ARCHER User Survey

2016

1. Document Information and Version History

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

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2017-02-24	Initial draft	Anne Whiting
0.2	2017-03-01	Version for internal review	Anne Whiting
0.3	2017-03-08	Review	Alan Simpson
0.4	2017-03-08	Review	Andy Turner
0.5	2017-03-09	Updates after review	Anne Whiting
1.0	2017-03-09	Final version	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, comparison 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 15 February 2017. 161 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 99% 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 2015 and a superset of those in the Annual Survey in 2014 to allow comparison between different periods. 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.

3. Selected Quotes

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

"If we did not have access to ARCHER, we would not be able to do our science! It is an essential service for the UK"

"It's a very good service. Generally reliable and when problems do occur the communication with users is very good."

"ARCHER is extremely consistent, and reliable - as an outsider user I see it as a resilient and a great system"

"ARCHER is an outstanding resource which I feel privileged to use."

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

"This is probably the best helpdesk I have ever come across"

"I've had lots of excellent interactions with the helpdesk. They respond quickly and are more than willing/able to go in-depth as needed."

"I find it very helpful and quick. Keep up the good work!"

"Surprisingly fast answers, always very clear"

There were a number of comments on the queue times and the work done to improve this:

"Archer's compute platform is excellent. The interconnects are great and keep everything moving. The file system remains much better than HECToR's, even with the increased load. Access to the RDF is a really big bonus. The only reason I haven't given Archer a 5 is because the queues have been abominable. This has remained a problem throughout the calendar year. There was a big improvement in from Dec 2016 to Jan 2017 and I was able to work effectively."

"I'm not sure, but I think that the queue wait times have been getting a little better lately. For a while, long queue times really hampered my efforts."

"The policy of giving higher priority to large jobs negatively affected my ability to use the ARCHER service this year, as my jobs are mainly longer jobs which use fewer cores. The changes to the policy implemented over the summer made some improvement to this."

"Queue times were a serious problem in 2016. Changes to how the queues work did improve things, but queue times remain long. Continued assessment of queues and changes that could be put in place should continue to be monitored"

4. 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 1 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 down from 230 in 2015 to 161 in 2016 but up from 153 in 2014. Table 2 shows the responses to the 2015 survey and Table 3 those for 2014 for comparison purposes.

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 1: 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 2: 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 3: Summary of scores for different aspects of the ARCHER Service 2014

Table 4 shows that the mean ratings for different aspects of the service are slightly higher in general for 2016 than the ratings in the previous year. All aspects of the ARCHER service continue to receive very high satisfaction ratings from the users. In particular, the Helpdesk continues to stand out as the highest rated aspect of the service in both surveys with an extremely high rating. 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)
Overall Satisfaction	4.4	4.3	4.3
Hardware	4.1	4.1	4.2
Software	4.0	4.0	4.2
Helpdesk	4.5	4.5	4.5
Documentation	4.1	4.1	4.2
Website	4.1	4.2	4.2
Training	4.1	4.1	4.2
Webinars	3.6	3.9	3.9
Online training	-	4.0	4.1

Table 4: Comparison of mean scores from 2014, 2015 and 2016 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 only 3 responders rating the service below 3 on a 1-5 scale from "Very Unsatisfied" to "Very Satisfied", i.e. 2 percent. The mean rating is 4.3 and the median rating is 4, the same as in 2015.

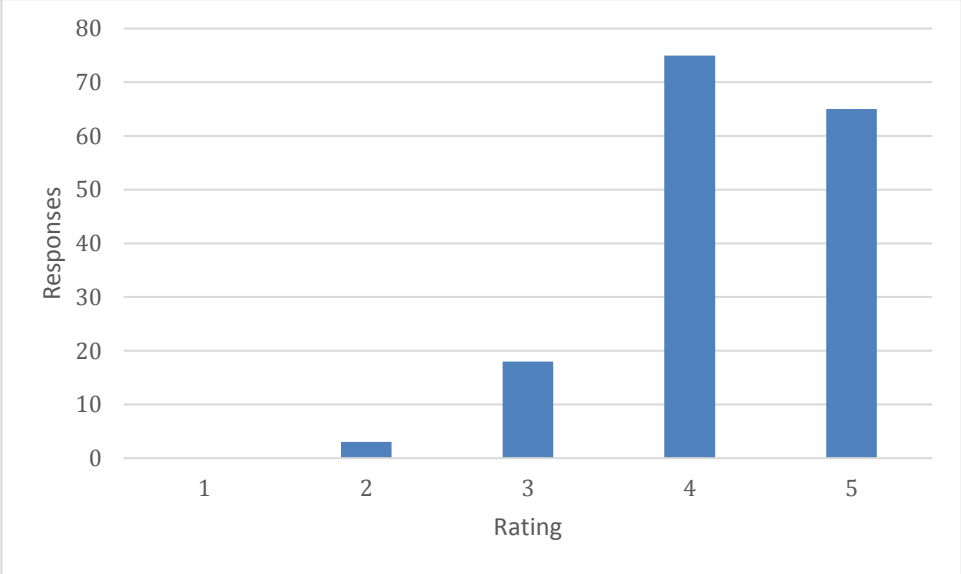

Figure 1: Distribution of scores for overall satisfaction with the ARCHER service (161 responses in total).

