



ARCHER SP Service Quarterly Report

Quarter 2 2017



Document Information and Version History

Version:	1.0
Status	Draft
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	14/06/17	Initial Draft	Anne Whiting
0.2	03/07/17	Added Graphs	Jo Beech-Brandt
0.3	05/07/17	Added phone and service metrics information	Jo Beech-Brandt
0.4	06/07/17	Added Infrastructure report	Mike Brown, Jo Beech-Brandt
0.5	06/07/17	Reviewed	Alan Simpson
1.0	10/07/17	Final version	Anne Whiting

1. The Service

1.1 Service Highlights

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

April 2017, May 2017 and June 2017.

- Utilisation on the system during 17Q2 was 89% as compared to 90% in 17Q1. The continued high utilisation of the service supports the need for ongoing investment in HPC.
- The CLE upgrade to 5.2 UP04 was successfully carried out on 12th and 26th April 2017. This will improve the user environment and provide access to current software and additional functionality.
- In collaboration with the CSE service and Cray we have used SAFE to combine multiple data streams providing information on resources used by applications running on ARCHER. Using SAFE we can now correlate application used, memory used, energy used, with standard SAFE properties such as user, project, research area. Combining data in this way allows us to answer questions about the use of the service that are not possible from the isolated data sources.
- On-site Cray personnel developed diagnostic software for the Lustre filesystem which records a high level of detail about application usage, to help in diagnosing and investigating potential performance problems. OSG integrated this into their monitoring and metrics systems and produced a dashboard which enables all service partners to see performance bottlenecks in real-time, so issues take less time to detect and less time to resolve.
- A joint Tier 1 / Tier 2 Champions event ARCHER Champions Workshop took place on Monday June 26th and Tuesday June 27th at Hartree, Daresbury.
- Tier 2 administration has been integrated into the ARCHER SAFE. The first Instant Access applications for Cirrus have been reviewed and projects set up. The Tier 2 Technical Working group has been established to guide the Tier 2 infrastructure across the UK.
- Weekly digest user mailings communicating relevant service highlights and upcoming training courses are now being sent out every Tuesday with the mailing content available in the ARCHER website News. The mailings help to ensure the user community has a consistent method of keeping up to date with upcoming events and news, and positive user feedback has been received.

1.2 Forward Look

- The HPC-UK website is under development, providing information for users and providers on UK HPC facilities. This includes information on facilities available, how to gain access, what training courses are available, and links to other useful resources.
- Work is ongoing with Cray and NCAS on the repurposing of a compute node as 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 Unified Model code. The functionality is currently being tested on the TDS.
- We will repeat the scheduler analysis that we undertook throughout summer 2016 to ensure that the scheduler configuration remains well tuned to the ARCHER workload.

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 17		May 17		Jun 17		17Q2
	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	99.6%	-2	98.6%	-1	99.8%	-2	-5
2.6.4 – UR	1 WD	0	1 WD	0	1 WD	0	0
Total		-7		-6		-7	-20

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 can be found in Section 2.3.2.

2.1.3 Service Credits

As the Total Service Points are negative (-20), no Service Credits apply in 17Q2.

2.2 Detailed Service Level Breakdown

2.2.1 Phone Response (PR)

	Apr 17	May 17	Jun 17	17Q2
Phone Calls Received	37 (8)	36 (5)	42 (5)	115 (18)
Answered 2 Minutes	37	36	42	115
Service Level	100.0%	100.0%	100.0%	100.0%

The volume of telephone calls remained low in 17Q2. Of the total of 115 calls received above, only 18 were actual ARCHER user calls that either resulted in queries or answered user questions directly.

2.2.2 Query Closure (QC)

	Apr 17	May 17	Jun 17	17Q2
Self-Service Admin	570	481	970	2021
Admin	141	128	142	411
Technical	14	18	19	51
<i>Total Queries</i>	725	627	1131	2483
<i>Total Closed in 2 Days</i>	722	618	1129	2469
Service Level	99.6%	98.6%	99.8%	99.4%

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 17Q2:

	Apr 17	May 17	Jun 17	17Q2
Change Requests	0	1	2	3

2.2.3 User Registration (UR)

	Apr 17	May 17	Jun 17	17Q2
No of Requests	94	99	112	305
Closed in One Working Day	92*	99	112	303
Average Closure Time (Hrs)	1.0	0.6	0.3	0.6
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 requests were not able to be processed in one day due to ARCHER EPSRC approved maintenance session on 26/04/17.

