ARCHER2
SP Quarterly Report

April - June 2025

EPCC
The University of Edinburgh

Document Information and Version History

Version:	1.0
Status	Release
Author(s):	Clair Barrass, Jo Beech-Brandt, Alan Simpson, Anne Whiting
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	16/06/2025	Template created	Jo Beech-Brandt
0.2	01/07/2025	Added narrative, graphs, service statistics	Jo Beech-Brandt
0.3	02/07/2025	Added critical success metrics	Lorna Smith
0.4	02/07/2025	Added ISO related information	Anne Whiting
0.5	09/07/2025	Reviewed	Alan Simpson
1.0	09/07/2025	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: 1 April – 30 June 2025.

1.1 Service Highlights

- Following EPSRC and UKRI discussions regarding the spending review outcome, a one-year extension of the ARCHER2 SP and CSE contracts was agreed and users were notified. The new end date for the service is 21st November 2026.

- Utilisation remains high on ARCHER2 with usage over the quarter at 91%. We are awaiting the approved applicants from the recent Access to HPC Call, and projects will then be set up and new users onboarded.

- EPCC have worked with HPE to test and demonstrate an alternative to the time-consuming standard software update process for ARCHER2, which would allow us to keep up to date for software patches while minimising downtime. Key users from the various Consortia have now been involved in testing the updated software on the main system.

- ARCHER2 staff attended the International SuperComputing Conference in Hamburg. An EPCC booth was able to promote the ARCHER2 service using the outreach Raspberry Pi cluster, Wee Archie. Staff were able to attend sessions, meetings and participated in various sessions.

- A User Forum took place within the ARCHER2 Celebration of Science. Users were able to hear about the latest news and developments from the service and participate in an open discussion with staff. Presentations were on: the extended shutdown which is required for the replacement of switchgear in the machine room; training provision; the operating system update; and the latest work the ARCHER2 team are leading with on green activities.

- At the request of the User Advisory Group, further Capability Days were planned for 17th – 19th June. The days were advertised to users and a failsafe measure was put in place to return the nodes to service if the nodes were idle for more than 6 hours. Unfortunately, this was the case and the nodes were returned to normal service at the end of the first day. We will continue to plan further Capability Days, endeavouring to increase the promotion of the days with the users and fully documenting the failsafe measure.

- We are pleased to announce we have been recommended for ISO re-certification after a very successful recent external audit. We had full re-certification audits for ISO 9001 quality service delivery and ISO 22301 business continuity and disaster recovery. We also had a surveillance audit for ISO 27001 information security. Our continued investment in ISO certification helps to ensure that we use best practice to deliver the highest quality of service to our user community and to the science they deliver.

- We successfully invoked the hot weather processes for the first heatwave of the summer on 20th June to protect both ARCHER2 and the plant infrastructure, and then wrote this up as a report.

1.2 Forward Look

- EPCC are working with HPE towards the deployment of a server running “View for Clusterstor” software. This should enable better understanding of work file system performance and a more immediate ability to identify the origin of problems on work file systems. EPCC have configured and provided a server for HPE to use for this deployment.

- EPCC will continue to work with HPE on the refinement and testing of a final software image to be deployed utilising the new process developed by HPE. In particular, we will help investigate any issues or performance problems identified during user testing.

- EPCC continue to work closely with contractors in planning the full site power shutdown at the Advanced Computing Facility.

- Given the current weather trends, we are expecting a hot summer and are continuing to improve our understanding and our processes around hot weather mitigation to the ARCHER2 service and supporting plant.

2 ARCHER2 Performance Report

This is the contractual performance report for the ARCHER2 SP Service for the Reporting Periods from 1 April until 30 June 2025.

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	Apr 2025		May 2025		Jun 202		Q2 2025	
	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)	0.05	-2	0.07	-2	0.07	-2	**0.06**	**-6**
SP_Level3 (MTR)	2.96	-2	0.00	-2	0.00	-2	**2.96**	**-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**		**-36**

2.1.2 Service Credits

As the Total Service Points are negative (-36), no Service Credits apply in 25Q2.

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.

