



# ARCHER User Survey

2015



## 1. Document Information and Version History

<b>Version:</b>	1.0
<b>Status</b>	Final
<b>Author(s):</b>	Anne Whiting
<b>Reviewer(s)</b>	Alan Simpson, Andy Turner

<b>Version</b>	<b>Date</b>	<b>Comments, Changes, Status</b>	<b>Authors, contributors, reviewers</b>
0.1	2016-03-07	Initial draft	Anne Whiting
0.2	2016-03-22	Further updates	Anne Whiting
0.3	2016-04-13	Updates	Anne Whiting
0.4	2016-04-18	Reviewed	Andy Turner
0.5	2016-04-19	Post-review updates	Anne Whiting
0.6	2016-04-19	Reviewed	Alan Simpson
1.0	1016-04-19	Version for EPSRC	Alan Simpson, Anne Whiting

## 2. Description of the Survey

The ARCHER User Survey closed on 17<sup>th</sup> February 2015. 230 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 responders, but over 93% of responders also provided feedback in at least 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.

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:

*"I find the ARCHER service to be responsive and helpful, easy to use, and flexible. Excellent service."*

*"The online driving test was how I obtained access to ARCHER, and I think it is a brilliant idea for widening access. This has allowed me to put together an eCSE application for the next call, and hopefully accelerate my research in future."*

*"Archer really is a great service. I have used other supercomputers but the assistance and the guidance on Archer is significantly better."*

*"The UK National Supercomputer service is an essential tool in support of scientific research within the UK."*

*"An excellent service! Thank you for all your hard work!"*

Quotes on the helpdesk (which also reflect on the centralised CSE team) echo the high ratings for this aspect in particular are shown below:

*"Just to re-iterate the helpdesk / support makes archer so beneficial and has really saved me an immense amount of time when trying to compile software, or when a problem has arisen in a calculation."*

*"Fantastic. They know when to bring experts in on particular package specific questions also."*

*"Always helpful, and I like the fact that they check with you before closing a job to make sure everything is fixed. "*

*"All help has been prompt and directed towards improving my experience. Some suggestions to improve performance have been offered without my request and that's also proved very useful."*

*"The helpdesk tends to be quick to respond, efficient, and helpful, even with fairly complex requests. Great job! It's always a pleasure to work with the helpdesk."*

There were only two comments on the file systems problems:

*"Disk failures and prolonged use of RaidCheck diagnostics greatly reduced the performance of the system and the throughput of my jobs over the past 12 months. "*

*"The issues with the filesystems have had an adverse effect on our research work, mostly in terms of a time lost. "*

A full list of the comments can be found in Section 5.

## 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 up from 153 in 2014 to 230 in 2015. Table 2 shows the responses to the 2014 survey for comparison purposes.

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

**Table 1: 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)
<b>Overall Satisfaction</b>	153	4.4	4
<b>Hardware</b>	153	4.1	4
<b>Software</b>	153	4.0	4
<b>Helpdesk</b>	198	4.5	5
<b>Documentation</b>	142	4.1	4
<b>Website</b>	144	4.1	4
<b>Training</b>	81	4.1	4
<b>Webinars</b>	41	3.6	4
<b>Online training</b>	-	-	-

**Table 2: Summary of scores for different aspects of the ARCHER Service 2014**

Table 33 shows that the mean ratings in 2015 for different aspects of the service are very similar to the equivalent 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)
<b>Overall Satisfaction</b>	4.4	4.3
<b>Hardware</b>	4.1	4.1
<b>Software</b>	4.0	4.0
<b>Helpdesk</b>	4.5	4.5
<b>Documentation</b>	4.1	4.1
<b>Website</b>	4.1	4.2
<b>Training</b>	4.1	4.1
<b>Webinars</b>	3.6	3.9
<b>Online training</b>	-	4.0

**Table 3: Comparison of mean scores from 2014 and 2015 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 8 responders rating the service below 3 on a 1-5 scale from “Very Unsatisfied” to “Very Satisfied”, i.e. 3 percent. The mean rating is 4.3 and the median rating is 4. These ratings are very similar to the rating from 2014 User Survey where the mean rating was 4.4 and the median rating was 4.




Figure 1: Distribution of scores for overall satisfaction with the ARCHER service (230 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 5 users rating the hardware below 3 and 9 users rating the software below 3. The single rating of 1 (“Very Unsatisfactory”) for the hardware on ARCHER was accompanied by a comment stating that the hardware is fine but the turnaround for NERC jobs was not good recently. The rating of 1 for the software on ARCHER is from the same user who gave the hardware a score of 1 and has a negative comment about recent job turnaround. The response has been submitted anonymously so we cannot follow up the comments with the responder. The mean rating for hardware is 4.1 (median is 4) and the mean rating for the software is 4.0 (median is 4). These ratings are exactly the same as those from the 2014 User Survey.




Figure 2: Distribution of scores for satisfaction with the ARCHER hardware (230 responses in total).




Figure 3: Distribution of scores for satisfaction with the ARCHER software (230 responses in total).

The satisfaction ratings for the ARCHER Helpdesk showed 4 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 2014 User Survey. Of the 198 responses 128, 65%, gave a score of 5 (“Excellent”). One of the users who left a score of 1 indicated that they did not wish to be contacted regarding their response to the survey and the other has been contacted for feedback as they did not leave a comment.




Figure 4: Distribution of scores for satisfaction with the ARCHER helpdesk (198 responses in total).

ARCHER documentation (fig 5, mean = 4.1, median 4) and website (fig 6, mean = 4.2, median 4) show the same high level of satisfaction as that shown for the overall service and have high respondent rates. The results for ARCHER training (fig 7, mean=4.1, 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. As for the other aspects, these scores are very similar to the scores from the 2014 User Survey.

The webinars and online training have a much lower responder 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 (figs 8 and 9, mean = 3.9 and 4.0 respectively, median = 4). There was some

apparent responder confusion between training and online training and webinars, and thus the questions in the survey for 2016 will be refined to avoid this.




