ARCHER SP Service Quarterly Report

Quarter 2 2015

Document Information and Version History

Version:	1.1
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	2015-06-22	Initial Draft	Anne Whiting
0.2	2015-06-30	Updates	Anne Whiting
0.3	2015-07-09	Added graphs	Jo Beech-Brandt
0.4	2015-07-14	Heat map added	Andy Turner
0.5	2015-07-14	Systems input added	Mike Brown
0.6	2015-07-15	Reviewed	Alan Simpson
1.0	2015-07-16	Version sent to EPSRC	Anne Whiting, Alan Simpson
1.1	2015-08-21	Updated version for web	Alan Simpson, Jo Beech-Brandt

1. The Service

1.1 Service Highlights

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

April 2015, May 2015, June 2015

- Utilisation on the system during 15Q2 was 87%, compared to 85% in 15Q1. Given the file system problems discussed below, this is remarkably positive and demonstrates that the mitigation procedures helped.

- Major Sonexion Lustre file system issues were experienced during May and June 2015 resulting in data loss on the /work area of file system 3 (fs3). Whilst this caused extensive downtime and user data loss, user impact was minimized by the provision of:
  - SAFE functionality to lock access to running jobs on ARCHER based on the file system that a project is using, forestalling the user frustration and confusion that would arise if jobs were allowed to submit but fail whilst users on other file systems continue to work.
  - Temporary project space on fs4 was provided on an unaffected file system for fs3 users. For the 3-day period from 8 May 2015 where only fs4 was available, the temporary project accounted for 79% of the usage of ARCHER.. For the 3-week period from 8 May 2015 to 22 May 2015 (when all three file systems were again available) the temporary project accounted for 46% of the usage of ARCHER. Total utilisation over the 3-week period from 8 May 2015 to 22 May 2015 was 83%; the temporary project played a key role in keeping utilisation on the system up throughout the turbulent period.
  - The move to a resilient model of package installation enabled users to access packages independently of any particular file system being available
  - The success of targetted and informative communications can be seen in the take up of the temporary project space and the small number of complaints received.

- All service levels were met during the period despite additional queries being raised due to the file system issues and it being not possible to complete a number of systems processes because of the file system issues.

- An end-of-lifecycle ARCHER policy for ending projects and former users has been published.

- A data management process has been published to facilitate the movement of data from /work to the RDF. Movement of data to the RDF increased after the file system problems.

- Benefits realisation work continues including the creation and submission of 5 case study pro formas to EPSRC.

1.2 Forward Look

- The CLE upgrade scheduled for 22 July has been deferred until after the completion of the disk replacement programme; we now expect that the upgrade will take place in late September. It was felt prudent to defer the upgrade since the timescales for the completion of the disk replacement are so tight. The session on 22 July has been offered to Cray so that they can upgrade the Sonexion software in advance of the disk replacement programme. Cray have confirmed that the existing version of CLE will be supported for critical security updates until CLE 5.2 is installed on the service during the Autumn:
    - The CLE on the TDS was upgraded during the week of 15 June 2015. Extensive testing to be carried out on the TDS to prepare support staff and to create any guidance information for users;
    - Advice will be provided to users on actions required following upgrade (such as recompiling applications and verifying output);

- RSIP node implementation will allow the use of applications requiring license servers on the compute nodes, for example, compilers and ISV applications;

- Investigation into file system problems:
    - Cray working together with Seagate and Western Digital will establish and document the root cause of the file system problems and high rate of disk failure:
    - A lessons learned analysis will be carried out involving EPSRC, Cray and EPCC to review the handling of the file system problem and to ensure that any recommendations for service improvement are identified and implemented;

- Replacement of the disks is planned to reduce the risk of further significant disk failures;

- There have been a number of PBS-related issues that the Service has raised with Cray and which Cray are taking forward directly (where appropriate) with Altair. These issues are principally with regard to scheduling parameters and it is hoped that improvements in job turnaround and better job placement will result;

- Planning is ongoing for the Workshop for ARCHER Champions scheduled for October;

- There is ongoing work on improving SAFE through the SAFE Development project that will improve the service for users. Upcoming targets include:
    - Providing greater support for Technical Assessments of new project applications within the SAFE;
    - Announcing new ARCHER project mailings automatically on twitter;
    - Updates and enhancements to the reports generated by the SAFE;
    - Improved automation for courses;
    - New forms for requesting access to packages with access restrictions;
    - An improved, more intuitive user interface.