Similarly, for the hardware and software (Figure 2 and Figure 3 respectively), the overall satisfaction with the service is high, with only 1 user rating the hardware below 3 and 5 users rating the software below 3. There were no ratings of 1 ("Very Unsatisfactory") for the hardware or software on ARCHER this year. The mean rating for hardware is 4.2 (median is 4) and the mean rating for the software is 4.2 (median is 4). These ratings are slightly up from the 2015 User Survey.

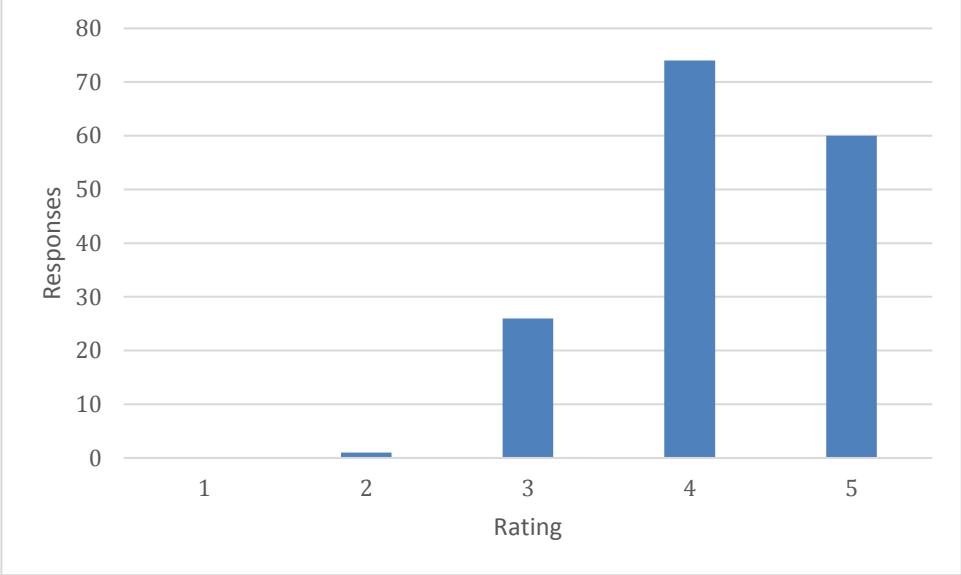

Figure 2: Distribution of scores for satisfaction with the ARCHER hardware (161 responses in total).

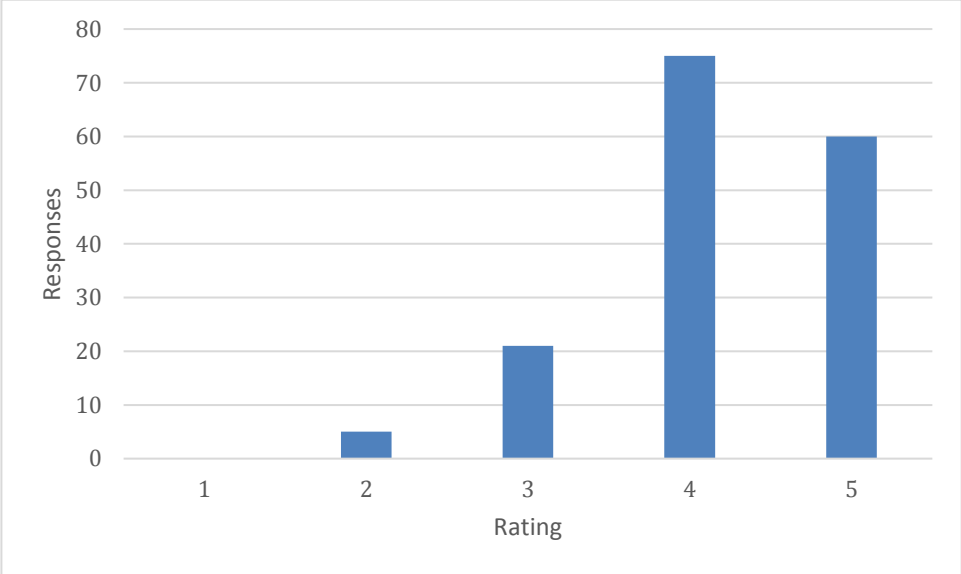

Figure 3: Distribution of scores for satisfaction with the ARCHER software (161 responses in total).

The satisfaction ratings for the ARCHER Helpdesk showed 2 responses with a score under 3 and a mean rating of 4.5 (median is 5). These are identical to the mean and median in the 2015 User Survey. Of the 136 responses 84, 62%, gave a score of 5 ("Excellent"). No users gave a score of less than 2.

