



ARCHER SP Service Quarterly Report

Quarter 4 2014



Document Information and Version History

Version:	1.0
Status	Final
Author(s):	Alan Simpson, Jo Beech-Brandt, Andy Turner, Mike Brown, Stephen Booth
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2015-01-05	Initial Draft	Liz Sim, Alan Simpson, Jo Beech-Brandt, Stephen Booth, Andy Turner
0.2	2015-01-12	Updates and additions	Jo Beech-Brandt
0.3	2015-01-13	Review	Alan Simpson
1.0	2015-01-16	Final version for EPSRC	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:

October 2014, November 2014 and December 2014

- Utilisation on the system during 14Q4 was 85%, compared to 74% in 14Q3.
- All service levels were met during the period.
- The System was upgraded during this period and the upgrade was completed on schedule. A photo blog was updated regularly to allow users to follow the progress of the upgrade.
- The batch system configuration was revised to include a short queue for short jobs such as debugging tasks. This uses a standing reservation so that the nodes can be guaranteed to be available during operational hours as they will not be allocated to other tasks this had not been possible with the previous debug queue.
- Public-facing web pages have been created to help demonstrate the impact of ARCHER and encourage new users to apply for access including a dedicated Industry section.
- Collaborated with CSE to produce the mailing list for the ARCHER Image Competition calendar; this was sent out to key ARCHER users and stakeholders before Christmas.
- We assisted in the promotion of *The Story of Hector* user survey.
- Two ARCHER SP staff attended the Machine Evaluation Workshop organised by Daresbury.
- SAFE functionality was extended to allow a Project Manager to be assigned to manage a sub-group project within a larger project. This had been requested by PIs of large groups.
- On EPSRC's request, automatic mailings from the SAFE to Heads of Consortia have been implemented. This includes a graph of allocated resources plotted against used resources to assist PIs and Consortia heads in managing their allocations.
- The ARCHER Annual User Survey was opened to users on 8 December 2014 and closed on 7 January 2015. 152 responses were received (more than 10% of the active user base):
 - Users were asked to rate and comment on the following aspects of the service: overall satisfaction, hardware, software, helpdesk, documentation, website, training, webinars.
 - An initial pass of the data reveals a high level of satisfaction with the service: the mean score for overall satisfaction is 4.4/5 and all other aspects show a similar high level of satisfaction. This supports the anecdotal evidence we currently have regarding user perceptions of the service.
 - A large number of positive comments and suggestions were also received revealing that users are engaged with the service and feel that they can contribute positively to its development.
 - A full analysis of the results will be undertaken (see Forward Look section below) to gain as much benefit from the survey as possible.

1.2 Forward Look

- Cray Linux Environment (CLE) will be upgraded from CLE 5.1 to CLE 5.2.
 - We will advise users to recompile their codes.
- RSIP node implementation will allow the use of applications requiring licence servers on the compute nodes, including compilations and ISV codes.
- At the request of SMB, we are planning to run a Workshop for ARCHER “Champions.”
- Implementation of Periodic allocations for consortia and large users.
- Progress the collection of data describing ARCHER users’ career stage
- Further analysis and compilation of results from the ARCHER Annual User Survey Report.
- Complete an end-of-lifecycle ARCHER policy for former users.
- There is ongoing work on improving SAFE through the SAFE Development project.
Upcoming targets include:
 - Enabling UK-Federation authorisation to the SAFE. This will give users the option of using the same authorisation credentials as they use in their home institution.
 - Providing greater support for Technical Assessments of new project applications within the SAFE
 - Announce 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
- ARCHER Annual User Survey results analysis:
 - SP will lead the analysis of the responses involving all Service Partners.
 - A short paper analysing the results will be provided along with the Annual Report and published on the ARCHER website
 - Comments and suggestions will be reviewed and a prioritised list of possible improvements presented to the ARCHER Operations Meeting for discussion and possible implementation.

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 receive requests which 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 14		Nov 14		Dec 14		14Q4
	Service Level	Service Points	Service Level	Service Points	Service Level	Service Points	Service Points
2.6.2 - PR	100.0%	-5	100.0%	-5	100.0%	-5	-15
2.6.3 - QC	98.8%	-1	98.9%	-1	98.8%	-1	-3
2.6.4 - UR	1 WD	0	1 WD	0	1 WD	0	0
Total		-6		-6		-6	-18

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

2.1.2 Service Failures

Type	Severity	Service Points	Notes
Unscheduled Maintenance	Sev-1	None	Maintenance started 30 minutes early on 5 th Nov

There was one Service Failure in the Service Quarter. A planned maintenance session started earlier than had been pre-approved by the Authority. Details are available in Section 2.3.2 of this report.