April 2025					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1157	1157	0	0.092	0:00:55
SP_Level2	78	77	17	7.169	0:16:37
SP_Level3	0	1	0	20	0:08:25
May 2025					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	774	774	0	0.112	0:00:59
SP_Level2	62	61	18	7.557	0:11:23
SP_Level3	0	0	0	0	0:00:00
June 2025					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1458	1458	0	0.055	0:00:38
SP_Level2	78	82	14	7.756	0:16:00
SP_Level3	0	0	1	0	0:00:00
Q2 2025					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	3389	3389	0	0.081	0:00:51
SP_Level2	218	220	14	7.495	0:14:56
SP_Level3	0	1	0	20	0:08:25

2.3 Query Resolution

Metric	**Apr 2025**		**May 2025**		**Jun 2025**		**Q2 2025**	
Service Level	**MTR**	**Resolved**	**MTR**	**Resolved**	**MTR**	**Resolved**	**MTR**	**Resolved**
SP_Level1	0:00:08	1157	0:00:05	774	0:00:20	1458	0:00:10	3389
SP_Level2	0:27:59	77	0:39:32	61	0:40:32	82	0:35:22	220
SP_Level3	29:38:03	1	0:00:00	0	0:00:00	0	29:38:03	1
Total		**1235**		**835**		**1540**		**3610**

A total of 3610 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 April, there were 17 feedback scores received during this period. 100% were Good, Very Good or Excellent with 88% given the highest score of Excellent.

During May, there were 20 feedback scores received during this period. 100% were Good, Very Good or Excellent with 90% given the highest score of Excellent.

During June, there were 17 feedback scores received during this period. 100% were Good, Very Good or Excellent with 82% given the highest score of Excellent.

£54 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
Partial	2025-06-30 1330	2025-06-30 2100	7 hrs 30 mins	Some compute nodes were removed from service to ensure adequate cooling was available. Possible longer queue times.	Unusually hot weather in Edinburgh area	Accommodation
Partial	2025-06-25 0800	2025-06-25 1710	9 hrs 10 mins	Increased queue times caused by reduced node availability. Possible intermittent issues with file system or internode communication due to change in interconnect topology while cabinets and switches are unavailable.	Pump replacement on 3 cabinets which were removed from service	HPE
Partial	2025-06-20 1100	2025-06-23 0830	2 days 21 hrs 30 mins	Some compute nodes were removed from service to ensure adequate	Unusually hot weather in Edinburgh area. Full return to service was	Accommodation

				cooling was available. Possible longer queues times.	delayed due to faulty valve which is awaiting replacement	

3 ARCHER2 Service Statistics

3.1 Utilisation

Utilisation from 1 April – 30 June is 91% which is the same as the previous quarter. Utilisation for April was 91%, for May 91% and for June 89%. The utilisation was slightly lower during June due to capability days and also the hot weather incident. During unusually high temperatures in Edinburgh, we removed some nodes from service to allow the cooling system to cope with the additional load.

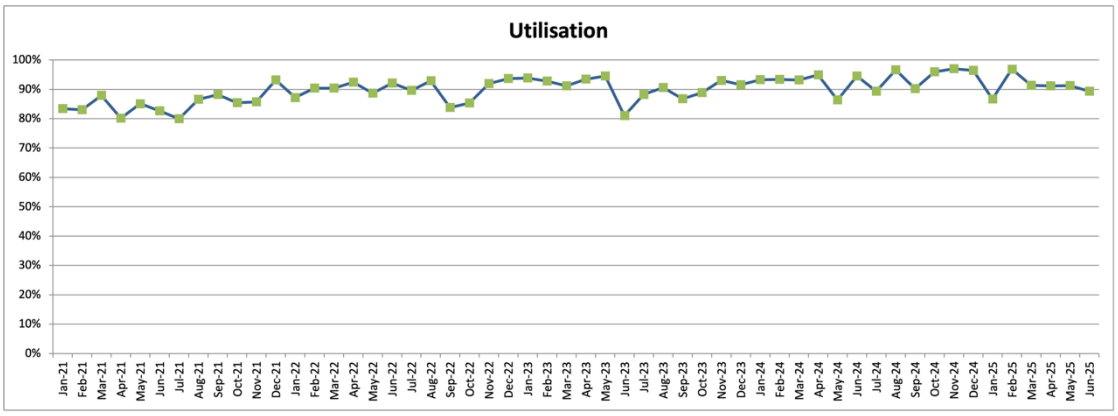

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER2 by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER2. It can be seen that EPSRC exceeded their target this quarter with their usage being at 78% (against their target of 66.8%). It should also be noted that the proportion of EPSRC's uncharged utilisation has increased further this quarter from 11% in 1Q25 to 12% in this quarter.

