2024 Minerva User Satisfaction Survey Results

**Scientific Computing and Data**

February 11, 2025

The 2024 Minerva User Satisfaction Survey (distributed in January 2025) solicited feedback from 1,100 active Minerva users. Of these, 65 users responded (6% response rate) with 37 comments in total.

We asked five questions:

Q1: Overall, how satisfied are you with LSF queue structure, storage, and compute resources (file system, GPUs, high-memory nodes, etc)? **Overall Results: 82.8% Good or Excellent.**
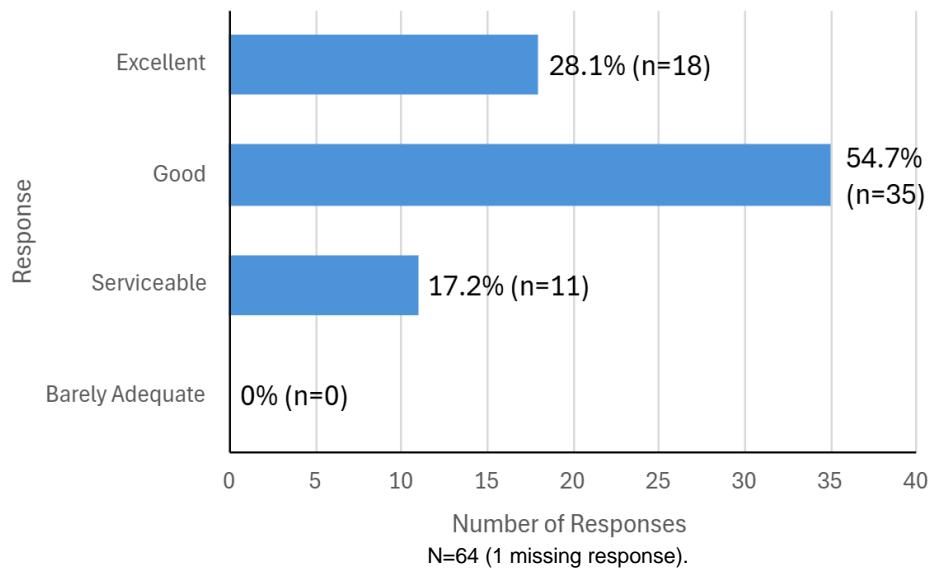

N=64 (1 missing response).

Q2: Please rate current software environment (packages and services such as database, data transfer, container etc). **Overall Results: 76.6% Good or Excellent.**
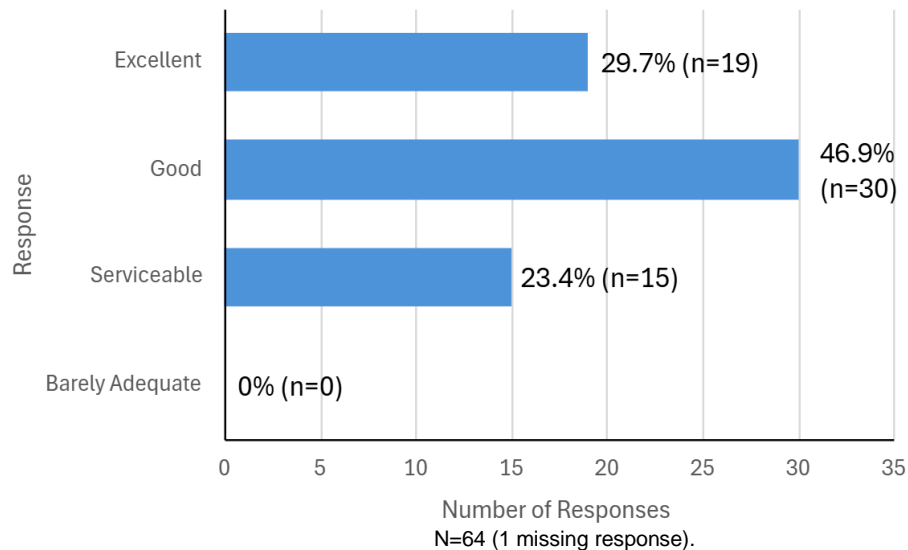
N=64 (1 missing response).