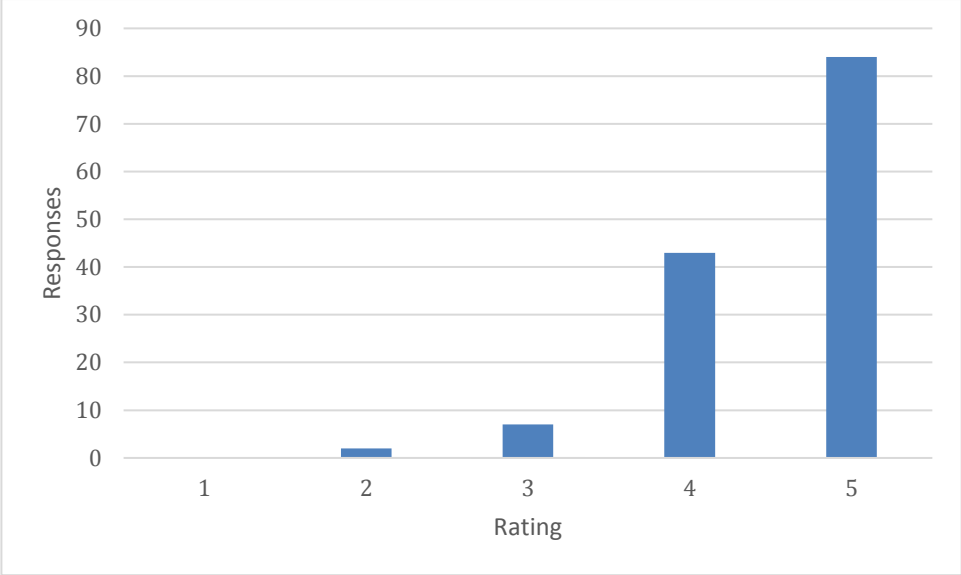

Figure 4: Distribution of scores for satisfaction with the ARCHER helpdesk (136 responses in total).

ARCHER documentation (Figure 5, mean = 4.2, median 4) and website (Figure 6, mean = 4.2, median 4) show the same high level of satisfaction as that shown for the overall service, as well as having high respondent rates.

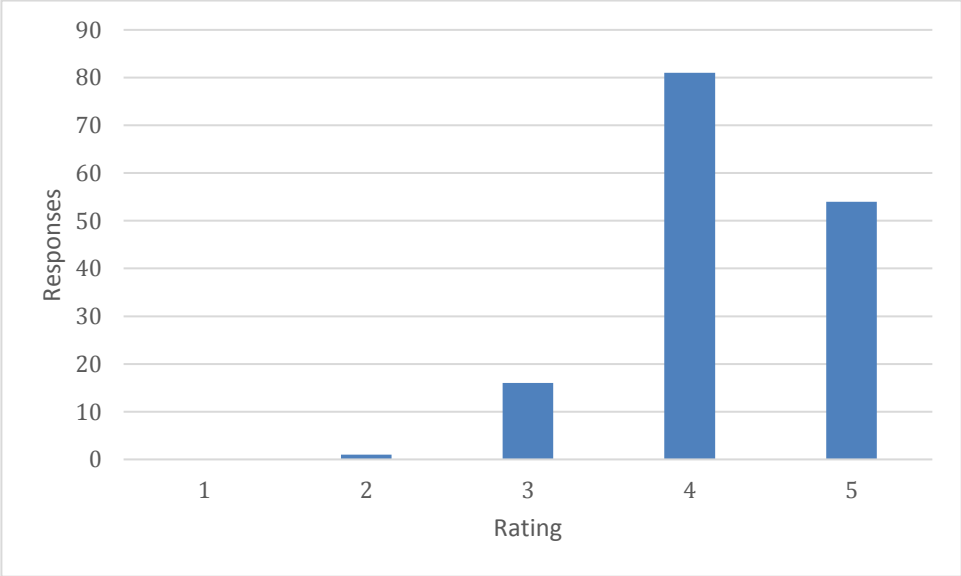

Figure 5: Distribution of scores for satisfaction with the ARCHER documentation (152 responses in total).

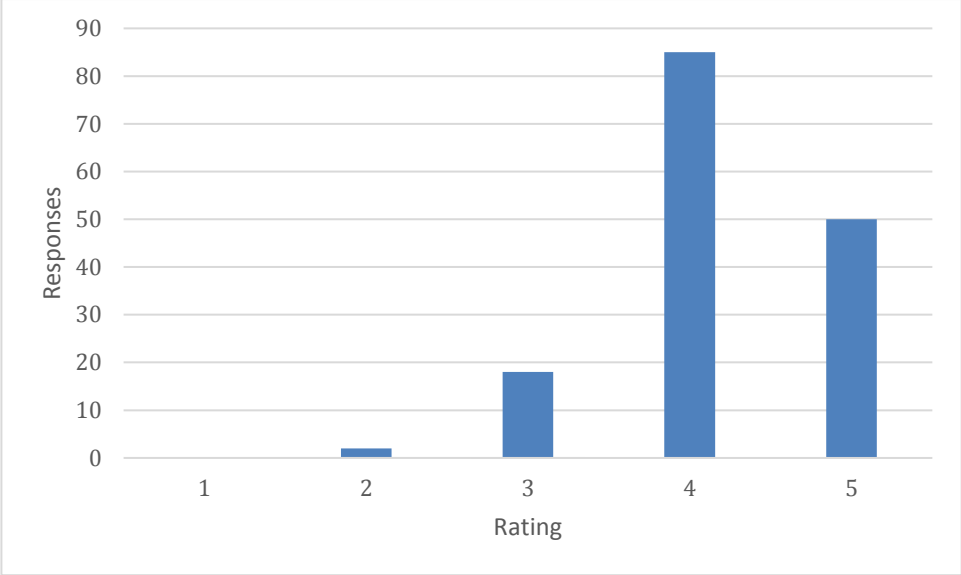

Figure 6: Distribution of scores for satisfaction with the ARCHER website (155 responses in total).

The results for ARCHER training (Figure 7, mean=4.2, median = 4) are high and consistent with the course survey results presented in the CSE Service quarterly reports. There are no comments from users with scores under 3 and a number of responders have scored for training with a comment that they had not attended any training. These scores are all also very similar to the scores from the 2015 User Survey.