NERC did not meet their target utilisation of 18.2% and it was significantly less at 10%. This is also lower than their usage in the previous quarter which was 17%. The higher utilisation by NERC at the end of Q1 may be due to the end of the annual allocation period on 31st March and users running a lot of jobs at the end of the period. It is disappointing to see the reduced utilisation for NERC in Q2.

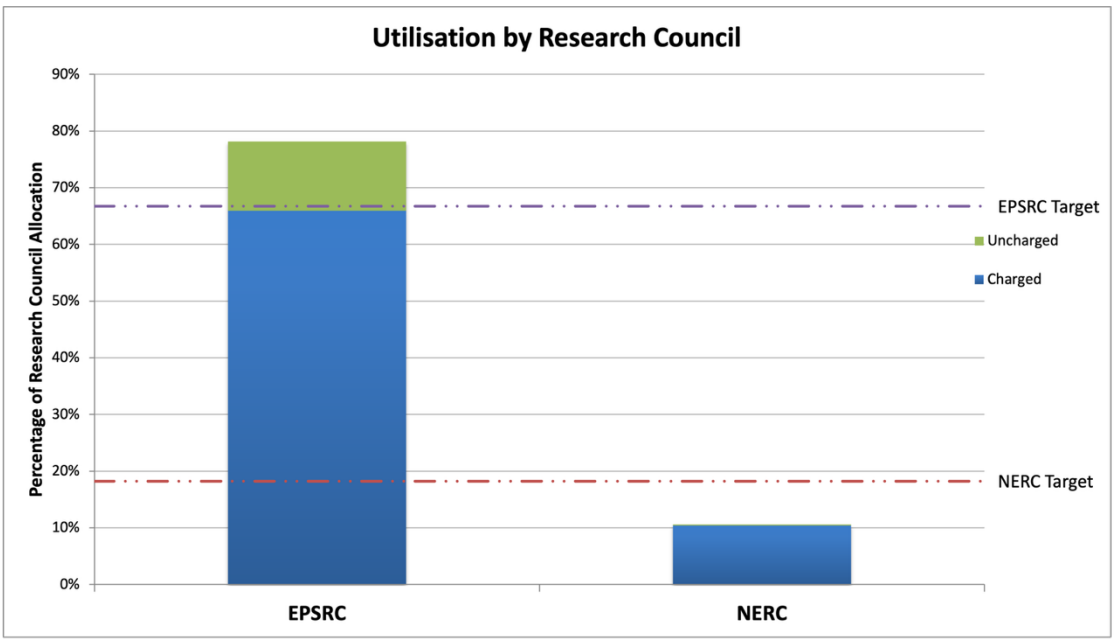

The stacked graph below shows the trend of charge and uncharged utilisation since the start of the service.

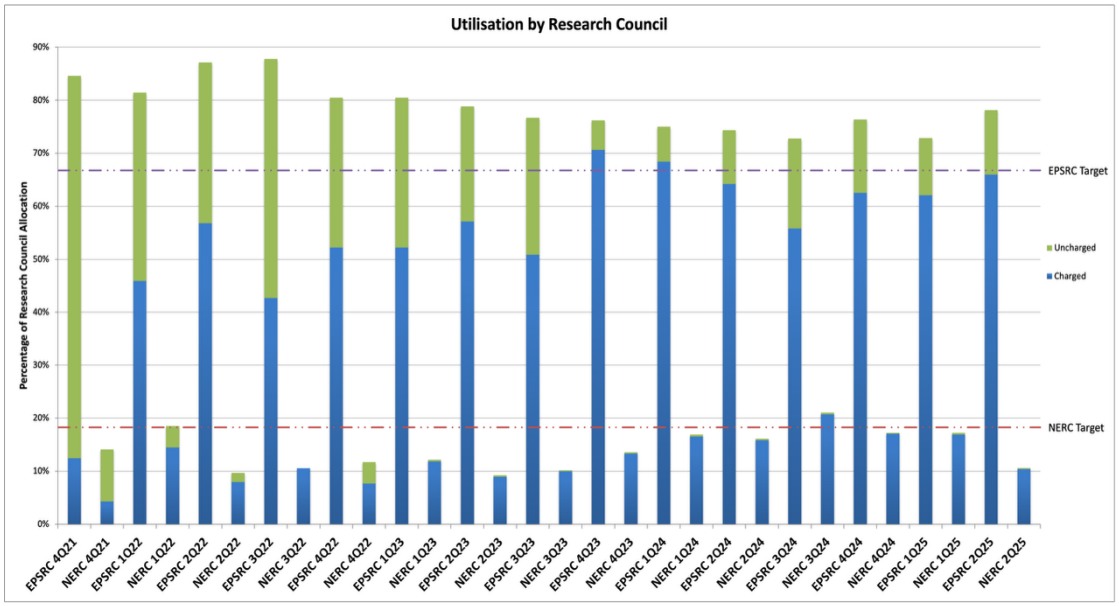

3.2 Scheduling Coefficient Matrix

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.