Figure 5: Distribution of scores for satisfaction with the ARCHER documentation (215 responses in total).




Figure 6: Distribution of scores for satisfaction with the ARCHER website (221 responses in total).






Figure 7: Distribution of scores for satisfaction with the ARCHER training (147 responses in total).




Figure 8: Distribution of scores for satisfaction with the ARCHER webinars (102 responses in total).




Figure 9: Distribution of scores for satisfaction with the ARCHER Online Training (104 responses in total).

## 5. List of comments

The number in brackets indicates the ID of the respondent.

### Hardware

- It is not possible to use part of processors on a node: this constrains scalability to a multiple of 24, i.e. the number of processors on a node (14)
- It would be extremely useful to have a few nodes with more memory (17)
- Less crashing (23)
- The focus on traditional CPUs, rather than accelerators, is great for scientists needing to spend time on science, not rewriting their software (where their algorithm is amenable) for GPUs. For my applications, interconnect is not the most important thing, and the fastest cores possible is the priority. (28)
- Would be nice to have access to a few GPUs for trial purposes. (32)
- Big memory nodes are very useful - more of these would be great. (34)
- Very pleasant to use, the system fast and the queues are generally quite quick under normal usage loads. (40)
- Please stop doing maintenance mid week... it causes huge queue back logs and destroys the flow of the week at times (48)
- The hardware and interconnect is excellent, and much better than the regional N8 machine (scalability is ~ 200 cores on N8 vs 6000 cores on ARCHER). (50)
- The issues with the filesystems have had an adverse effect on our research work, mostly in terms of a time lost. (55)
- There is insufficient storage space to post-process any data. Data transfers are slow. (58)
- Same architecture and software across compute and serial queues. Longer access (>12hrs) and better performance of serial queue nodes. In future in memory computing would be useful to reduce data size and output to disk. (60)
- RAM on individual nodes are on lower side. Most codes scales well with 8 GB per core RAM. Also, if general queue allow 72 hour jobs, it will be helpful. (62)
- Any way to better handle many small files? (66)
- Long queue and waiting time for the interactive sessions (67)
- unknown, yet. (68)
- I often use the large memory CPUs available on Archer and due the smaller number of those specific CPUs the waiting times are much longer compared to the standard CPUs. Possibly I would like to see an increase of the number of such CPUs. (73)
- Nice and fast! (74)
- It would be nice if there was less lag when accessing file. (77)
- To date, I have only used ARCHER for training courses. (80)
- Currently trying out DDT reverse connection. It worked briefly which was good but only a toy code for testing. Its great that we have this facility (89)
- It would be nice if there was a default 6 month extension of CPU hours beyond the end of a grant. Many papers are submitted at the end of a grant and extra time for reviewer comments would be useful. (91)
- I have produced a substantial number of simulations on ARCHER this year. Mostly I have been very impressed with the speed of jobs going through the queue. However, there have been instances where jobs have waited for several days before being processed and this has held up my research. (102)
- The jobs I run on Archer typically perform best with numbers of cores that are powers of 2. With 24 cores per node, it is not possible to reach a power of 2 and use all available cores at the same time. It would be more useful if future systems had something like 16 or 32 cores per node. It

would also be useful to increase the number of high memory nodes to a total of 512. There are some jobs I would like to run that would benefit from this number of high memory nodes. (104)

- ARCHER can provide more ability for the post-processing nodes since I found that the speed for my post-processing with OpenFOAM utilities is always very slow. These jobs are submitted to serial queue in Archer. (108)
- Very reliable, good performance (109)
- Great for running large jobs. However because the queues are often long it isn't that practical to debug smaller jobs etc. (119)
- None at present as I have only just been trained on how to use it (121)
- Transfer speeds between /work disk to /nerc and /nerc to JASMIN even using lightpath are quite slow. That has been the major limitation of how much simulations we can perform within given days rather than the computation speed or length of queue. That's quite a shame isn't it? (126)
- The compute times and queue times are manageable. However, the file system is often extremely sluggish during weekdays when people use ARCHER. Jobs that normally take ~2 hrs can sometimes hit the 5hr walltime due to slow read/write speeds. (127)
- Excellent hardware for my Group's needs! (128)
- For big data tasks more use of the MOM nodes for a mother program (i.e. for extremely parallel quantum monte carlo) would be nice in the future. Otherwise great. (131)
- It's be helpful to have a larger /work space (132)
- The only issue I have is that reading and writing files seems quite slow. Particularly when I want to tar up a set of results so I can copy them elsewhere. It's a fairly minor point. (133)
- Stable, fast, good amount of memory per core. Only slight caveat is the lack of diskspace. (138)
- Stable, fast, good amount of memory per core. Only slight caveat is the lack of diskspace. (143)
- Availability of >2GB/core memory as well as the existence of large memory node has been fundamental to sustain our research. (146)
- The hardware is fine, when jobs run. However, over the past few months the turnaround for our NERC jobs has become terrible. With decent turnaround Archer would be excellent for our work. (151)
- Faster cores and faster networking (probably lower latency rather than bandwidth) would give better performance for our massive atomistic simulations. Don't get me wrong, ARCHER is fantastic but our 40 million atom simulations scale linearly to 3072 cores but only 80% efficient at 6144. So I cant give "5 Exceeded requirements" although I love the machine. (152)
- You may want to experiment adding accelerators on some nodes. I can see that you may gain experience on this field, which may become more important in the future. My suggestion however, is to keep the ARCHER as the very large computer cluster that it is right now, and upgrade it with faster memory and more powerful processors. (153)
- Good performance per node (155)
- Disk failures and prolonged use of RaidCheck diagnostics greatly reduced the performance of the system and the throughput of my jobs over the past 12 months. Additionally, the scheduler's strong preference for large (>200 node) jobs means that users who, like me, require smaller number of cores for longer periods of time are at a disadvantage. Many of my simulations require weeks to run, but use only 16-64 nodes. Throughput for these jobs is poor. (160)
- It may be a software issue rather than hardware. The file system slows down frequently and once several files are in a folder (about 1000 or so) the filesystem becomes difficult to work with (161)
- More storage space would be even better. (171)
- I am very happy with ARCHER and have not had a problem yet. (178)
- The hardware is very good, my down-voting is more in terms of availability. To work efficiently I'd need very quick access to a development queue available for debugging, can't wait for 2h in the development queue to "sort out the edges"...Also for production runs queue times have been ridiculously long this year. (182)
- I was wondering, whether Archer is able to have a large memory login node for some simple post-processing use if possible. (183)

