


2022 Minerva User Survey Comments and Response
Scientific Computing and Data
Feb 16, 2023

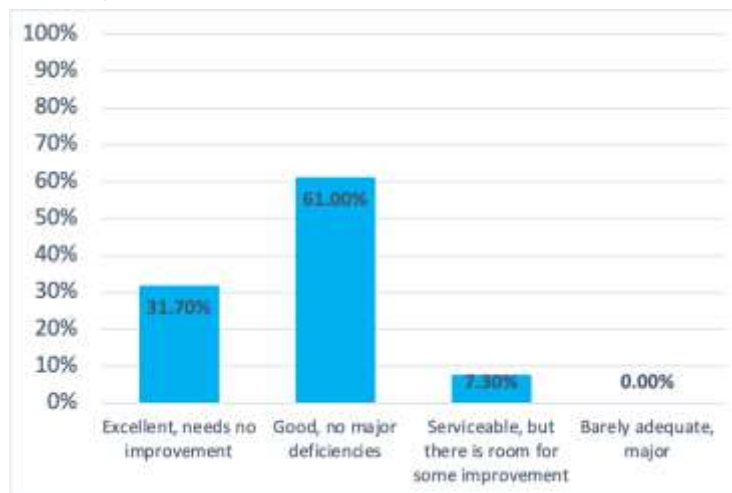
The 2022 Minerva user survey —distributed in January 2023—solicited feedback from 845 active Minerva users. Of these, 44 users responded (5.2% response rate). 31 comments were received.

We asked five questions:

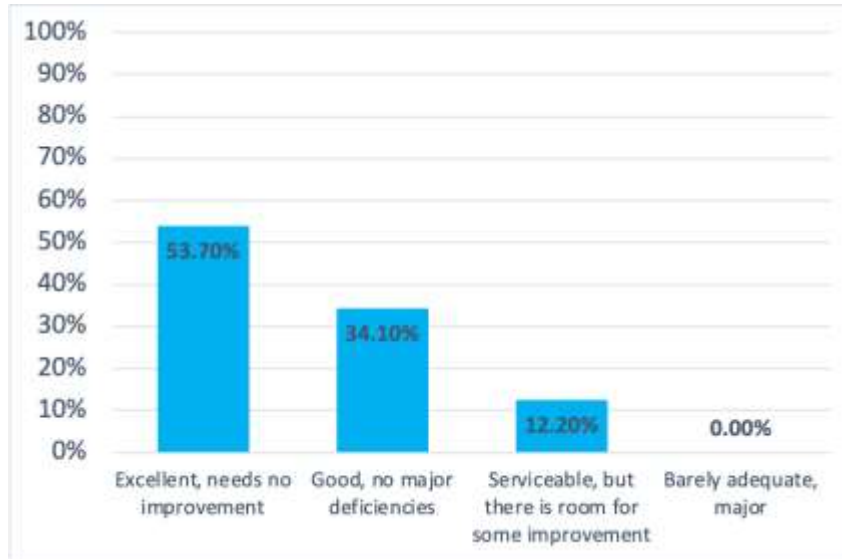
Q1: Overall, how satisfied are you with the LSF queue structure, compute and storage resources (GPUs, high-memory nodes, TSM, etc.)?



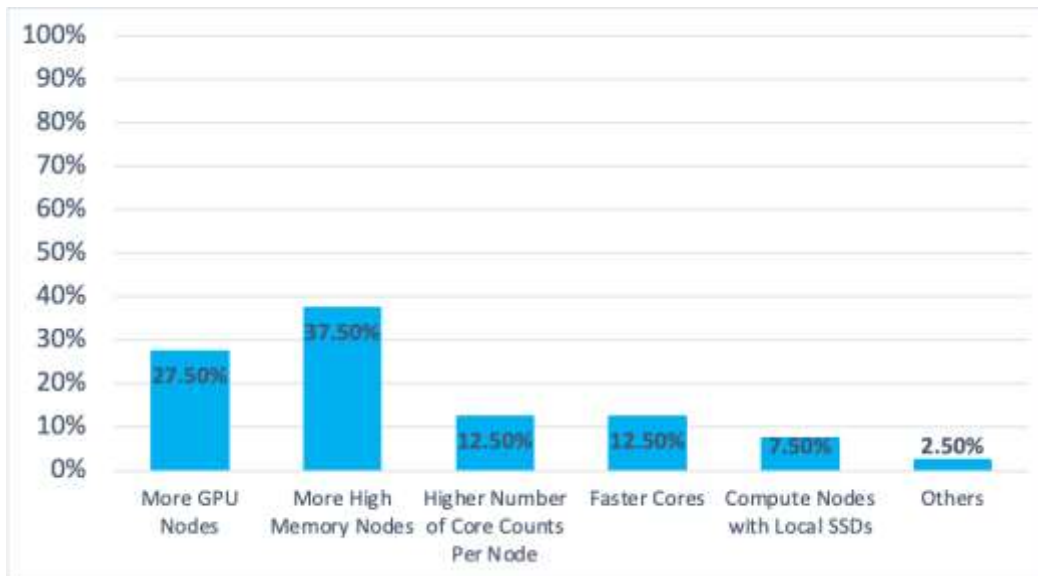
Q2: Please rate current software environment (packages and services such as database, data transfer, container etc).



