



ARCHER2

SP Quarterly Report

October - December 2023
EPCC
The University of Edinburgh



Document Information and Version History

Version:	1.0
Status	Release
Author(s):	Clair Barrass, Jo Beech-Brandt, Stephen Booth, Paul Clark, Daniel Gleed, Kieran Leach, Alan Simpson, Anne Whiting
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	29/12/2023	Template created	Jo Beech-Brandt
0.2	04/01/2023	Added heatmap and usage graphs	Clair Barrass
0.3	06/01/2023	Added narrative, metrics and graphs	Jo Beech-Brandt
0.4	05/01/2024	Added narrative	Anne Whiting
0.5	08/01/2024	Added Critical Success Factors	Alan Simpson
0.6	09/01/2024	Added comments	Paul Clark
0.7	12/01/2024	Reviewed	Alan Simpson
1.0	15/01/2024	Version for UKRI	Alan Simpson, Jo Beech-Brandt

1 The ARCHER2 Service

This is the report for the ARCHER2 SP Service for the Reporting Period: 1st October – 31st December 2023.

1.1 Service Highlights

- The AMD GPUs for ARCHER2 have arrived and upcoming maintenance sessions have been advertised to users to allow their integration during January.
- To enhance security, multi-factor authentication (MFA) was introduced on ARCHER2. User documentation was updated, and users are now asked to connect using both a SSH key passphrase and a time-based on time passcode (TOTP).
- Due to contention on the existing three file systems, a decision was taken to deploy the final file system fs4, as an additional file system rather than a /scratch file system. The NVMe file system, is now available to all projects as a /scratch file system when requested and the team are also proactively working with user groups to migrate to the /scratch file system if projects have a heavy i/o load.
- A number of ARCHER2 staff undertook ITILv4 Foundation training. The ITIL Framework is used to ensure good practise across service delivery and service management.
- ARCHER2 was represented at both SuperComputing '23 in November and CIUK '23 which took place in Denver and Manchester respectively.
- All the ARCHER2 chillers, located in in plant room C (prC), were cleaned over the quarter.
- All transformer meters swapped out in prC with additional training required to allow EPCC staff to interrogate new information on power quality.
- Computer Room 3 (cr3) Air Conditioning (CRAC) units recommissioned to balance the load in the room and improve efficiency.
- An Annual Change Freeze was agreed via the Change Advisory Board (CAB) and was in place during the festive period, this was noted as a heightened period of scrutiny and awareness. The Service remained open during the festive period with the exemption of the agreed public holidays (25th, 26th December and 1st January).

1.2 Forward Look

- The GPU nodes will be integrated into ARCHER2 and users are expected to have access to the GPU nodes during Q1 2024. Documentation has been produced and a user forum is planned at the end of January, with both HPE and EPCC staff to advise and answer any questions regarding the GPUs.
- Auto-deletion of files older than 28 days will be enabled on the /scratch file systems for the scratch NVMe file system. Users will be informed of this policy; documentation will be updated, and it will then be deployed on the service.
- Preparation is underway for the migration of our Information Security Management System from the previous version of the ISO 27001 standard to the new one. This will help ensure that the latest best practice is applied to the secure handling of user data as the information security landscape has changed somewhat since 2013, and the new version places a much greater emphasis on cyber security risks and preventive measures.
- Mechanical and Electrical design received to address multi point plan developed for prC in Q4 2023. Programme planning ongoing for changes throughout 2024.

2 ARCHER2 Performance Report

This is the contractual performance report for the ARCHER2 SP Service for the Reporting Periods from 1st October 2023 until 31st December 2023.

2.1 Service Points and Service Credits

The Service Levels and Service Points for the SP service are defined by EPSRC in Schedule 2.2 of ARCHER2 SP Service Contract.

The Working Day (WD) for the ARCHER2 Service is 10 Working Hours (WH) as the Service operates from 0800-1800. The Median Time to Resolution is measured in WD.

