ARCHER2
SP Quarterly Report

January - March 2024

EPCC
The University of Edinburgh

epcc

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	01/04/2024	Template created	Jo Beech-Brandt
0.2	02/04/2024	Added heatmap and usage graphs	Clair Barrass
0.3	05/04/2024	Added narrative, metrics and graphs	Jo Beech-Brandt
0.4	09/04/2024	Added ISO details	Anne Whiting
0.5	10/04/2024	Added hosting update	Paul Clark
0.6	11/04/2024	Updated Critical Success Factors	Alan Simpson
0.7	12/04/2024	Reviewed	Alan Simpson
1.0	15/04/2024	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 January – 31st March 2024.

1.1 Service Highlights

- The AMD GPU nodes were integrated into ARCHER2. A user forum took place with both HPE and EPCC staff, and users had access to the nodes following this. The status page now includes a graph showing the usage of the GPU nodes.

- On the scratch NVMe file system, a purge policy was implemented and users were notified that auto-deletion of files older than 28 days was being enabled. Documentation was updated and a sample command included so users can check which files will be purged.

- 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 in recent years, and the new version places a much greater emphasis on cyber security risks and preventive measures.

- A Capability Day took place which allowed users to run large scale jobs that were uncharged. This event was oversubscribed and a report detailing the day and recommendations for future days was created and sent to UKRI.

- A hybrid User Forum took place recently as part of the 2-day Celebration of Science in-person event. This focused on some of the recent developments including GPU arrival, NMVe file system purge, Capability Day and ARCHER2 Training.

- To gain a more detailed understanding of the IO patterns in practise on the service, IO profiles were captured by the Darsham IO profiling tool for all parallel jobs launched using 'srun' for a 24-hour period. Users were notified and informed how to opt their jobs out of the log collection if they wished. The logs will be analysed and allow us to work on improving user experience with IO on the /work lustre file systems and provide better advice on improving IO performance to users of the service.

- ARCHER2 staff attended the Digital Research Infrastructure (DRI) Federated Compute NetworkPlus meeting and are now engaged within the project team submitting a proposal for the NetworkPlus.

- Technical free-cooling report produced from AR2-RFC-0189 in February. Survey completed of the Dry Air Coolers and deep clean planned prior to summer 2024.

1.2 Forward Look

- A webinar is being planned especially aimed at PIs and the use of SAFE to manage their projects. This is in addition to the specific documentation for PIs.

- A Business Continuity Disaster Recovery (BCDR) test is being planned for the forthcoming quarter.

- Future Capability Days are planned with adopted amendments following the previous Capability Day.

- Analysis is currently taking place for the fair-share policy on ARCHER2. A proposal is being prepared to change parameters and then to analyse the impact of the changes especially in respect to the usage for the HEC consortia.

- ARCHER2 staff will participate in the Digital Research Infrastructure (DRI) Cybersecurity workshop engaging with the wider DRI community.

- ARCHER2 staff are preparing for the annual ISO external audit, which takes place in June this year. We are upgrading to the latest version of ISO 27001, the Information Security standard. This standard had not been updated for 10 years, and the emphasis has changed a lot to reflect the increased emphasis on cyber security risks. The audit includes information security, quality service delivery and business continuity and disaster recover, to ensure that we continue to apply best practice in our management of ARCHER2 and the user science it hosts.

- Deep clean of Dry Air Coolers planned before warmer temperatures.

2 ARCHER2 Performance Report

This is the contractual performance report for the ARCHER2 SP Service for the Reporting Periods from 1st January 2024 until 31st March 2024.

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 2024		Feb 2024		Mar 2024		Q1 2024	
	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)	0.04	-2	0.07	-2	0.06	-2	**0.06**	**-6**
SP_Level3 (MTR)	8.00	-0.5	0.00	-2	0.35	-2	**4.17**	**-4.5**
Initial Response (%)	100%	-1	100%	-1	100%	-1	**100%**	**-3**
Query Satisfaction (%)	100%	-2	100%	-2	100%	-2	**100%**	**-6**
Total		**-10.5**		**-12**		**-12**		**-34.5**

2.1.2 Service Credits

As the Total Service Points are negative (-34.5), no Service Credits apply in 24Q1.