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	Apr 15		May 15		Jun 15		15Q2
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	98.3%	0	97.4%	0	97.8%	0	0
2.6.4 – UR	1 WD	0	1 WD	0	1 WD	0	0
Total		**-5**		**-5**		**-5**	**-15**

The details of the above can be found in Section 2.2 of this report.

2.1.2 Service Failures

There were no SP service failures in the period, as defined in the contract. However outages were approved by EPSRC as required during the process of working to resolve the file system issues.

2.1.3 Service Credits

The total Service Credit applicable for each Service Quarter is calculated in the following way:

[redacted]

Where:
"Applicable Charge" = the relevant Annual Maintenance Charge divided by four (4) (to form the Maintenance Charge relevant for the Service Periods being assessed)
"SC" = Service Credit
"TSP" = Total Service Points for the Service Quarter

As the Total Service Points are negative (-21), no Service Credits apply in 15Q2.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Apr 15	May 15	Jun 15	15Q2
Phone Calls Received	40 (10)	32 (7)	28 (5)	**100 (22)**
Answered 2 Minutes	40	32 (7)	28 (5)	**100 (22)**
Service Level	**100.0%**	**100.0%**	**100.0%**	**100.0%**

The volume of telephone calls remained low in 15Q2. Of the total of 100 calls received above, only 22 were genuine ARCHER user calls that resulted in queries or answered user questions directly.

2.2.2 Query Closure (QC)

	Apr 15	May 15	Jun 15	15Q2
Self-Service Admin	347	413	412	**1172**
Admin	174	200	242	**616**
Technical	19	29	43	**91**
Total Queries	*540*	*642*	*697*	***1879***
Total Closed in 2 Days	*531*	*625*	*682*	***1838***
Service Level	**98.3%**	**97.4%**	**97.8%**	**97.8%**

In addition to the Admin and Technical queries, the following Change Requests were resolved in 15Q1.

	Apr 15	May 15	Jun 15	15Q2
Change Requests	3	4	0	**7**

2.2.3 User Registration (UR)

	Apr 15	May 15	Jun 15	15Q2
No of Requests	80	60	74	**214**
Closed in One Working Day	80	60	74	**214**
Average Closure Time (Hrs)	0.82	1.42	0.92	**1.03**
Average Closure Time (Working Days)	0.09	0.15	0.10	**0.11**
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:

	Apr15	May15	Jun15	15Q2
CRAY	18	13	4	**35**
ARCHER_CSE	149	89	85	**323**
ARCHER_SP	920	1120	1001	**3041**
Total Queries Assigned	1087	1222	1090	**3399**
Total Assigned in 1 Hour	1087	1222	1090	**3399**
Service Level	**100%**	**100%**	**100%**	**100%**

2.3.2 Maintenance

SP is allowed to book a maximum of two maintenance occasions in any 28-day period, and these shall last no longer than four hours; these are defined as Permitted Maintenance. Such Maintenance Periods are recorded in the Maintenance Schedule. A 6-month forward plan of maintenance has been agreed with the Authority.

It has been agreed with the Authority that SP may combine the hours normally allocated for two consecutive maintenance sessions into a single session with a maximum of eight hours and this has been the normal mode of operation as recorded in the table below. This reduces the number of sessions taken, which reduces user impact since the jobs running on the service have to be drained down once and not twice.

If greater than 4 hours downtime is required for maintenance, 20 days prior approval is required from 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
22nd April 2015	0900	1628	7:28	Pre-Approved	EPSRC Approved 0900 - 1700	Planned Maintenance
20th May 2015	Postponed due to file system problems			Pre-Approved	EPSRC Approved 0900 - 1700	Planned Maintenance
24th June 2015	0900	1613	7:13	Pre-Approved	EPSRC Approved 0900 - 1700	Planned Maintenance

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 87%.