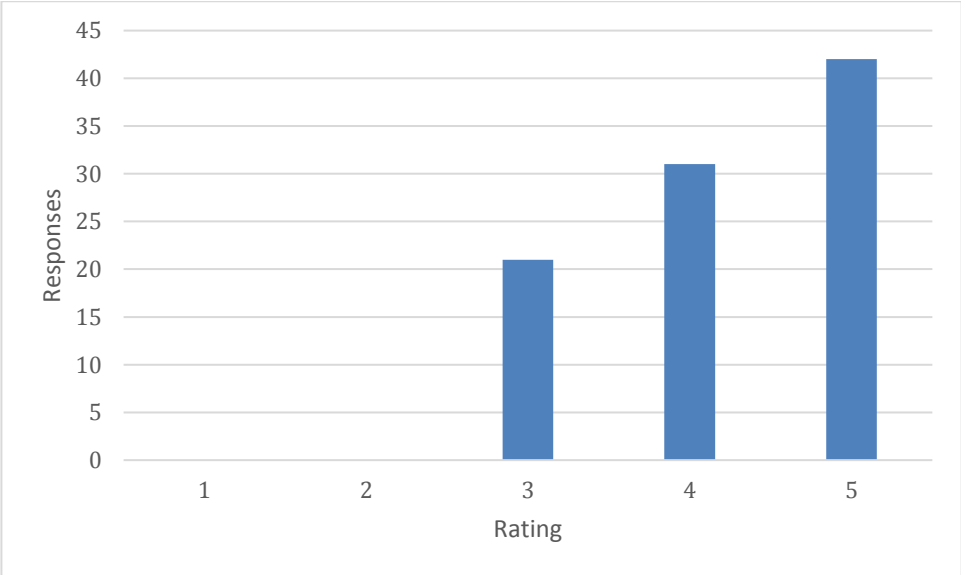

Figure 7: Distribution of scores for satisfaction with the ARCHER training (94 responses in total).

The webinars and online training have a lower respondent rate (possibly due to the fact that the technical nature of the webinars is of interest to a subset of ARCHER users) but show a high satisfaction rating (Figures 8 and 9, mean = 3.9 and 4.1 respectively, median = 4).

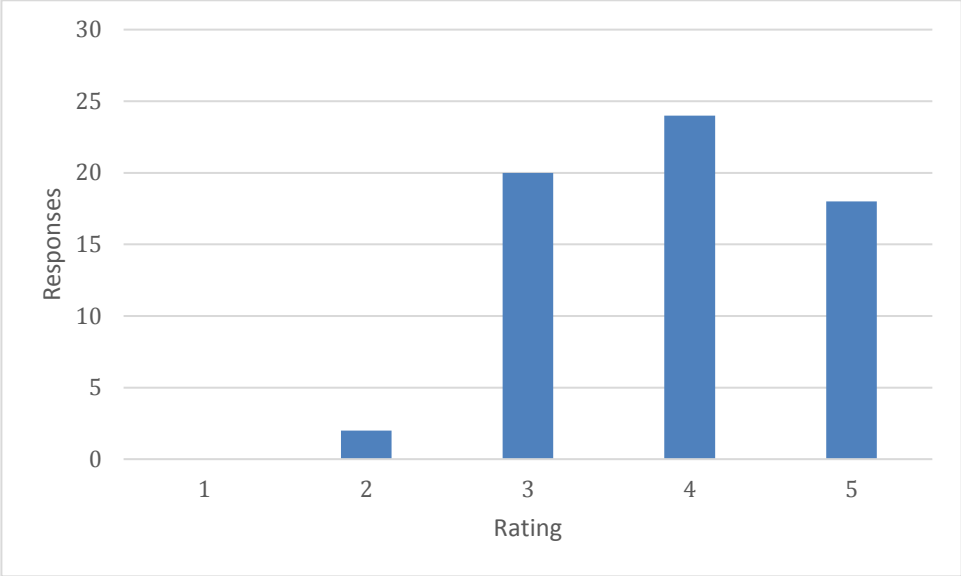

Figure 8: Distribution of scores for satisfaction with the ARCHER webinars (64 responses in total).

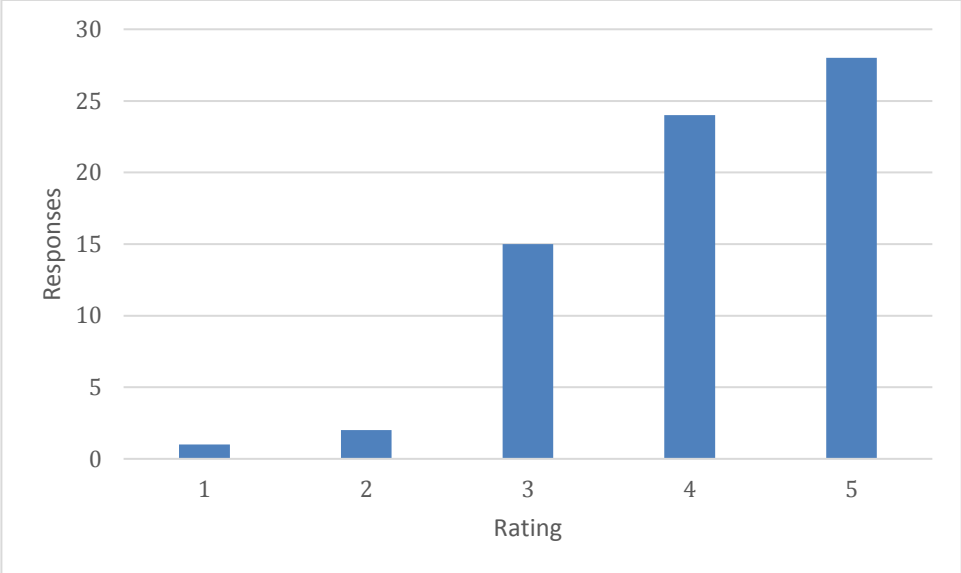

Figure 9: Distribution of scores for satisfaction with the ARCHER Online Training (70 responses in total).

