ARCHER2
SP Quarterly Report

April – June 2024

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	01/07/2024	Template created	Jo Beech-Brandt
0.2	02/07/2024	Added heatmap and usage graphs	Clair Barrass
0.3	03/07/2024	Added narrative, metrics and graphs	Jo Beech-Brandt
0.4	04/07/2024	Added narrative	Andy Turner
0.5	05/07/2024	Added hosting update	Paul Clark
0.6	09/07/2024	Reviewed and updated CSFs	Alan Simpson
1.0	10/07/2024	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: 1 April – 30 June 2024.

1.1 Service Highlights

- The AMD GPU nodes were integrated into ARCHER2, and their usage is now reported on the ARCHER2 status page.

- We are now certified for the latest version of ISO 27001, the Information Security standard. The standard was updated after 10 years, and the emphasis has changed significantly to reflect the increased emphasis on cybersecurity risks. The audit included information security, quality service delivery and business continuity and disaster recovery, and ensured that we continue to apply best practice in our management of ARCHER2 and the research it hosts.

- A Business Continuity Disaster Recovery (BCDR) test was performed across the service. This took place ahead of the annual ISO audit for the ARCHER2 service. The test focused on communication escalation following a simulated incident.

- Capability Days 2 took place from 4-6 June 2024 with improvements following feedback from Capability Days 1. Again, demand was high with more work scheduled by users than could be run during the session. We are now planning Capability Days 3 for Sep 2024.

- Further analysis has taken place on the Slurm scheduler fair-share configuration which is part of the priority calculation for jobs on the service. A short report has been circulated to UKRI for comment. Changes to fair-share that reflected the relative size of different projects (in terms of users) were found to make wait times more equitable across projects on ARCHER2. We are putting together a proposal for further tests of the fair-share configuration to aim to refine and further improve the service for all users.

- We have added support for functional accounts[1] to allow users to manage long-running services (such as licence servers) to support a wider variety of software and research workflows on ARCHER2. User documentation was added to the User Guide.

- ARCHER2 staff participated in the Digital Research Infrastructure (DRI) Cybersecurity strategy workshop engaging with the wider DRI community. The workshop was very interesting with positive discussions. The programme is ongoing with the next meeting focussed more on the technical aspects. We look forward to continuing our involvement in this collaborative endeavour.

- Completed improvements to the ARCHER2 Disaster Recovery (DR) infrastructure at Kings Buildings in Edinburgh via AR2-RFC-0128. This moved the connectivity from 2 x 10Gbps connections to 2 x 100Gbps and to newer hardware.

- Had positive discussions with Scottish Power Energy Networks (SPEN) at the ACF regarding the growth of site and the need for quality power distribution across their network.

[1] https://docs.archer2.ac.uk/user-guide/functional-accounts/

- Following preparation works in April 2024 to support major works in another part of site, we put a number of improvements in place to provide better long-term resilience to the ARCHER2 service. This included adding diverse power supplies to the outgoing Ciena network connections to the JaNET backbone (ACF-RFC-0269); the migration of infrastructure hosts to multiple locations within the ACF (ACF-RFC-0258) and diversification of power supplies for our main comms areas (ACF-RFC-0261).

- LDAP configuration changes (AR2-RFC-0199) to reduce the level of logs gathered. The settings had been retained since early in the service handling initial integration work and were now of little benefit. This reduction means an improvement of the server performance without negatively impacting the team's ability to debug issues.

- Resolved an issue with one of the prC transformers via AR2-RFC-0201. The transformer was isolated and the underlying issue resolved over a 3-day period and brought back into service for resiliency to the power infrastructure.

- A part of the continual site service improvement, following on from the chiller cleaning in prC at the end of 2023 / early 2024, additional works were carried out over the last period to complete deep cleaning of the Dry Air Coolers (DACs) on the roof of prC that support ARCHER2, along with changes to the controls strategy for the plant. This has improved the efficacy of the plants cooling capability and will improve heat rejection over the hotter summer months.

- All ARCHER2 login nodes have been integrated with the Globus Online data transfer infrastructure and this is currently under testing; appropriate documentation has also been prepared.

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.

- EPCC are engaged in high-level discussions with HPE about the future of the system software. The aim is to ensure good information security throughout the service, while minimising disruption to users.

