ARCHER2
SP Quarterly Report

January - March 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	26/03/2025	Template created	Jo Beech-Brandt
0.2	04/04/2025	Added narrative, graphs, service statistics	Jo Beech-Brandt
0.3	05/04/2025	Added critical success metrics	Lorna Smith
0.4	05/04/2025	Added ISO related information	Anne Whiting
0.5	08/04/2025	Reviewed	Alan Simpson
1.0	09/04/2025	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 January – 31 March 2025.

1.1 Service Highlights

- During March, the ExCALIBUR-ECP Hardware and Enabling Software Knowledge Exchange Programme provided financial support for an exchange visit for three members EPCC staff: the Director of Research; the Director of the Systems team; and a HPC systems specialist who leads on ARCHER2. They visited the National Renewable Energy Laboratory (NREL) and the National Centre for Atmospheric Research (NCAR) in Colorado and Wyoming. The primary focus of the visit to see data centres and data centre technology and how these US organisations approach hosting large scale systems in as sustainable a manner as possible. The visit allowed us to compare and contrast our relative approaches to HPC system hosting and data centre policies and procedures. It also enabled us to kick start longer term collaborations on technology innovations; policy and procedure standardisation; and best practises for hosting and operating HPC systems. This is invaluable knowledge for ARCHER2 and future services.

- After working closely with the EPCC training team, the ARCHER2 Systems team are now delivering a new MSc module in System administration. This will support students within industry and academia to build core skills in HPC system support.

- Utilisation remains high on ARCHER2 with usage over the quarter at 91% and a new Access to HPC Call recently opened.

- EPCC have worked with HPE to test and demonstrate an alternative to the time-consuming standard upgrade process for ARCHER2, which would allow us to keep up to date for software patches while minimising downtime. This has been included in the options paper recently sent to the ARCHER2 Management Board.

- Module software was moved from the heavily utilised work1 file system to the work4 file system which is not as heavily utilised, and we hope this will provide better resiliency for users.

- ARCHER2 staff continue to engage with the wider UKRI DRI projects and workshops, and helped organise the recent launch of the National Federated Compute Services NetworkPlus event as well as attending the Data Interoperability workshop in Birmingham.

- Accommodation works to complete the replacement and renewal of blockwork at the cr3 door entrance completed successfully with minimal impact on site stakeholders.

- New ARCHER2 livery installed on security fence in prC.

- Successful change freeze / period of additional scrutiny completed successfully with normal change management processes back from 6th January.

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. Server hardware has been built and configured to support View and will be handed over to HPE for View installation and integration shortly.

- ARCHER2 SP staff are attending and presenting at the ISC High Performance Conference in Hamburg in June 2025.

- A User Forum event is planned within the ARCHER2 Celebration of Science event that will take place in Edinburgh in May 2025.

- The User Advisory Group met recently and spoke positively about Capability Days. The ARCHER2 team were asked to plan further capability days and these will be advertised once scheduled.

- EPCC are working towards their annual ISO external audits for ISO 9001 (quality service delivery), ISO 27001 (information security) and ISO 22301 (business continuity and disaster recovery) in June. EPCC uses these standards to ensure that they are meeting user’ requirements, applying best practice and identifying areas for improvement. The annual external audit is an opportunity to take stock and ensure we are still complying with the requirements of the standards.

2 ARCHER2 Performance Report

This is the contractual performance report for the ARCHER2 SP Service for the Reporting Periods from 1 January until 31 March 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	Jan 2025		Feb 2025		Mar 202		Q1 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)	0.06	-2	0.05	-2	0.06	-2	**0.06**	**-6**
SP_Level3 (MTR)	4.54	-2	0.00	-2	0.00	-2	**5.40**	**-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 25Q1.

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.