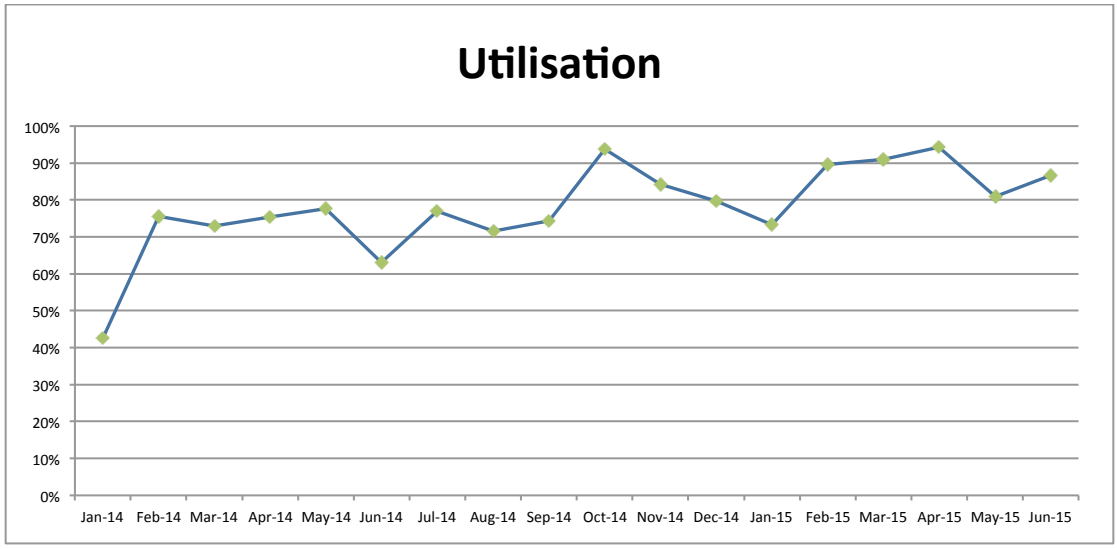

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. The large amount of uncharged time this quarter is primarily from the temporary project (v01) that was put in place during the filesystem issues.