5. 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

- Queue times were a serious problem in 2016. Changes to how the queues work did improve things, but queue times remain long. Continued assessment of queues and changes that could be put in place should continue to be monitored (8)
- The hardware seems appropriate for the type of DFT calculation I'm running. (18)
- There has been some downtime every now and then. Failing disks, nodes etc.
  I'm not sure whether the impact on users can be reduced, but it is rather problematic when it happens. Also, the use of PP nodes I find confusing. Users can submit jobs to them, or simply log in and use them. This, I think, means they can overburdened with jobs, explaining why computations there are often really slow. (21)
- Sometimes the memory requirements force me to depopulate the nodes. Maybe more large memory nodes would be helpful (22, 28)
- I wish more GPUs were available. Even the cheaper series of the NVidia 10 series would significantly help with the computation in the case of molecular dynamics. The power consumption per computation is also significantly lower. (30)
- Very fast nodes, sometimes the queue is very long though. (31)
- Nodes with more memory...and more capacity! Jobs queue for quite a while these days.(38)
- Don't waste money on underused GPUs/Accelerators - more x86.(43)
- Can sometimes take a while to do runs, but this is because ARCHER is so popular and has got better more recently. I have done some 32k core runs which all went through fine, I would ideally like to do some runs in the 65k-100k core range but think it would probably take up too much of the machine.... so it would be nice if "ARCHER 2" had many more cores ;) (46)
- Faster disk read/write. (48)
- Please increase the number of computational cores available by at least one order of magnitude!!!
  It is not possible at the moment to do on ARCHER production runs of highly-resolved simulations with dozens of thousands cores. We are trailing behind France and Germany in Europe and behind the US and China worldwide in terms of number of AUs available for academic research. (50)
- More cores please. The queue times are too long.(51)
- Looking forwards to having a play with Knights Landing nodes.(52)
- software is more important than hardware (56)
- The queuing times for UM execution have been poor this year and this has adversely affected my ability to develop the model. The ability to use e.g. github repositories is also poor due to the reluctance to allow the use of ssh-agents on ARCHER. (57)
- Memory remains an issue in many different of calculation (hybrid DFT) (59)
- It's a good mechanism for making novel architectures like the Knights Landing nodes widely available. I haven't had occasion to use the x86
  nodes as my high performance codes are written for GPUs or Knights Landing.(62)
- A GPU cluster would be a useful potential future development. Higher memory nodes (even higher than the 128GB ones currently available) would also help to meet growing computational needs in memory-bound problems. Current Intel chips are beginning to show their age in terms of performance. (63)
- It mostly works very well for what I want. I do like the experimental KNL cluster. It seems to work well. (64)
- I would suggest to include a GPU-based resource. (65)
- Running MPI with a fixed number of cores. Straightforward. (67)

- ARCHER is extremely consistent, and reliable - as an outsider user I see it as a resilient and a great system (72)
- It is great to have access to nodes with different hardware to port and test code (e.g.KNL) (84)
- ARCHER is overloaded; increases the availability/number of cores. Demonstrate a convincing set of scheduling tests showing that you have evaluated different scheduling algorithms (e.g.slurm) which can schedule large numbers of small jobs as well as small numbers of large jobs, the current ARCHER preference (87)
- Improved visualisation and data analysis hardware directly attached to the rdf (88)
- Over the last year, there seems to have been more unscheduled outages than previously - perhaps related to an ageing system? (92)
- I'm not sure whether the system is over-subscribed or what, but the waiting times over the last year have been ridiculous. (94)
- More hardware is always better. Would be nice to reduce wait times and allow for bigger jobs. A filesystem that can deal with many files would be useful in some cases. (96)
- The National Service needs to provide access to GPUs (97)
- It's been OK but looking a bit long in the tooth now. A lot of our codes run very well on GPUs, it would be nice to have a significant amount of these available. (105)
- The queuing time for more than half a year has been unbearable, not sure if this was a management issue or due to the cluster not being big enough for the request (128)
- more taskfarming, i.e. single core jobs (134)
- An expansion on the KNL nodes might be a useful stopgap measure before retiring the machine to allow for optimisation of codes and such. (135)
- It serves us very well (142)
- Archer's compute platform is excellent. The interconnects are great and keep everything moving. The file system remains much better than HECToR's, even with the increased load. Access to the RDF is a really big bonus. The only reason I haven't given Archer a 5 is because the queues have been abominable. This has remained a problem throughout the calendar year. There was a big improvement in from Dec 2016 to Jan 2017 and I was able to work effectively. (148)
- Could be made easier to run small-scale jobs. (151)
- It could do with having more nodes (10%?). It appears to be running near capacity, and queuing times have increased significantly over the last 18 months. (160)