- We are awaiting details of the latest round of the Access to HPC Call applications. Applicants will be asked to use the new process for Technical Assessments by directly submitting their details into the SAFE which will hopefully streamline the application process. User documentation has been prepared for this new process.

- ARCHER2 service staff will be attending the annual RSE Conference and early plans are taking place for the next ARCHER2 Celebration of Science event.

- During August 2024 we plan to test a new Slurm scheduler fair-share configuration based on project allocations to provide an improved balance for access to resources with size of allocation.

- Globus Online data transfer capacity will be advertised to users.

2 ARCHER2 Performance Report

This is the contractual performance report for the ARCHER2 SP Service for the Reporting Periods from 1 April 2024 until 30 June 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	Apr 2024		May 2024		Jun 2024		Q2 2024	
	Perf	Points	Perf	Points	Perf	Points	Perf	Points
Availability	100%	-3	99.8%	-3	100%	-3	**99.9%**	**-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	0.09	-2	**0.08**	**-6**
SP_Level3 (MTR)	0.00	0	0.00	-2	0.00	-2	**12.18**	**-4**
Initial Response (%)	100%	-1	100%	-1	100%	-1	**100%**	**-3**
Query Satisfaction (%)	100%	-2	100%	-2	100%	-2	**100%**	**-6**
Total		**-10**		**-12**		**-12**		**-34**

2.1.2 Service Credits

As the Total Service Points are negative (-34), no Service Credits apply in this Quarter.

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 2024					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	2227	2227	0	0.115	0:02:52
SP_Level2	97	102	23	8.216	0:21:33
SP_Level3	0	1	0	37	0:04:38
May 2024					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1144	1144	0	0.03	0:00:59
SP_Level2	87	86	24	7.651	0:16:58
SP_Level3	1	0	1	0	0:00:00
June 2024					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	1726	1726	0	0.079	0:01:06
SP_Level2	113	109	27	6.587	0:14:40
SP_Level3	1	0	2	0	0:00:00
Q2 2024					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
SP_Level1	5097	5907	0	0.093	0:01:51
SP_Level2	273	276	25	7.793	0:17:56
SP_Level3	2	1	2	37	0:04:38

2.3 Query Resolution

Metric	Apr 2024		May 2024		Jun 2024		Q2 2024	
Service Level	**MTR**	**Resolved**	**MTR**	**Resolved**	**MTR**	**Resolved**	**MTR**	**Resolved**
SP_Level1	0:01:03	2227	0:00:30	1144	0:00:52	1726	0:00:54	5097
SP_Level2	0:49:34	102	0:33:28	86	0:53:39	88	0:45:13	276
SP_Level3	121:46:24	1	0:00:00	0	0:00:00	0	121:46:24	2
Total		**2807**		**1453**		**1320**		**5375**

A total of 5375 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 38 feedback scores received during this period. 100% were Good, Very Good or Excellent with 92% given the highest score of Excellent.

During May, there were 27 feedback scores received during this period. 100% were Good, Very Good or Excellent with 88% given the highest score of Excellent.

During June, there were 21 feedback scores received during this period. 100% were Good, Very Good or Excellent with 95% given the highest score of Excellent.

£86 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	2024-05-08 0900	2024-05-08 2100	12 hrs	Users unable to connect to ARCHER2. Compute jobs will run and users cannot access data.	Replacement of operating system certificates	HPE
Partial	2024-04-15 1530	2024-05-08 1640	70 mins	Users cannot connect to ARCHER2	Physical moving of the server hosting the ARCHER2 LDAP server	Accommodation
Partial	2024-04-25 0930	2024-04-25 1040	70 mins	Serial nodes, DVN01 and DVN02	Heavy load on metadata server impacted Slurm controller and caused the Slurm daemon to fail on the nodes.	SP
Partial	2024-05-27 0300	2024-05-27 1300	10 hrs	Compute nodes	Power incident on UK national grid in Edinburgh area resulted in loss of power to ARCHER2 compute nodes	Accommodation

3 ARCHER2 Service Statistics

3.1 Utilisation

Utilisation from 1 April – 30 June is 92% which is slightly decreased from 93% the previous quarter. Utilisation for April was 95%, for May 86% and for June 95%.