January 2025					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1908	1908	0	0.102	0:01:29
SP_Level2	111	95	23	7.505	0:19:09
SP_Level3	1	1	1	13	0:00:00
February 2025					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1624	1624	0	0.126	0:01:35
SP_Level2	108	116	15	6.655	0:11:41
SP_Level3	0	0	1	0	0:00:00
March 2025					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1159	1159	0	0.173	0:01:54
SP_Level2	111	110	16	6.691	0:16:09
SP_Level3	0	0	1	0	0:00:00
Q1 2025					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	4691	4691	0	0.128	0:01:40
SP_Level2	330	321	16	6.919	0:15:26
SP_Level3	1	1	1	13	0:00:00

2.3 Query Resolution

Metric	**Jan 2025**		**Feb 2025**		**Mar 2025**		**Q1 2025**	
Service Level	**MTR**	**Resolved**	**MTR**	**Resolved**	**MTR**	**Resolved**	**MTR**	**Resolved**
SP_Level1	0:00:43	1908	0:00:14	1624	0:00:19	1159	0:00:24	4691
SP_Level2	0:36:35	95	0:30:21	116	0:35:54	110	0:32:53	321
SP_Level3	45:24:18	1	0:00:00	0	0:00:00	0	45:24:18	1
Total		**2004**		**1740**		**1269**		**5013**

A total of 5013 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 January, there were 28 feedback scores received during this period. 100% were Good, Very Good or Excellent with 86% given the highest score of Excellent.

During February, there were 34 feedback scores received during this period. 100% were Good, Very Good or Excellent with 88% given the highest score of Excellent.

During March, there were 30 feedback scores received during this period. 100% were Good, Very Good or Excellent with 87% given the highest score of Excellent.

£92 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
Full	2025-03-26 0700	2025-03-27 1410	31 hrs 10 mins	Users cannot connect to login nodes, access data and jobs will not run	Replacement of operating system certificates	HPE
Partial	2025-03-25 1015	2025-03-25 1135	1 hr 20 mins	Compute nodes are unavailable. Users can still connect to login nodes and access data	Failed server	HPE
Partial	2025-02-03 0800	2025-02-03 1600	8 hrs	Users cannot access data on the RDFaaS : /epsrc and /general	Essential work on the RDFaaS hardware	Not within scope
Partial	2025-01-24 1245	2025-01-27 1130	70 hrs 45 mins	Compute nodes are unavailable. Users can connect to login nodes and access data. Jobs will not run and jobs running will have failed.	Extreme weather in Edinburgh area caused interrupted power supply at ACF data centre. Restoration took place once weather alert had passed.	Accommodation
Partial	2025-01-22 1445	2025-01-22 1800	3 hrs 15mins	Users cannot connect to login nodes, jobs running at time	Failure on solid state lustre file system	HPE

				of failure will have failed		
Partial	2025-01-15 1145	2025-01-16 0920	21 hrs 35 mins	Users cannot connect to login nodes, jobs running at time of failure will have failed	Failure on solid state lustre file system	HPE
Partial	2025-01-09 1230	2025-01-10 0810	19 hrs 40 mins	No user jobs can start for projects hosted on this file system	Hardware issue on work2 lustre file system	HPE

3 ARCHER2 Service Statistics

3.1 Utilisation

Utilisation from 1 January – 31 March is 91% which is slightly lower than the 96% from the previous quarter. Utilisation for January was 87%, for February 97% and for March 91%. The utilisation was slightly lower during January due to some interruptions to service caused by severe weather alert in Edinburgh and also issues with the solid-state file system.