2.1.3 Service Credits

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



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 (-18), no Service Credits apply in 14Q4.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Oct 14	Nov 14	Dec 14	14Q4
Phone Calls Received	35 (10)	36(6)	26(5)	97(21)
Answered 2 Minutes	35	36	26	97
Service Level	100.0%	100.0%	100.0%	100.0%

The volume of telephone calls remained low in 14Q4. Of the 97 calls received in total above, only 21 were genuine ARCHER user calls that resulted in queries.

2.2.2 Query Closure (QC)

	Oct 14	Nov 14	Dec 14	14Q4
Self-Service Admin	373	352	246	971
Admin	168	146	138	452
Technical	56	25	36	117
<i>Total Queries</i>	<i>597</i>	<i>523</i>	<i>420</i>	<i>1540</i>
<i>Total Closed in 2 Days</i>	<i>590</i>	<i>517</i>	<i>415</i>	<i>1522</i>
Service Level	98.8%	98.9%	98.8%	98.8%

In addition to the Admin and Technical queries, the following Change Requests were resolved in 14Q4.

	Oct 14	Nov 14	Dec 14	14Q4
Change Requests	6	4	2	12

2.2.3 User Registration (UR)

	Oct 14	Nov 14	Dec 14	14Q4
No of Requests	95	86	67	248
Closed in One Working Day	94	80	67	241
Average Closure Time (Hrs)	1.30	3.08	1.48	1.95
Average Closure Time (Working Days)	0.14	0.32	0.16	0.21
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:

	Oct 14	Nov 14	Dec 14	14Q4
CRAY	13	9	8	30+
ARCHER_CSE	151	79	73	303
ARCHER_SP	953	870	664	2487
Total Queries Assigned	1117	958	745	2820
Total Assigned in 1 Hour	1116	958	740	2814
Service Level	100.0%	100.0%	99.0%	99.6%

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.

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
22/10/14	0900	1617	5 hrs 17mins	Permitted	EPSRC Approved 0900-1700	Planned Maintenance
05/11/14	*0830	1305	4 hrs 35mins	Permitted	EPSRC Approved 0900 - 1700	Phase 2 Preparation

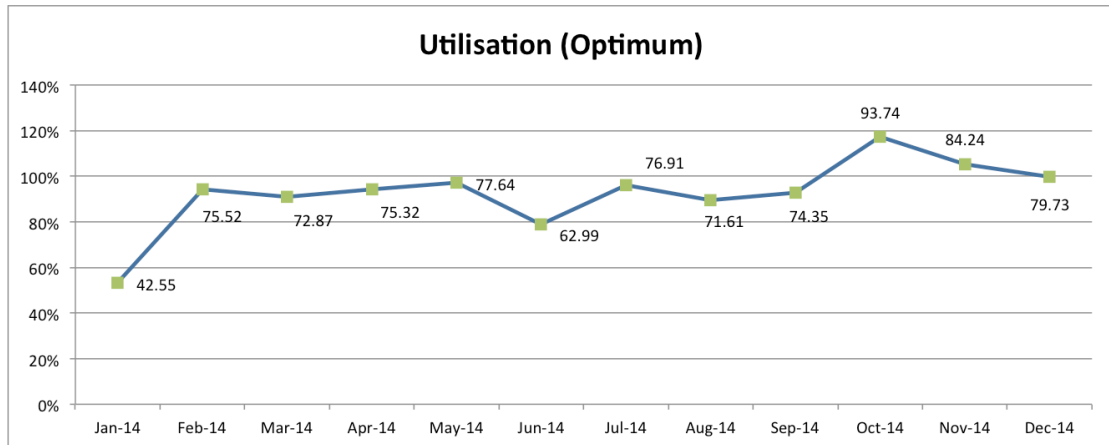
*The maintenance for the Phase 2 Preparation work was approved by the Authority to start at 0900 but work started at 0830. It was completed well within the approved 8 hours.

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 84.8% (or 106.0% of optimum).




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 slightly overused their target allocation, whereas EPSRC underused their allocation.

