
# ARCHER SP Service Quarterly Report

Quarter 2 2019

# 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 | 28/06/19 | Initial Draft | Anne Whiting |
| 0.2 | 08/07/19 | Graphs added | Jo Beech-Brandt, Anne Whiting |
| 0.3 | 10/07/19 | Added highlights | Anne Whiting |
| 0.4 | 10/07/19 | HPC Systems input | Linda Dewar |
| 0.5 | 10/07/19 | Reviewed | Alan Simpson |
| 0.6 | 11/07/19 | Updated post review | Anne Whiting |
| 1.0 | 16/07/19 | 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:

April 2019, May 2019 and June 2019.

- Utilisation over the quarter was 93%, an increase from the previous quarter where the utilisation was 86%. Despite the machine being so busy, the positive impact of the work carried out by EPCC to modify the job priority formula to balance queue times across different job sizes is helping to ensure that no job has to wait too long to run. This is reflected in the scheduling coefficient matrix on p11 of this report.

- EPSRC and EPCC reached agreement to extend the ARCHER SP service to end on 18 February 2020 to allow additional time for the procurement process for ARCHER 2 to be completed. ARCHER2 is expected to enter service in May 2020

- After user feedback received, two new reports have been added to SAFE to enable users and group leaders to run the LASSi reports to analyse and show parallel IO performance by user and by group. The reports include a summary of data written and read and statistics from individual jobs. The LASSi information can be used by users to assist in the tuning of their I/O use to try and avoid file system performance reductions. The LASSi tool was developed by Cray and integrated into SAFE by EPCC.

- We are keen to minimise the impact of service outages for full maintenance sessions on the user community and aim to carry out maintenance tasks during at risk sessions wherever possible. We are pleased to be able to say that we have taken no maintenance outages this quarter.

- An HPC Champions event was held in May in Edinburgh which was co-located with the launch of the Women in High Performance Computing (WHPC) Chapter at EPCC. Sessions included site updates from all Tier2 centres, the current HPC Landscape and the UKRI roadmap, the transition between ARCHER and ARCHER2, an update on HPC Europa, and developments at the ACF. Lightning talks given were very varied and included sessions by academic users, system staff, vendors and also a UKRI research fellow. The day ended with discussion groups focusing on new technologies, training and future Champions events.

- Members of the Systems and User Support teams continue to represent ARCHER within the community, sharing ideas and expertise by attending and presenting at regional, national and international meetings and workshops. This quarter these have included the Scottish Regional HPC technical meeting in Dundee, HPC-SIG in Belfast, ISC in Frankfurt and Women in HPC in Edinburgh. Sessions across these events have ranged from Machine Learning, common storage issues and new technologies. Talks were given by our team at the Women in HPC launch including the benefits of taking paternity leave by our Director of HPC Systems and a 20-year career as a female systems administrator by our Principal Systems Administrator.

## 1.2 Forward Look

- EPSRC have announced that there will be a break in service of up to 11 weeks from 18 February 2020 between the ARCHER and ARCHER2 services. EPCC is working with EPSRC and NERC to provide the user community with the necessary information and support to be able to plan projects and to mitigate the impact of the break in service. Updates will be provided as they become available.

- Plans are underway to hold a Tier1/Tier2 half day workshop at the RSE Conference 2019, which is being held from 16 – 19 September in Birmingham. This Champions event will build on the recent Champions event held in May 2019 in Edinburgh. New Champions are always welcome, and interest can be expressed by sending in a message to the ARCHER Helpdesk.

- In order to facilitate faster data movement to the RDF GPFS filesystems from the ARCHER login nodes, plans are being finalised to upgrade the existing bonded pair of 10gbit links from the ARCHER core switches to the RDF sitewide network a pair of 40gbit links. The increase in transfer speed will assist users in transferring data from ARCHER to the RDF.

- EPCC are currently updating the Data Management Guide to provide guidance on the transfer of data from ARCHER in preparation for the end of service. A webinar on this topic is planned for September 2019.

- Work is underway to prepare for a combined ISO 9001 Quality Management and ISO 27001 Information Security external audit in the autumn. Moving to a combined system and audits leverages the strengths of the process-based quality management approach with the controls provided by the information security management system to deliver the best and most secure service to our users.

# 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 19 | | May 19 | | Jun 19 | | 19Q2 |
|---|---|---|---|---|---|---|---|
| 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 have been no unplanned outages this quarter.

### 2.1.3 Service Credits

As the Total Service Points are negative (-21), no Service Credits apply in 19Q2.