- In the last years the Barcelona Supercomputing Center has maintained simultaneously a CPU based machine (Marenostrum) and a GPU based machine (Minotauro). With the growing popularity of GPU based simulation codes, this may be a suitable option to explore in the future. (186)
- As always, we could use a bigger and faster computer... (187)
- Prioritise disk space before number of cores. It is very easy to generate TBs of data with the current size of the machine, but there isn't the disk space available to hold it. The shared nodes are utterly useless for compiling code, and make me not want to use the service at all. As the shared nodes are exactly that, and are often overloaded, what should take ~10 minutes to compile times out at an hour or more. (188)
- Matlab on the few post-processing nodes would be quit helpful. (195)
- Archer hardware excellent, especially largemem cores. (198)
- just we wait sometimes for long time in the queue, if there is any solution for this problem we will be grateful. (208)
- Very powerful. (209)
- Just as I need it to be. (216)
- I use the RDF a lot at the moment. It's very useful for model data storage. I will use the CPUs later this year again, for UKESM. (218)
- Good architecture, but could do with being a bit larger. I/O can be extraordinarily slow, and there's a lot of noise in run times still (+/- 20% is typical, but sometimes much larger) and it would be good to solve this. (225)
- Please open some queue with longer walltime. For long (10 days) simulations, the 24h walltime means wasting our time in restart (229)
- Hardware is good. The number of processors per node (24) can make domain decomposition awkward, since a power of 2 is usually preferred, but this is a small matter. (230)

## Software

- Installation of often required Matlab (1)
- Sometimes package upgrades (which often break code) are not as well documented as removals. (13)
- The linear equation solver that I used seemed slow in comparison to other systems, but I only found this out after the event as it wasn't a major part of my code. Some documentation pointing to performance of different linear algebra tasks on different compilers would be useful. (18)
- You could add BerkeleyGW as a preinstalled module. (19)
- A permanent test queue is needed with short wall-times. With the current setup, testing is very difficult as even small jobs may queue for in excess of 24 hours. (22)
- The Cray compilers have often caused trouble, and I have instead had to rely on the GNU programming environment. Libraries and packages however have never caused me problems. (28)
- I'd be happy if Midnight Commander (unix file manager) operates in a full regime, i.e. with the subshell which actually works. (29)
- Can take a very long time for smaller jobs to start running - sometimes up to 5 days. Would be good to tweak the scheduling system to make this wait a bit shorter. (34)
- Paraview runs slow with large datasets, but this may be inevitable. (36)
- The only comment that I have is that I would like automatic highlighting on the head node for file types e.g. tar.gz and for directories to make it a little simpler to distinguish between different items. (40)
- ARCHER rdf could benefit from having a few more software. For example, NCL could be installed as well as a software (other than gs) for viewing pdf documents. (44)
- I would need to use software to perform QM calculations like Guassian, QChem, Orca. I understand these programs do not parallelize very efficiently but it is what I need for doing my research (48)
- Generally the compiler is good, but the current default version of the Cray C compiler is buggy and doesn't compile my code (it compiles on all other platforms/compilers without issue). The next version compiles the code OK. (50)
- Some software was out of date but once I emailed the helpdesk, they updated it to the latest version quickly enough. (53)
- The python changes have been a bit of a pain - hopefully it'll be stable now. (56)
- The code I'm using is LAMMPS (not the provided version but one which includes modifications). It is a pity that LAMMPS does not compile properly with the default compiler, and it is recommended to switch instead to the GNU compiler as this surely affects performance. I don't know whether this is a compiler problem, a LAMMPS problem or something else. (57)
- Same software/compilers for compute and post processing nodes. Graphical/visual parallel debugging + performance analysis would be useful on compute nodes - could be done on interactive nodes if same architecture+software environment is used. It can be difficult to debug/test using small non-production type runs. (61)
- Is it possible to have an estimation the waiting time after submitting a job? (67)
- Compilers and environment modules system is excellent. ARCHER is by far the easiest environment I have used to compile scientific software, and a good selection is available (74)
- To date, I have only used ARCHER for training courses. (80)
- Maybe it is already possible, and if so please accept my apologies for this comment. I think that it would be useful for some people (I know actually quite a few of them I work with, particularly in the bio-physics, photo-physics, molecular dynamics sectors) to get access to very-long queues. Something like lasting a week. They do not need large systems, or large nodes (materials scientists like me do) just a few CPUs working uninterruptedly for a week or two. They are not willing or it is not inherently possible to implement middle check points in their software.