Q3: Please rate your satisfaction with operations (ticket system, responsiveness of staff, documentation, user support etc).



Q4: Which of the following would you most prefer for future Minerva expansion?



Q5: What suggestions do you have for improving our service?

All comments and responses are as following:

System Related Issues

- Sometimes the file system slows down to the crawl making it difficult/impossible to access any Minerva mounted data. • Apart from the occasionally slow file system, Minerva has been better lately.
 - Considering the fact that the overall file system of Minerva has been pretty stable over the past few years, we believe the slowness users sometimes suffer from was mostly due to the misused login nodes by some users running compute and/or I/O intensive production jobs on them. We keep monitoring and removing these offending processes to resolve the problem as quickly as possible. We will continue to improve our monitoring strategy on the usage of the login nodes and I/O. We will also emphasize the shared use of login nodes with our users.

- MPI jobs with more than 20 often die due to socket communication problems.
 - We noticed that sometimes MPI processes on the compute nodes were not cleanly removed after the job completion, causing unexpected issues to the next jobs landing on them. We believe this issue has already been taken care of. If you are still experiencing the problem please contact us at hpchelp@hpc.mssm.edu and provide a bit more details about the issues you encountered so that we can take a look.
- Need more GPU nodes. • Need more GPUs, this is a rate limiting step for my research. • Significantly more GPUs are needed. • Please get more GPU nodes or even the new intel Gaudi2 processors (<https://habana.ai/training/gaudi2/>) if GPUs are expensive. This would really help train much larger models that could be very useful to the genomics community. Having more GPUs available for training could dramatically speed this up and enable research that isn't currently being done at research universities.
 - We added 2 more GPU nodes (2TB RAM & 4 *80GB A100 on each) in 2022 to serve more runs, especially heavy-memory jobs.
 - We plan to expand and update our compute resource this year with new hardware including extra GPU nodes to provide more compute capacity and boost up productivity for users.

Job Scheduler & Queues

- At times, I had experienced my submitted job in a pending status for more than 24 hours, while the same job script could be scheduled for running by a different colleague. To this day, it is still why it is so, given the same resources requested. • Occasional long lags in job assignment to queues (which otherwise appear to have plenty of room for quick job assignment - based on 'bqueues').
 - There could be many reasons for a job to stay in the queue for longer than usual; The whole cluster or the compute partition your job is supposed to be dispatched to may be all busy; Your job requests a large amount of compute resources such as amount of memory and compute cores so that no compute node can satisfy the resource requirement at the moment; Your job would overlap with a scheduled PM. To see the pending reason use the “bjobs” command with the “-l” flag:

```
$ bjobs -l <JobID>
```

- Occasionally there are users running 1000s of jobs at once and making it difficult to secure resources. • It would be useful to explore resources such as the LSF job scheduling could be fairly or equally assigned to users/groups, rather than just adding more resources to the current system, which will unfortunately exacerbate resource hoarding by some users/groups.
 - We set the global job slot limit to 4,000 (the maximum number of running jobs foreach user is 4,000 by default) to avoid the whole cluster being taken by the jobs of a few users as well as considering the throughput of the whole cluster. We have been also putting extra restrictions on certain users who use heavy resources such as memory in order to better balance the resources among

users. You can check the slot limit by "`blimits -u userID`" while you have running jobs.