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 2024					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	2651	2651	0	0.104	0:01:23
SP_Level2	93	94	13	6.374	0:15:21
SP_Level3	1	1	1	15	3:08:53
February 2024					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1,343	1,343	0	0.13	0:00:56
SP_Level2	108	110	23	7.555	2:51:04
SP_Level3	0	0	1	0	0:00:00
March 2024					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1210	1210	0	0.121	0:01:11
SP_Level2	113	109	27	6.587	0:48:07
SP_Level3	1	1	1	9	0:03:00
Q1 2024					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	5204	5204	0	0.115	0:01:12
SP_Level2	377	374	27	6.783	1:10:42
SP_Level3	2	2	1	12	1:35:56

2.3 Query Resolution

Metric	**Jan 2024**		**Feb 2024**		**Mar 2024**		**Q1 2024**	
Service Level	**MTR**	**Resolved**	**MTR**	**Resolved**	**MTR**	**Resolved**	**MTR**	**Resolved**
SP_Level1	0:01:05	2651	0:00:36	1343	0:00:49	1210	0:00:52	5204
SP_Level2	0:26:26	155	0:43:33	110	0:37:51	109	0:35:26	374
SP_Level3	80:00:00	1	0:00:00	0	3:27:04	1	41:43:32	2
Total		**2807**		**1453**		**1320**		**5580**

A total of 5580 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 50 feedback scores received during this period. 100% were Good, Very Good or Excellent with 92% given the highest score of Excellent.

During February, there were 32 feedback scores received during this period. 100% were Good, Very Good or Excellent with 84% given the highest score of Excellent.

During March, there were 38 feedback scores received during this period. 100% were Good, Very Good or Excellent with 82% given the highest score of Excellent.

£120 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	2024-03-24 2300	2024-03-25 1100	12hrs	New jobs could not run and jobs running at the time of the incident failed	Compute nodes lost power	Accommodation

3 ARCHER2 Service Statistics

3.1 Utilisation

Utilisation from 1st January – 31st March is 93% which is slightly increased from 91% the previous quarter. Utilisation for January was 93%, for February 93% and for March 93%.

Please note the first Capability Day took place during this period on March 12th but this did not impact the overall utilisation for March which remained high at 93%.

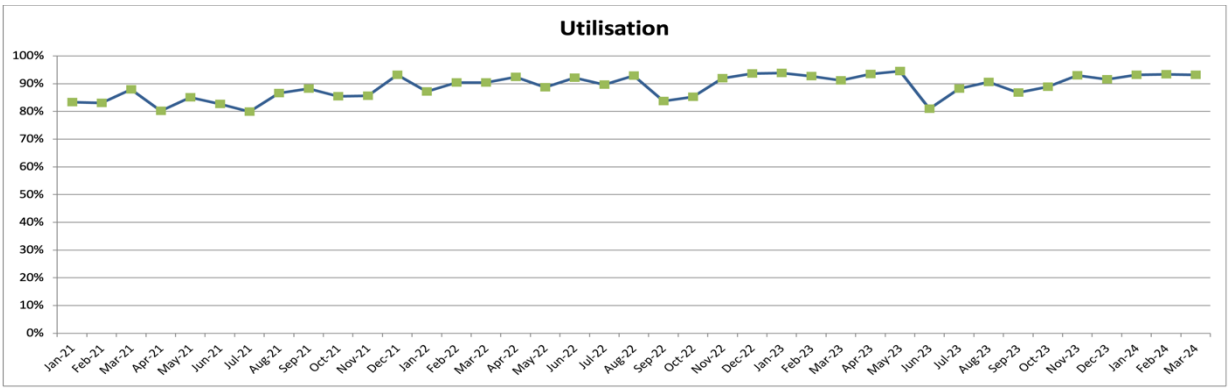

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER2 by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER2. It can be seen that EPSRC exceeded their target this quarter with their usage being at 75% (against their target of 66.8%). It should be noted that the proportion of EPSRC's uncharged utilisation has reduced and is 6.5%. The use of the Capability Day was uncharged.