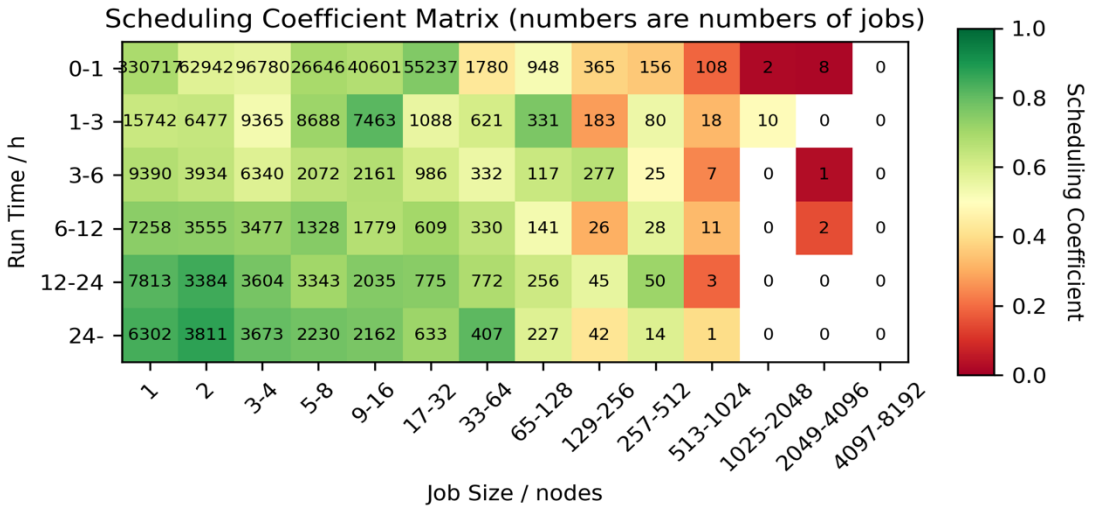

The usage heatmap below provides an overview of the usage on ARCHER2 over the quarter for different job sizes/lengths. The colour in the heatmap indicates the number of CUs expended for each class, and the number in the box is the number of jobs of that class.

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.

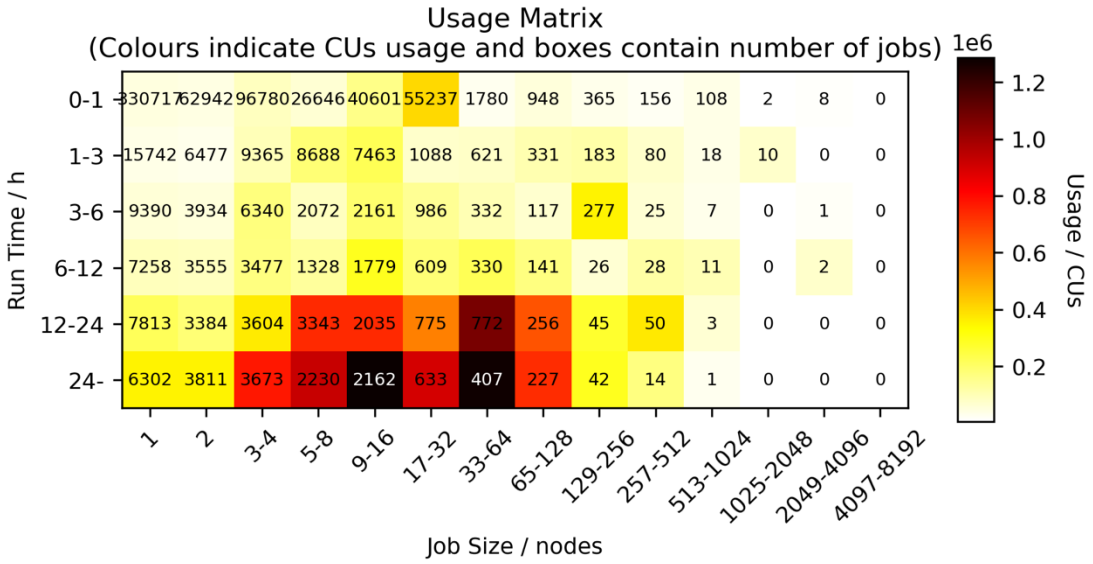

Appendix: Critical Success Factors

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: 4

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]
- Developed a set of new tools to help users estimate the environmental impact of their computing simulations and workloads [November 2024]