- **Availability:** *Service Threshold: <=96.5%; Operating Service Level: >98.0%, ≤ 98.5%.*
- **ARCHER2_SP_Level1 (MTR):** The Median Time to Resolution, of all SP queries falling within Level 1 resolved by the Contractor in the Reporting Period. *MTR Service Threshold: >1 WD; Operating Service Level: >0.3 WD, ≤ 0.45 WD.*
- **ARCHER2_SP_Level2 (MTR):** The Median Time to Resolution, of all SP queries falling within Level 2 resolved by the Contractor in the Reporting Period. *MTR Service Threshold: >8 WD; Operating Service Level: >2 WD, ≤4 WD.*
- **ARCHER2_SP_Level3 (MTR):** The Median Time to Resolution, of all SP queries falling within Level 3 resolved by the Contractor in the Reporting Period. *MTR Service Threshold: >25 WD; Operating Service Level: >12 WD, ≤16 WD.*
- **Initial Response to Queries (%):** The percentage of the total number of SP queries assigned to the Contractor in the Reporting Period responded to within 3 Working Hours. *Service Threshold: <96.00%; Operating Service Level: 98.00 – 98.99%.*
- **Query User Satisfaction (%):** The percentage of the total number of query satisfaction surveys completed in each Reporting Period, rating the quality of the resolution of Queries by the Contractor as “Good”, “Very Good” or “Excellent”. *Operating Service Level: 82.00 – 87.99%*

2.1.1 Service Points

Metric	Oct 2023		Nov 2023		Dec 2023		Q 2023	
	Perf	Points	Perf	Points	Perf	Points	Perf	Points
Availability	100%	-3	100%	-3	100%	-3	100%	-9
SP_Level1 (MTR)	0.00	-2	0.00	-2	0.00	-2	0.00	-6
SP_Level2 (MTR)	0.07	-2	0.09	-2	0.03	-2	0.06	-6
SP_Level3 (MTR)	0.00	-2	6.36	-2	2.20	-2	4.28	-6
Initial Response (%)	100%	-1	100%	-1	100%	-1	100%	-3
Query Satisfaction (%)	100%	-2	100%	-2	100%	-2	100%	-6
Total		-12		-12		-12		-3

2.1.2 Service Credits

As the Total Service Points are negative (-36), no Service Credits apply in 23Q4.

2.2 SP Query Statistics

The metrics were specified by EPSRC in Schedule 2.2 of ARCHER2 SP Service Contract.

- **Assigned:** The number of SP queries assigned to the Contractor within each query resolution category in the Reporting Period.
- **Resolved:** The number of SP queries resolved by the Contractor within each query resolution category in the Reporting Period.
- **Backlog:** The number of SP queries assigned to the Contractor that remained unsolved within each query resolution category in the Reporting Period
- **Correspondence:** The average number of pieces of correspondence generated for SP queries in each query resolution category.
- **First Response:** The average time taken for the Contractor to first respond to the Originator of the SP query.

October 2023					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
SP_Level1	859	859	0	0.19	0:00:52
SP_Level2	93	94	13	7.80	0:16:33
SP_Level3	0	0	1	0	0:00:00
November 2023					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
SP_Level1	1,157	1,157	0	0.13	0:01:45
SP_Level2	78	76	15	8.11	0:21:33
SP_Level3	1	1	1	41	0:06:43
December 2023					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
SP_Level1	688	688	0	0.11	0:00:50
SP_Level2	107	98	24	6.714	0:13:36
SP_Level3	1	1	1	14	0:21:24
Q4 2023					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
SP_Level1	2704	2704	0	0.15	0:01:12
SP_Level2	278	268	24	7.59	0:16:53
SP_Level3	2	2	1	27.5	0:14:03

2.3 Query Resolution

Metric	Oct 2023		Nov 2023		Dec 2023		Q4 2023	
Service Level	MTR	Resolved	MTR	Resolved	MTR	Resolved	MTR	Resolved
SP_Level1	0:00:09	859	0:00:58	1157	0:00:17	688	0:00:33	2704
SP_Level2	0:44:13	94	0:55:41	76	0:19:47	98	0:34:23	268
SP_Level3	0:00:00	0	63:34:21	1	22:02:34	1	42:48:27	2
Total		953		1234		787		2974

A total of 2974 queries were resolved by the ARCHER2 SP Service in the Reporting Period. The percentage of user queries responded to within 3 hours was 100%.