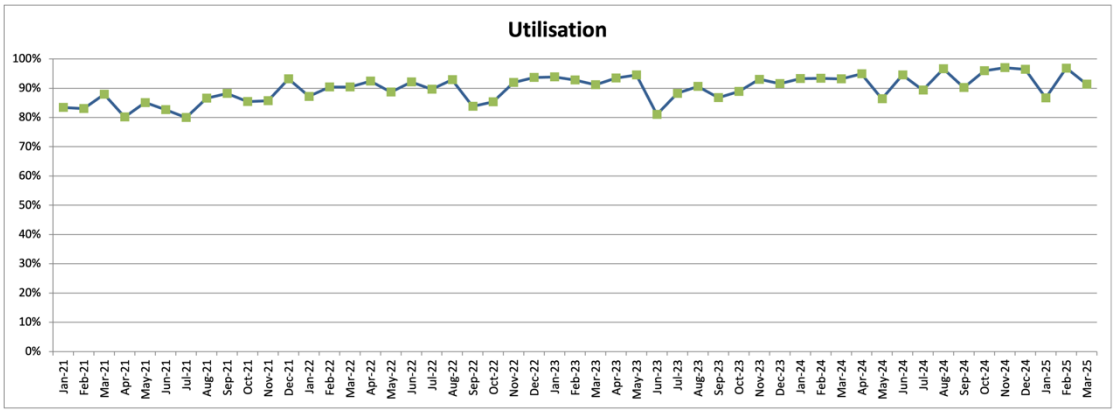

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER2 by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER2. It can be seen that EPSRC exceeded their target this quarter with their usage being at 73% (against their target of 66.8%). It should also be noted that the proportion of EPSRC's uncharged utilisation decreased further this quarter from 14% in 4Q24 to 11% in this quarter.

NERC almost met their target with utilisation being 17% (against their target of 18.2%) which was the same usage in the previous quarter.

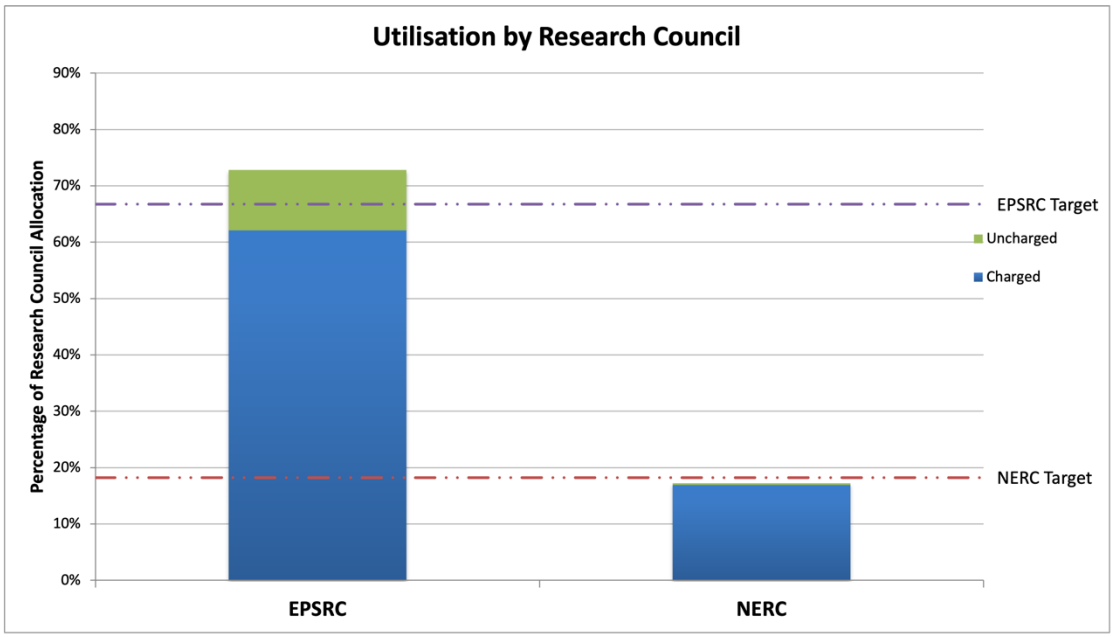

The stacked graph below shows the trend of charge and uncharged utilisation since the start of the service.