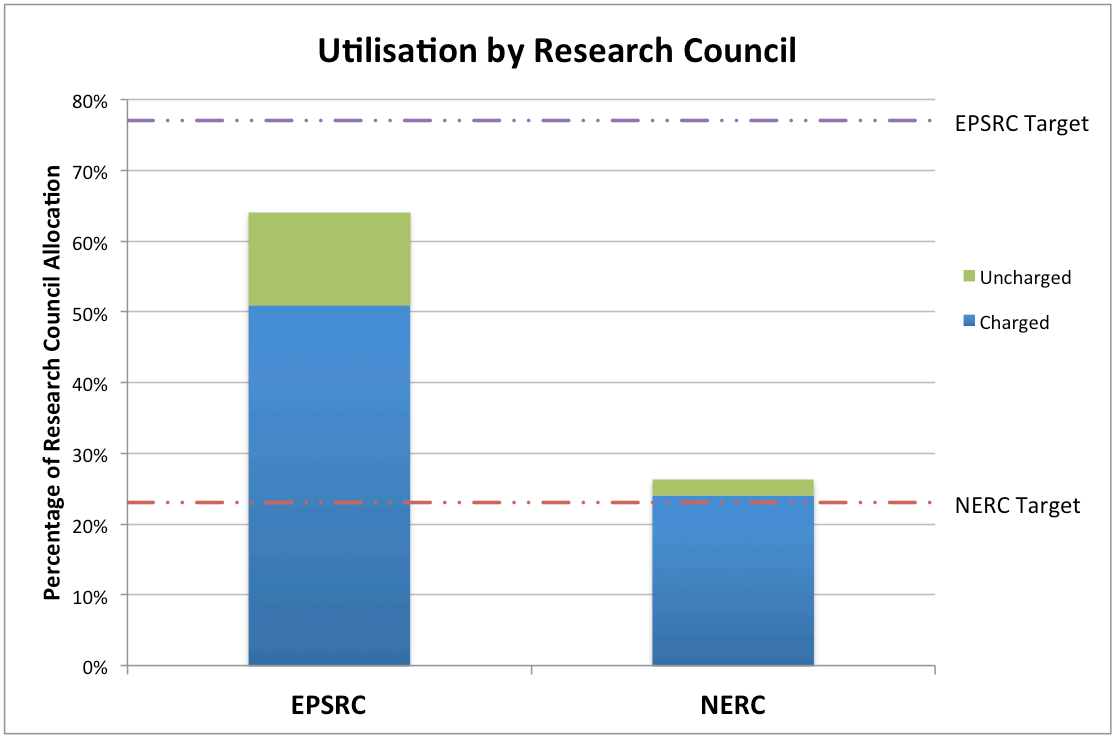

3.2 Scheduling Coefficient Matrix

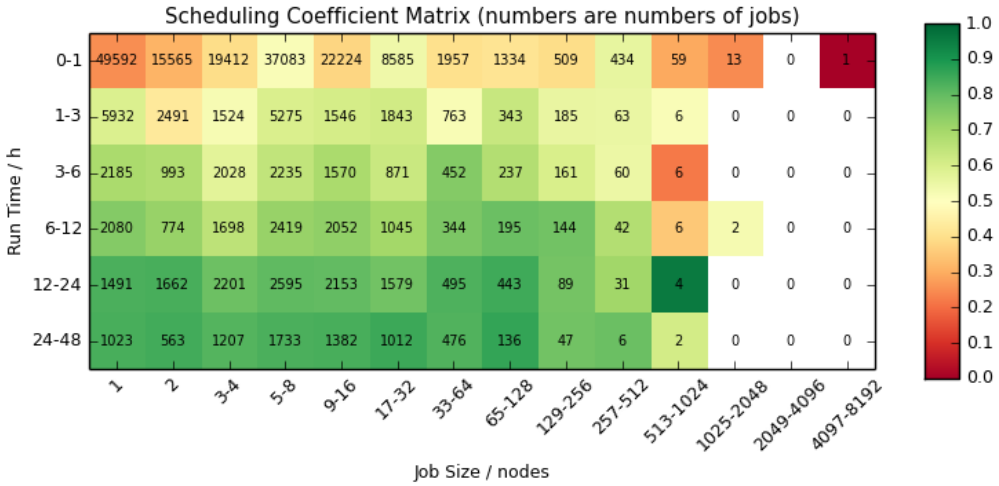

The colour in the matrix indicates the value of the Scheduling Coefficient. This is defined as the ratio of runtime to runtime plus wait time. Hence, a value of 1 (green) indicates that a job ran with no time waiting in the queue, a value of 0.5 (pale yellow) indicates a job queued for the same amount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

The matrix shows that generally queuing times are short. The only cases where longer wait times than runtimes are encountered are either for very short jobs (as there is always a scheduling overhead) or for very large jobs (where the system has to drain compute nodes to make space for the jobs).

3.3 Additional Usage Graphs

The following charts provide different views of the distribution of job sizes on ARCHER.