The slight decrease in the utilisation in May was due to a few factors. Firstly, a full-service maintenance for the renewal of the operating system certificates. A rolling reboot of all of the compute nodes took place over a week to update the Cray Programming Environment (CPE). A power incident on 27 May also impacted the utilisation, and finally some large-scale testing of the new LFRic (Met Office) package took place.

Please note the second Capability Days took place during this period on 4 - 6 June but this did not impact the overall utilisation for June which remained high at 95%

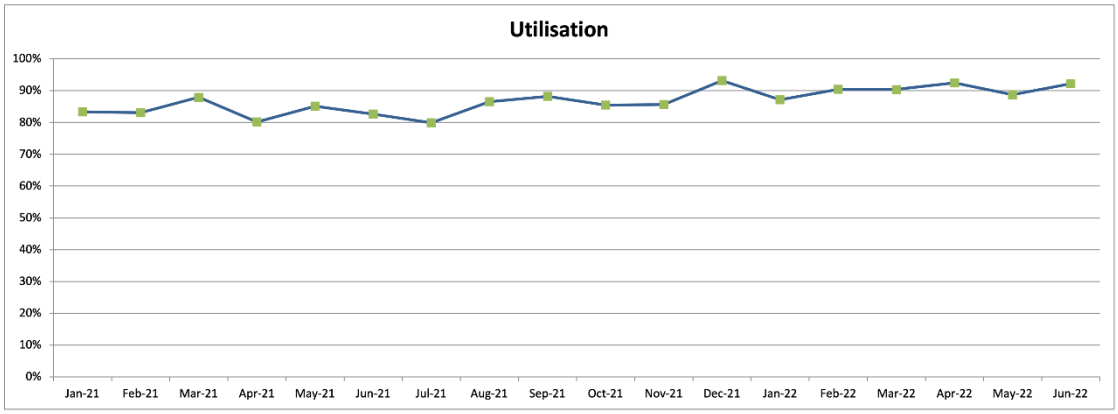

The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER2 by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER2. It can be seen that EPSRC exceeded their target this quarter with their usage being at 74% (against their target of 66.8%). It should be noted that the proportion of EPSRC's uncharged utilisation increased this quarter and is 10.15%. Capability days jobs are uncharged.

NERC missed their target with utilisation being 15.8% (against their target of 18.2%) which was a slight decrease from 16.5% in the previous quarter.

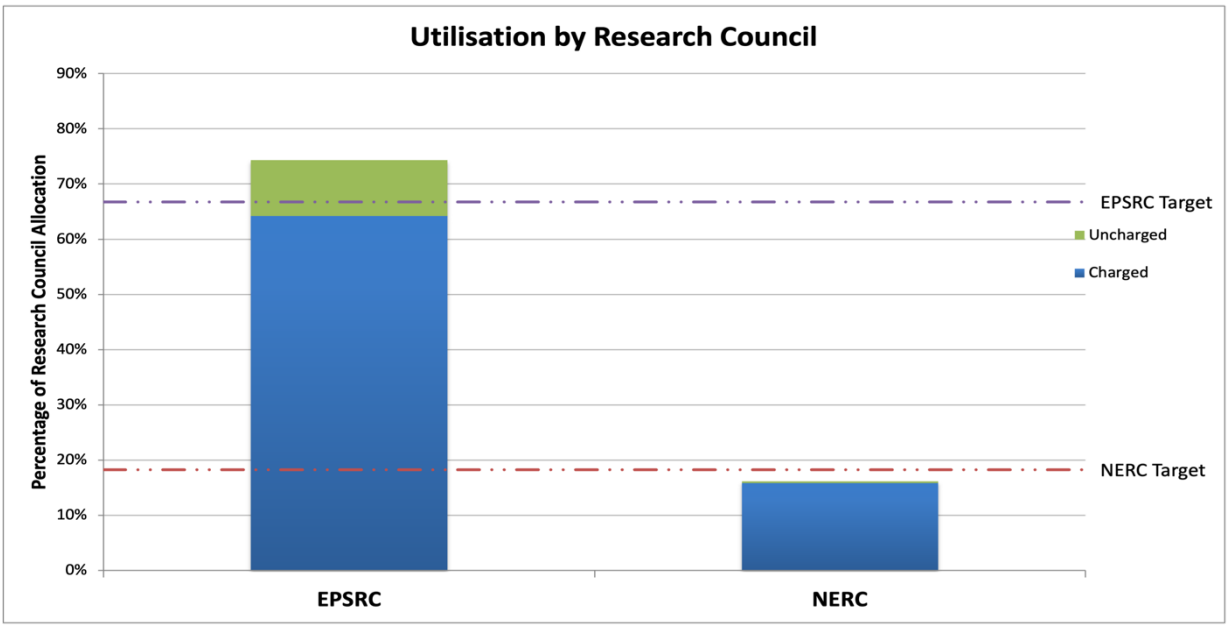

The stacked graph below shows how the trend of charge and uncharged utilisation since the start of the service. There is a trend of uncharged time decreasing and the NERC usage increasing.