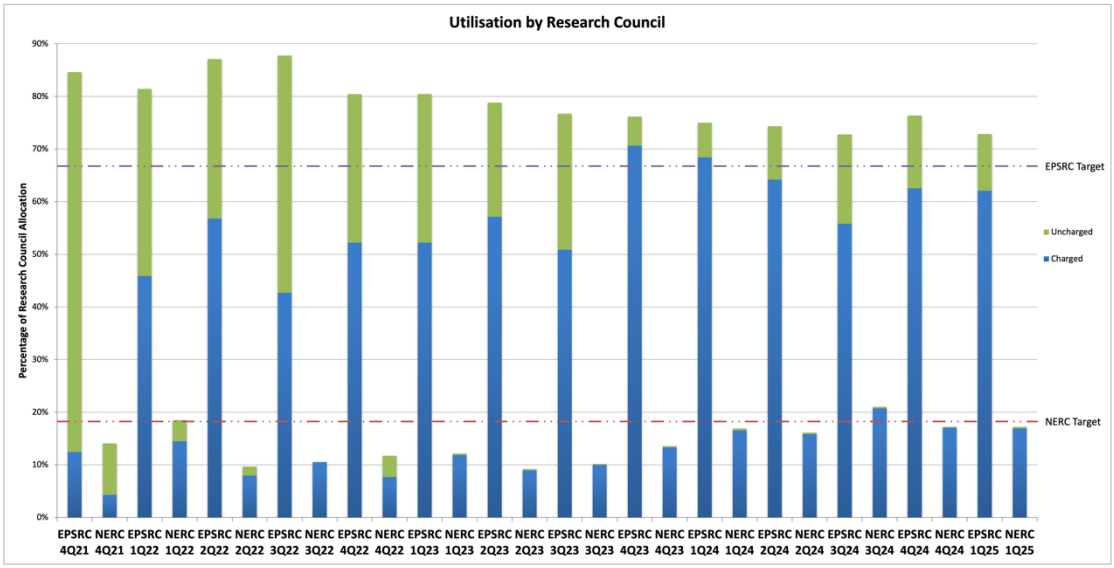
Utilisation by Research Council
Percentage of Research Council Allocation
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
EPSRC Target
NERC Target
Uncharged
Charged
EPSRC 4Q21
NERC 4Q21
EPSRC 1Q22
NERC 1Q22
EPSRC 2Q22
NERC 2Q22
EPSRC 3Q22
NERC 3Q22
EPSRC 4Q22
NERC 4Q22
EPSRC 1Q23
NERC 1Q23
EPSRC 2Q23
NERC 2Q23
EPSRC 3Q23
NERC 3Q23
EPSRC 4Q23
NERC 4Q23
EPSRC 1Q24
NERC 1Q24
EPSRC 2Q24
NERC 2Q24
EPSRC 3Q24
NERC 3Q24
EPSRC 4Q24
NERC 4Q24
EPSRC 1Q25
NERC 1Q25

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.