Power Usage

	4Q 21*	**1Q 22**	**2Q 22**	**3Q 22**	**4Q 22**	**1Q 23**	**2Q 23**	**3Q 23**	**4Q 23**
Average Power	3.31	3.16	3.15	2.86	2.90	2.51	2.56	2.46	2.53
	1Q 24	**2Q 24**	**3Q 24**	**4Q 24**	**1Q 25**	**2Q 25**			
Average Power	2.58	2.54	2.64	2.57	2.55	2.55			

* 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 on and off ARCHER2 each Quarter.

Data Transferred…	**4Q21***	**1Q22**	**2Q22**	**3Q22**	**4Q22**	**1Q23**	**2Q23**	**3Q23**	**4Q23**
…to ARCHER2 (TB)	534	163	68	220	44	67	42	65	99
…**from** ARCHER2 (TB)	236	582	667	822	834	1231	1022	1472	1771
Data Transferred…	**1Q24**	**2Q24**	**3Q24**	**4Q24**	**1Q25**	**2Q25**			
…**to** ARCHER2 (TB)	108	93	98	228	114	150			
…**from** ARCHER2 (TB)	2056	2443	3956	2227	2915	1815			

* Partial

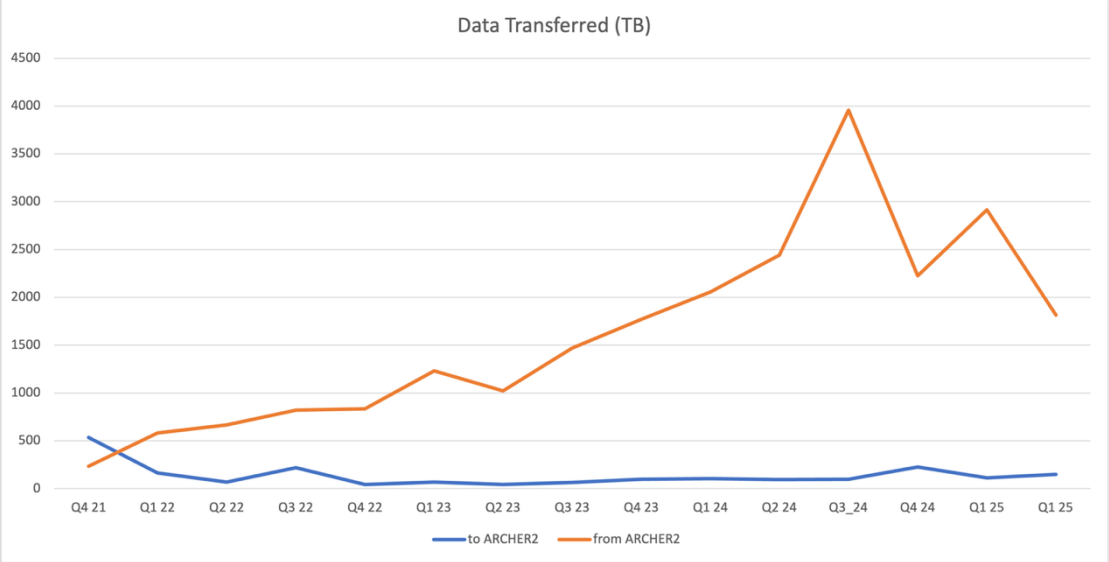

The amount of data moved off ARCHER2 has reduced from the previous quarter, however it remains higher than any point previous to 2024.

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[1] and report the MPI-IO write performance with the following settings:

- Global data structure of 2048^3: 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 and provide boxplots illustrating the performance variation. (On the boxplots, the green triangles mark the mean value and the whiskers extend to the last datapoint within the range 1.5 x IQR.)

During Q1 2025, we have worked to remove data from a2fs-work1 as much as possible. With the usage of the file system now below 80% we see a significant improvement in performance compared to Q4 2024 where the usage on the file system was over 80%. We are working to ensure that usage on any of the ARCHER2 Lustre file systems does not go above 80% to try and maintain good performance for users on the service.