NERC missed their target with utilisation being 16.5% (against their target of 18.2%) but this is nevertheless an increase from 13.6% in the previous quarter.

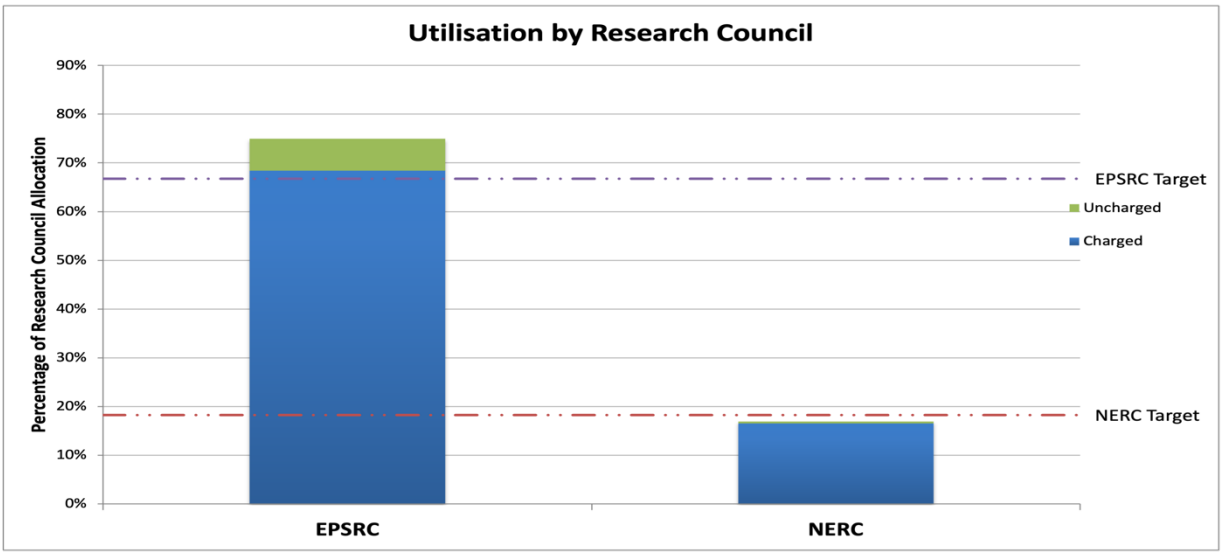

3.2 Scheduling Coefficient Matrix

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.