3.2 Scheduling Coefficient Matrix




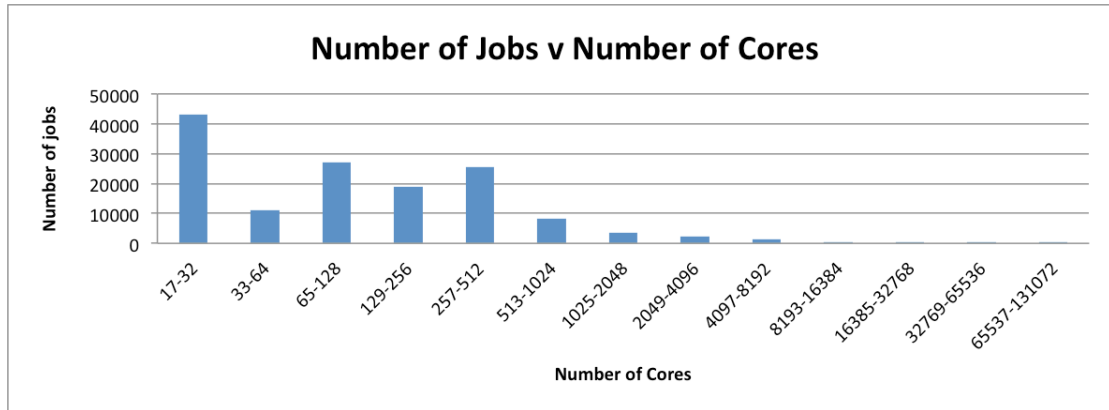
The colour in the matrix indicates the value of the Scheduling Coefficient. This is defined as the ratio of runtime to runtime plus wait time. Hence, a value of 1 (green) indicates that a job ran with no time waiting in the queue, a value of 0.5 (pale yellow) indicates a job queued for the same amount 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). The green values for the very largest jobs (4097-8192 nodes) arise from the Special User Testing period during the Phase 2 upgrade where only a single user was allowed access to the queues at any one time and so the wait to start was effectively zero.

3.3 Additional Usage Graphs


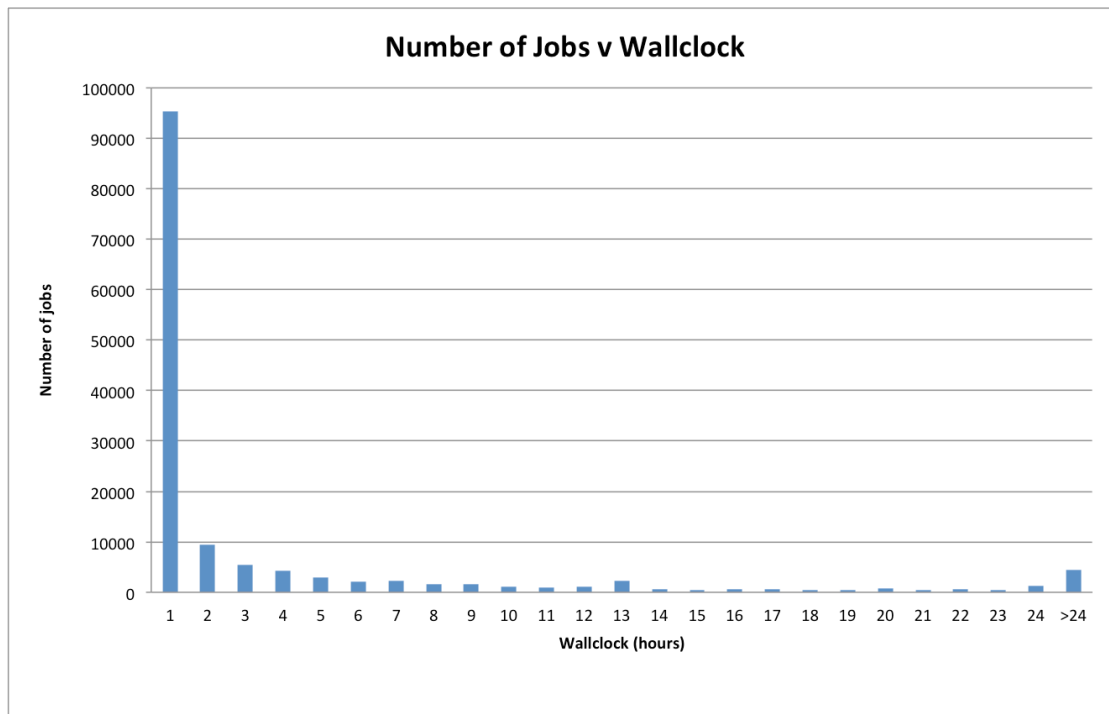
The following charts provide different views of the distribution of job sizes on ARCHER.

