



ARCHER User Survey

2018



1. Document Information and Version History

Version:	1.0
Status	Draft
Author(s):	Anne Whiting
Reviewer(s)	Alan Simpson, Andy Turner, Lorna Smith

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2019-03-13	Initial draft	Anne Whiting
0.2	2019-03-13	Added survey analysis data	Clair Barrass
0.3	2019-03-14	Added analysis text	Anne Whiting
0.4	2019-05-16	Final check of language	Lorna Smith
1.0	2019-05-16	Update after review and pdf format	Anne Whiting

Structure of this paper

Section 2 provides a description of the survey, its questions, the scoring and how it was constructed.

Section 3 gives some highlights of the comments provided by responders to the survey.

Section 4 provides an analysis of the responses received, comparisons to previous years and graphical distributions of the scores.

Section 5 lists the comments received in full and unedited form by question together with the ID of the anonymous respondent.

2. Description of the Survey

The ARCHER User Survey closed on 20 February 2019. 188 responses were received from ARCHER users. The survey asked for ratings (on a scale of 1 to 5) with the following questions:

1. Please rate your overall experience of the ARCHER Service (required) [Very Unsatisfied (1) – Very Satisfied (5)]
2. Has the ARCHER hardware configuration met the requirements of your research? (required) [Not met any requirements (1) – Exceeded requirements (5)]
3. Has the software on ARCHER met the requirements of your research? (required) [Not met any requirements (1) – Exceeded requirements (5)]
4. If you have used the ARCHER helpdesk, please rate your experience [Very Unsatisfied (1) – Very Satisfied (5)]
5. If you have used the ARCHER documentation, did it provide the information you required? [Did not provide the information I required (1) – Provided all the information I required and more (5)]
6. If you have used the ARCHER website, please rate the quality of the content and ease of navigation [Very poor (1) – Excellent (5)]
7. Please rate your experience of any ARCHER Training you have used (either online or face-to-face)? [Very Unsatisfied (1) – Very Satisfied (5)]
8. If you have attended any ARCHER webinars or virtual tutorials, did you find the session worthwhile? [A complete waste of time (1) – Extremely interesting and useful (5)]
9. If you have attended any ARCHER online training material (e.g. Online Driving Test, screencasts), how useful did you find the material? [Of no use (1) – Extremely useful (5)]

Only the first three questions were compulsory for all survey respondents, but 97% of respondents also provided feedback to some of the optional questions. Users were also provided with the opportunity to offer comments or suggestions under all of the above headings, and provided with space for any other comments or suggestions at the end of the survey. These questions are the same as those in the Annual Survey in 2017, 2016 and 2015 and a superset of those in the Annual Survey in 2014 to allow comparison between different periods. As previously, user feedback received will be reviewed to identify opportunities for improvement.

The survey was constructed using Google Forms and embedded directly into the ARCHER website.

Executive Summary

The results of the 2018 annual ARCHER User Survey have been analysed. 188 responses were received with the mean results shown below (scores 1 representing “Very Unsatisfied” and 5 representing “Very Satisfied”):

Service Aspect	2014 Mean Score (out of 5)	2015 Mean Score (out of 5)	2016 Mean Score (out of 5)	2017 Mean Score (out of 5)	2018 Mean Score (out of 5)
Overall Satisfaction	4.4	4.3	4.3	4.4	4.5
Hardware	4.1	4.1	4.2	4.3	3.9
Software	4.0	4.0	4.2	4.1	3.8
Helpdesk	4.5	4.5	4.5	4.6	4.5
Documentation	4.1	4.1	4.2	4.2	4.0
Website	4.1	4.2	4.2	4.2	4.0
Training	4.1	4.1	4.2	4.1	4.3
Webinars	3.6	3.9	3.9	4.2	3.9
Online training	-	4.0	4.1	4.2	3.9

As can be seen, users have generally provided very positive feedback with the overall service receiving its highest rating since the start of service. The number of responses received was higher than in 2017. The rating for hardware and software was slightly down on previous years perhaps reflecting the age of the service and the desire of users for newer functionality. Many users provided suggestions and comments on all aspects of the service. Where users have rated elements of the service under 3 or provided negative comments, the users will be contacted to obtain further details if they have provided contact information.