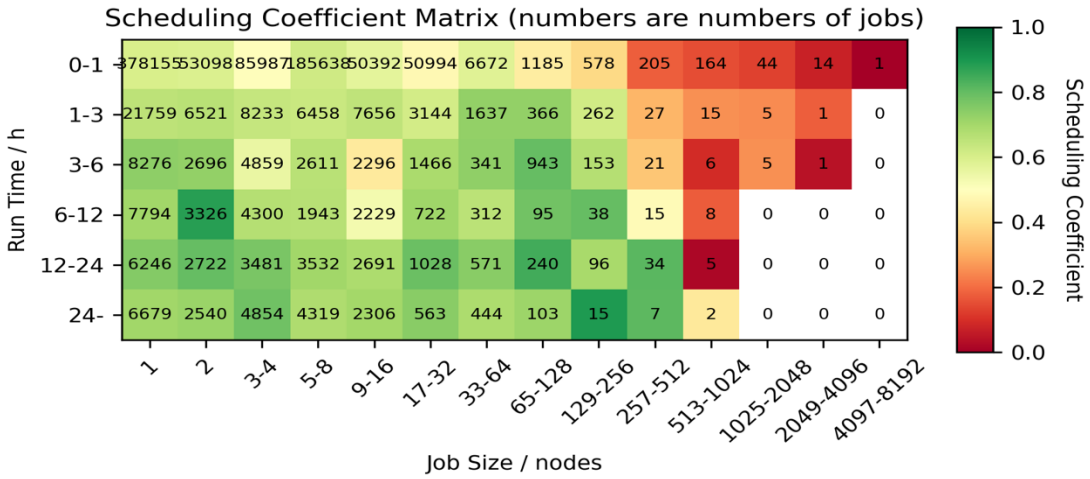

The usage heatmap below provides an overview of the usage on ARCHER2 over the quarter for different job sizes/lengths. The colour in the heatmap indicates the number of CUs expended for each class, and the number in the box is the number of jobs of that class.

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.

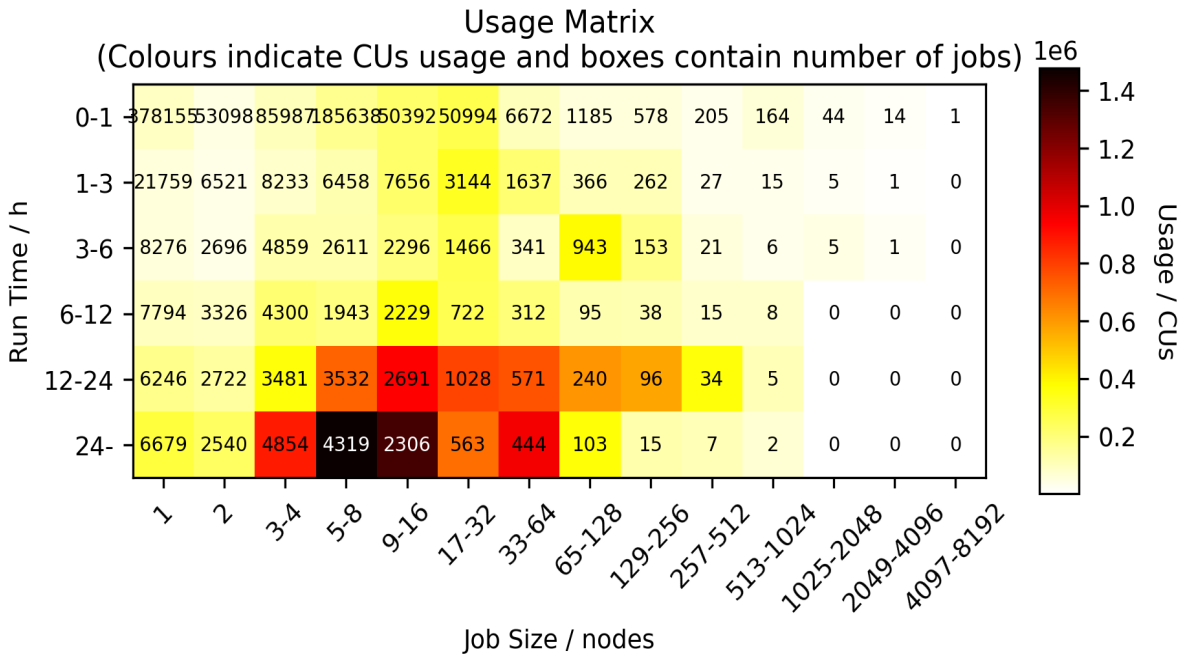

Appendix: Critical Success Factors

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	1Q24
Average Power	3.31	3.16	3.15	2.86	2.90	2.51	2.56	2.46	2.53	2.58