Furthermore, you would get more users from Engineering background by adding some more multi-physics software or finite elements software. (86)

- I use only my software but I have access to all the libs I need to compile it (91)
- support a bit slow, python modules struggling to keep up with e.g. releases on Ubuntu LTS? (92)
- More complete python compute set of modules for scientific computing would be useful and save some time installing them locally. (101)
- It would be very useful if users were allowed to run screen on the login nodes. Also, I was asked recently to stop using "ls --color" as it slows down the filesystem. This is a bit sad considering how useful it is and that this is allowed on every other supercomputer I have ever used. (104)
- I'm glad that ARCHER has Python! (105)
- Very good. (108)
- Most of the software we use is our own but all libraries we need were available. (109)
- tmux! (115)
- tmux! (116)
- Whenever anything hasn't been available it has been added quickly. (119)
- it is good to have the makefiles for some of the installed software packages available at the web page; this serves as a starting point for own compilations where changes in the source code have to be accounted for (123)
- I spent 3 months trying to discover what feature of my code lead to some nodes running more slowly than others, only to have that erratic performance behaviour vanish as a result of a 'system upgrade'. (124)
- Lack of IDL, ncdump, ncvview, nco have been the limitation of ARCHER. This has been a reason why I need to transfer data to JASMIN or local computer which takes time. (126)
- A few days ago, I had an incident when I could not compile my code due to problems with the path to a library. On the following day the problem seemed to have been resolved as the same code could compile without any problems (although I did not notify anyone from the ARCHER Support Team). (128)
- One small thing would be if your personal qstat listed how many jobs out of your total you have submitted so far - i.e. 14/16 jobs in queue. (131)
- The error messages that you get when a job fails are rather impenetrable - particularly when it is job control fault. (133)
- We normally use R and rgdal package to process geographical data but it requires external libraries such as GDAL and PROJ4. Would it be possible to have these libraries installed on the Archer? (134)
- I tend to self-install all necessary software. (138)
- no comments. All available as needed. (139)
- Ideally I would like to be able to use the post-processing software Paraview remotely in parallel for visualizing results from OpenFOAM CFD simulations. (141)
- Would be useful to be able to run Paraview on Archer for post-processing. This would save having to regularly download large quantities of data. (142)
- Currently Archer has limited support for Intel compiler tools. It would be great if the latest Intel Compiler Suite and performance analysis tools are also supported. (147)
- Fortran coarray corner cases need more research. Some unexplained behaviour encountered. Cray support has been excellent, but this seems an area where improvements can be made. (148)
- Again the machine is great when jobs run. However, the turnaround has become terrible. Even quite small jobs (64 cores) queue for many days (including weekends). Not sure if this a software issue or hardware or priority.... (151)
- There is a great up-to-date amount of software installed. Thank you! (153)
- Providing custom modules to other ARCHER users is still rather cumbersome. (154)

- The software is updated too frequently, particularly the compiler suites and the system libraries (e.g., MPICH), requiring recompilation of code. No reasons are provided to users for the updates except to increase version numbers. I would prefer a more stable software environment. (160)
- The messing with the python stack has been a bit frustrating, but otherwise it's great. (173)
- I use gromacs and it is working very well (178)
- Frequently broken modules after update cycles, issues with python libraries, gui libraries... (182)
- great! (183)
- I am generally happy with the software on ARCHER, and any that isn't available has been installed reasonably well. I would say that sometimes only the gnu version is installed, when it would be good to have a version compiled under the cce instead. (188)
- Module conflicts between scientific libraries and MPI libraries are unfortunate (but difficult to avoid).  
There is no tool that can produce a communication timeline a la Vampir. (201)
- Some sort of more automated archiving from work to the RDF would be useful. (203)
- Really satisfied with coherent packages of MPI, FFTW and other libraries for different programming environment. If only this was available on smaller university clusters.  
Unfortunately, the parallel profilers turned out to be buggy and they didn't work with my code in eCSE 05-14. We have spent a lot of time with them. (205)
- More job scripts for those scientific softwares under individual webpages. (209)
- I'm not an advanced user, don't have intensive demands. So yes, has met my expectations. (211)
- It would be nice to be able to run small short scripts to move files around. (215)
- In my experience, ARCHER is stripped back to just what is needed, and I am happy to analyse elsewhere (220)
- Having a graphic session and some tool to visualise atomi structures would help our jo (229)



## Helpdesk

- The helpdesk is exceptional, all my interactions with them have been prompt and very helpful. (13)
- Fantastic. They know when to bring experts in on particular package specific questions also. (22)
- Always helpful, and I like the fact that they check with you before closing a job to make sure everything is fixed. (34)
- support was very helpful in making necessary scripts to run jobs on Archer and in testing the scripts to understand optimal scalability (35)
- I only ever had to use the help desk once as the documentation on the website covered almost everything that I have needed so far. The time that I did contact the helpdesk the response was quick and it solved the issue that I was having. (40)
- Very fast response and very helpful at solving the problem. (45)
- Very helpful staff and quick reply. (49)
- Did what they said they would within the time frame they said they would. (53)
- Great responses and response time. (56)
- Fast and helpful. (58)
- Online videos/webinars useful, please keep these going. (60)
- Make it possible to reply through the web page for helpdesk queries. (66)
- very efficient (68)
- The responses are always very fast and helpful. (73)
- Actually nothing to say...you are great! The query system works perfectly and efficiently as never. Since I am asked and since I did it recently, I would standardize the procedure of buying computational time for external users and centers, to be able to access Archer by paying (I mean, other than applying for grants). (86)
- Was waiting an additional (beyond due date) two months to have query response. (92)
- The support has been and continues to be excellent. Without the support given to me I would definitely have struggled to conduct my research. (97)
- All help has been prompt and directed towards improving my experience. Some suggestions to improve performance have been offered without my request and that's also proved very useful. (102)
- I have used the help desk a lot, with anything from quick 1 line queries, to more complicated compiling and debugging questions. The team have been invaluable to my research (103)
- The helpdesk tends to be quick to respond, efficient, and helpful, even with fairly complex requests. Great job! It's always a pleasure to work with the helpdesk. (105)
- The help desk can always solve my problems. Very helpful and professional. (108)
- Some of my students have complained that response time has been slow in some instances and no real solution to a problem was provided. (109)
- Helpdesk always reply very quickly and are very friendly and helpful. (111)
- Very helpful. (119)
- quick and kind help! Thanks! (125)
- Helpdesk has been extremely helpful. I really appreciate it. (126)
- Often quick and helpful replies. (127)
- This works well. (133)
- As a recent ARCHER user, I've asked a number of questions and required considerable support from the ARCHER helpdesk and the responses have always been timely and incredibly helpful, thank you. (141)
- Have been very helpful when answering all of our queries. (142)
- Helpdesk queries are answered quickly and informatively. The ARCHER helpdesk is one of the most useful support services I have used. (160)
- They have replied promptly and helped me when required (178)
- Usually good, in one case didn't receive further feedback and after a good while (a week or so) found the issue myself, fed back the solution but never heard back... (182)
- great! (183)
- It seems to work fine. (188)