Selected Quotes

The following quotes reflect the tone of the majority of responders to the survey with regard to the ARCHER service:

- *Fantastic service, best HPC system I have used.*
- *A very good service. If only all computing facilities were run by ARCHER.*
- *The current Cray system 'just works' for me; more so than any other type of cluster.*
- *I would just like to thank the ARCHER Service for their high level of provision.*
- *Thank you - without Archer, I would not be making as good progress with my research.*

Quotes on the helpdesk (which also reflect on the centralised CSE team) echo the particularly high ratings for this aspect of the service:

- *The help desk is amazing. They have helped me with multiple difficult requests (e.g. installing custom software) and really went above and beyond what I expected them to help me with. They always respond quickly and are determined to see the problem through.*
- *Fantastic. The support is also absolutely amazing.*
- *We had to install a dedicated software and got great support in compiling it!*
- *The helpdesk provides outstanding service with great courtesy and in a timely fashion.*

Changes have been made in response to previous user feedback and requests. These include modifying the development short queue so that it is available 24x7 rather than just within working hours. To encourage use of ARCHER at the weekend when it is quieter a discounted weekend queue has been introduced. Users have provided feedback on these changes:

- *Thank you for opening up the short (debug) testing queue as a 24/7 queue. It is very helpful.*
- *not sure if this is the best section, but I am a huge fan of the new short queue.*
- *I have always found the ARCHER service to be very good. I appreciate the initiatives that the service has taken to keep up to date with user usage, e.g. extending the debug queue to be available 24 hours, and starting a new 'weekend' queue.*

With ARCHER heading towards the end of its life, users providing comments highlighting the need for a replacement service:

- *An upgrade seems necessary within the next year or two in order to compete with other world-leading machines.*
- *Could be upgraded as it started years ago.*
- *Ivy Bridge CPUs are quite old now.*
- *From my experience the file-system is slow compared with other machines that I am using for example Hazel Hen.*

3. Ratings

All questions asked responders to rate their satisfaction with each particular aspect of the survey on a scale of 1 to 5 with 1 representing “Very Unsatisfied” and 5 representing “Very Satisfied”. Table summarises the ratings for each aspect and reveals the all aspects of the ARCHER Service are rated highly by users. The number of responses was 188, up from 164 in 2017. Table 2 shows the responses to the 2017 survey, Table 3 from 2016, Table 4 those for 2015 and Table 5 those for 2014 for comparison purposes.

Service Aspect	Total Responses	Mean Score (out of 5)	Median Score (out of 5)
Overall Satisfaction	188	4.5	5
Hardware	188	3.9	4
Software	188	3.8	4
Helpdesk	144	4.5	5
Documentation	176	4.0	4
Website	183	4.0	4
Training	103	4.3	4
Webinars	64	3.9	4
Online training	80	3.9	4

Table 1: Summary of scores for different aspects of the ARCHER Service 2018

Service Aspect	Total Responses	Mean Score (out of 5)	Median Score (out of 5)
Overall Satisfaction	164	4.4	4
Hardware	164	4.3	4
Software	164	4.1	4
Helpdesk	132	4.6	5
Documentation	156	4.2	4
Website	161	4.2	4
Training	98	4.1	4
Webinars	65	4.2	4
Online training	74	4.2	4

Table 2: Summary of scores for different aspects of the ARCHER Service 2017

Service Aspect	Total Responses	Mean Score (out of 5)	Median Score (out of 5)
Overall Satisfaction	161	4.3	4
Hardware	161	4.2	4
Software	161	4.2	4
Helpdesk	136	4.5	5
Documentation	152	4.2	4
Website	155	4.2	4
Training	94	4.2	4
Webinars	64	3.9	4
Online training	70	4.1	4

Table 3: Summary of scores for different aspects of the ARCHER Service 2016

Service Aspect	Total Responses	Mean Score (out of 5)	Median Score (out of 5)
Overall Satisfaction	230	4.3	4
Hardware	230	4.1	4
Software	230	4.0	4
Helpdesk	198	4.5	5
Documentation	215	4.1	4
Website	221	4.2	4
Training	147	4.1	4
Webinars	102	3.9	4
Online training	104	4.0	4

Table 4: Summary of scores for different aspects of the ARCHER Service 2015

Service Aspect	Total Responses	Mean Score (out of 5)	Median Score (out of 5)
Overall Satisfaction	153	4.4	4
Hardware	153	4.1	4
Software	153	4.0	4
Helpdesk	129	4.5	5
Documentation	142	4.1	4
Website	144	4.1	4
Training	81	4.1	4
Webinars	41	3.6	4
Online training	-	-	-

Table 5: Summary of scores for different aspects of the ARCHER Service 2014

Table 6 shows a comparison of mean scores for the different questions across Annual Surveys since the service began. This comparison shows that the mean ratings for hardware and software are slightly lower than in previous years, perhaps reflecting the age of the service. All aspects of the ARCHER service continue to receive very high satisfaction ratings from the users with the overall service rating for 2018 being the highest over the lifetime of the service. The Helpdesk continues to stand out as the highest rated aspect of the service in both surveys with an extremely high mean score. This is testament to the hard work of all service partners (SP, CSE and Cray) in ensuring that responses to the users through the helpdesk are timely, accurate, useful and polite.