Number of Cores

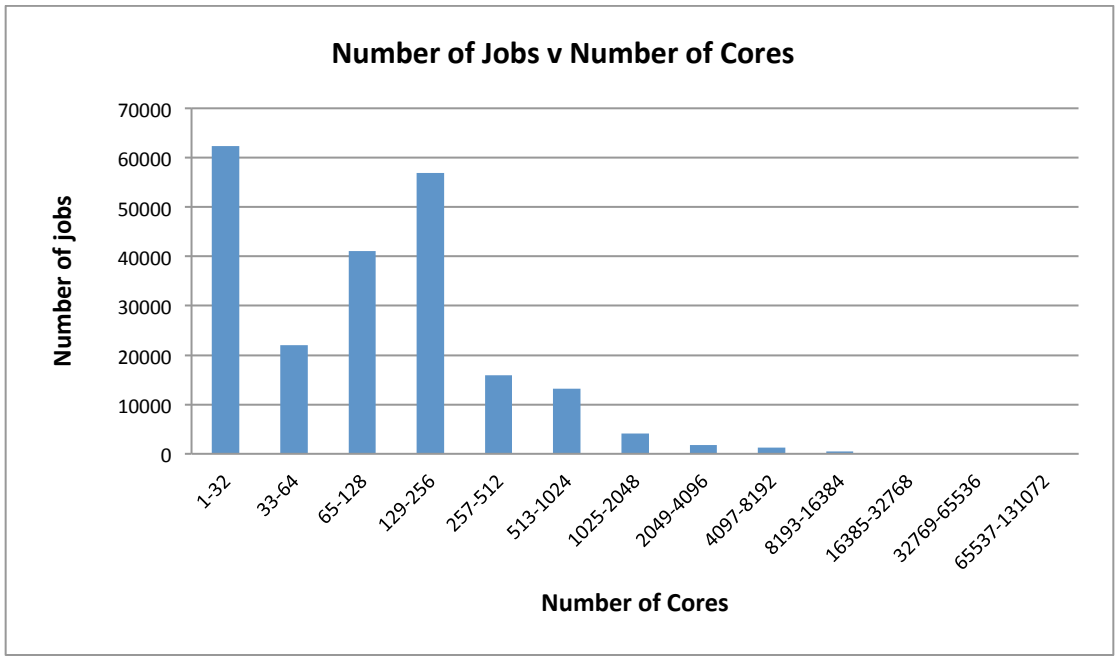

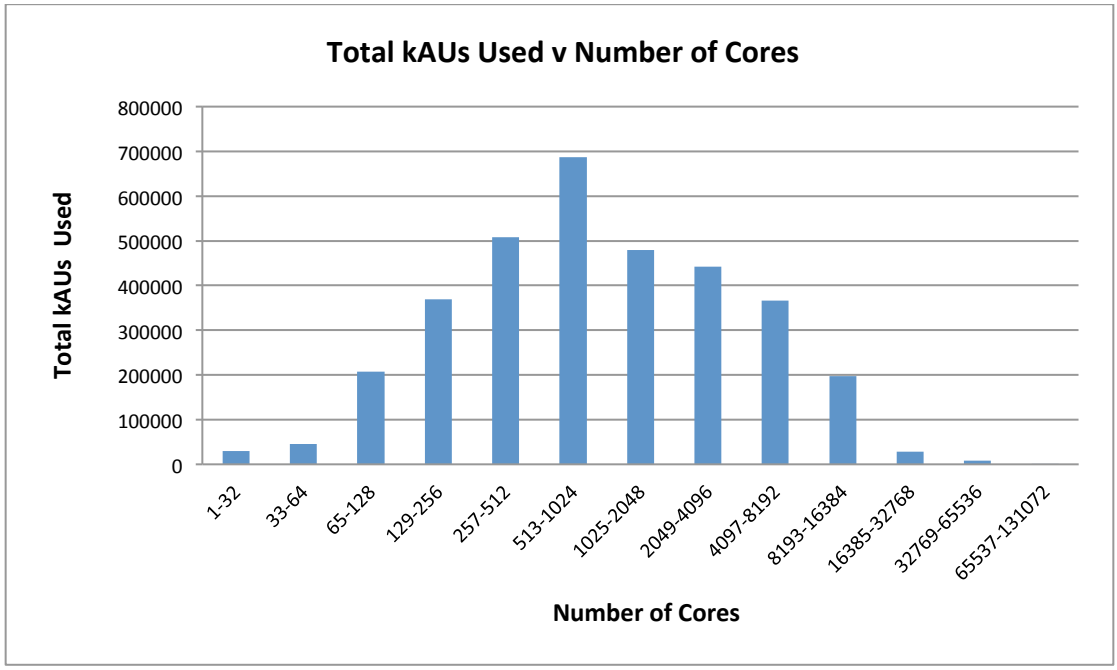

The first graph shows that, in terms of numbers, there is a significant number of jobs using no more than 512 cores. However, the second graph shows that most of the kAUs were spent on jobs between 257 cores and 8192 cores. The number of kAUs used is closely related to money and shows better how the investment in the system is utilised.