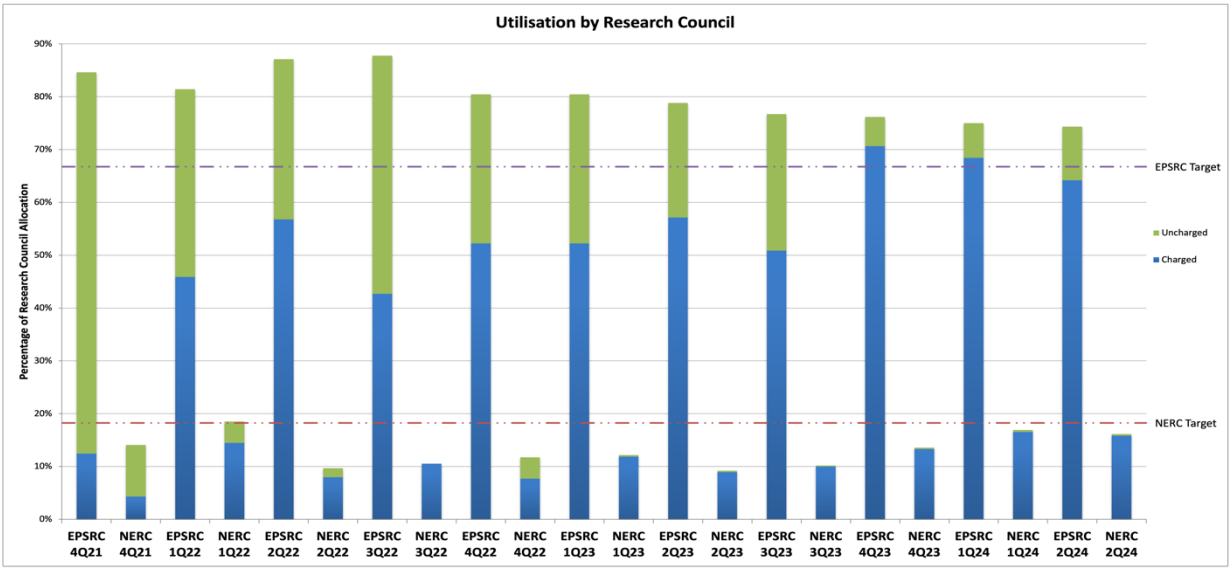

3.2 Scheduling Coefficient Matrix

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.