Q3: Please rate your satisfaction with operations (ticket system, responsiveness of staff, documentation, user support, etc). **Overall Results: 75.4% Good or Excellent.**

Number of Responses
N=65.
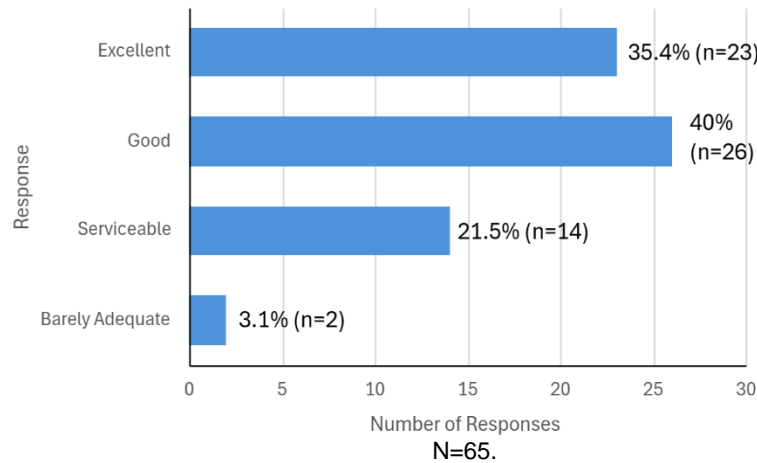
Q4: Is there any other new training session you would prefer adding to 2025 Minerva HPC training serials?

Number of Responses
N=62 (3 missing responses).
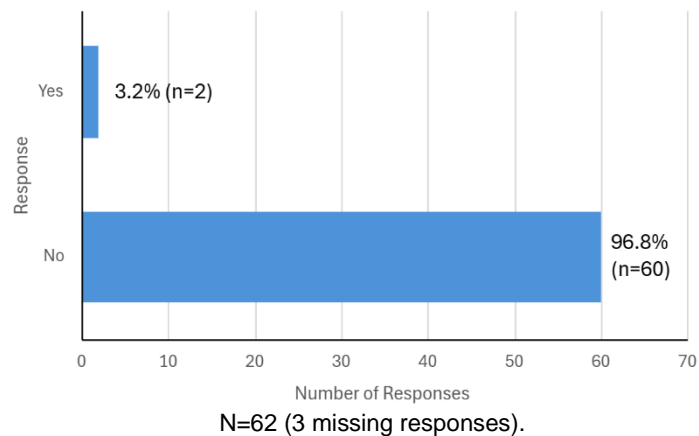
Q4 *(Continued)*: If yes, please specify any new training session topic you are interested in. Currently, the following sessions are available: *1. Minerva Introduction. 2. LSF Job Scheduler. 3. Open OnDemand. 4. Introduction to GPU resources. 5. Accelerating Biomedical Data Science with GPUs: Practical Approaches and Tools. 6. Leveraging Large Language Models in Biomedical Research. 7. How to Accelerate Genome Analysis Toolkit (GATK) by using Parabricks. 8. Data Ark.*

- No responses.

Q5: What suggestions do you have for improving our service? (If you wish to be contacted, please leave your contact information here).

All comments and responses are as follows:

Category	Comments	Responses
Compliments	- Everything about Minerva and the team is amazing! - As someone who works/has worked with multiple different HPCs, I find the Minerva management to be excellent. Documentation is good and support is helpful, making it simple for a moderately experienced researcher to make use of the service. - Year after year, HPC has always been able to help me with whatever needs I had and offered computational resources in a timely manner. Thank you! - Love working with the HPC team and for the most part, the services offered are great. - I want to start by saying that I think the HPC department is great! Minerva is awesome! And I appreciate everything you have done to make it as good as it is right now. - Helpful during the transition time. - There is always a place for improvement but excellent. - The HPC team is incredibly helpful when it comes to tickets. - The staff is responsive and helpful, thank you for all your help in troubleshooting!	The you for the positive feedback.
System Performance Issues	- Love working with the HPC team and for the most part, the services offered are great. Minor hiccups related to the substantial update to the system. - Increased downtime post major update. To be expected, but mildly disruptive. - Overall, your service is excellent. Responses are fast and very helpful. However, there were more issues with disk operations this year compared to the previous year, so I can't rate the first question as "excellent." That said,	We acknowledge there was more maintenance scheduled after the Minerva upgrade. This maintenance is necessary for keeping the system secure and stable with improved features. This amount of extensive maintenance work has not been done recently, resulting in a large amount of outdated hardware. To reduce downtime in the future, we will schedule routine maintenance for system upgrades, send announcements ahead of time, and send repeated reminders so you can plan in advance.