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:

	Apr 17	May 17	Jun 17	17Q2
CRAY	8	11	13	32
ARCHER_CSE	188	202	121	511
ARCHER_SP	1192	1454	1599	4245
Total Queries Assigned	1388	1667	1733	4788
Total Assigned in 1 Hour	1388	1667	1729	4784
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 a single day each month (fourth Wednesday of each month). This is marked as a full 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 once per month and not twice. It also eases the planning for training courses running on ARCHER.

Such Maintenance Periods are defined as 'Permitted Maintenance ' and recorded in the Maintenance Schedule. A 6-month forward plan of maintenance has been agreed with 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
12/04/17	0900	1259	3 hrs 59 mins	Full Outage	EPSRC Approved 0900 – 1700	System Set Backup in preparation for CLE5.2UP04 Upgrade
26/04/17	0800	1848	10 hrs 48 mins	Full Outage	EPSRC Approved 0800 – 2000	CLE5.2UP04 Upgrade
24/05/17	0900	1359	4 hrs 59 mins	Full Outage	EPSRC Approved 0900 – 1700	CLE5.2UP04 Upgrade
28/06/17	0900	1404	4 hrs 4 mins	Full Outage	EPSRC Approved 0900 – 1700	CLE Field Notices

2.3.3 Quality Tokens


No quality tokens have been 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 89%. 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 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 both Research Councils did not meet their respective targets this quarter with EPSRC being at 71.6% (against their target of 77%) and NERC's utilisation being 21.6% (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.



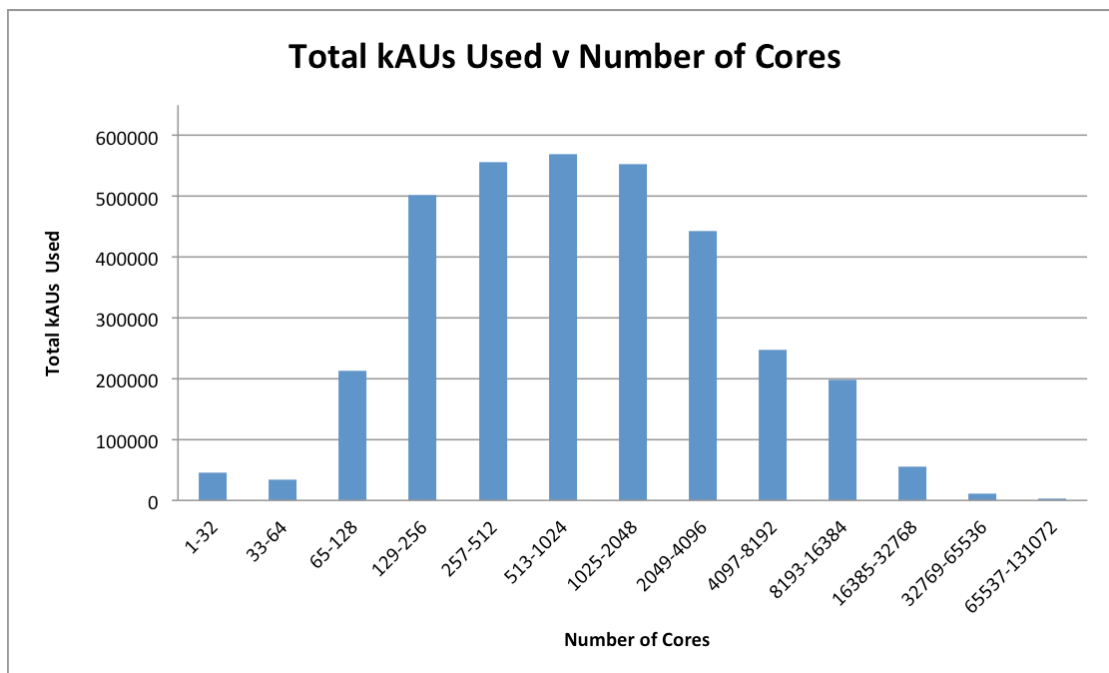

3.3 Additional Usage Graphs

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 kAU 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.

Appendix – Infrastructure report

The University of Edinburgh has nothing to report under the site Lease Agreement regarding infrastructure work within this quarter.

The University is planning a full site power outage for High Voltage trip testing and the other statutory safety checks. The Authority will be notified of the planned date and duration with appropriate notice.