2.4 Query Feedback

During October, there were 38 feedback scores received during this period. 100% were Good, Very Good or Excellent with 74% given the highest score of Excellent.

During November, there were 35 feedback scores received during this period. 100% were Good, Very Good or Excellent with 69% given the highest score of Excellent.

During December, there were 32 feedback scores received during this period. 100% were Good, Very Good or Excellent with 75% given the highest score of Excellent.

£105 donation was made to our chosen charity Save the Children with £1 donated per query feedback item received.

2.5 Maintenance and Outages


Type	Start	End	Duration	User Impact	Reason	Attributable
Unplanned	2023-10-19 07:45	2023-10-19 10:30	2 hours 45 mins	Three cabinets lost power and some user jobs failed. Jobs were not charged.	Power incident	Accommodation
Unplanned	2023-11-06 09:20	2023-11-06 11:40	2 hours 20 mins	All ARCHER2 cabinets lost power. All running jobs failed.	Power incident	Accommodation

3 ARCHER2 Service Statistics

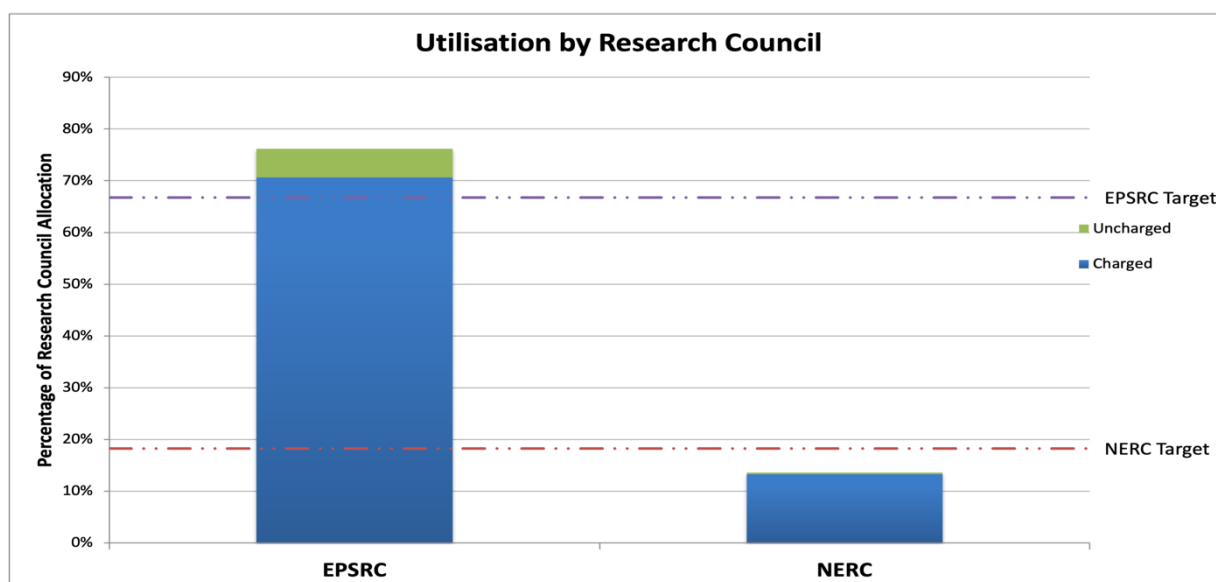
3.1 Utilisation

Utilisation from 1st October – 31st December is 91% which is slightly increased from 89% the previous quarter. Utilisation for October was 89%, for November 93% and for December 92%.

Please note the festive period took place during this period. This had little impact on the utilisation as changes were made to the QoS layout on the service to allow users to submit a larger number of jobs to the service.




The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER2 by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER2. It can be seen that EPSRC exceeded their target this quarter with their usage being at 76.2% (against their target of 66.8%). It should be noted that the proportion of EPSRC's uncharged utilisation has reduced and we expect that this is due to the increase in allocation of ARCHER2 with the uncharged portion being 5.5% during this quarter compared to 25.9% in the previous quarter. NERC missed their target with utilisation being 13.6% (against their target of 18.2%) though this is still an increase from 10.2% in the previous quarter.