		I truly appreciate all your efforts and support. Thank you!	● We also acknowledge that there are more unexpected issues with the home directory and Arion file system.
		● Since the major update, maintenance and failures at the service are constant.	● As announced, we hit an extremely rare hardware failure which caused the home /hpc file system to be unavailable in December 2024. An unexpected power module failed which caused both battery-driven storage controllers to reboot. We have been working with the vendor to understand the root cause of the power module failure to prevent this from happening again. ● Also, in late January 2025, there were two unexpected outages on Arion (a brief outage and a longer one lasting around 80 minutes). We apologize for the interruption to your research work. In short, the cause of the issue was the online upgrade of the Minerva Arion file system. We completed the upgrade during a scheduled maintenance window in February 2025. In the future, we will also schedule regular windows for these types of challenging upgrades instead of online rolling upgrades to avoid interruptions to your work. We had not rebooted the storage system for a long time due to downtime required, which sometimes makes online rolling upgrades unpredictable. ● We also understand that the recent hardware and OS update has caused incompatibility issues on software modules. Please see details below in the software-related section. ● We appreciate your understanding concerning these issues, and we are sorry for the interruptions to your work. We are working to improve and reduce these disruptions and keep the system available as much as possible. Our team is closely monitoring the system and resolving problems as they arise.
		● There need to be fewer days of downtime. The frequency is causing major disruptions to work. ● Minerva is down a lot, and it does interrupt our jobs, costing labs a lot of money.	● We understand the impact of downtime on your work and lab resources. We are actively working to minimize it. Downtime is often necessary for critical system updates, security patches, and

		hardware maintenance to ensure the stability and performance of Minerva. We appreciate your patience and will continue to communicate planned outages well in advance. If you experience unexpected downtime, please reach out to us at hpchelp@hpc.mssm.edu. ● For unexpected downtime, please see responses above.
	● Sometime file systems errors pop up and limit the usability of Minerva.	● File system errors are rare but there might be hiccups due to high I/O loads or misconfigured jobs. We are actively monitoring the file system and addressing issues as they arise. We are continually refining our monitoring strategies to minimize such disruptions. ● Also, in late January 2025, there were two unexpected outages on Arion as mentioned above.
Job Scheduler & Queues	● Queue times can be quite long. I prefer SLURM over LSF for job scheduling efficiency.	● We understand that queue times can sometimes be long, and we continue to work on improving them with the current available resources. Regarding your preference for SLURM, we have received mixed feedback on switching from LSF to SLURM over the years. Most users prefer to stick with LSF to avoid disruptions to existing pipelines. However, we remain open to revisiting this option in the future as needs evolve. For now, you can use LSF commands like `bqueues` to identify less crowded queues or `bjobs -l <jobID>` to check why your job might be pending. If queue times remain a bottleneck for your work, please contact us at hpchelp@hpc.mssm.edu with specifics so we can explore tailored solutions.
	● The queue priority could be improved as usually a few users take most of the resources available.	● Concerning queue priority, we have implemented a global CPU job slot limit of 4,000 per user, 60 GPU job slots per user and a total memory cap of 15TB per user to prevent resource monopolization and balance the whole throughput of the cluster. We are also adjusting these global limits as needed. Our LSF fairshare scheduling policy also adjusts job priority based on