Service Aspect	2014 Mean Score (out of 5)	2015 Mean Score (out of 5)	2016 Mean Score (out of 5)	2017 Mean Score (out of 5)	2018 Mean Score (out of 5)
Overall Satisfaction	4.4	4.3	4.3	4.4	4.5
Hardware	4.1	4.1	4.2	4.3	3.9
Software	4.0	4.0	4.2	4.1	3.8
Helpdesk	4.5	4.5	4.5	4.6	4.5
Documentation	4.1	4.1	4.2	4.2	4.0
Website	4.1	4.2	4.2	4.2	4.0
Training	4.1	4.1	4.2	4.1	4.3
Webinars	3.6	3.9	3.9	4.2	3.9
Online training	-	4.0	4.1	4.2	3.9

Table 6: Comparison of mean scores from 2014, 2015, 2016 and 2017 User Surveys for different aspects of the ARCHER Service

As can be seen from Figure 1, the overall satisfaction with the ARCHER service is extremely high with no responders rating the service below 3 on a 1-5 scale from “Very Unsatisfied” to “Very Satisfied.” The mean rating is 4.5, up from 4.4 in 2017. The median rating is up from 4 in 2017 to 5.



Figure 1: Distribution of scores for overall satisfaction with the ARCHER service (188 responses in total).

For the hardware and software (Figure 2 and Figure 3 respectively), the overall satisfaction with the service is high, though lower than for 2017, with 3 users rating the hardware below 3 and 4 users rating the software below 3. There was one rating of 1 (“Very Unsatisfactory”) for the software and none for the hardware on ARCHER this year. The mean rating for hardware is 3.9 (median is 4) and the mean rating for the software is 3.8 (median is 4).

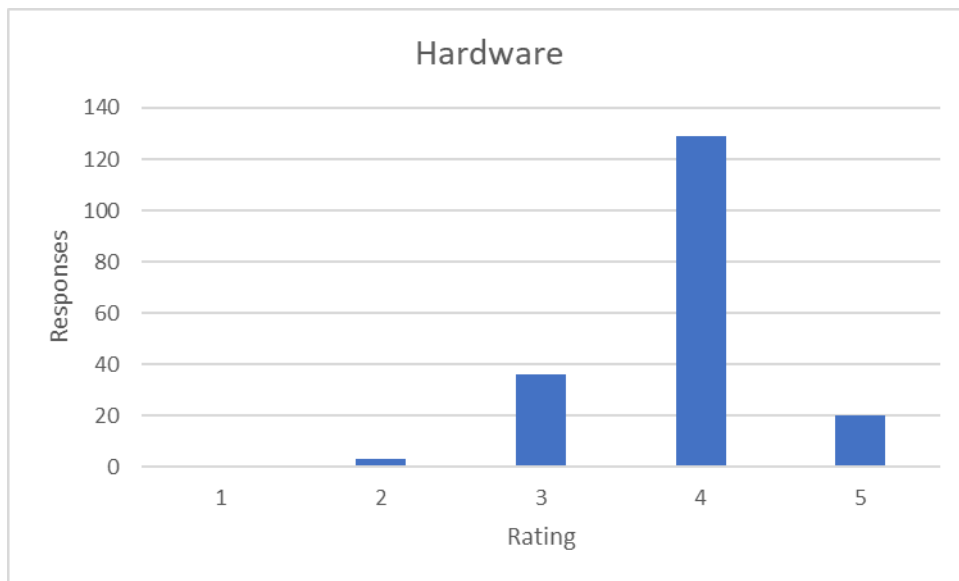


Figure 2: Distribution of scores for satisfaction with the ARCHER hardware (164 responses in total).