* 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…	4Q 21*	1Q 22	2Q 22	3Q 22	4Q 22	1Q 23	2Q 23	3Q 23	4Q 23	1Q 24
… **to** ARCHER2 (TB)	534	163	68	220	44	67	42	65	99	108
…**from** ARCHER2 (TB)	236	582	667	822	834	1231	1022	1472	1771	2056

* Partial

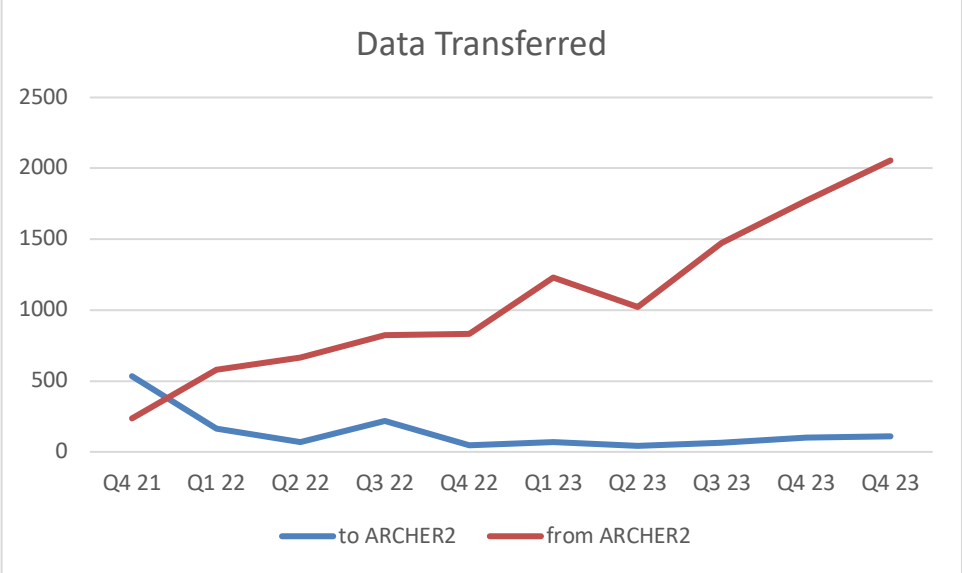

The trend of more data being moved off each quarter continues. Indeed There was almost as much data transferred off ARCHER2 in the last quarter as there was in the first year of the ARCHER2 service.

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.

Note: Recent, in-depth IO performance analysis undertaken by the CSE team into IO patterns on the ARCHER2 Lustre file systems has revealed that the IO performance test reported in previous quarters does not seem to provide as robust a measure of IO performance as we would want. We are currently investigating the performance variability shown by the original test and comparing to potential improved tests, and we will provide an update by the next Quarterly Review meeting for discussion with UKRI.

Benchio MPI-IO medium (GiB/s)	1Q23	2Q23		3Q23	4Q23	1Q24
a2fs-work1	8.2	7.6±0.5		10.5 (8.8:11.8)	10.9 (8.3:12.5)	N/A
a2fs-work2	8.5	7.3±0.6		10.4 (7.2:12.4)	10.4 (7.7:13.0)	N/A
a2fs-work3	8.3	9.6±0.7		10.0 (8.2:11.6)	10.7 (8.1:11.9)	N/A
a2fs-work4					9.7 (9.1:10.2)	N/A
a2fs-nvme				10.1 (9.6:11.5)	10.1 (9.5:12.4)	N/A

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	1Q24
Relative Research Output per kWh	100	109	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*
Energy per CU (kWh)	0.719	0.713	0.728	0.715	0.650	0.545	0.669	0.590	0.568	0.552

*partial

Energy Cost per CU Delivered

	4Q21*	1Q22	2Q22	3Q22	4Q22	1Q23	2Q23	3Q23	4Q23	1Q24*
Cost per CU (£)	£0.089	£0.090	£0.098	£0.096	£0.088	£0.074	£0.162	£0.143	£0.136	£0.133

*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.