## 2.2 Detailed Service Level Breakdown

### 2.2.1 Phone Response (PR)

| | Apr 19 | May 19 | Jun 19 | 19Q2 |
|---|---|---|---|---|
| Phone Calls Received | 24 (4) | 22 (4) | 17 (2) | **63 (10)** |
| Answered in 2 Minutes | 24 | 22 | 17 | **63** |
| **Service Level** | **100.0%** | **100.0%** | **100.0%** | **100.0%** |

The volume of telephone calls remained low in 19Q2. Of the total of 63 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)

| | Apr 19 | May 19 | Jun 19 | 19Q2 |
|---|---|---|---|---|
| Self-Service Admin | 426 | 613 | 322 | **1361** |
| Admin | 105 | 107 | 72 | **304** |
| Technical | 23 | 24 | 19 | **64** |
| *Total Queries* | 554 | 744 | 431 | **1729** |
| *Total Closed in 2 Days* | 544 | 739 | 429 | **1712** |
| **Service Level** | **98.2%** | **99.3%** | **99.5%** | **99.0%** |

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 19Q2:

| | Apr 19 | May 19 | Jun 19 | 19Q2 |
|---|---|---|---|---|
| Change Requests | 1 | 0 | 1 | **2** |

### 2.2.3 User Registration (UR)

| | Apr 19 | May 19 | Jun 19 | 19Q2 |
|---|---|---|---|---|
| No of Requests | 71 | 78 | 63 | **212** |
| Closed in One Working Day | 71 | 78 | 63 | **212** |
| Average Closure Time (Hrs) | 0.7 | 0.5 | 0.6 | **0.6** |
| Average Closure Time (Working Days) | 0.07 | 0.06 | 0.06 | **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:

| | Apr 19 | May 19 | Jun 19 | 19Q2 |
|---|---|---|---|---|
| CRAY | 4 | 4 | 8 | **16** |
| ARCHER_CSE | 177 | 139 | 79 | **395** |
| ARCHER_SP | 868 | 1061 | 716 | **2645** |
| Total Queries Assigned | 1049 | 1204 | 803 | **3056** |
| Total Assigned in 1 Hour | 1049 | 1204 | 803 | **3056** |
| **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 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 at most once per month. 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.

No planned maintenance outages were taken this quarter.

### 2.3.3 Quality Tokens and query feedback emails

One 5 star positive quality token was received communicating how happy the user is with the service and giving a special mention for a member of the CSE team who carried out a review of a technical assessment for the user at very short notice.  No negative tokens were received.

5 very positive feedback emails were received from users upon closure of their queries.  No negative feedback emails were received.

# 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 93%, up from 86% the previous quarter. Utilisation for April was 88%, for May 95% and for June 96%. 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.

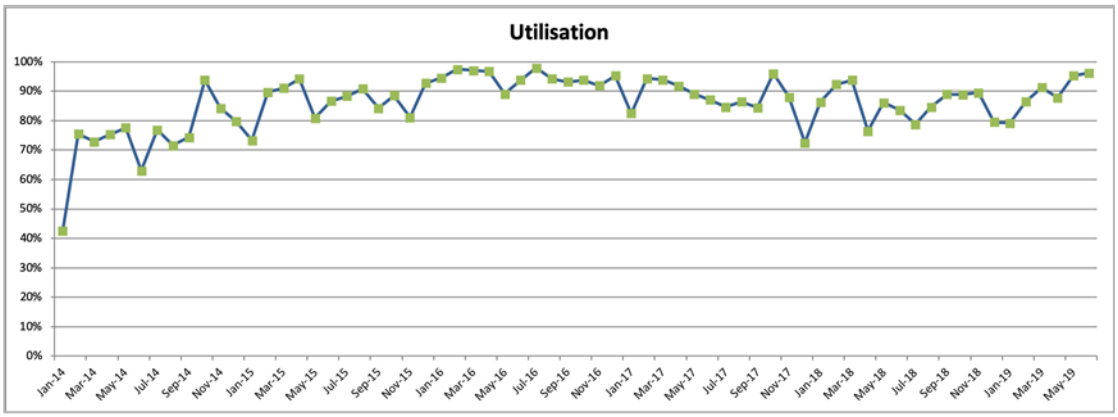


The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER. It can be seen that EPRSC almost met their target this quarter with their usage being at 76% (against their target of 77%) whereas NERC exceeded their target with utilisation being 25% (against their target of 23%). This compares with 63% for EPSRC and 25% for NERC for the previous quarter.