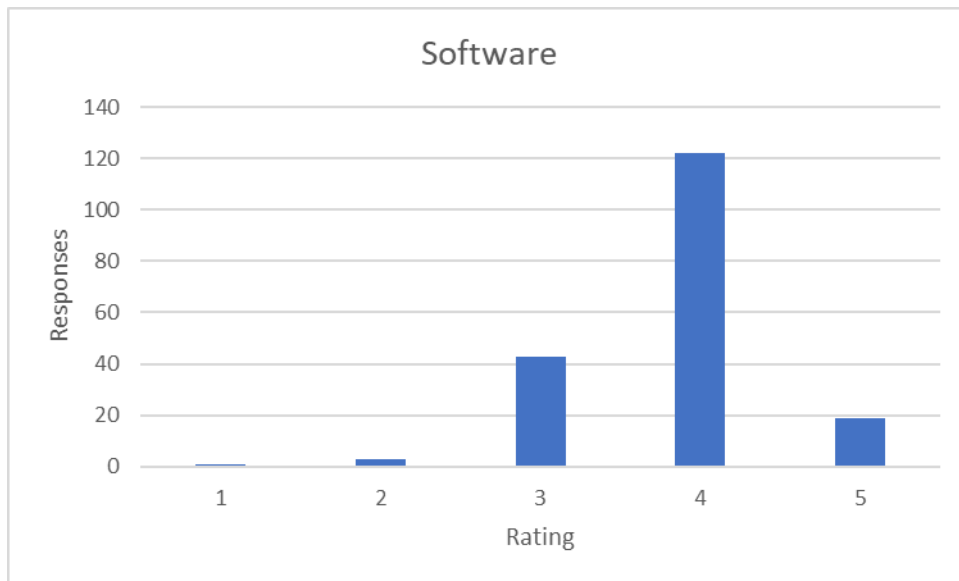


Figure 3: Distribution of scores for satisfaction with the ARCHER software (164 responses in total).

The satisfaction ratings for the ARCHER Helpdesk showed no responses with a score under 3 and a mean rating of 4.5 (median is 5). Of the 144 responses, 85 (60%) gave a score of 5 (“Excellent”). No users gave a score of less than 3.




Figure 4: Distribution of scores for satisfaction with the ARCHER helpdesk (144 responses in total).

ARCHER documentation (Figure 5, mean = 4.0, median 4) and website (Figure 6, mean = 4.0, median 4) show the same high level of overall satisfaction as that shown for the overall service, as well as having high respondent rates. The 6 users who gave a score of 2 will be contacted where they have provided their email addresses to ask for further details.

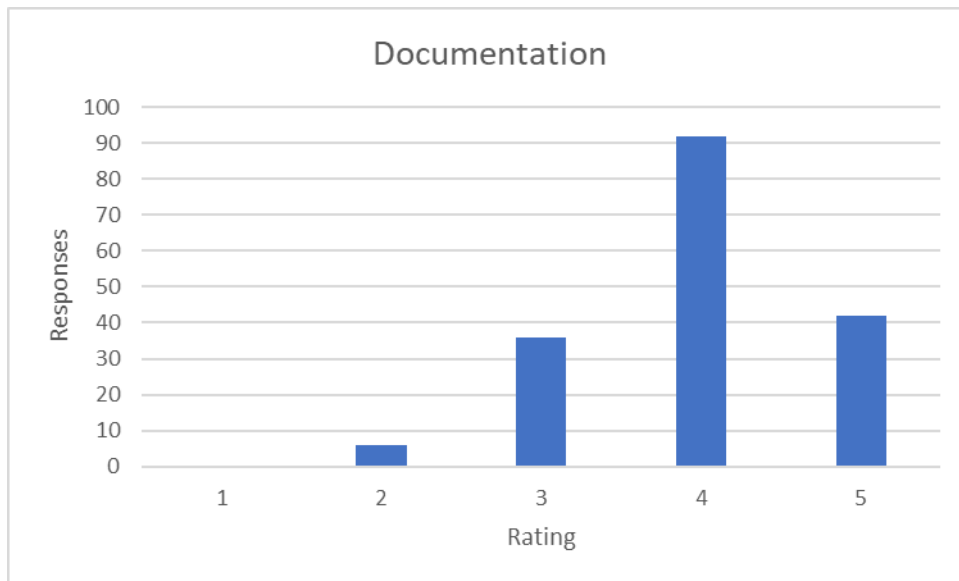


Figure 5: Distribution of scores for satisfaction with the ARCHER documentation (176 responses in total).



Figure 6: Distribution of scores for satisfaction with the ARCHER website (183 responses in total).

The results for ARCHER training (Figure 7, mean = 4.3, median = 4) are high and consistent with the course survey results presented in the CSE Service quarterly reports. There were 15 users with a rating of 3 and no responses with a rating under 3.



Figure 7: Distribution of scores for satisfaction with the ARCHER training (103 responses in total).

The webinars and online training have a lower respondent rate (possibly due to the fact that they are of interest to a subset of ARCHER users) but show a high satisfaction rating (Figures 8 and 9, mean = 3.9, median = 4 for both).



Figure 8: Distribution of scores for satisfaction with the ARCHER webinars (64 responses in total).



Figure 9: Distribution of scores for satisfaction with the ARCHER Online Training (80 responses in total).