- The helpdesk has always been quick to respond to my queries (203)
- Very efficient. (209)
- Not used sufficiently to provide useful commentary (220)
- I haven't really used it that much. (225)
- Response was very quick and helpful. (230)

## Documentation

- More documentation on optimising code performance would be useful. Particularly for estimating the amount of AUs to be used prior to requesting time. (9)
- Need to improve in some aspects such as there is no detailed documentation for submitting jobs with dependency and restart jobs. (10)
- I think that the web links sometimes led me in circles. (18)
- I very much enjoyed the introductory tutorial videos. (19)
- Possibly a bit more information on the big memory nodes. (34)
- I have only had to use the help desk once to find the answer to a problem as the documentation on the website is very good. (40)
- The structuring of the documentation is a bit confusing at times: some information is in the user guide, some in the best practice guide etc and it is not clear where the best place to look is (41)
- The Youtube videos are helpful and there is a lot of written information but it can be hard to find what you're looking for in the written documentation. A tips and tricks section could be helpful. (43)
- I have found it hard to find information on the queueing system and what queue accepts certain jobs. Particularly recently when the long queue changed to greater than 25 hours - I found this by trial and error. (45)
- The documentation is excellent. (50)
- Mostly very clear. (60)
- Add how the used AUs are computed (my calculations, using what I think how it is done, are factor 2 off). have not discussed with helpdesk so far. (68)
- The documentation of Globus and Grid-FTP in the Data Management Guide would benefit from an update. More detailed instructions on installation of the required tools and certificates on the client system would especially be helpful. (71)
- I recently found the task farm jobs instructions extremely helpful. (73)
- It could be a little easier to find things, but the information is well-written and offered at about the right depth (74)
- It is very difficult to find out which libraries are available. It would also be useful to have some more information on how things are laid out (e.g. location of libraries and include paths). (84)
- Just perfect. What you need in a clear and concise style. (86)
- I know it is difficult to keep it up to date but referencing a version of software that is no longer available could be better. If stating version numbers and modules one should really provide the date that it was valid.i.e. the default version as of .... (89)
- This has been for the bps scripts mainly (91)
- There is not a specific documentation for libraries and modules installed on archer. Sometimes one want to know specific environmental variables and paths for specific modules, but there is no documentation for that. (94)
- it would be nice to have it organized in a more user-friendly way, with easy access to the important information, but maintaining high-level help, possibly updated with the help of users. For instance, a special CASTEP or VASP section with benchmarks, suggestions and utilities from the community will be great. I think the section with help for windows user to connect and backup files could be upgraded with suggestions for speeding up connection and data transfer as well as links to useful tools such as xming. (96)
- I've looked at the documentation for submitting multiple jobs using arrays and found it baffling. There were very few clues as to how to start with submitting in this fashion. (102)
- There are a few too many guides available, such as the Quick Start Guide, the User Guide, and the Best Practices Guide. The division of information between these three is not at all clear, especially between the User and Best Practices Guide. These would be better off combined. (104)
- The documents are very helpful and can solve most of my questions. (108)
- The documentation is good - clear, right level of detail, wide ranging. However, it's not always easy to find what you need - it's not clear what is in the user guide and what's in the best practice guide, so it's easy to miss things that are in the best practice guide. (114)

- Sometimes a bit hard to find what you're looking for. e.g. running the serial nodes, accessing the post processing bit of archer. (119)
- Not easy to get the information... a search system or a better organisation as a wiki could be very good. (122)
- The documentation on the web can be better. There have been a few occasions when I followed the documentation, didn't work, consulted the helpdesk, helpdesk offered help and updated the documentation. That means the documentation was not initially correct or not updated. (126)
- I can find most things OK. The only thing that I find hard to find is the details of the currently available queue limits, and the recommended sizes of jobs to maximise utilisation. (133)
- I don't rely on the documentation often anymore these days, but it's to the point and useful when I do. (138)
- Availability of wiki containing up to date information on compilation of supported modelling packages is a well-defined merit of the ARCHER documentation (146)
- Simplify documentation so that it's easy to pull out specific tasks (e.g. compilation against specific libraries etc) (149)
- The documentation is scattered throughout several "guides", but it is not clear why. The information could be more usefully organised into a single document. Recently I encountered an issue with running mixed-mode MPI/OpenMP jobs when compiling software with the Intel compiler, in which the job would use only one OpenMP thread no matter how many were specified through the 'aprun' command. The solution to this problem was in the documentation, but listed only in the section on "pure OpenMP" jobs, which I had skipped over because my job was a mixed-mode one. Including this information in the correct section would have saved me two weeks of testing and debugging my code. (160)
- Again, usually good but not always catching up with the update cycles (182)
- The documentation is good, it would be better to provide some key word sorting function. (183)
- Whenever I've needed to use it, I've found what I need. (188)
- The use of the queues is still not entirely clear to me. Is the allocation of nodes dependent on total resource request (i.e. node hours) or just number of nodes used in a job. i.e. for long runs that need to be run in sequence, is it better to run these in the long queue, or in the standard? This is mainly related to a problem for when you need one experiment to run in consecutive dependent chunks (i.e. climate model experiments), rather than having many independent experiments that can be submitted to the standard queue. Sometimes, it feels like other users are at an advantage as they can have many experiments running at once, while a single run ages in the queues. (203)
- I found a lot of useful things in the documentation. Still I had to ask the user support for some details, but most users probably don't need those. (205)
- The documentation is very useful for me, especially new to ARCHER. (209)
- Generally helpful so far, only been a user for about a month. (211)
- It didn't mention the work directory! (215)
- Not used sufficiently to provide useful commentary (220)
- The guidance on using DDT was a little unclear, not helped by there being both "ddt" and "allinea" modules. (225)