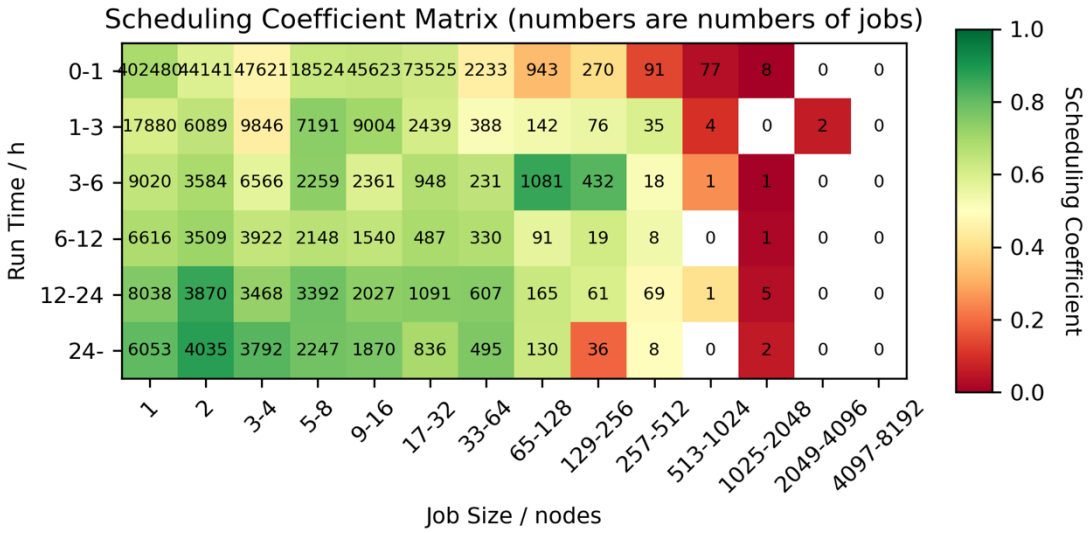

The usage heatmap below provides an overview of the usage on ARCHER2 over the quarter for different job sizes/lengths. The colour in the heatmap indicates the number of CUs expended for each class, and the number in the box is the number of jobs of that class.

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.

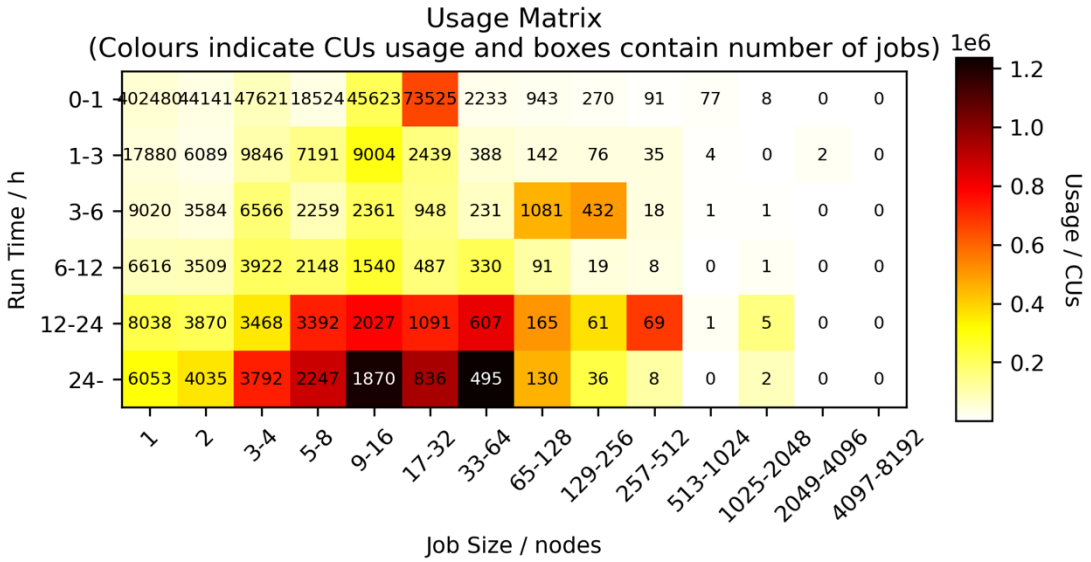

Appendix: Critical Success Factors

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**				
Average Power	2.58	2.54	2.64	2.57	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**				
…**to** ARCHER2 (TB)	108	93	98	228	114				
…**from** ARCHER2 (TB)	2056	2443	3956	2227	2915				