4. List of Comments

The comments shown are all the comments received for each question in an unedited form. The number shown in brackets at the end of each comment represents the ID of the anonymous respondent.

Hardware

- Could be upgraded as it started years ago (1)
- The hardware is excellent, and provides the performance that I have been looking for at small and large scales. I think that appropriate hardware has bought and installed for this system (16)
- It is almost satisfy my research usage. (22)
- I use simulation software that occasionally requires very large output files (several gigabytes each) I find that very often writing these files takes upwards of 10-20 minutes each. If multiple files are required on a large job this can inflate the cost of running the jobs somewhat. I have not attempted to quantify or diagnose the cause, but it seems far slower than writing large files on any other machine I use. (27)
- coupled model climate simulations are quite slow on ARCHER compared to the MONSOON HPC platform (32)
- Addition to GPUs (33)
- My work is better served by Tier 2 resources, but it is not possible to directly apply for these in EPSRC proposals. (38)
- Increased RAM on normal nodes are needed for high level GW calculations (39)
- Fantastic. The support is also absolutely amazing.(41)
- Would be great if there were nodes with large (>12GB) memory per core. 43)
- GPU units would be very useful. Sometimes I require more hard disk for temporary scratch space(52)
- Less bespoke, more commodity would reduce the learning curve. (54)
- Access to a larger percentage, or expansion of ARCHER urgently required. (55)
- Conventional configuration nodes make porting and maintaining code easy. The Xeon Phi nodes have been interesting, and work to make use of them has resulted in better vectorisation which also benefitted runs on standard nodes. The overall performance on the Xeon phi nodes was comparable to standard nodes. (70)
- Very good. (73)
- Excellent (83)
- Some of my calculations would benefit from more memory (I underpopulate nodes to get round this) (87)
- Not sure this is a hardware issue. I have a serious problem with memory management of my jobs. Some of my jobs failed with no specific reason and the information about memory usage was not advisable at all. I still don't have a clue how to estimate the memory usage of my jobs, especially for the failed ones. (89)
- Slow to compile in comparison to other machines I have used. Though I understand there might be a trade off between compile and compute speed?(92)
- Might be useful to have a few more high memory nodes, but great otherwise! (95)
- Increase capacity rather than capability (97)
- not sure if this is the best section, but I am a huge fan of the new short queue. (99)
- Ivy Bridge CPUs are quite old now (101)
- Access to GPU resources will cover more areas of research. (109)
- Same executable should work on login, post processing and compute nodes (111)
- At the moment ARCHER hardware configuration works well with applications which use massively MPI and are poorly memory bounded. Unfortunately, I wasn't able to get good performances for the applications I've been working with that don't use MPI and require an high memory bandwidth. Most of them are applications related to Computer Vision. (115)
- Incorporation of GPU nodes will speed up many calculations (126)
- More RAM per node might be nice (127)

- Bigger! (135)
- Hardware failures need to more clearly notified to users. (138)
- Shared queues need to be able to be reservable. For training the shared queue is the weak point as compilation times are exceptionally variable. (146)
- From my experience the file-system is slow compared with other machines that I am using for example Hazel Hen.(148)
- The budgets should be updated more frequently. So jobs can stop immediately when the budget is consumed instead of charging other research groups in the same consortium or project. (149)
- ARCHER needs to be expanded. The UK is lagging behind the international competition in computational resources. (151)
- I'd prefer a DAC that was more interactive and easier to use (155)
- the job queue system can be updated (167)
- I wish there were more nodes in the 'bigmem' queue (more nodes with larger amounts of RAM memory) (168)
- An upgrade seems necessary within the next year or two in order to compete with other world-leading machines. (174)
- Would be good to have other testbeds like the Xeon Phi - or a formal mechanism for using kAUs on other Tier 2 systems. Instead I'm having to invest grant income on my own local testbed may not always be value for money for the taxpayer. (182)
- all is fine at the moment (186)