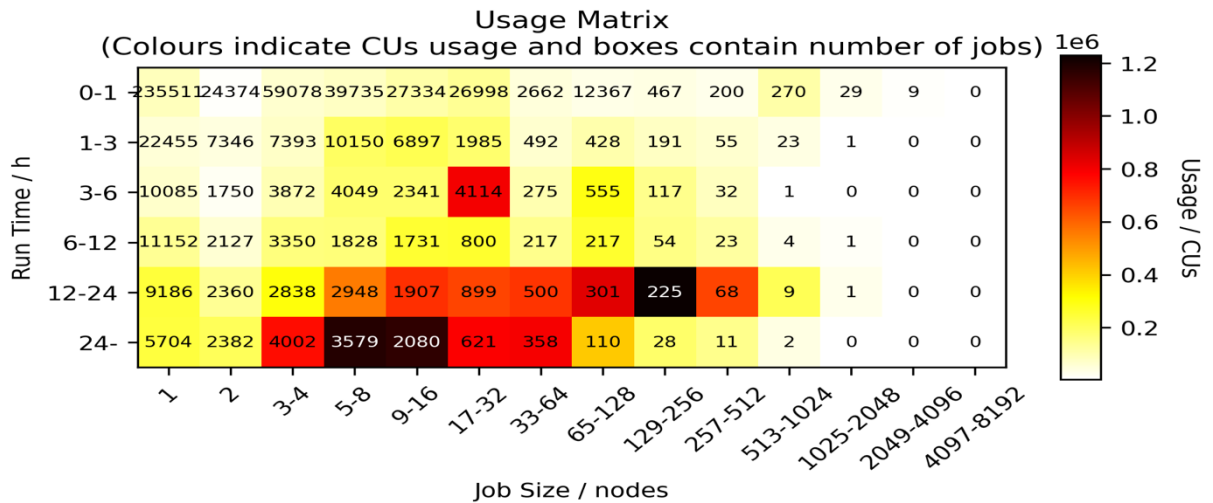
3.2 Scheduling Coefficient Matrix

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 usage heatmap below provides an overview of the usage on ARCHER2 over the quarter for different job sizes/lengths. The colour in the heatmap indicates the number of CUs expended for each class, and the number in the box is the number of jobs of that class.

It should be noted that there was an increase in the number of larger sized jobs during this quarter as users were encouraged to submit larger jobs during the data centre network maintenance session.



Appendix: Critical Success Factors

1. Context

EPCC have been asked by UKRI to provide quarterly data for a number of critical success factors:

- CSF04 Implementation of environmentally considerate energy policies
- CSF07 Deliver and maintain a reliable data I/O function
- CSF08 Be cost-effective, cost-efficient and drive towards lowering of operational costs

In the sections below, please find the relevant metrics and data.

2. CSF04 Implementation of environmentally considerate energy policies

Implementation of environmentally considerate energy policies with a drive to reducing costs and environmental impacts.

All electricity provided to the ACF and ARCHER2 is on a 100% green, renewable energy tariff.

Environmentally considerate policies: 3

Since the start of full Service, EPCC have worked on implementing the following policies:

- Move from High Performance Mode to Low Power Mode: reduced average power draw from 3.2 MW to 2.9 MW (9%) with negligible input on performance [May 2022]
- Reduced default processor frequency: further reduced average power to around 2.5 MW (19%) [December 2022]
- Increase in coolant temperatures: this will result in an increase in passive cooling (“free cooling”) [ongoing]

Power Usage

	4Q21*	1Q22	2Q22	3Q22	4Q22	1Q23	2Q23	3Q23	4Q23
Average Power	3.31	3.16	3.15	2.86	2.90	2.51	2.56	2.46	2.53

* Partial

So far, the average power draw has been reduced by around 0.7MW (21%) which will reduce electricity usage by up to 6M kWh per annum, significantly reducing annual running costs.

3. CSF07 Deliver and maintain a reliable data I/O function


The compute resource will deliver and maintain an efficient, effective and reliable data I/O function which meets the requirements of users and their software. It will evolve and expand to accommodate new software or hardware architectures as required by the Service or its user base.