Wallclock

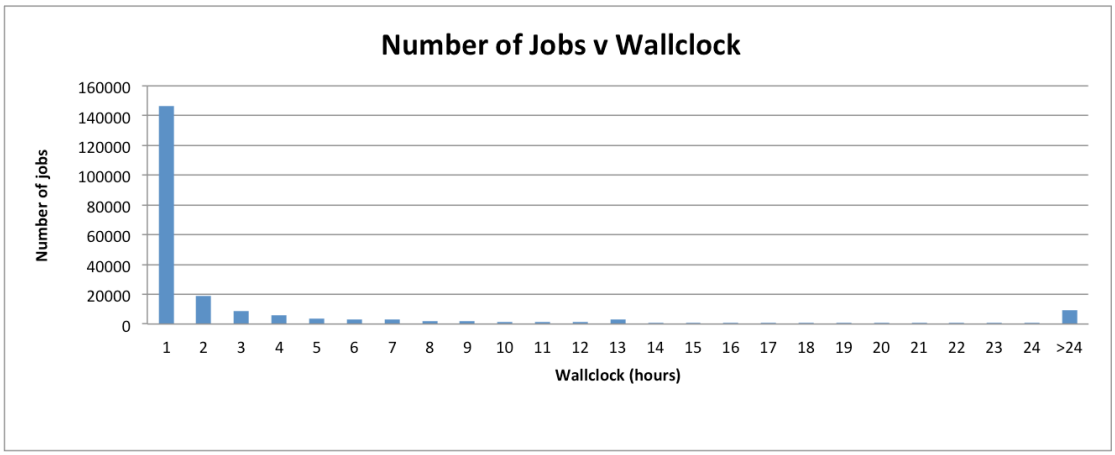

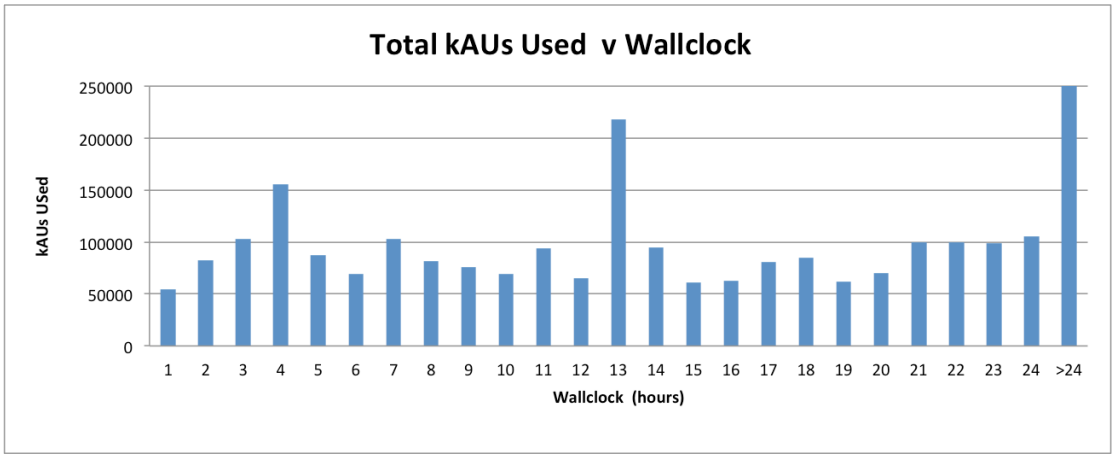

From the first graph, it would appear that the system is dominated by short jobs. However, the second graph shows that actual usage of the system is dominated by jobs of more than 12 hours.