Software

- Small updates to the software stack (mostly Cray environments, in this case) have had a large negative effect on the performance of programs that had already been compiled. In some cases, recompilation fixed the issue, in others a deeper investigation was required, detracting from the amount of time able to be given to science performed on Archer (17)
- Perfect (24)
- the JASMiN VM for data analysis is useful, but at times significantly slower than working on JASMIN and the user cannot run scripts in batch mode (32)
- Optimised compilations of common chemistry software, with appropriate licenses could save a bunch of hassle and kAUs (39)
- Takes a long time for software to be updated/installed. Why isn't LAMMPS on Archer? (40)
- Most packages I needed were already installed, there were a few lesser-known ones that I needed to install myself (41)
- The compiler and MPI toolchains usage differs too much from non-Cray equivalent tools commonly found on HPC clusters, making the port of software on ARCHER excessively painful (57)
- The current Cray system 'just works' for me; more so than any other type of cluster (84)
- Update gromacs to 2019 version (86)
- The TSK field for output from the qstat command only gives three digits, so jobs with 1000 processors or more are not displayed correctly. (95)
- Very difficult to install and run python software. Need a standardised python environment. (97)
- It would be nice to have more MPI implementations to choose from, e.g. Intel MPI or Open MPI. Cray MPICH2 does not support ILP64 interface, while the other two can. (101)
- Same executable should work on login, post processing and compute nodes (111)
- I wasn't able to find always the release version: intel compiler, chapel compiler for example (115)
- Dynamic Process Management might be very helpful for some users in the future (133)
- make scp work (135)
- Need CASTEP 18.1 (142)
- Would be good if python 3 libs (iris) were easier to install on compute nodes and python3.6 or 3.7 were supported. (150)
- data analysis and data transfer/management software could be improved (151)
- I run a model written in Matlab (not my choice!!), which does not exist on Archer. I understand this is a difficult software to support due to the licensing complications, but I have worked on

other supercomputers that have figured out a system for users to add their own licenses. Right now I am using Matlab Compiler to compile the model on my local machine and Matlab Runtime to run it on Archer. This works, but adds complexity to my workflow and the compiler license was expensive. (152)

- latest mpi4py and ase (atomic simulation environment) modules would be very useful to have pre-installed (154)
- I wish the experience across compute / log-in nodes was more similar and log-in nodes that didn't kick you off after a while (155)
- Update NAMD to latest version (2.13) please (158)
- I have no particular comment on software. I am fully satisfied. (174)
- Older compilers are removed, which creates enormous amounts of work when finicky older codes are not supported by newer compilers. This has happened with HiFi, an MHD solver built on SEL. I don't have the time to research and develop the changes needed to make it work with newer compilers. (176)
- Sometimes have issues with OpenFOAM on Archer, especially with decomposition (181)
- Had to install some software (182)
- all is fine, slight extension the accessibility to depreciated software versions would be useful (186)
- Write a tool (showbf) that tells which resources are immediately available, namely how many nodes are free and for how much time. (I have access to machines where there is this showbf script, and it is enormously useful) (188)

Helpdesk

- Excellent
- I did not have recourse to go straight to the Archer helpdesk, but colleagues at EPCC did, and received prompt responses. From what I can tell they are doing their job well. (16)
- Good enough (22)
- Outstanding help. (33)
- When I have needed any help and advice replies have been speedy and very informative (36)
- The ARCHER helpdesk has always provided outstanding assistance with my queries. (39)
- Quick and helpful responses, although actual work to change any software (ex enable an internal module on an available module) took months (40)
- 10/10 really fantastic support (41)
- They attended all my questions properly. (52)
- The helpdesk provides outstanding service with great courtesy and in a timely fashion. (61)
- Fast response and detailed answer which was very helpful. (70)
- very helpful (73)
- No it is fine (83)
- Prompt and helpful! (95)
- They were prompt and helpful. (97)
- We had to install a dedicated software and got great support in compiling it! (109)
- Your helpdesk should be more responsive when it comes to different software versions that need to be installed. (121)
- Quick helpful responses to any issues (125)
- Very helpful resolving an issue (127)
- Very helpful. Respond within quick time too! (132)
- Very friendly and fast replies. (133)
- Very helpful in setting up training accounts for UKCA training (146)
- they are the best (151)
- The help desk is amazing. They have helped me with multiple difficult requests (e.g. installing custom software) and really went above and beyond what I expected them to help me with. They always respond quickly and are determined to see the problem through (152)
- Thank you so much for your excellent support throughout the year. (165)
- I had very fast responses from the helpdesk which was appreciated. (171)
- All queries have been answered promptly, and staff have been most helpful (173)
- I have always been very happy with their service and the speed of response. (174)
- They are very quick to respond and helpful (181)
- all was helpful (186)