Software

- I have enquired about installing Gaussian09 on the cluster, which admittedly does not scale well at enormous nodecounts, however works well up to a dozen nodes. The majority of G09 users would require maximum 4 nodes but for longer times rather than the large corecounts required for solid state jobs. I was hoping this could be looked into again. (12)
- Maybe you could consider to install a scientific software in between ab-initio and molecular mechanics, such as tight-binding DFT. (20)
- I find the scheduler and its rules confusing. In ocean modelling, the jobs we want to run often need ~20-30 nodes for ~2 hrs, and the jobs are often tuned and hardwired to use a certain number of nodes.
  It is therefore annoying that there is no way of running jobs of this size on the "debug" queue. Instead, we have to use the "standard" queue, in which the job can easily sit for a day. If we are testing something new, it is very impeding on the work flow to wait for a day before the job runs for 30s and crashing due to a minor bug. Also, why the debug queue is only open during weekday work hours, I don't understand. I, for instance, sometimes work on weekends, and to not be able to use the debug queue is quite annoying. (21)
- All software required is on Archer which is great (22)
- The prioritisation of large jobs is surprisingly good, which encourages full usage of the system as a capability machine.(24)
- More profiling tools (25)
- I hope ATAT code can be installed for all users. (26)
- Cray compilers do not work on our code, we have to use intel (28)
- The new configuration and (fortran) compilers are very aggressive and vectorize strongly (previous DECI iterations on ARCHER worked like a charm). This is hard to control, and batch runs which are fine on other machines will diverge on ARCHER (I run a development version of ABINIT mainly - the existing module you have suffers from the same issues). I have also tried other compilers, to no avail, and do not want to sacrifice performance completely. This is also a problem on other machines with ifort > 14 or so, but I have not been able to solve it fully on ARCHER. I should probably open a new ticket, but reproducing the results is not always simple or quick.
  In short, it would be _very_ beneficial to have a tool to tune/limit this kind of optimization on ARCHER, or to check executables for "vectorizability", memory alignment, etc... The issue has become critical for many applications and MIC will be even worse! At any rate, your operation is truly professional and strongly appreciated. (31)
- improve the data virtualization with parallel mode working with multi-cores .
  For example, VisIt cannot work using multi-cores job submit. (45)
- Decent balance of range, stability and timeliness (52)
- Would have loved to use screen (but this may be a safety issue?) (54)
- Fortran coarrays work uncovered several CCE bugs which took a very long time to fix or work around. (56)
- See below about making the Intel performance evaluation and tuning tools available.(62)
- The only thing I find harder to do that I would like is simple profiling. I don't really need all the detail that I can get from the complex profiling (because I'm not going to change my code depending on the cache size) but simply finding out how much time is spend in different parts of my code is not as easy as it could be (it ought to be no more complicated than adding a compiler switch, or changing a runtime module). (64)
- Everything perfect for my needs. (65)
- Experience only with compilers. (67)
- It would be great to get the pathways plugin for vmd (70)
- The only thing I cannot do on ARCHER is use matlab for pre- and post-processing data (72)
- I rate the Cray XC / slurm environment on ARCHER very highly: it makes my life easier supporting an application that depends on c++ and fortran and mpi and hdf5 and python and ... when those components are so well supported by the system. As far as I am concerned, *this* is the standard by which other clusters are measured.(81)
- Everything we need has been available (84)

- You MUST support all relevant codes; comments made by ARCHER representative at recent materials chemistry consortium (MCC) meeting indicated 'streamlining' of code support; this is not appropriate to a scientific community (87)
- If more space was available, it would be great. (91)
- Clear, easy to use (92)
- I would like to use quantum chemistry software like Gaussian or Qchem but they are not available in archer (102)
- We are making increasing use of python. The Archer environment is OK, but could probably do with a clean-out and refresh - some of the system-level packages and modules are a bit idiosyncratic.(105)
- I am aware that processes that detach from the shell are disallowed on login nodes (they are automatically killed). I find terminal multiplexers like GNU screen and tmux to allow me to work much more efficiently, but these get killed due to this policy. If it were possible to run something like tmux or GNU screen on the login nodes, it would be very useful to me (and I imagine many other users). (110)
- It would be useful to be able to run Matlab on ARCHER for basic post-processing. I have an individual license.That being said, I should probably just go ahead and learn Python! (111)
- Making it easier to submit smaller jobs, perhaps by giving them a separate queue, would be very useful. (121)
- As above, we might need to do some work to be able to run codes efficiently on new hardware. (135)
- Update the intel compilers to a new version >=17 (138)
- Python 2.6 on compute nodes, outdated packages. (139)
- Install the latest version of NAMD as the current version is quite old now (141)
- All the software we use and want to use is on Archer. The NAG teams that optimise the most popular softwares that we use is also very useful. (142)
- The software itself is fine. Having multiple version of everything can be very helpful when compiling new codes. (148)
- I am doing very heavy post-processing, the limits set for the use of serial processors are very restricting for my work. Particularly the 24 hour limit to run jobs, and that you can only use 5 processors at the time and que 10 other jobs. (149)
- Intel compilers are slightly old but this is outweighed by the amazing profiling software (153)
- No, it's very well set up. (160)