		historical usage to balance access. If you are still seeing delays due to resource-heavy users, please run `blimits -u <userID>` to check your limits, and contact us at hpchelp@hpc.mssm.edu with details of your experience so we can investigate further.
	● The queuing system is weird. Sometimes a job is waiting in a queue for several days, but a newly submitted job gets run immediately.	● This situation can occur for several reasons. A new job might be scheduled quickly if there are available resources that meet its request, while other jobs may have "pending" status due to resource constraints (e.g., available memory or cores). The dispatch of the jobs depends on the queues and resources requested. We encourage you to use the `bjobs -l` command to check the pending reason for your job. If you notice any recurring issues, feel free to contact us, and we can assist you with a more detailed analysis.
	● Slow running and delays, shuts off in the middle of running code.	● We are sorry to hear about the slow performance and job interruptions you have faced. Jobs shutting off mid-run could be due to exceeding walltime (the actual time it takes for a computer program to run from start to finish) limits, resource overuse triggering a kill signal, or others. To troubleshoot, check your job's status and termination reason with `bjobs -l <jobID>` after it stops. Ensure your resource requests (e.g., memory, cores, walltime) match your code's needs—test runs with smaller datasets can help refine this. If the issue persists, please reach out to hpchelp@hpc.mssm.edu with your job ID and script details so we can pinpoint and resolve the problem.
	● Recently improved but previously had times when I would wait hours in queue and never connect in the fall of 2024.	● We are glad that you have noticed recent improvements, and we are sorry for the long queue waits you experienced earlier in fall 2024. As announced, in mid-July, we decommissioned 90 Minerva compute nodes (4,320 cores) for facility construction work for the major Minerva upgrade, which makes the queue busier.

		● We are continuously working to optimize resource allocation and reduce wait times. However, we understand that each case is unique. If you are still facing delays or want help optimizing your submissions, please contact us at hpchelp@hpc.mssm.edu with more details about your job (e.g., job ID, resource requests). We will see if there is a way to improve it further for you.
	● Some compatibility issues after the major hardware upgrade and change to Rocky9.	● We appreciate your feedback and are sorry for the issues you have encountered. Before the upgrade, we tried to rebuild all major modules with Rocky9 and set up some test nodes for users to run test and report issues on missing libraries and compatibility issues. However, there are still packages and libraries that haven't been rebuilt since there are over 1,000 packages with multiple versions. We will continue to address any reported issues. ● Please use the new version of the software as much as possible, as these have been specifically installed for Rocky9. Please share specific details about the compatibility issues you face by contacting us at hpchelp@hpc.mssm.edu. We will work with you to identify and resolve these problems as quickly as possible.
Software / Packages	● Can be improved. The R package installation system is still a bit difficult for some packages and there are incompatibilities with the new LSF migration. ● I prefer to be able to install R packages freely, which is not possible right now.	● Thank you for your feedback. We understand that R package installation can be challenging, especially after the LSF migration, and that some users prefer more flexibility in installing packages. Currently, users can install R packages locally by specifying the "R_LIBS_USER" environment variable to a desired location, rather than the default "~/.Rlib" directory. For more detailed instructions, please refer to our documentation: https://labs.icahn.mssm.edu/minervalab/documentation/r/. If you encounter any issues or need further assistance, feel free to reach out to us at hpchelp@hpc.mssm.edu, and we will be happy to help.

● Many packages used for R (especially those requiring github access) are not compatible with R webstudio. ● The new implementation of rstudio is much improved, but loading + installing new packages in this environment could be easier. e.g. workarounds for loading git2r and its dependencies are currently necessary.	● We understand that many R packages, especially those requiring GitHub access or dependencies like `git2r`, may not be fully compatible with the current RStudio setup, and that workarounds are sometimes necessary for installing and loading these packages. While R WebStudio is a lightweight tool designed for initial development and testing, it may not support all packages out-of-the-box, especially those with external dependencies. As a solution, you can install these packages locally by setting the `"R_LIBS_USER"` environment variable to a custom directory and configuring your session accordingly. We are also continually improving the RStudio environment and will consider your feedback for future updates. If you encounter specific compatibility issues or need assistance with package installations, please do not hesitate to reach out to us at hpchelp@hpc.mssm.edu.
● Running python is honestly a huge hassle. Every time I need to install a new package, it feels like everything breaks. Maybe it would be possible to provide a written guide to this, if venv vs conda is recommended, etc.	● We have installed most of the popular Python packages. If a specific Python package is not available, you can install it yourself. Please check the following links for instructions on how to use `Conda`, `venv`, and `pip`. If you need any help, you can open a ticket at hpchelp@hpc.mssm.edu. ● Minerva Python Instructions: https://labs.icahn.mssm.edu/minerva lab/documentation/python-and-jupyter-notebook/ ● Minerva Conda Instructions: https://labs.icahn.mssm.edu/minerva lab/documentation/conda/
● I recently was unable to apply my conda environment to my Minerva-based jupyter notebook that I launched with minerva-jupyter-module-web.sh. I needed to customize this script in order to successfully apply my full conda environment. Instructions online for facilitating this more easily	● Thank you for bringing this to our attention. You can activate a `Conda` environment using the Jupyter on-the-fly script as shown below: ● `sh minerva-jupyter-module-web.sh -mm anaconda3/2024.06 -env envName`