Documentation

- Sometimes it takes time to find the necessary information, especially when it comes to craypat tools, but also with simpler things like resubmitting jobs, creating array of jobs and having dependencies on them. (15)
- Excellent comprehensive documentation (30)
- Not much (or any) documentation on molecular dynamics modules available (gromacs, Tinker, Tinker-HP, etc) (40)
- Documentation was generally good. It would be nice to have more examples on creating PBS submission scripts for various cases such. There was one situation where I needed to run on few cores and more nodes for memory reasons and struggled to get the code running. The support quickly addressed the issues though. (41)
- More detail about the compiling options for the compilers, when to use them and what they actually do. (52)
- Would need an update on the style (61)
- Need more information on python @ archer (73)
- There could be more on the nitty gritty on the parallelization of sequential code. For example, how exactly a loop do or for i=1 to N is parceled up and sent to the different processors. I know the code to do that but it seems to work magically, and I prefer to understand what is going on more precisely. (83)
- There is a summary of the different queues, but it always takes me a little while to find it - it's a little hidden (half way down a page). Perhaps there could be a separate page outlining it? (87)
- The various user guides (intro guide, full guide, best practices) are a little overlapping. It's not always clear which one is the best to consult. I think it would be more helpful to condense them. (95)
- Concise and comprehensive (97)
- when i was signing up via the EPSRC system and so forth I found it quite difficult to find the appropriate documentation, but that's probably more of a them problem than a you problem. (99)
- The command for cancelling a job, qdel, and using the term 'delete' for stopping a running job is not necessarily intuitive for an engineer using the system. I recommend using keywords and like 'stopping', or 'cancel' around the description of the qdel command could help people find it quicker. (104)
- Intel Phi thread/help thread placement unclear (111)
- I was very glad to find the documentation spanned introductory as well as more advanced material, I found it well layed out and easy to use (114)
- some times it is obsolete and with broken links (115)
- Might be more clear on different python versions and associated libraries (119)
- I found not too much info on using qstat and related commands for monitoring jobs (not a problem though, but this could be placed more prominently) (133)
- Has not changed much, could be organised better. (138)
- needs some update in content and style (151)
- Update how to run NAMD (I only managed to get it to work after direct communication & help from James Gebbie as the online instructions were out of date) (158)
- I found the documentation and the user guide are a little confusing and I have do several clicks to find the information (167)
- Perhaps a simple COMMAND LINE script that reminds the user the different available queues and their resource limits (max wall time, max memory, max nodes, etc). In the same spirit as those command line tools at '/home/y07/y07/cse/epcc-archer-tools/bin/' (168)
- Documentation on website is set out helpfully and is easy to query which is good. (171)
- all was fine. Scheduling option (batch job sequences submission) examples for commonly used packages (like WRF or similar) would be very useful. (186)

Website

- the embedded Google search bar was a simple feature to have in the website, however if that comes with the usual "I give you this in exchange of all your information" ethos typical of Google I'm happy as it is (4)
- Search has got harder & finding eCSE reports is not easy as no obvious links. (14)
- Website too convoluted, I would appreciate more direct access to documentation. (34)
- Some more current webinars would be nice (36)
- sometimes difficult to find documentation (40)
- I could find everything I was looking for quite easily. The website is very functional. (46)
- Could be easier to get to the online training information (both pdfs and recorded lectures) (50)
- I really miss a "search" engine (52)
- Also in need of updating and redesigning (61)
- Search option could be improved (69)
- Font and layout is quite cramped and not very easy to navigate for visually impaired people (88)
- Provide a live feed of archer energy consumption (97)
- Training events are added and not emailed through to your contact list immediately, has caused me to miss a course I was eagerly waiting for. (103)
- ARCHER calculator and information were really helpful, especially when supporting students for HPC3-Europe applications. (109)
- see above. it looks old fashioned ((152)
- Can be very slow to generate reports (156)
- There is quite a lot of duplication of content from all of the various course materials. It would be great to have a clearly labelled "master" copy of the most current documentation and the rest just in an archive.e (164)
- May be a forum for user groups to post and answer questions (165)
- Maybe provide a user portal to track the HPC time, etc. (167)
- I have had no issue with the ARCHER website. (174)
- I find I have to talk students & pdras through it (182)
- all is rhte (186)

Training

- All the course leaders were very knowledgeable and while sometimes the course was a miss, then were able to adequately answer my questions after the training so I suggest always leaving a section for user issues, etc. (15)
- You can tell that the training provided comes from professionals that have had time to tune their teachings. The delivery of content, and the content itself is very good. At times, EPCC training materials and presentations are amongst the first search results when looking for information on technical concepts online (e.g. Cray's Aries network and its topology) (16)
- Good for the very beginning level users but I would like more immediate level HPC theory/practicals and programming training (50)
- Increase online/distance learning and training providing some feedback (61)
- Need for more exercises (73)
- Attended MPI training at EPCC, which was excellent. (82)
- More sessions (97)
- The more the merrier! (103)
- I did not received training, but visitors to my lab did: they found it extremely helpful. (109)
- suggest to provide more training opportunities for beginners. (113)
- some times the material is obsolete. For example the Advanced OpenMP doesn't treat enough the offloading that plays an important role in the modern applications (115)
- Intermediate course on using HPC required - courses go from beginners to very very specific. There doesn't seem to be an in-between course. Also a general course on how HPC work. (141)
- training in data management, workflow, new software (151)
- A more in-depth python course would be much appreciated (153)