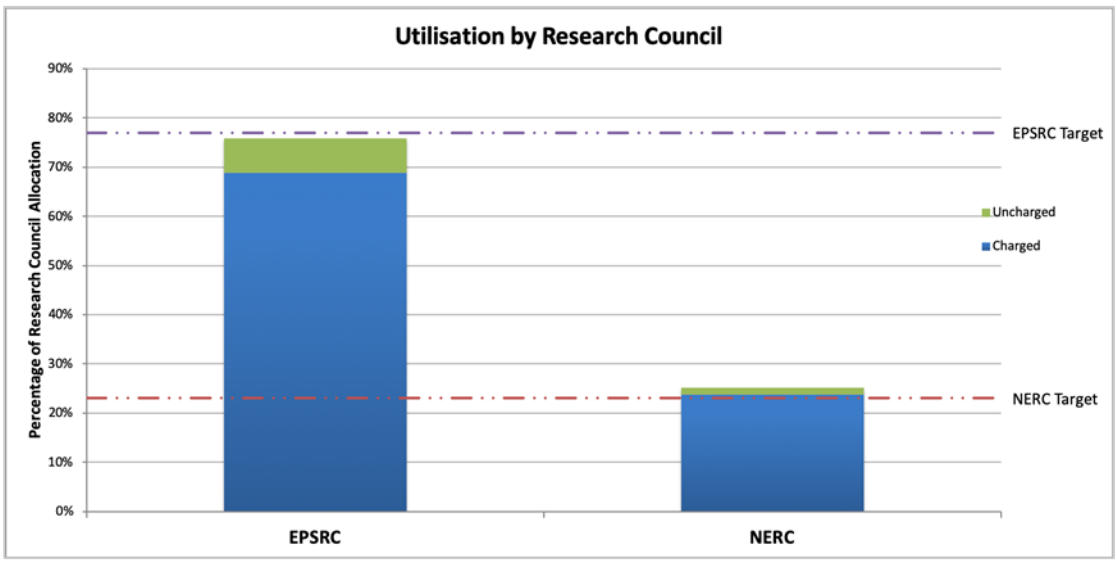

The cumulative allocation utilisation for the quarter by the Research Councils is shown below:

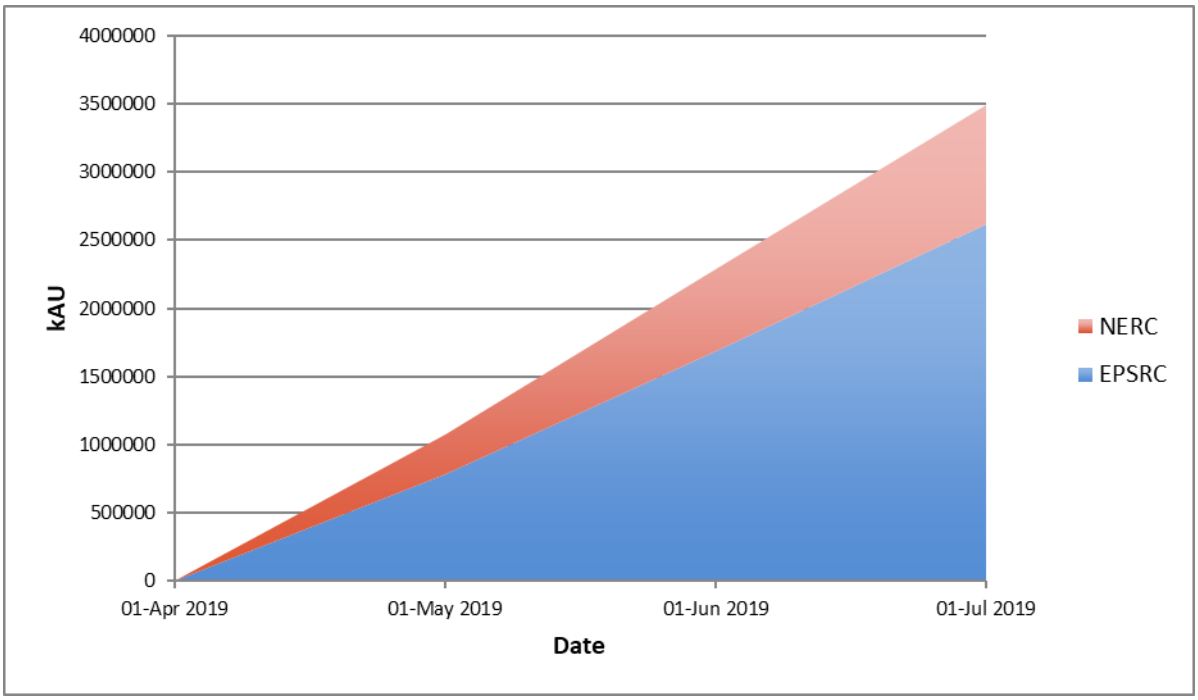


The cumulative allocation utilisation for the quarter by EPSRC broken down by different project types (see below) shows that the majority of usage comes from the scientific Consortia (as expected) with significant usage from research grants, CSC (the Finnish IT Center for Science) and ARCHER RAP projects. EPSRC have agreed some time for CSC on ARCHER whilst CSC are awaiting their new HPC system. The times used by Instant Access projects is very small.

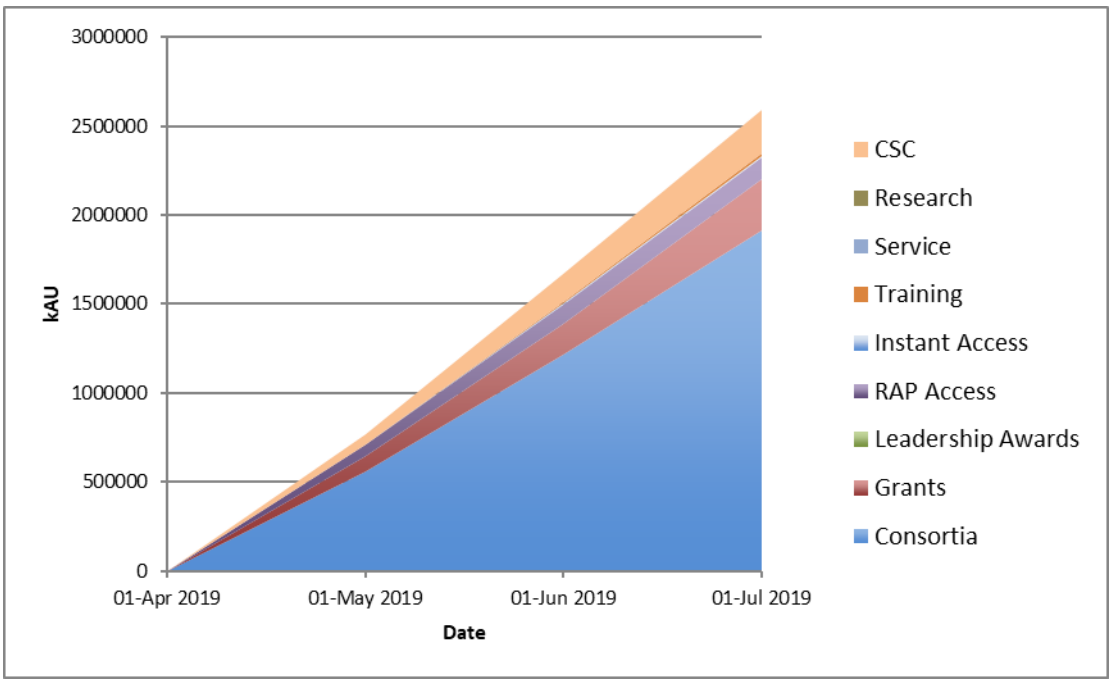

## 3.2 Scheduling Coefficient Matrix

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.

Despite the high utilisation this quarter, the heatmap shows that wait times have remained modest.