	would be appreciated for other users, I think.	● To see the Jupyter on-the-fly script usage instructions, run the following command: ● `sh minerva-jupyter-module-web.sh --help` ● We will update our documentation to include instructions on how to use your `Conda` environment with the Jupyter on-the-fly script. If you need help, please contact us at hpchelp@hpc.mssm.edu.
Tickets / Staff Response	● Sometimes ticket response time is way too long. Please hire more people. The staff is great, but it would be helpful to receive faster responses to the tickets.	● We acknowledge that response times can sometimes be longer than expected, and we apologize for any inconvenience this may have caused. Thank you for your understanding of the staff bandwidth. Yes, it was quite challenging late last year with limited staff. As announced, there was a hiring freeze at the end of 2023, so we must extend the response time. In 2024, the HPC team handled 5,740 tickets and major upgrades on hardware and OS/software. ● We will improve our ticket response time. We finally got the positions reopened and hired 1 more computational scientist in 2024 and 2 more HPC admins in 2025. We are also actively hiring more staff to refill some positions.
	● My main complaints are all due to a lack of staff/resources. I would prefer more documentation, and the user experience is pretty bare bones compared to HPCs I've used at other institutions. A few specifics that come to mind are the log/error file system defaults, the fact that jobs requesting impossible resources pend forever instead of throwing an error, and the lack of backups for user dirs. However, I know it's only a handful of people running this crucial resource that I use every day. I can't expect them to make many major improvements when they're already actively maintaining the system and putting out fires daily. That's why my main recommendation is just to hire more staff. I have had wonderful experiences with the	● We greatly appreciate your patience and feedback. We are actively expanding our team and hiring more staff to better support our growing user base. Additionally, we are updating and improving our documentation to make the user experience smoother. For jobs pending due to impossible resources, we will explore adding error notifications in a future update. We provide a Tivoli Storage Manager (TSM) service for backup and archiving user data. However, we do not back up user data. Users are responsible for their own backups. Thank you for recognizing the hard work of our team.