- We also have the LSF fairshare scheduling policy implemented. This approach schedules jobs from each user with weights applied to their priority, based upon the user's historical usage information, but does not restrict the number of jobs a user can submit or run. We set the maximum number of pending jobs for each user to be 20,000 by default to regulate the overall load on the LSF scheduler.
 - We will keep adjusting our queue structure and dedicated nodes to reduce the turnaround time and increase the throughput along with the dynamic job needs.
- Sometimes certain groups hog a lot of the resources leaving others with long wait times - including for short/small resource jobs. Is there a way to guarantee a small or certain amount of resources for groups so that we are not waiting a week for small jobs (4 cores/8mb)? ● Sometimes there is a long wait for small compute resources, even though not many resources are being used by others. Not sure why this happens but would be helpful to look into when it does.
 - You may want to use the LSF command "`bqueues`" to check the load status of the queues, which would be helpful for you to determine which queue is less crowded and where your jobs should be submitted to.
 - To find out the pending reason, use the "`bjobs`" command with the "`-l`" or "`-p`" flag. In case of short/small jobs the most common cause of delay in scheduling is a large amount of memory requested. Please adjust the required amount of memory for your job accordingly so that LSF can schedule jobs effectively and avoid running them on the high-memory node unnecessarily, which is the most crowded compute partition on Minerva lately.
- Balance the scheduler for shorter GPU jobs and distribute over more users to prevent single users hugging GPUs for days or weeks. ● The gpu queue is often not efficient. It seems to be based on chronological criteria rather than splitting the resources between users (eg. few users are using most of the resources, while the others are waiting for days).
 - We have set job limits on gpu queues for heavy users to prevent the resources solely occupied by a handful of users.
 - Our recent analysis on the usage statistics indicates that we still have a satisfactory amount of GPU resources based on the average usage percentage of the whole GPU partition but its occupancy pattern is rather sporadic and occasional. We also see that a greater portion of jobs on the GPU partition consists of fairly long workflows, which leads to a rather long wait time for jobs in the queue once the nodes are occupied by those long jobs. To alleviate the pain we have implemented the new "`gpuexpress`" queue that can accommodate short and small jobs more effectively.
 - We plan to expand and update our GPU compute resource this year. We hope this will ease all the inconvenience eventually.

Software/Packages

- Some packages need to be updated. ● The version control (specifically R) is incomplete: many of the newer versions of R do not have packages present on the older versions, and sometimes the locally install packages do not play well with the environment (example "snow" package in 4.2.1 vs 4.2.0)

- We support multiple versions of packages on Minerva via the module system (<https://labs.ica hn.mssm.edu/minervalab/documentation/software-environment-lmod/>). You can see what versions are available and which one is the default using the module command “ml avail” or “ml spider”. The version number has to be specified explicitly to pick a specific version if not the default. E.g.


```
$ ml avail star
$ ml star/2.7.10a
```
 - If you need a specific (or updated) version of a package please contact us by opening a ticket at hpchelp@hpc.mssm.edu. We cannot really remove all the old software and keep the latest version, since many users still rely on those for reproduction and consistency of their publications.
 - Each version of the R module on Minerva comes with more than 1,000 R packages that are continuously updated and added in response to users’ requests. It is not practically feasible to keep all versions consistent. In case you need a specific package installed please let us know.
 - If you need help on a package install in general please open a ticket at hpchelp@hpc.mssm.edu.
- Pretty well kept up, usually the staff is responsive in installing upgrades when requested. There will be major upgrades to the deep learning environment (see <https://www.semianalysis.com/p/nvidiaopenaitritonpytorch>). It would be good for the staff to think about upgrading the software to enable these things as it will enable much faster model training given the hardware currently available.
 - Thanks for the suggestion. We will take a look and see if we can set up and implement the desired environment for faster model training on Minerva.
 - It would be useful to explore and support more container technologies such as rootless docker, Shifter, Charliecloud and Podman.
 - Because of the security concerns we do not support Docker. Instead you can use the Singularity tool on Minerva. Docker container images can easily be pulled and run as Singularity container images, which are safe to run on the shared cluster. See <https://labs.ica hn.mssm.edu/minervalab/documentation/running-container-singularity/> for more information.