- (1) More courses on [near] Warwick University! (2) Wish the courses lead to an official certificate, such as that in Canada's SCINET <http://bit.ly/2LZIQ5R> (168)
- My PhD students and post-docs who attended the ARCHER Training sessions have said that they were very pleased. (174)
- Add an N/A option here on the survey (176)
- Adrian Jackson has a charisma in explaining HPC and MPI concepts. (178)
- Training events especially with on-line remote connection are very useful (186)

Webinars

- 'case studies' from researchers who have used ARCHER to discuss work that has resulted in recent publication (36)
- I like it when they are made available (as an archived resource) for people who weren't able to attend the webinar live (40)
- Need for more in depth webinars on HPC (73)
- I'd be interested in more tutorials on parallel profiling on ARCHER/Cirrus. (75)
- Attended one on git version control. Did not find it extremely useful as it just taught the basics, rather than aspects like collaboration on code. (82)
- I'm going to start going through some of these RSE ones as they look like they might be of the right level to get me interested in the nitty gritty. Will report back in the next survey. (99)
- (based on one sample) (100)
- Thanks for all the efforts to make these facilities available to users. Very useful and inspiring. (165)
- More pin-point (and concise) virtual tutorials would be very useful. (174)
- Attendance appears low (182)
- all I attend were great (186)
- useful (187)

Online Training Material

- I am not aware of this material (40)
- Can you please record more lectures with online slides to allow for easier self-study for those who can't attend in-person training sessions (50)
- Well explained (52)
- Not enough actual coding, too much focussed on hardware. (83)
- Might be worth summarising some of the training into summary online multiple choice questions to allow us to revisit and refresh important points (103)
- I wish there was more online video training materials with good production value (151)
- Thanks for all the efforts to make these facilities available to users. Very useful and inspiring. (165)
- This is in line with the above (14 and 15). (174)
- very good (186)

Other Comments

- The only gripe I would have with the Archer service is that the updates that Cray provide to their environment can subtly change the configuration in non-obvious ways. This can cause issues, which do not have a fast turnaround time for being fixed by a large corporation such as Cray, for an issue that may only affect a single application. This makes it difficult to plan longer-term studies of performance of applications on the system. (16)
- overall very satisfactory (24)
- Would be nice to have an overview of training material provided to people when they first set up an Archer account. I found it difficult to navigate through and have later found stuff that would have been very helpful to me at the beginning. (40)
- Fantastic service, best HPC system I have used. (41)
- I would just like to thank the ARCHER Service for their high level of provision. (51)
- The scratch space can be very limited for large groups (52)
- Report generator - although many options are provided it might be useful to have an initial page of reduced options that are typically chosen (and a link to the page of more extensive options). (55)
- A very good service. If only all computing facilities were run by ARCHER... (70)
- Excellent (83)
- Thank you - without Archer, I would not be making as good progress with my research (87)
- Would be very useful if the different accounts owned by the same person could have permission to read and write within the folder tree each account freely. It would save a lot of time transferring the data (sometime order of 100GB) between the accounts. (90)
- Thank you for opening up the short (debug) testing queue as a 24/7 queue. It is very helpful. (94)
- It would be really useful if a user could have a single account associated with multiple projects, rather than being required to have a different account for each project. I currently have 5 accounts and keeping them straight is a real headache. (95)
- Many steps to profile a code, not enough time to get working, could be scripted? (111)
- I just ran out of kAUs on the project I'm allocated (n02-chem), it would have been great to be emailed in advance to be warned that the budget was running low (or it would have been great if this hadn't happened at all!) (112)
- When there are unplanned outages, these are announced on Twitter (which is then visible on the website) but there is no email to all users to let them know. I would really appreciate receiving such emails, as I'm not in the habit of checking the Archer website or Twitter. (152)
- I'd like to have more information about what's going to happen at the ARCHER end of life. (164)
- I have always found the ARCHER service to be very good. I appreciate the initiatives that the service has taken to keep up to date with user usage, e.g. extending the debug queue to be available 24 hours, and starting a new 'weekend' queue. (171)
- More testbeds and a refresh - h/w is old and outdated (182)
- Applications are very difficult and time consuming. an easier access route, possibly a more capped service, would be really useful. (185)
- extend jobs duration limit to say one week (186)