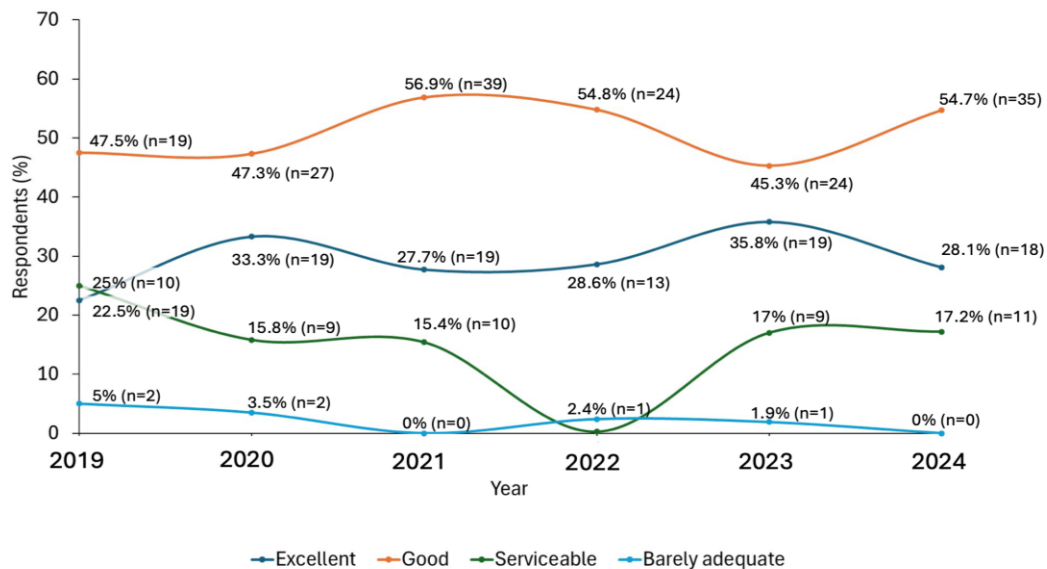|		
	HPC staff in my couple of years here at MSSM. I watched them work very hard to prepare for the OS upgrade this fall/winter, and they are impressively responsive to my tickets & questions considering all the other responsibilities they have. I have attended their trainings multiple times and always found them very helpful. Specifics: I use the RStudio web server every day, and Jielin has been awesome about keeping it working even through the OS changes. Lili is always great about giving users plenty of advanced warning before any service outages or major changes, so I've always been able to plan around them and have minimal disruption to my research. Thanks to them, my coworkers & I were able to test out our scripts on the new computing nodes a few weeks before the upgrade officially went through, so we worked out all the new bugs before they could become an issue. I haven't interacted as much with Wei or Hyung Min yet, but they have always been helpful answering my questions during training and via email. In summary, Minerva is outdated/basic, but the HPC team has managed to make it functional through a lot of hard work and communication. I am impressed with them all. The real need here is for more staff and funding.	
	● Also, a new ticket system where it is possible to follow up on a question would be helpful. Right now, each email opens a new ticket which can be confusing. ● As a user with multiple open tickets at any given time, it would be helpful to include some additional information from the initial email with the responses from HPC staff.	● We understand the challenges with the current ticketing system, particularly regarding the status of the tickets. We are investigating our current free OS ticketing system to see whether we can improve the status check so you can see all the previous replies with your ticket number and follow up directly with the user interface. ● We are also exploring the UI interface that will allow users to submit tickets, track their status, and receive answers—all in one place. In addition, we are also investigating other ticketing systems for more

	friendly service management and user interface.
● We understand staffing is tight now, but some more information about the status of requests, where requests are in the queue and the ability to withdrawal requests, if necessary, would be helpful. Also, I would encourage your team to consider if there should be a separate mechanism for requesting additional feature, rather than submitting as an incident ticket. Thank you.	● Thank you for your feedback. Your suggestions are valuable. We will discuss them within our team and work on improving our ticketing system. Please also see answers above regarding the ticketing system.
● Overall fine, we only had one ticket where staff did not want to solve the issue, but another one was willing to help even though this situation should not have been created in the first place.	● We are glad to hear that your overall experience has been good, and we appreciate your feedback on this situation. While we work to resolve every issue well, we are sorry that this happened. We will use this as an opportunity to improve staff training and communication to avoid similar situations in the future. Please do not hesitate to submit a follow-up ticket.
● The ticket system works okay, but problems were not always solved when I opened a ticket in the past.	● We are sorry to hear that some of your issues were not fully resolved. We aim to solve every problem, but we admit that mistakes can happen. If a ticket slips through or remains unresolved, please send us a reminder at hpchelp@hpc.mssm.edu. We have also expanded our team to improve our ability to address issues effectively.
● Staff responsiveness is often slow, and responses are not entirely helpful, especially for new users. It would be helpful to have more in-person times to talk to staff to get quicker responses and more detailed answers / feedback.	● We are sorry that our response time has not fully met your expectations. We are working to improve this by hiring more staff to answer requests faster. We offer Digital Concierge sessions every Wednesday from 3:30 to 4:30 PM on Zoom and encourage you to join. Detailed information and registration links are emailed weekly. ● You can also and always request one-on-one consultations by emailing hpchelp@hpc.mssm.edu. We will schedule the session with you soon.
● Need quality staff.	● We are committed to hiring and retaining highly skilled staff to meet your needs. Our team is continuously being trained and expanded to ensure we provide the