Helpdesk

- Always respond promptly and are very helpful (6)
- This is probably the best helpdesk I have ever come across (19)
- I think the service status page is really well done. (21)
- Very helpful and great help!! (22)
- Very helpful and courteous. (24)
- Very long turn around on Helpdesk inquiry. (32)
- Service generally great, but this seems like a suitable place to complain about the frustrating restrictions on password choice that result in having to either write it down or reset every time a change is forced. Why do I need to recall the existing password for use with passwd when forced to change it? I *don't use it*, I use SSH keys to login. Password reset through SAFE was slow last time I tried. (52)
- I always get a quick response from ARCHER Helpdesk. I am grateful for the Helpdesk team for their support. (53)
- Thanks for responding to my questions so quickly. (54)
- Occasionally there's confusion with reply emails, I think some Cray support mails get straight to me, with no copy to helpdesk. (56)
- Turn around times poor. (57)
- I totally understand the password policy, but sometimes the password criteria are negative (cfr. https://xkcd.com/936/) (59)
- Regarding my enquiry Q808817 I think it was unhelpful to reconfigure the Knights Landing queues so that the example submission script (still) in the documentation stopped working, because the maximum wall time had been lowered from 30 minutes to 10 minutes. (62)
- I find it very helpful and quick. Keep up the good work! (63)
- Works very well. (64)
- Everything perfect for me. (65)
- surprisingly fast answers, always very clear (66)
- I have used the helpdesk mainly when installing new software, and found them very helpful and quick to respond. They also do not shy away from tricky and involved problems. (72)
- Always very quick response (84)
- Recent experience with polling algorithm proved to me that parts of the best practice guide were obsolete. You MUST at least keep the 'best practice guide' up to date !!! (87)
- Very quick response to several queries I had during the last year; comprehensive, accurate answers too. (92)
- Always reply very quickly. (94)
- Always prompt and helpful. (96)
- Helpdesk always responds quickly to enquiries. (103)
- The Helpdesk continues to be really good. (105)
- I appreciated the speed with which the helpdesk responded to my requests. (110)
- I've had lots of excellent interactions with the helpdesk. They respond quickly and are more than willing/able to go in-depth as needed. (111)
- Always got prompt response and useful advice, very satisfied. (122)
- One agent should deal with a request as it can be confusing talking to different people about the same problem (130)
- Not really. Help is usually prompt and efficient. (135)
- Very prompt service and helpful staff. (137)
- Always very helpful (142)
- The helpdesk staff are generally helpful, but the response time is too slow. I usually solve any issues myself before the helpdesk responds. (148)
- Always helpful and courteous. (160)

Documentation

- Gromacs compilation documentation does not seem to work. (4)
- Not updated for sometime, for example how to compile certain software (9)
- It's quite hard to find help on which libraries to use and compile with - the information is present, but is hard to locate. (13)
- A little more introduction on code compilation would be helpful. (15)
- References to full PBS Pro docs could be better; tend to find command-line flags and queue features by browsing relevant topics which feels inefficient (52)
- Very good documentation. Can't think at this point what's missing, but I think there were moments when I need to go beyond the archer docs. (56)
- The "instruction" for rdf could be improved. I guess the main reason that most people do not use it, it is the complexity of login in and out from such nodes. (59)
- The documentation for using the Knights Landing nodes could have addressed confusing issues for former Knights Corner users, e.g. that -mmic is now not the flag to use when compiling (despite its name and apparent purpose) and that using it will lead to error messages implying a misconfigured system. It would be useful to document how to use numactl within the Cray programming environment. Most generic Knights Landing documents assume one can get a shell on the execution node itself to invoke numactl. (62)
- Some of it needs updating. Most of it is fine, but some sections are a little sparsely worded and could use more examples. (63)
- It is basically very good. A few more real-world examples would always be nice. (64)
- It was clear and effective. (65)
- The variety of compilers available for MPI was confusing. We ended up doing test runs to decide how to compile our code for efficient running. (67)
- Some of the info hasn't been updated for new versions of the software (gromacs) (71)
- See my comments to previous question about the best practice guide (87)
- Comprehensive. (96)
- There is a good amount of documentation which is very helpful. Only problem is that the information is spread over several different guides, so that it sometimes not clear where to look for a piece of documentation. (103)
- It is pretty good, clear and easy to understand. (109)
- The documentation is very comprehensive and has been a great use to me. (110)
- The manuals are comprehensive and easy to follow when using the system. (117)
- Seems to cover most things. (135)
- Generally very good, although took me some time to find out how to run long jobs. Wasn't very clear from what I read. (137)
- A cheat sheet could be created. (143)
- Quite comprehensive with good links to further reading (153)
- The KNL documentation was a bit patchy, but the main service is well documented. (160)

Website

- Better guidance on how to use safe to request project access and access to correct budgets. (14)
- Easier means of viewing/managing time. (19)
- Status page is very useful (25)
- More information regarding the recent queueing time would be helpful, the graphical information available at the moment is not much helpful. (27)
- Navigation could be further improved by making content accessible via search topics that group related content, rather than the tradition, linear layout, which requires the user to navigate in and out of topics. (48)
- Takes quite a few clicks to see (e.g.) individual usage report for last week/month. Possible to put a summary/overview on SAFE login page? (61)
- Calculator and status pages are really useful. (65)
- Missing/unable to find reports for completed eCSE projects. (73)
- SAFE password algorithm is too stringent (87)
- In SAFE, I would like an overview page that shows remaining kAU for all my accounts on the same page. (96)
- Web page works fine. (103)
- The site generally contains excellent content and is easy to browse, on a large screen. It is much harder to navigate on a small screen (e.g. on a smartphone). I suggest that the layout of the website for small screen devices could be improved. (110)
- Beginners like me can find the resource needed to start using the HPC. (117)
- The website is very easy to navigate and well laid out. Its search function always finds the most useful page. (148)
- Tabs in SAFE could be made a little clearer - initially difficult to locate where information about credits associated with accounts is. (151)
- I have always found what I am looking for although it isn't the simplest site. (153)