Tickets

- HPC team is very responsive, but I cannot view my ticket. ● In the ticketing system emails, link saying " view this ticket's progress online." doesn't work.
 - The check status feature of the ticketing system (osticket) is disabled due to security reasons. We are waiting for a reliable fix from the developers.
- Sometimes it takes a week or more to receive an answer to a ticket. ● I have tried to contact the department several times, the waiting time is in the order of weeks and eventually the email gets picked up by the wrong department and the issue remains.
 - We handled more than 3,240 tickets in 2022 and tried our best to reduce the time to respond to users’ tickets. If you do not hear from us in a couple of days since you opened a ticket at hpchelp@hpc.mssm.edu it is most likely slipped out

mistakenly. Sorry for that. We have to admit that we made mistakes sometimes. Please send us a note and remind us about your open ticket.

- Staff is responsive but often unaware of any previous history of dealing with a specific user which leads to endless repetitions of assessments already performed. This in turns leads to reworking details and repeating things in a circle that slow down the process toward achieving a solution of the actual problem. • For users that require more consistent interactions with the help service, dedicated support that is consistently the same and can build upon previous development of tools/services.
 - Please understand that since each staff member of our HPC team has to take care of a myriad of requests/issues from many users everyday it is pretty hard to track the history of dealing with a specific user. It would be very helpful for us if you let us know who you worked with before about your issue so that we can assign the right person to your case and process it more effectively.

Documentation on Website

- Using jupyter notebooks on Minerva requires better documentation for dummies like me. Simply providing the bash script for launching it does not provide enough guidance on ideal usage / pipeline • Maybe it's me, but I find access to documentation, example scripts, case studies a bit cumbersome. • More detailed documentation regarding software, databases, supported web servers and various advanced operations with the job scheduler would cut down the need to contact HPC for enquiries.
 - The documentation section on Minerva's website (<https://labs.icaahn.mssm.edu/minervalab/>) was restructured to improve its accessibility.
 - We also performed major updates on its contents to supply more details about the system and usage instructions. We are continuously working to improve our documentation. We have been updating the changes on the system and adding new services.

Others

- Ability to use Globus to transfer directly to/from MSSM OneDrive accounts would be very useful.
 - To enable direct data transfer between Globus and OneDrive we are required to have a subscription to the Globus premium connectors service, which is not free and doesn't seem to justify the cost of purchasing yet.
- It is quite challenging to provide service to a diverse community. Minerva seems also mainly geared toward the LSF queue service that limits its reach. • Lay language for everybody. Often time users balk at the slurry of acronyms used and the taken-for-granted technical language used by Minerva in communicating with users. This approach is likely to turn down people and send them to outside services where everything is pre-packaged and clearly explained. The damage is twice, Minerva gets less users and the chance of providing the users with more tailored and effective service is lost. There is a need for better communication!
 - You can always request 1-1 sessions in case you need further assistance by opening a ticket at hpchelp@hpc.mssm.edu

- If your research group would like to schedule an additional tutorial, or training sessions tailored to meet the needs of your group or focused on special HPC topics, you can reach out to our computational scientists (hpchelp@hpc.mssm.edu)
 - We have several training classes a year and announce these through the Minerva mailing list and post them on our Minerva website (<https://labs.icaahn.mssm.edu/minervalab/>).
- More attention to the record of individual users. i.e. mistakes made by users that tend to be consistently within the rules should be dealt with trying to understand what's going on instead of blindly applying policies. While everybody has, no doubt, to abide to the policies instituted by Minerva, "unexpected" breaches by otherwise compliant users are more likely to be related to Minerva's issues than to unruly users.
 - Minerva is a shared compute resource bound by Mount Sinai regulations, and it is expected that all Minerva users, regardless of their titles and ranks at Mount Sinai, must abide by our rules and policies and stay within the boundary. The rules and policies are the minimum requirement for the protection of the system and users, ensuring "security", "stability" and "fairness", not for the convenience in terms of management. Please understand that any system-related issue cannot be taken to be an acceptable excuse to ignore the policy. We are trying our best to accommodate the needs of users but unfortunately we cannot track all individuals practically. If you have something to discuss please don't hesitate to reach out to us and let us know what your concerns and issues are.

Thank you for your nice words!!!

- I work with Minerva since 2017. I witnessed major improvements in service over the last few years as well as a substantial increase in resources. Fees for space usage are also very affordable. Communication has never been a problem, the team replies quite fast and, more importantly, always tries to accommodate our needs. Thank you.
- Excellent, needs no improvement
- The staff is quite responsive to requests.