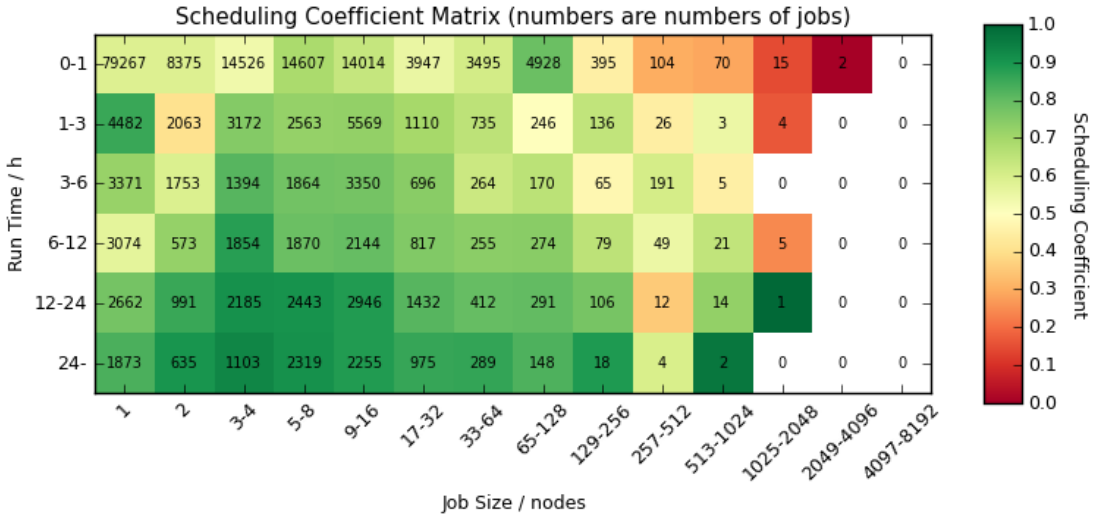


## 3.3 Additional Usage Graphs

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.