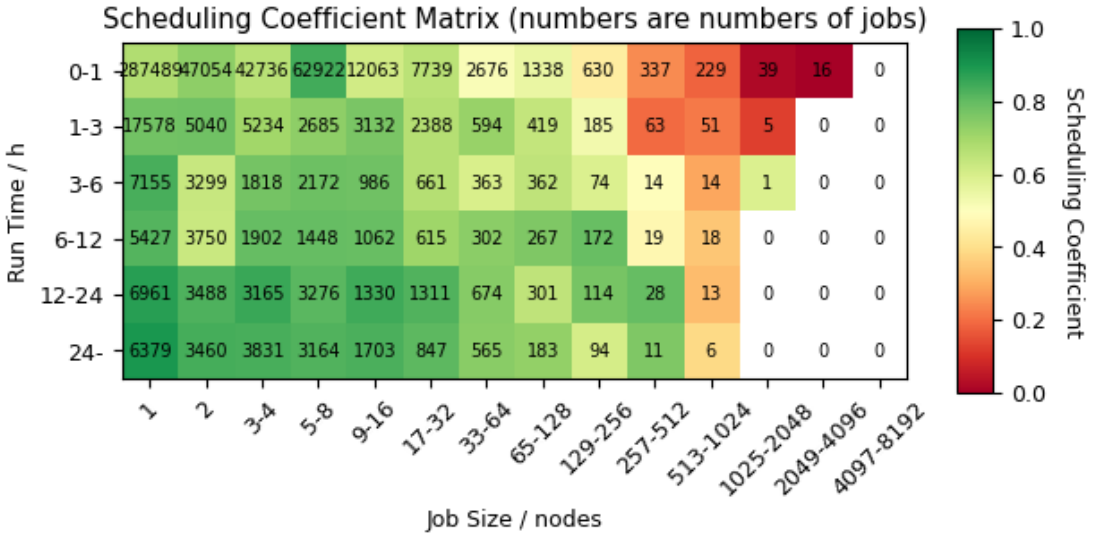

The usage heatmap below provides an overview of the usage on ARCHER2 over the quarter for different job sizes/lengths. The colour in the heatmap indicates the number of CUs expended for each class, and the number in the box is the number of jobs of that class.

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.

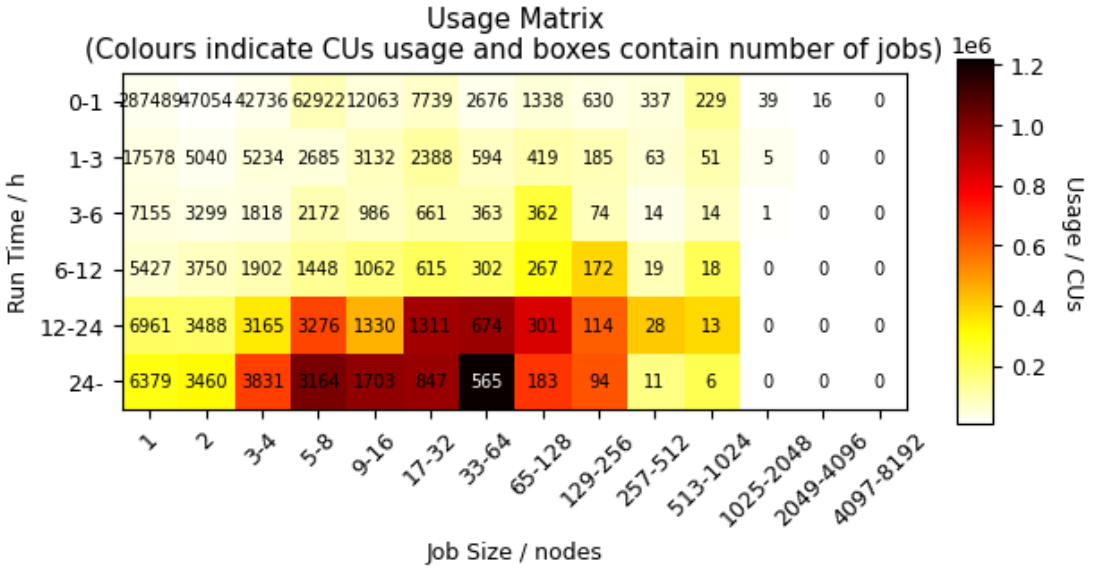

Appendix: Critical Success Factors

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 could result in an increase in passive cooling ("free cooling") [ongoing]

Power Usage

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

* 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	2Q 24
… **to** ARCHER2 (TB)	534	163	68	220	44	67	42	65	99	108	93
…**from** ARCHER2 (TB)	236	582	667	822	834	1231	1022	1472	1771	2056	2443