Core Hours

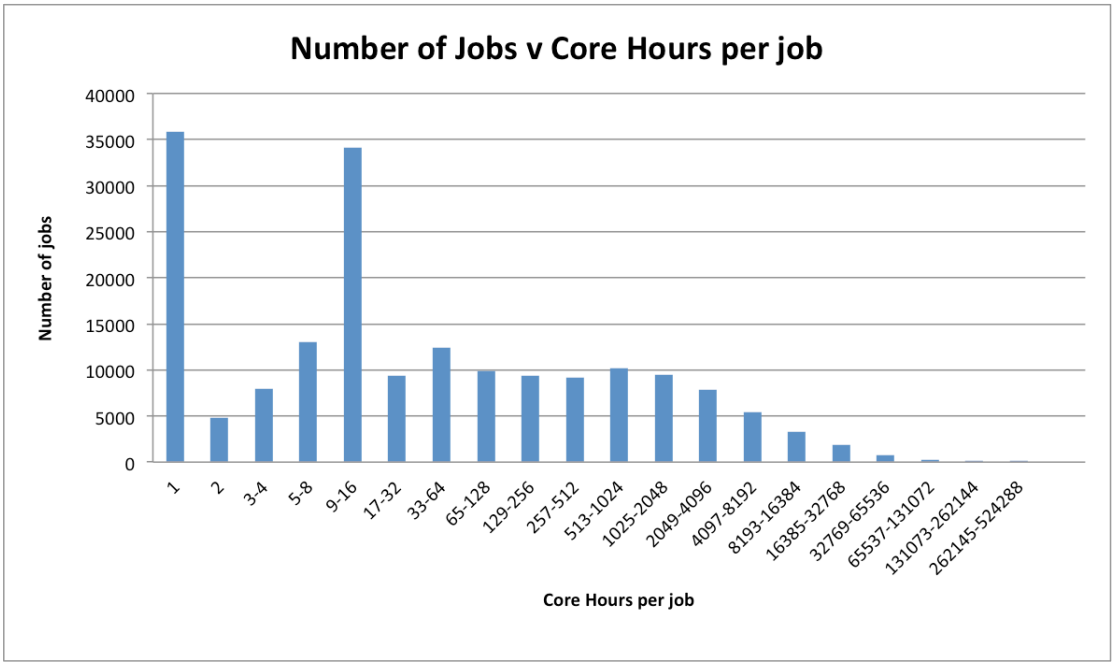

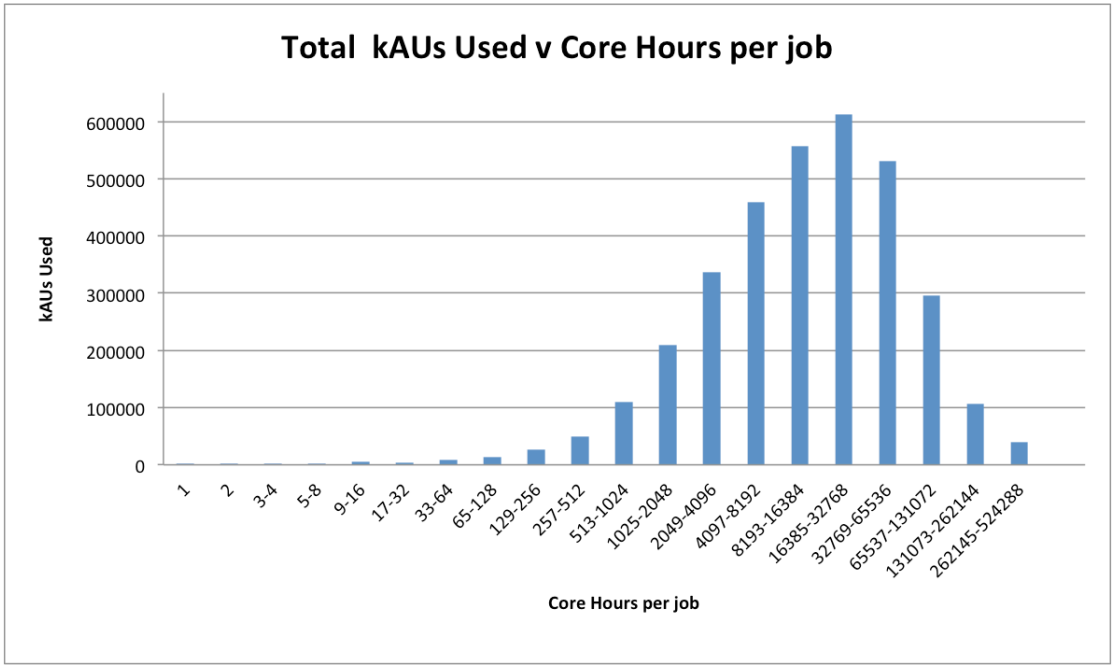

The above graphs show that, while there are quite a few jobs that use only a small number of core hours per job, most of the resource is consumed by jobs that use tens of thousands of core hours per job.