* Partial

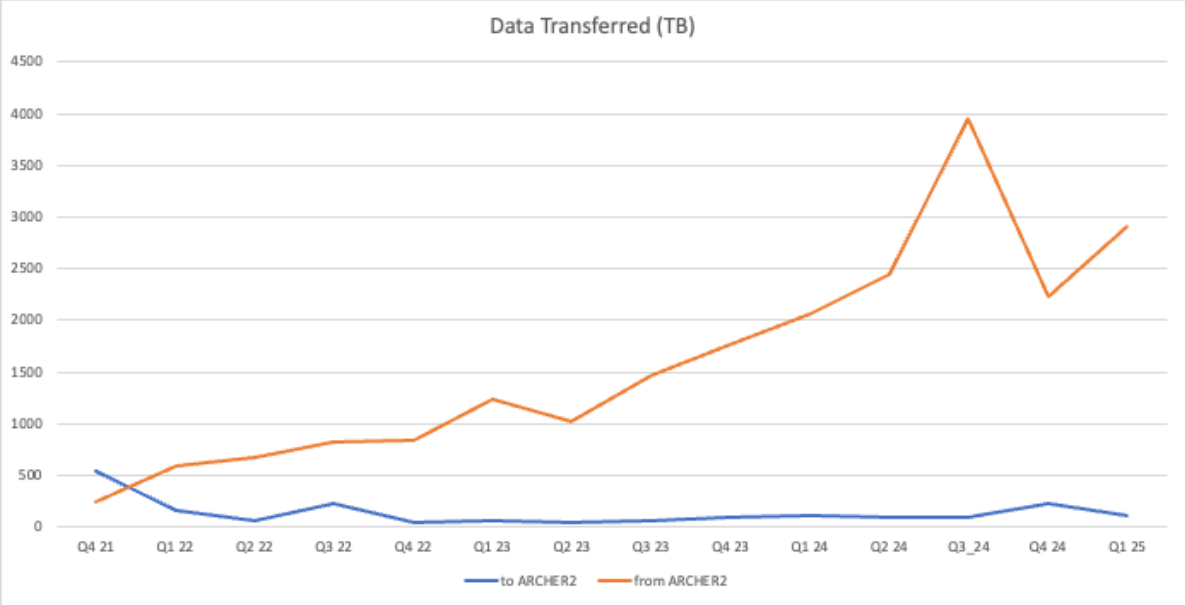

The amount of data being moved off this quarter has increased from the previous quarter.

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.

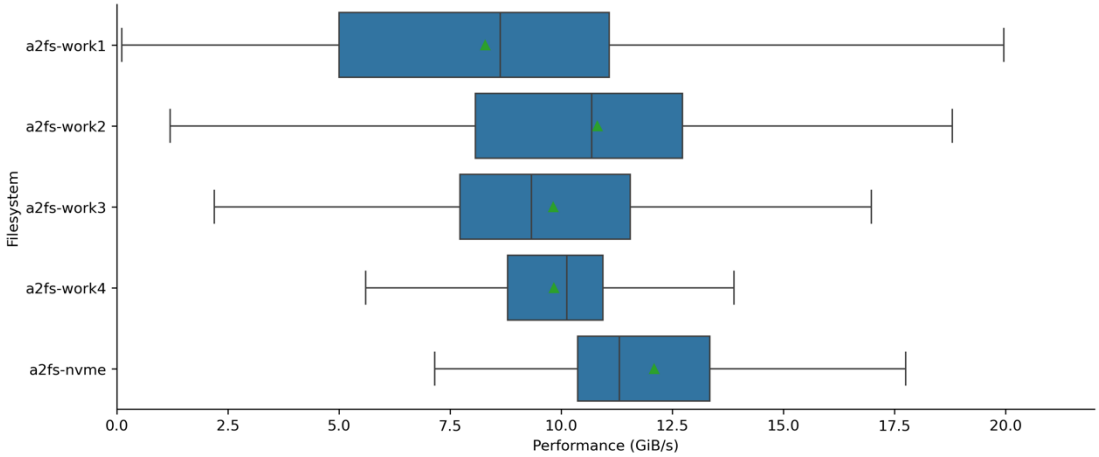

[1] https://github.com/davidhenty/benchio

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
a2fs-work1	8.6 (5.0:11.1)
a2fs-work2	10.7 (8.1:12.7)
a2fs-work3	9.3 (7.7:11.6)
a2fs-work4	10.1 (8.8:10.9)
a2fs-nvme	11.3 (10.4:13.3)

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
Relative Research Output per kWh	100	109	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
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

Energy Cost per CU Delivered

	4Q21*	1Q22	2Q22	3Q22	4Q22	1Q23	2Q23	3Q23	4Q23	1Q24	2Q24	3Q24	4Q24
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

* 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.