## Website

- The only way I can find things is through search, which I suppose is okay. (16)
- The service status information is not as detailed as it was on HECToR, which provided a count of which queues were busiest. This was useful when considering whether to split large jobs or run whole. (28)
- I was trying to find information on how many AUs used certain software package modules (e.g. CASTEP) on ARCHER, to get a rough estimate on how much ARCHER time goes towards CASTEP, LAMMPS, etc (being aware that people may forgo the ARCHER modules in favour of their own compiles. Anyway, I could not find such a statistic on module use. (37)
- Safe could be improved. The right button can be hard to find to perform an operation. For new users to gain project access it is not obvious what the user or the project manager will see/do. (60)
- Navigation could be easier, but a good range of functions. (74)
- It's usually easy enough to find the material I need via google/ search on webpage. (75)
- Finding information is complicated. Archer SAFE used to be particularly non-trivial to locate. (77)
- It is more difficult than necessary to access the course descriptions from the course/event scheduling page. Description/registration/scheduling information should be included in the same place. (80)
- it would definitely benefit of an update in style (a bit too much old-fashioned) and content. (96)
- Sometimes the compilation / jobscribs are out of date but help was always at hand when contacting the helpdesk. (97)
- The text is a little confrontational, but navigation is good. (102)
- I find the overall status / live reports very helpful (103)
- It would be very useful to show information on completed or expired allocations on the user-specific page. Once allocations are fully used or expired, they disappear completely as if they never existed. Such information is useful when writing proposals and is also useful when an allocation stops being active before a user expects. (104)
- Sometimes a bit hard to find what you're looking for. e.g. running the serial nodes, accessing the post processing bit of archer. (119)
- A text book has an index. Websites require you to know where the info is before you start looking for it. Archer's web site is no different. (124)
- Service status does not always shows the problem. By looking at the graphs down the page I can see it's quite crowded, but it does not provide any clue about how much longer my job needs to be on the queue. (126)
- This can be tricky to navigate - especially the SAFE tool ... not entirely clear how certain processes (e.g. get a new password) are meant to happen. (137)
- The website is very useful. However, the remaining budget section only seems to update occasionally which has led to over-using the available budget on occasion. (142)
- Navigating between the main site and archer-safe could be easier. (152)
- Some useful links and information, sometimes hard to find certain pages for example, if there was a link in an email and I wanted to look the page up again later without the link from the email, using the search tool on the archer website it can be hard to find that page sometimes. (181)
- great! (183)
- Some things hard to find because not sure what words you use. Once accustomed to the language all was fine. (187)
- Nice layout. (209)
- Good (211)
- Could be made clearer, especially the documentation, current format looks messy (214)
- I've always found the website to have what I need. In part, this probably stems from the limited way in which I need to interact with ARCHER (e.g. running large ocean jobs on it). (220)
- Up to date. cool! =D (229)

## Training

- I appreciate the high frequency of webinars. (19)
- not used, but now I cannot unclick the answer. (37)
- I have not completed any of these, of my own fault, they are well advertised through the mailing lists. Only given a 3 because I have never taken advantage of this facility. (40)
- Audio quality can be an issue when following video guides. (43)
- The range of training available is great. (50)
- Online videos/webinars useful, please keep these going. (60)
- I attended the Scientific Python course particularly for the section on integrating Python with other languages, e.g. C, Fortran. However, as the course had run over time this section was almost completely skipped. I would be interested in more medium-level course, or hands-on workshops for people with some HPC/coding experience but not advanced users, on things like data management/analysis using batch processing, particularly how to use different software packages to complete different parts of the analysis. (I am aware of how to do this with bash + gnuplot, but have no idea how to integrate say, data generated in fortran with analysis done in MatLab/Mathematica/Origin.) (80)
- Unfortunately I did not use them, but I will give it a go in the future. I believe, looking at the trends around, that delivering lectures on basic and advanced FORTRAN (or C and C++) and on the use of GPGPU and CUDA will be enough for normal users. Furthermore (and I am not sure here again already you offer this possibility) you can increase the interest of programming courses (sometimes boring) if you add specific ones on applications (I am thinking about fluid-dynamics, astrophysics, high-performance molecular dynamics etc...) (86)
- Have not attended some as they are based in Edinburgh. (89)
- I would have loved to participate to the computational software training, but it is difficult for PIs to attend a full day (or more) of training. Events in London or events online work better for me. I would incredibly appreciate if the material, examples, tutorials and presentations videos and slides would be made available to the ARCHER community, especially to users that could not attend the sessions. (96)
- I attended the Python and Fortran Training session and they were very useful. Maybe running different levels of training might be worthwhile in the future. (97)
- They look very interesting although I haven't actually attended any. More on python. (119)
- Very informative (121)
- I attended one. it was very good. I just felt not enough days and practical works. (122)
- I have only watched one Tutorial from the ARCHER website. The rest information I needed I found in written form (on the ARCHER website). (128)
- I wish some of these were easier to get to... (133)
- MPI in Python might be a useful addition to the MPI course (which currently includes C and Fortran) (137)
- I have not received recent training, so I give a neutral score. (138)
- I have attended several Training sessions and they have been really useful to me. (147)
- I have not attended to any of the webinars, as I have long experience in HPC and supercomputing, but have looked at the website and the materials look great. However, you could take a step forward and do some workshops on numerical approaches, specific libraries such as PETSc or Trilinos... You can find inspiration in the workshops and training sessions that are offered in the SuperComputing conference in the states. (153)
- I found the Advanced MPI course a bit too fast paced + too much new information covered. This was fed back on the feedback form for the course along with some suggestions for improvement etc. (156)
- It is very useful indeed! thanks! (183)
- - For your info, I haven't attended any ARCHER training so far, but it is useful for ARCHER to provide these. (184)
- Didn't do this myself, but heard good stories from others. (187)
- Could you improve your announcements about your training courses to be more international. I mean there are many international students need your courses but they don't know about your services. (208)