Number of Cores



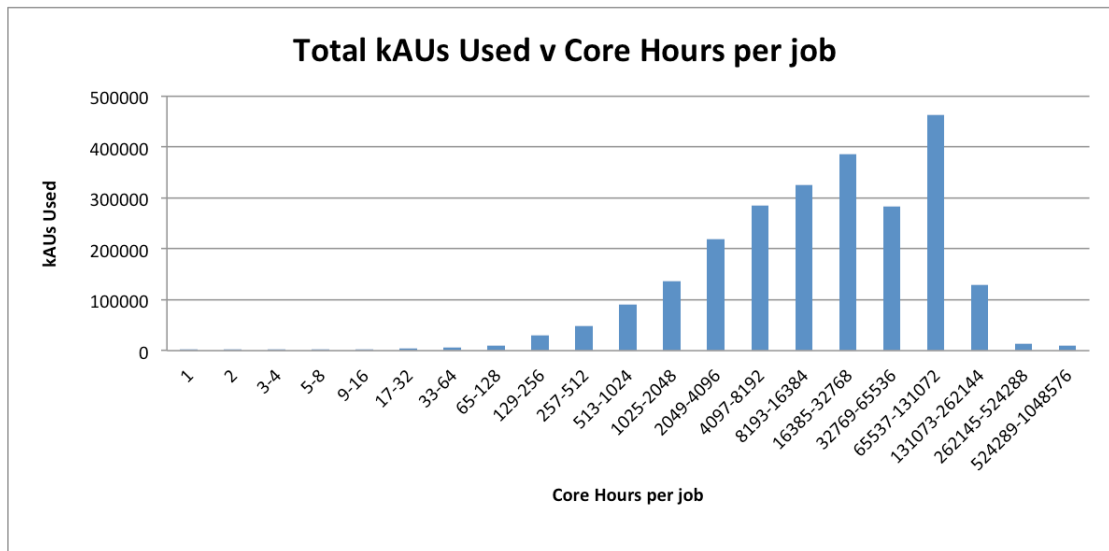
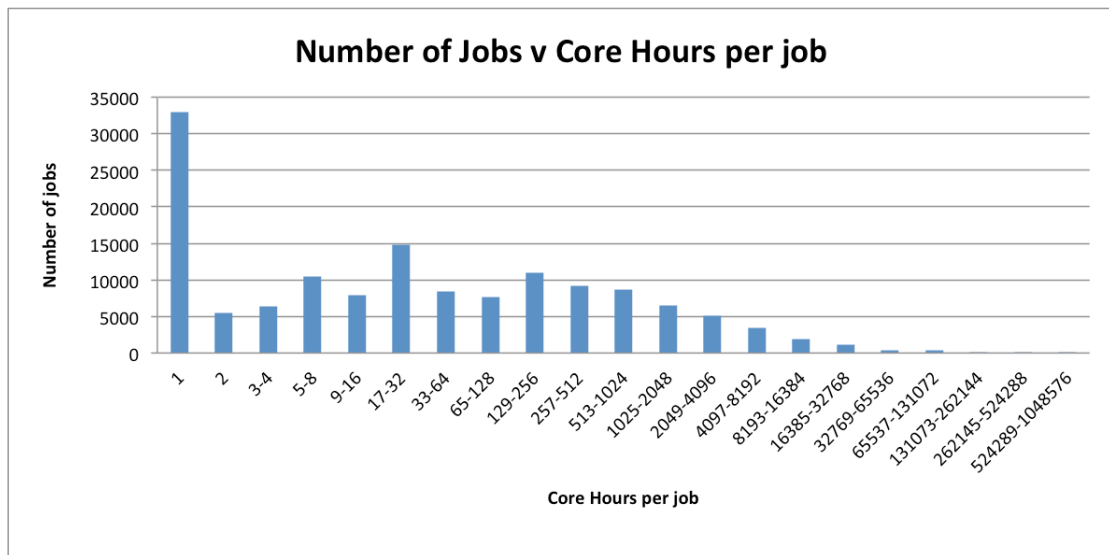
The first graph shows that, in terms of numbers, there is a significant number of jobs using no more than 512 cores. However, the second graph shows that most of the kAUs were spent on jobs between 257 cores and 8192 cores. The number of kAUs used is closely related to money and shows better how the 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 the kAUs are fairly evenly distributed across jobs of all lengths. (Due to testing of the Phase 2 system, a number of longer test jobs were run contributing to the spike in jobs of >24 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.