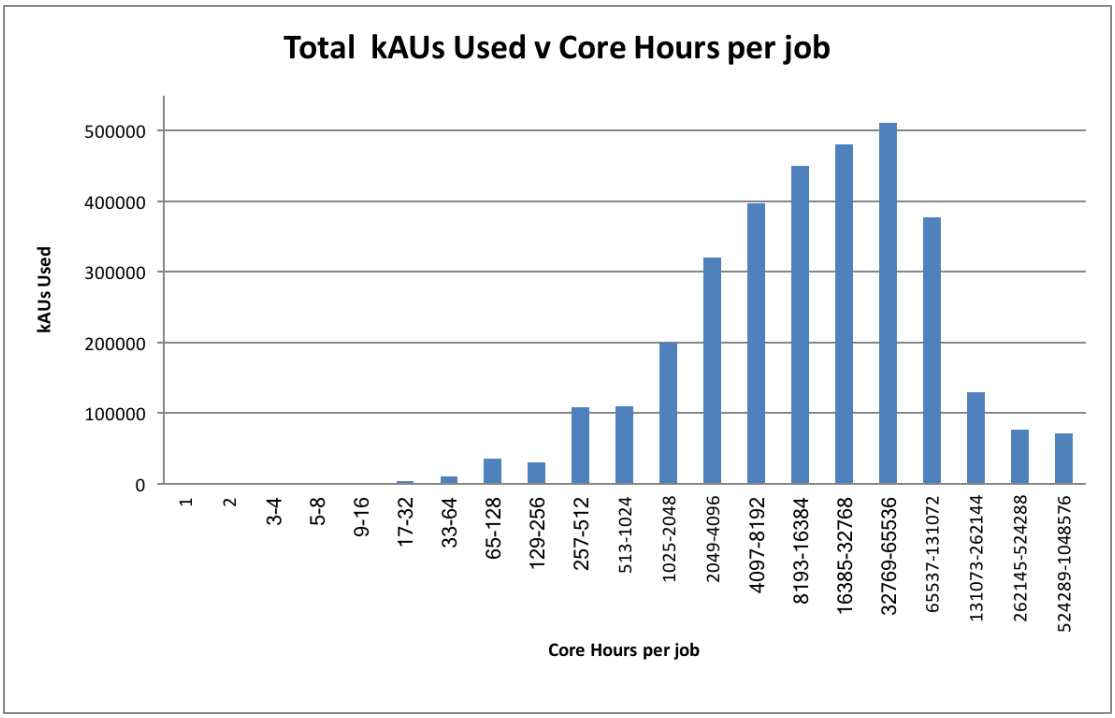

The above graphs show that, while there are quite a few jobs that use only a small number of core hours per job, most of the resource is consumed by jobs that use tens of thousands of core hours per job.