- It covers many useful training courses for me. (209)
- I haven't been to any, but I've heard good reports. (225)

## Webinars

- This was very useful. (18)
- Some of the introductions to a given topic were a bit too basic for me, but I suppose the clue was in the title... (53)
- The webinar software is a bit clunky, a simple web interface would be nice. (74)
- The HPC overview course (recommended as the starting point for new users planning to take the HPC driving test) was hard to focus on remotely. There were technical difficulties at the beginning (videos only showed presenter with slides in the background, rather than a video of the presentation next to the slides themselves). The presenters did a good job of informing participants of the issues and updating the slides via the course mailing list, but I found it difficult to set aside the time during the work day to watch the 3 hour presentation after the ineffective first session. (80)
- Often cannot attend as they clash with other work..Have used YouTube channel. Worked well for me.  
Although on one occasion the presentation was difficult to follow due to font size. (89)
- Good for asking questions but generally I tend to switch off half way through. (97)
- Some were useful, some a bit too basic. The archive of materials from past tutorials is very useful. (114)
- It's great that the webinars are available on youtube after the event date. (158)
- Useful and interesting! (183)
- Didn't do this myself, but heard good stories from others. (187)
- Late afternoon timeslots for many of them is somewhat inconvenient. Otherwise, excellent (198)
- Very useful (209)



### Online Training Material

- Honestly, I only did the driving test to get the extra kAUs, level was very basic. (20)
- Possibly have extensions to the driving test for more complex tasks. (34)
- One of my students recently worked through the initial training for Archer's new user and he was very satisfied by the instructions. (73)
- Need to be more careful when referencing specific versions (89)
- Online driving test is a great way to get people involved in Archer and HPC! (115)
- Online driving test is a great way to get people involved in Archer and HPC! (116)
- I took the Online Driving Test, and think that it is a great idea to as it introduces the future user nicely while motivating him/her to learn how the machine is configured before using it.(153)
- It is extremely useful to my coding and parallelism implementation. (183)
- I did take the driving test I think, but I can't remember it being particularly useful I'm afraid. (188)
- extend award limits (twice - triple kAu allocations) if resources would allow that. (195)
- Need more examples for explanation. (208)
- It is very useful to provide step-by-step information. (209)
- Driving test was good. I repeated it till I got 100%, couldn't figure out whether this was required or not, don't think I read it anywhere (211)

## Other comments

- Overall my experience with the ARCHER service has been very pleasant. (13)
- I've had a few MAJOR problems with ARCHER over the past year. First, the queue system is absolutely disgraceful. Having to queue more than 48 hours for a small, 3-minute run is ridiculous. I found that submitting these short runs to the long queue solved my problem, which I realize is not what you had in mind for the long queue, but when you are under time constraints, sometimes this is necessary. The fact that ARCHER was down for the better part of two months is a national disgrace. You are supposed to be the UK's best and fastest supercomputer. That was the key time for me to run simulations for my PhD. Having my models postponed for two months nearly cost me my PhD as I had no option of extending my funding. I hope you have learned from those problems and have a system in place to prevent this from happening in the future. (16)
- It would be great to be able to use the short queue all day long and not only between 09:00-16:00. (19)
- Archer really is a great service. I have used other supercomputers but the assistance and the guidance on Archer is significantly better. (27)
- The standard queues are often painfully long. My jobs take 6 hours to run if they are successful, but in the last few months have typically queued for 12-24 hours before starting to run. 24 hours is a long time to wait for something to crash as soon as it starts to run, which is a common occurrence when testing a new setup. Is there any way that the queuing system could be changed to reduce waiting times for users with repeatedly failing jobs? (38)
- It is a terrific resource (42)
- The only frustrating thing is when the queues always become blocked just before a renewal of budget. I know there is not much that can be done to prevent this but its the only frustrating thing about using Archer. (45)
- The online driving test was how I obtained access to ARCHER, and I think its a brilliant idea for widening access. This has allowed me to put together an eCSE application for the next call, and hopefully accelerate my research in future. (50)
- Archer should move towards high-throughput computation, much more useful for the vast majority of research, than the current single mpi process preference. (76)
- Please give me the prizes of 2000 kAU on Archer ;-)!!! (86)
- Just to re-iterate the helpdesk / support makes archer so beneficial and has really saved me an immense amount of time when trying to compile software, or when a problem has arisen in a calculation. (97)
- I have a few critical comments.

1) Why is the 'short' testing queue only open from 9-5 UK time? I am not always able to do my testing work during "normal business hours". I have to be flexible as to when and where I work. Sometimes I need to run tests while working from some other time zone - why should I be penalised for working from abroad? Having a testing queue that is only open from 9-5 hampers my progress and is very frustrating. It is also very unusual - I have worked on a lot of HPC systems, and ARCHER is the only machine I have encountered with such a strange time-limited queue. It really is bizarre.

2) The attitude that "ARCHER is not a testing platform" does not respect the needs of ocean/atmosphere/cryosphere modellers. Our models are \*always\* in development. We typically have to compile and run \*each\* case (e.g. different compile-time physical packages, adjoint setups) and test it. Why shouldn't I be able to test my setup with a short run? My typical "test run", just to make sure that my inputs are configured correctly, is less than 1000 cores for about 10 minutes. I shouldn't have to wait on the queue for hours just to perform this small test. Why doesn't ARCHER respect our needs for this very minimal kind of "testing"?

In short, I really need a larger "short" testing queue that is open 24 hours a day.

3) What can be done about the "rush to burn hours" in February and March? ARCHER is nearly unusable for those months because both NERC and EPSRC users are trying to use up all their

hours for the year. Why can't those allocation years start and end at different times?