Training

- Completed the ARCHER driving test and found all the online resources very helpful and informative (17)
- I did not follow any of them, due to shortage of time; nothing related to quality or proposed subjects. The thing is that people tend to learn new techniques when they are in need. I am not sure you can access those webinar after the scheduled broadcast. It could be a good idea to have a repository, so one can access those in the case they need that. (20)
- I expect more advanced not basic training. (26)
- Unfortunately, you don't offer training for beginners from zero and this does not meet the ambitious ! (34)
- They are really good. Language-specific courses would be good, i.e. Fortran 2003/8, C++ in built parallelism in C++11/14. (63)
- I've not been able to get to any of the ones that I thought were interesting or relevent. (64)
- I really like the residential courses - more please!!! (87)
- Would be nice to have more courses on advanced / specialist topics. (103)
- Face-to-face is usually more efficient that on-line. (135)
- An improvement would be to replace the "Driving Test" with an interactive learning course (1-3 hours long). Some of the questions at the driving test are meaningless for actual usage. (143)
- I attended a python course. It was a mixture of very basic stuff escalating too quickly to very complex stuff. Since it was a short workshop providing more materials to study a posteriori: code, explanations, exercises of tool commonly used in geoscience/oceanography would have been make it much more useful. (149)
- Extremely good, especially for free training. Some content is hard to follow but overall good. (153)
- Not attended, but I cannot uncheck the box (156)

Webinars

- Have attended several webinars, unfortunately some technical troubles made things hard. (6)
- I did not follow any of them, due to shortage of time; nothing related to quality or proposed subjects (20)
- I haven't attended any, partly because system requirements seem too much. It never works on my FreeBSD laptop. (56)
- The webinar environment works very well. (62)
- Did not attend virtual tutorials. Had difficulties setting up software on my computer. (103)
- I love the ARCHER virtual tutorials which include thousands of instructional videos. (109)
- It was easy to follow helped me get started using ARCHER. (117)
- There were some technical issues during the session, with the slides not working. (135)
- An improvement would be to replace the "Driving Test" with an interactive learning course (1-3 hours long). Some of the questions at the driving test are meaningless for actual usage. (143)

Online Training Material

- When starting any professional development IDE very often the "did you know" window is presented with some interesting act. Possibly a similar thing could happen upon logging into archer. (30)
- Driving test information seemed over complicated at times - it was understandable only after hands-on use of archer. (48)
- The Driving-test was a nice idea. (59)
- I'd rather this effort went into better regular documentation. (64)
- Useful but overshadowed by the residential courses (87)
- The screencasts were an ideal way to learn how to use the system initially. (94)
- Online driving test is a good idea - learned something new when taking it. (103)
- I think it is pretty good. I can find all my problems in these materials. (109)
- The material helped me a lot in determining my needs and making my experiments suitable to the system. (117)
- An improvement would be to replace the "Driving Test" with an interactive learning course (1-3 hours long). Some of the questions at the driving test are meaningless for actual usage. (143)
- Training manual is good. The driving test is largely irrelevant. There are a number of questions that bare little relevance to using ARCHER or best practices and the questions with multiple selections fail you if you miss one. To me it doesn't matter if you know every reason for doing something, only that you know at least one and that you do it. (153)

Other Comments

- Make ERC grant CPU cost same as EPSRC, if possible (9)
- Very satisfied with ARCHER; except the queues can be a pain sometimes, especially coming up to the expiry of one's allocation! (37)
- ARCHER is an outstanding resource which I feel privileged to use. The supporting services are quite adequate and allow ARCHER to be used to a very substantial part of its full potential. (52)
- Very good overall! (56)
- Very silly note. It is painful to avoid to use the -color option for ls, any workaround for that? (59)
- I'm not sure how one addresses this, but there's a tension between learning the Cray tools for code performance evaluation and tuning on ARCHER, and learning the Intel tools that are available at my institution. Ideally the Intel tools would be available on ARCHER too.(62)
- If we did not have access to ARCHER, we would not be able to do our science! It is an essential service for the UK (72)
- Would it be possible to stagger the allocations, such that there is not such a rush at the end of allocation blocks - this would avoid big queues and make it more feasible for users to manage their final calculations. (76)
- It's a very good service. Generally reliable and when problems do occur the communication with users is very good. (83)
- PLEASE make it easier for small jobs to execute - and be more polite when people e-mail to complain that their jobs are not executing - telling me 'you should not leave all your simulations to the end of the grant time allocation' is not useful when I've just spent 14 days waiting for simulations to launch (87)
- As a group leader, I still find the tracking of usage non-intuitive via the website at the moment. I (90)
- Live long and prosper. (103)
- Thanks for a generally great service. (110)
- I'm not sure, but I think that the queue wait times have been getting a little better lately. For a while, long queue times really hampered my efforts. (111)
- The policy of giving higher priority to large jobs negatively affected my ability to use the ARCHER service this year, as my jobs are mainly longer jobs which use fewer cores. The changes to the policy implemented over the summer made some improvement to this. (113)
- Better communication about system issues. (162)