* Partial

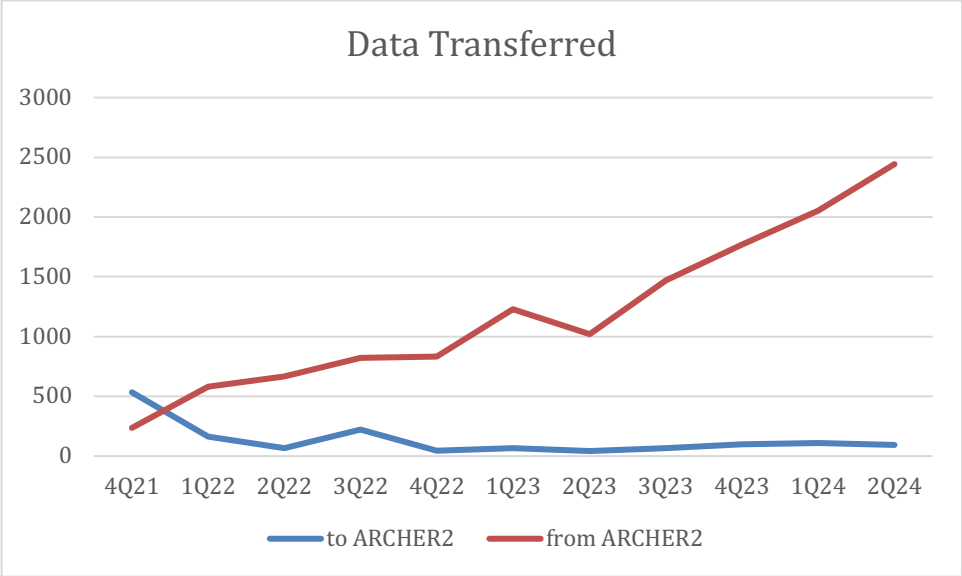

The trend of more data being moved off each quarter continues. Indeed, there was more data transferred off ARCHER2 in 2Q24 as there was in the first full year of the ARCHER2 service.

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 and provide boxplots illustrating the performance variation. (On the boxplots, the green triangles mark the mean values and the whiskers extend to the last datapoint within the range 1.5 x IQR.)

Note: For 1Q24, the analysis exposed a problem with the tests which has now been addressed.

Benchio MPI-IO Medium (GiB/s)	1Q23	2Q23		3Q23	4Q23	1Q24	2Q24
a2fs-work1	8.2	7.6±0.5		10.5 (8.8:11.8)	10.9 (8.3:12.5)	N/A	10.1 (7.0:11.8)
a2fs-work2	8.5	7.3±0.6		10.4 (7.2:12.4)	10.4 (7.7:13.0)	N/A	11.1 (8.0:12.5)
a2fs-work3	8.3	9.6±0.7		10.0 (8.2:11.6)	10.7 (8.1:11.9)	N/A	9.6 (8.4:11.8)
a2fs-work4					9.7 (9.1:10.2)	N/A	10.0 (9.2:10.8)
a2fs-nvme				10.1 (9.6:11.5)	10.1 (9.5:12.4)	N/A	11.1 (10.5:12.4)

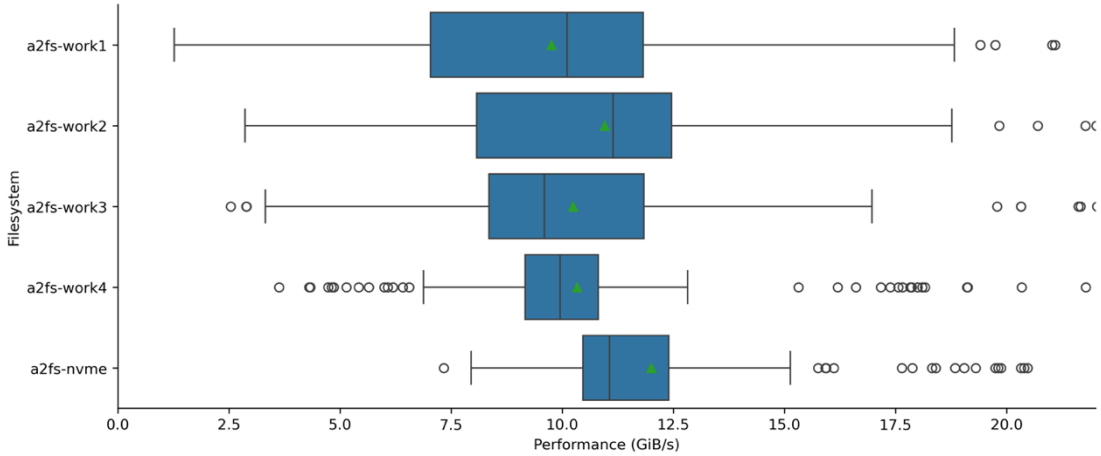

4. CSF08 Be cost-effective, cost-efficient and drive towards lowering of operational costs

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	1Q 23	2Q 23	3Q 23	4Q 23	1Q 24	2Q 24
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*
Energy per CU (kWh)	0.719	0.713	0.728	0.715	0.650	0.545	0.669	0.590	0.568	0.582	0.583

*partial

Energy Cost per CU Delivered

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

*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.