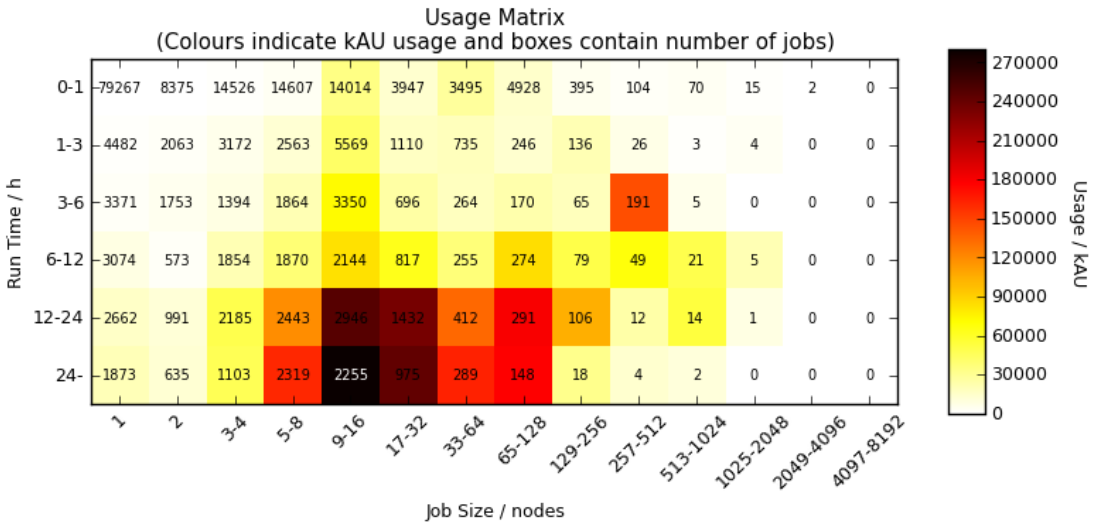

## Analysis of Job Sizes

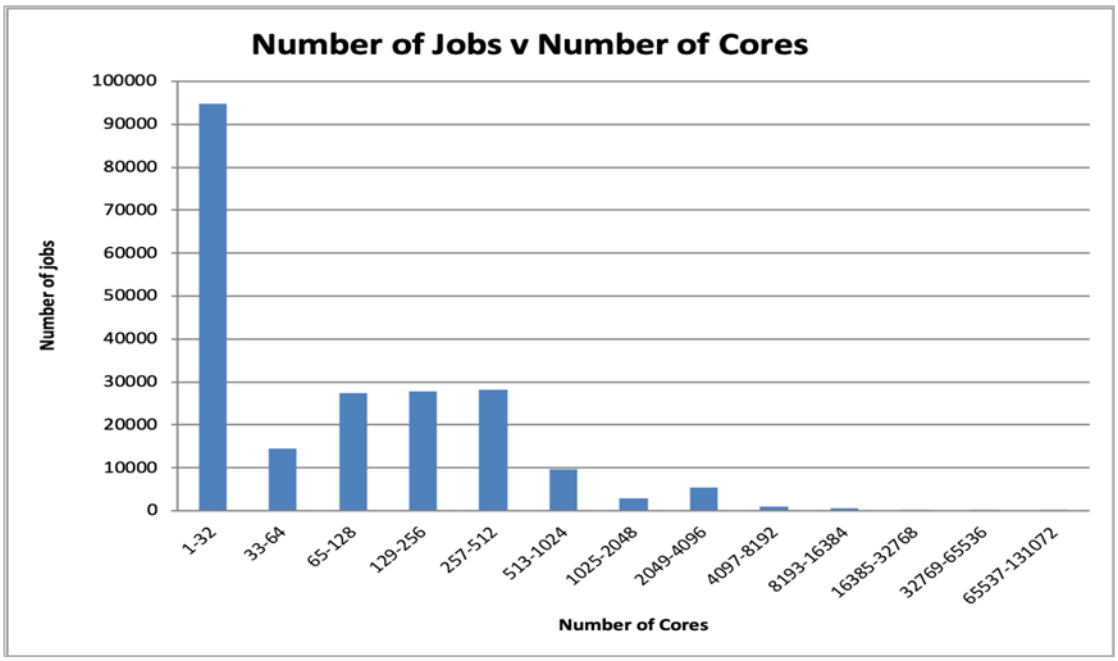


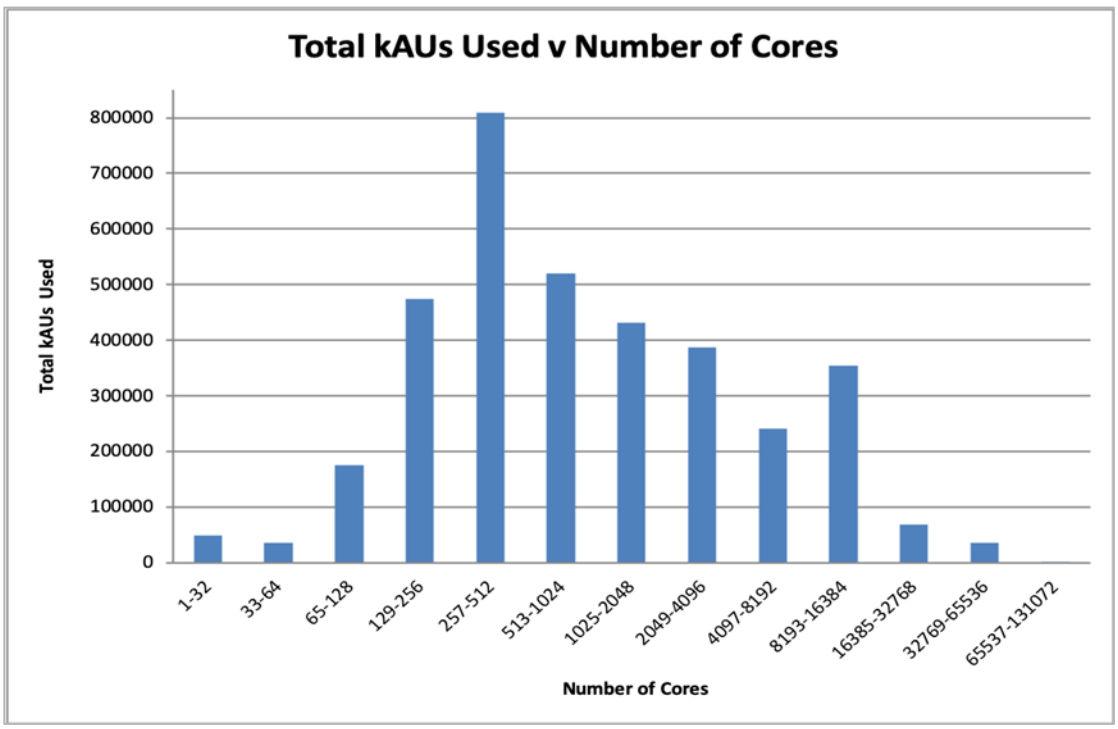


The first graph shows that, in terms of numbers, there are a significant number of jobs using no more than 1024 cores. However, the second graph reveals that most of the kAUs were spent on jobs between 65 cores and 16384 cores. The number of kAUs used is closely related to money and shows better how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