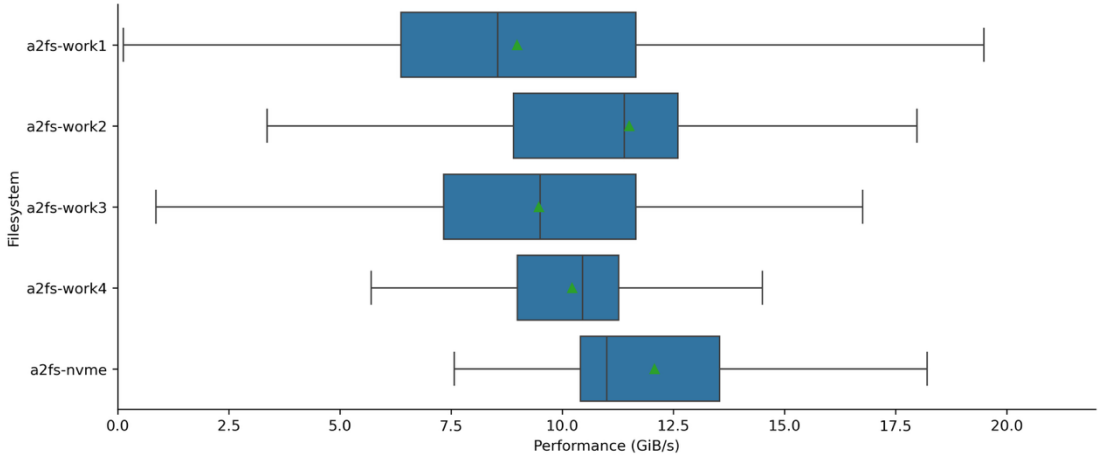

[1] https://github.com/davidhenty/benchio

Benchio MPI-IO medium (GiB/s)	1Q23	2Q23		3Q23	4Q23	1Q24	2Q24	3Q24	4Q24
a2fs-work1	8.2	7.6±0.5		10.5 (8.8:11.8)	10.9 (8.3:12.5)		10.1 (7.0:11.8)	9.7 (6.7:11.9)	4.0 (2.0:9.6)
a2fs-work2	8.5	7.3±0.6		10.4 (7.2:12.4)	10.4 (7.7:13.0)		11.1 (8.0:12.5)	11.1 (8.1:13.1)	11.3 (7.1:15.0)
a2fs-work3	8.3	9.6±0.7		10.0 (8.2:11.6)	10.7 (8.1:11.9)		9.6 (8.4:11.8)	9.6 (7.5:11.8)	9.3 (7.0:11.6)
a2fs-work4					9.7 (9.1:10.2)		10.0 (9.2:10.8)	10.6 (9.4:11.6)	10.7 (9.5:11.7)
a2fs-nvme				10.1 (9.6:11.5)	10.1 (9.5:12.4)		11.1 (10.5:12.4)	11.6 (11.1:12.7)	10.7 (10.0:11.8)

Benchio MPI-IO medium (GiB/s)	1Q25	2Q25
a2fs-work1	8.6 (5.0:11.1)	8.6 (6.4:11.7)
a2fs-work2	10.7 (8.1:12.7)	11.4 (8.9:12.6)
a2fs-work3	9.3 (7.7:11.6)	9.5 (7.3:11.7)
a2fs-work4	10.1 (8.8:10.9)	10.5 (9.0:11.3)
a2fs-nvme	11.3 (10.4:13.3)	11.0 (10.4:13.5)

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	01/2023 – 12/2023	1Q 24	2Q 24	3Q 24	4Q 24	1Q 25	2Q 25
Relative Research Output per kWh	100	109	115	115	115	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	1Q24	2Q24	3Q24	4Q24	1Q25
Energy per CU (kWh)	0.719	0.713	0.728	0.715	0.650	0.545	0.669	0.590	0.568	0.582	0.585	0.595	0.546	0.518

Energy Cost per CU Delivered

	4Q21*	1Q22	2Q22	3Q22	4Q22	1Q23	2Q23	3Q23	4Q23	1Q24	2Q24	3Q24	4Q24	1Q25
Cost per CU (£)	£0.089	£0.090	£0.098	£0.096	£0.088	£0.074	£0.162	£0.143	£0.136	£0.140	£0.160	£0.164	£0.149	£0.142

* 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. There was also an additional increase in the unit cost of electricity from April 2024.