		best service possible. If you have specific feedback or suggestions regarding staff performance, please share them with us.
Documentation on Website	● Documentation needs to be updated to reflect hardware and software changes. ● Documentation on new implementations of rstudio etc could be improved + better pointed to on the Minerva website.	∉ We are currently working on updating the documentation to reflect the latest hardware and software changes, including new implementations of RStudio. Additionally, we are improving our website to ensure that these updates are easy to find and access.
Other Comments	● Maybe more beginner friendly. Also work with MSSM VPN to develop a more feasible and easier way to use the login method. The vpn doesn't work outside of the US, and there is no way to connect to hpc when you are traveling.	● We are committed to making Minerva more accessible to beginners. Our training sessions and Digital Concierge Zoom meetings (Wednesdays, 3:30–4:30 PM) aim to support new and existing users. Please attend the above meetings as needed. ● Mount Sinai cybersecurity policy requires Minerva login nodes only be accessed within the Mount Sinai network. The HPC team does not handle VPN related issues. Please contact ASCIT (ASCIT@mssm.edu) for support. We know that some users outside the country can access Minerva. The issue that you experienced may be related to a specific location.
	● I think monitoring user activity can be improved. I will explain below. At the moment, from a user's perspective, the only issue that is affecting the quality of HPC and our experience is how other users seem to be utilizing these resources. And I think that issue can be fixed by communication and activity monitoring. For both points I am making in the following, I have witnessed instances of them firsthand. So, I know they are somewhat anecdotal; so, feel free to ignore both if there is already action being taken. First, is the issue of running code on login nodes, which can be solved through monitoring users. Sometimes I cannot run a simple submit command but could see in 'top' results that other users are running memory and CPU heavy	● Thank you for all the good suggestions. We are monitoring the login nodes and queue loads more actively these days and communicating with the users regarding resource misuse. Hopefully you will notice some improvement. We are also working to develop automated ways of monitoring and user education. ● If you observe specific instances of misuse, please also report them to us to help at hpchelp@hpc.mssm.edu so we can address them promptly.

	tasks on the login node. The second issue is the utilization of special nodes for unrelated tasks. For example, using HIMEM nodes when the tasks do not require large amounts of memory; or using GPU nodes when the tasks do not use GPUs at all simply due to these nodes having "better" hardware; or the queues for them being shorter at that moment. I know in some of the training sessions, users are advised not to use the system in this way. However, not everyone attends these sessions. Again, I appreciate all you do and the opportunity to provide feedback.	

Trends in Minerva User Satisfaction Survey Results (2019–2024)

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

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

Respondents (%) vs **Year**

- Excellent
- Good
- Serviceable
- Barely adequate

Good: 57.5% (n=23) [2019], 51.8% (n=30) [2020], 57.8% (n=39) [2021], 61.0% (n=27) [2022], 54.7% (n=30) [2023], 46.9% (n=30) [2024]

Excellent: 17.5% (n=7) [2019], 28.6% (n=16) [2020], 23.4% (n=16) [2021], 31.7% (n=14) [2022], 30.2% (n=16) [2023], 29.7% (n=19) [2024]

Serviceable: 15.0% (n=6) [2019], 12.5% (n=7) [2020], 15.6% (n=11) [2021], 7.3% (n=3) [2022], 11.3% (n=6) [2023], 23.4% (n=15) [2024]

Barely adequate: 5.0% (n=2) [2019], 3.6% (n=2) [2020], 1.6% (n=1) [2021], 0% (n=0) [2022], 3.8% (n=2) [2023], 0% (n=0) [2024]

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

Respondents (%) vs **Year**

- Excellent
- Good
- Serviceable
- Barely adequate

Excellent: 45.0% (n=18) [2019], 52.6% (n=30) [2020], 61.5% (n=42) [2021], 53.7% (n=24) [2022], 35.2% (n=19) [2023], 35.4% (n=23) [2024]

Good: 32.5% (n=13) [2019], 38.6% (n=22) [2020], 24.6% (n=17) [2021], 34.1% (n=15) [2022], 46.3% (n=25) [2023], 40.0% (n=26) [2024]

Serviceable: 15.0% (n=6) [2019], 5.3% (n=3) [2020], 10.8% (n=7) [2021], 12.2% (n=5) [2022], 18.5% (n=10) [2023], 21.5% (n=14) [2024]

Barely adequate: 7.5% (n=4) [2019], 3.5% (n=2) [2020], 3.1% (n=2) [2021], 0% (n=0) [2022], 0% (n=0) [2023], 3.1% (n=2) [2024]