## Analysis of Jobs Length

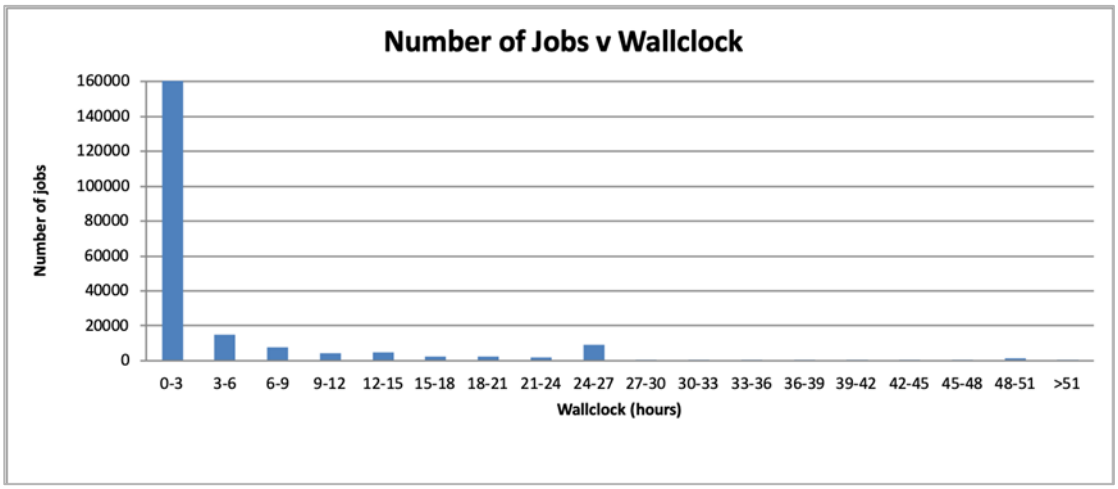


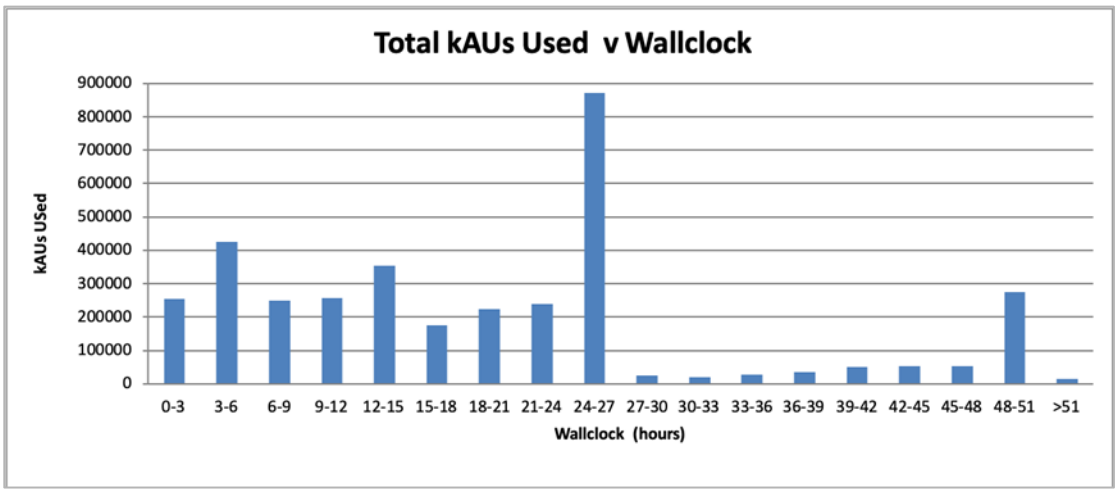


From the first graph, it would appear that the system is dominated by short jobs. However, the second graph shows that actual usage of the system is more spread and dominated by jobs of up to 27 hours with a second peak for jobs at 48-51 hours.