Data Transferred

EPCC monitor the data transfer rates in and out of the ARCHER2 system. Based on this, we now estimate the total amount of data transferred to and from ARCHER2 each Quarter.

Data Transferred...	4Q 21*	1Q 22	2Q 22	3Q 22	4Q 22	1Q 23	2Q 23	3Q 23	4Q 23
... to ARCHER2 (TB)	534	163	68	220	44	67	42	65	99
...from ARCHER2 (TB)	236	582	667	822	834	1231	1022	1472	1771

* Partial



Parallel IO Write Performance


We regularly monitor the parallel write performance between the compute nodes and the parallel Lustre (/work) file systems. We use the benchio synthetic IO benchmark application (<https://github.com/davidhenty/benchio>) and report the MPI-IO write performance with the following settings:

- Global data structure of 20483: writes a single file of 65,536 MiB (64 GiB).
- Uses 16 compute nodes and 128 MPI processes per node.
- Uses UCX as the MPI transport protocol.
- Sets the following environment variables:
 - FI_OFI_RXM_SAR_LIMIT=64K
 - MPICH_MPIIO_HINTS="*:cray_cb_write_lock_mode=2,*:cray_cb_nodes_multiplier=4

These settings have been found to maximise the IO performance for parallel writes using MPI-IO on the ARCHER2 file systems. Writes using the default settings on ARCHER2 typically have median write values 2-3 GiB/s lower than the optimised values.

Original reporting of this data (Q1 and Q2 2023) used the means from a small number of runs on the HDD-based Lustre file systems. From Q3 2023 onwards we have been monitoring performance regularly on both HDD and NVMe-based Lustre file systems throughout the quarter and report median (Q2) and lower (Q1) and upper quartile (Q3) performance, as well as providing a quarterly boxplot that illustrates the performance variation.

Benchio MPI-IO medium (GiB/s)	1Q23	2Q23	3Q23	4Q23
a2fs-work1	8.2	7.6±0.5	10.5 (8.8:11.8)	10.9 (8.3:12.5)
a2fs-work2	8.5	7.3±0.6	10.4 (7.2:12.4)	10.4 (7.7:13.0)
a2fs-work3	8.3	9.6±0.7	10.0 (8.2:11.6)	10.7 (8.1:11.9)
a2fs-work4				9.7 (9.1:10.2)
a2fs-nvme			10.1 (9.6:11.5)	10.1 (9.5:12.4)



4. CSF08 Be cost-effective, cost-efficient and drive towards lowering of operational costs

The Service shall be cost-effective and cost-efficient across its elements during its lifetime and drive towards lowering of operational costs by seeking efficiencies in delivery such that TCO presents an acceptable and cost-effective solution for the public. The Service will monitor and report its Power Usage Effectiveness (PUE) and strive to make efficiency savings where possible.

Relative Research Output

Measure	11/2021 – 5/2022	5/2022 – 12/2022	1Q23	2Q23	3Q23	4Q23
Relative Research Output per kWh	100	109	115	115	115	115

We define the initial measure of research output per kWh on ARCHER2 to be 100, and then estimate how this has changed with the introduction of the various environmentally considerate policies discussed under CSF04. This is estimated using applications benchmarks similar to those defined by UKRI for the procurement.

Energy Used per CU Delivered

	4Q21*	1Q22	2Q22	3Q22	4Q22	1Q23	2Q23	3Q23	4Q23*
Energy per CU (kWh)	0.719	0.713	0.728	0.715	0.650	0.545	0.669	0.590	0.558

*partial

Energy Cost per CU Delivered

	4Q21*	1Q22	2Q22	3Q22	4Q22	1Q23	2Q23	3Q23	4Q23*
Cost per CU (£)	£0.089	£0.090	£0.098	£0.096	£0.088	£0.074	£0.162	£0.143	£0.134

*partial

The two tables above are calculated using the total CUs delivered by ARCHER2, the total kWh of electricity consumed, and the unit cost for kWh. The increase in “Energy Cost per CU Delivered” from 2Q23 is caused by a significant increase in the unit cost of electricity from April 2023. For 2Q23, there is also an impact on the “Energy Used per CU Delivered” from the major software upgrade that took 3 weeks.