4) It's too cumbersome to share info between projects (e.g. n01, n02). I have collaborators in both n01 and n02, and sometimes sharing files on ARCHER between projects is a real headache. There are too many walls.

5) Please clearly publish any changes in your SSH/SCP policies. Sometime in the past year, a decision was made to disallow outgoing SSH connections. This made it impossible for me to do my work, since I use TAF (a FastOpt.com product), and TAF uses outgoing SSH connections for server-to-server communication. I wasted a lot of time wondering why TAF wasn't working for me. The ARCHER helpdesk did respond to my request fairly quickly, and the issue was fixed over the next couple of weeks. If possible, please be more explicit about these things in the future. For example, perhaps instead of just killing the SSH sessions, you could have an error message in the terminal explaining that the session is being killed (and why)? (105)

- On the SAFE webpage, it is very useful to have the table which summarises the projects you are involved in. Remaining budget is reported on this table, and it would be really useful to have reported how much of your total allocated time this was, and when it runs out. (111)
- It would be very helpful to get a warning message if you try to submit a job to the short queue out of the 9-5 hours; also if you submit a job before 5 pm which is likely to be still running at 5pm. It's easy to forget about the limited hours for this queue. (114)
- It would be invaluable to have a separate resource with better queues for smaller runs. (119)
- If there is a way to increase bandwidth for the filesystem, please do so. I've had several crashes due to extremely slow read/write speeds for my jobs, and post processing data is much slower on ARCHER than on other machines. (127)
- I find the ARCHER service to be responsive and helpful, easy to use, and flexible. Excellent service. (132)
- I have found that managing the budget of projects has been a little difficult due to a delay (~24 hours) in updating the remaining budget. I have found that the queuing periods can sometimes be quite long. I understand that reducing the walltime of simulations gives priority but in my work I tend to run relatively few long jobs so end up with long queue times. The upload and download speed of files can be incredibly slow as can deleting files. For simulations generating TBs of data this can take a really long time. (141)
- Downloading/uploading data from Archer can seem very slow or even stall at times. When this happens it is often quicker to stop the download and restart it again. Furthermore, deleting/copying data on Archer seems to be slow. (142)
- It would be great if there is further support for shorter duration debug queues (147)
- well done! (148)
- The service is worse now than ever in the past due to the job turnaround. I really hope that can be sorted soon. It is ironic as Archer has been upgraded, but the service is worse. The rest of the facility (machine and staff) is excellent. (151)
- Keep up the excellent work & Thanks (152)
- Please communicate system issues to users more frequently and in a more timely manner. I understand that there is a balance to be struck between "bombarding" users with emails and keeping users informed, but ARCHER already has different levels of mailing lists that could be used for this purpose. Over the past year, there were several instances of overnight system failures / machine reboots that were not communicated to users. Those of us who had jobs running at the time received cryptic error messages, which initially led us to blame software errors rather than a hardware crash. Only after emailing the ARCHER helpdesk to inquire did we discover that the entire machine had gone down overnight! Hardware errors (particularly disks) and subsequent maintenance have marred performance of the system over the past year. (160)
- 16 job limit on queue is frustrating -- I understand the need for throttling/controlling the number of jobs/cores active per user, but the queue should handle that, not reject jobs based on a hard (and quite small) limit. (167)
- Wonderful service, my research wouldn't be the same without it! (169)
- An excellent service! Thank you for all your hard work! (172)

- The time taken to queue jobs to run is quite frustrating. A 1.5 hr job often needs to be set up the day before to run. (175)
- some of the more useful workshops fill up really quickly, e.g. python related ones- would it be possible to have these classes more than once in a year? (181)
- Great - Thanks. Keep going. (183)
- - We run OpenFOAM simulations that produce a few thousands of files. Unfortunately this is a problem of OpenFOAM. However when a few researchers on a project are using ARCHER intensively, we find that the limit of allowable files (3 million) on ARCHER is sometimes exceeded and brings a halt to all our simulations. This could be improved, by say allowing these special cases an increase in this limit. (184)
- My research field is data assimilation, in which huge geoscience models of e.g. the atmosphere, with up to  $10^9$  variables, are coupled to huge data sets, up to  $10^8$  every 6 hours. Since these problems are nonlinear Monte-Carlo methods are needed for this. Present-day ARCHER is not fit for purpose, and only centres like the Met Office can do this at the moment. However, they need the academic community to push the quality of their methods forward, both in terms of accuracy and efficiency. In the USA the academic community does have access to machines that can do this, we in the UK are behind, while at the same time the newest and most exciting new ideas are coming from the UK, through close collaboration of mathematicians, statisticians, and geoscientists. This is not only for weather prediction and better understanding of the weather, but also for oceanography, biology, ecology, land-surface science, climate modelling, etc. There are very exciting and innovative things we could do with a bigger system. (187)
- Change the way the shared nodes work, or at least make some non-shared serial nodes that can be used for compiling. It is very difficult to compile when you don't know how long things will take to complete. This is particularly problematic when running training courses on ARCHER, which I do every year (UKCA). While we can reserve parallel nodes, I first get the students to compile. While this should take 10 minutes, they all then time-out after an hour as there are 20+ people hitting the shared nodes. It doesn't just waste their time and mine, but also we're paying for parallel nodes that we can't use. This is my biggest problem with the service as it is currently configured. If I could compile in the parallel queue, or reserve serial nodes, this would improve things greatly. When I called the helpdesk about this there was nothing they could do. (188)
- The UK National Supercomputer service is an essential tool in support of scientific research within the UK. (201)
- Please provide data analysis tutorials (for atmospheric model outputs e.g UKCA) (207)
- Good experience. (209)
- Please increase user's quota and give us more space. Much time is wasted in backups (229)
- Standard queue waiting times can be excessive: I have recently been submitting jobs of about 11 nodes for 12 hours, and these have quite regularly spent 2-3 days in the queue. I realise larger jobs are prioritised, but a faster turnaround on these smaller runs would be useful. (230)