## Core Hours per Job Analysis

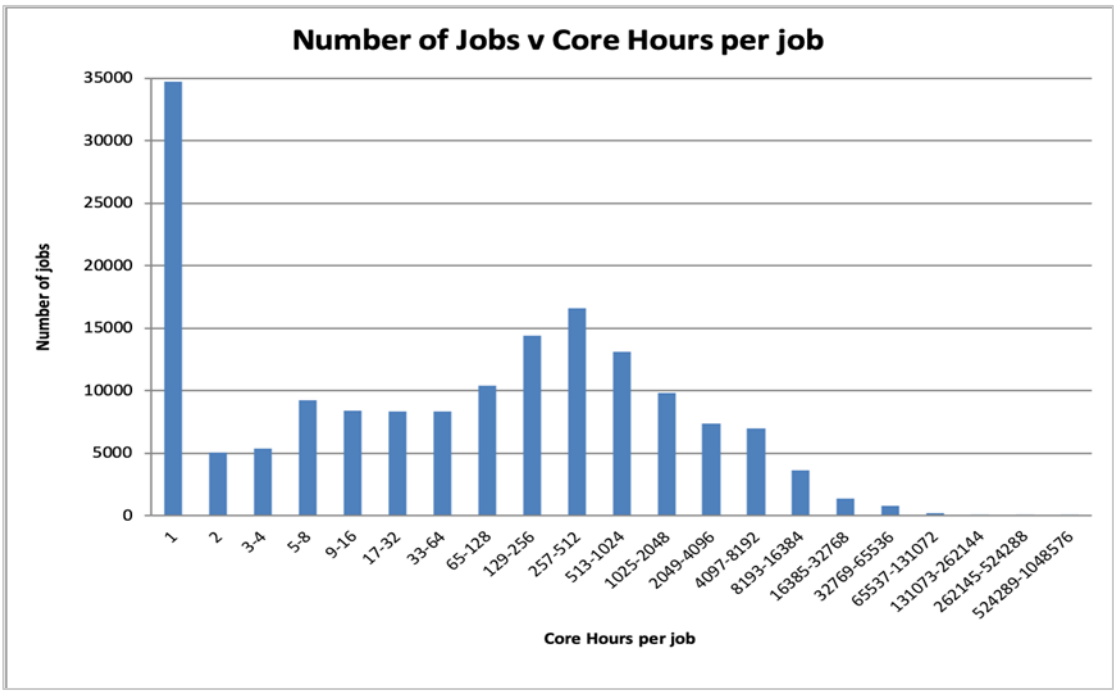


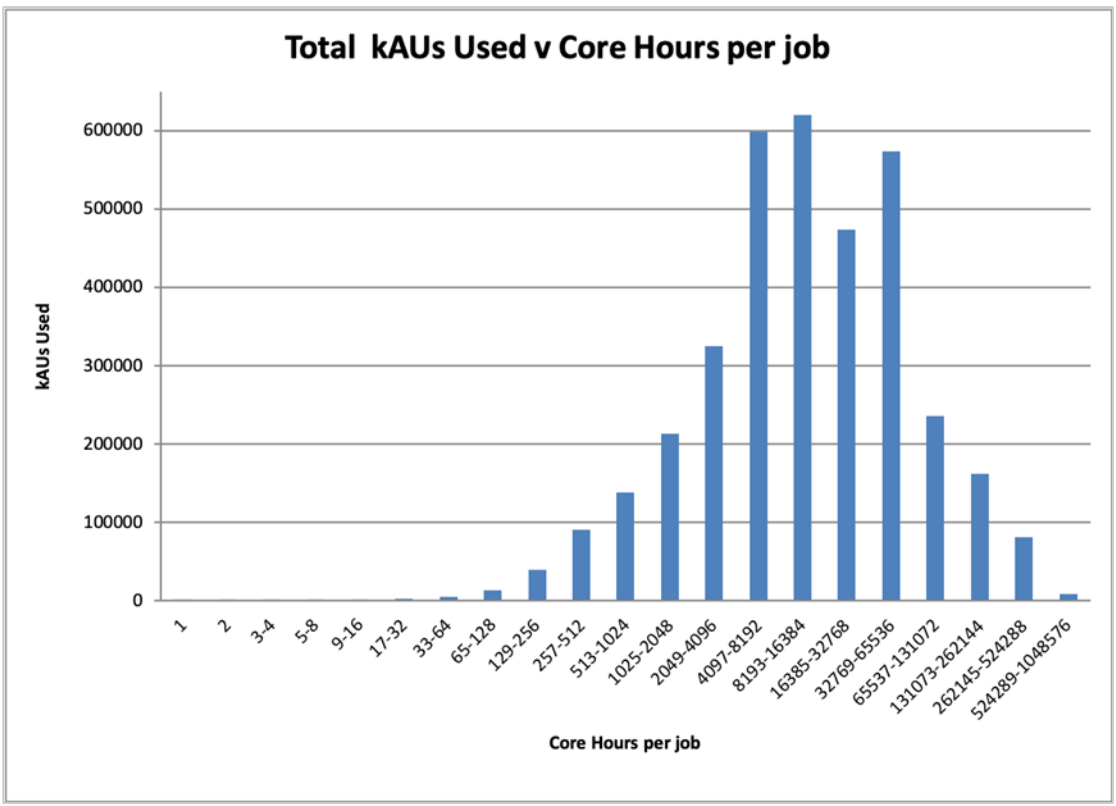


The above graphs show that, while there are quite a few jobs that use only a small number of core hours per job, most of the resource is consumed by jobs that use tens of thousands of core hours per job.