2020 Minerva User Survey Comments and Response
Scientific Computing and Data

System Issues

- The file system changed so fast, from minerva to hydra then to arion. I have to change my scripts all the time.
    - In 2019-2021, as both Orga and Hydra were out of warranty by end of 2020,, we did a complete refresh of the Minerva cluster, and a subsequent expansion of the cluster.. The file system name has changed from Orga -> Hydra-> Arion. The legacy hardware is now retired and Arion is the only online storage.The systemwill be stable for next 5 years.
    - To assist in the conversion of LSF job scripts, you can use the system environment variable ($GPFS) in your script. Please use $GPFS instead of /sc/arion or /sc/hydra.
      For example, /sc/arion/projects/myproject-dir/files -> $GPFS/projects/myproject-dir/files.

- Login to multiple and changing servers (hydra, orga, minerva, etc) makes it difficult to keep track. Need a single login that will redirect to home and links to stored data that will not change if data is moved.
    - We have retired all the old hardware and Arion is now the only online storage and will be stable for 5 years (see answers above).
    - The Minerva login is ssh userid@minerva.hpc.mssm.edu, which will direct you to the nearest login nodes, and stand at your home directory. You can make symbolic link to your data on Arion if you want to access from home directory (for example, ln -s /sc/arion/projects/myproject-dir/files files or ln -s $GPFS/projects/myproject-dir/files files).

- Sometimes login nodes are misused as compute nodes by users which could be resolved quicker. Also sometimes the disk access is jammed by massive I/O-operations. ; My work is heavily reliant on minerva. Although it's rare, minerva sometimes suffers from connectivity issues and unusable because the login node is not responding or slow. Reliability should be the number one priority in my opinion.
    - We already have some limitations on the login nodes (you can check by ulimit -a).

    - We will continue to improve our monitoring strategy on the usage of the login nodes and I/O.
    - We will emphasize the shared use of login nodes with our users.
- More temp disk space please.
    - Our current scratch has a quote of 100 TB for intermediate data during your pipeline processing. We also have set a per-user space limit on scratch to 15 TB. In addition, we have a policy running daily to purge the files every 14 days.
    - If you have a special situation, please contact us at hpchelp@hpc.mssm.edu to discuss the needs.
- Please send more notice and give more time before major change on minerva system. For example, /sc/hydra/work retired on 1/15. Only one notice is given on 1/13. People could lost import files if they miss the only one notification.
    - Our general policy is to give two-week notice, except for urgent security patches and major system issues. We don’t want to spam the users’ mailbox with too many notifications.. Going forward, we will give more advance notice on the system changes.
    - For hydra retirement, we have already sent notification to users on 10/23/2020 and 11/08/2020 indicating “Hydra will not be available after Dec 5th, 2020).” However, the hydra retirement is delayed until ?.
- Universal data transfer would be useful (MountSinai data point for example).
    - The recommended data transfer tool on Minerva is Globus. Due to security reasons, we cannot make data connected universally. Please send us an email at hpchelp@hpc.mssm.edu to discuss the specific questions or suggestions on other data transfer tools.

- Locking project directories based on the HIPAA form slowed productivity. It can be difficult to get busy PIs to sign forms - and it was confusing to them because they are not storing PHI. I suggest it would have been better to require a form to store PHI, but allow users who have no intention to store PHI to continue without interruption.
    - Mount Sinai Compliance and the IT Security departments has approved storage of Protected Health Information (PHI) data on Minerva to accommodate researchers’ request of processing PHI data.. This requires that all users understand and follow the new Minerva HIPAA policy and Minerva User Agreement. You need to agree to these new policies whether you plan to use PHI on Minerva or not.

    - We have sent several announcements starting from 09/2020, and the lock was implemented in late 10/2020 (about two month period).

- Need (and users are willing to pay for) more comprehensive and user-friendly backup solutions.
    - Please let us know if you have a specific backup solution in mind. . e currently use TSM system as our backup solution. You can setup a cron job on one of the login nodes to run backup automatically following instructions at https://labs.icahn.mssm.edu/minervalab/tsm-faqs/
    - If you need help on this, please send us ticket at hpchelp@hpc.mssm.edu to discuss your specific requirement

Connection Issues

- The login procedure with VIP tokens is very annoying.
    - For security reasons, we need the 2-factor authentication for Minerva login.
    - You can set ControlMaster to reuse ssh connections for all hosts and set alias for hostname to make it easier for multiple ssh connections. The setting details is at page 10 of slide https://labs.icahn.mssm.edu/minervalab/wp-content/uploads/sites/342/2020/10/Minerva_Intro_-2020-09-16.pdf
- VIP token lock
    - Usually, the token will be locked after too many times of failure(> 3 times). The 6 digit code generated from the VIP token changes every 30 seconds, and each code can only be used once for authentication. If you use it multiple times, this will result in the token lock. Also please do not use" the save password function" with your app.
    - The VIP token is managed by school IT. If you still have issue with your token, please contact school IT at ASCIT@mssm.edu

- Remote connection issue; connection drop; broken pipe.
    - Please use F5 VPN Tunnel for the campus network connection while off-campus. Details can be found here https://itsecurity.mssm.edu
    - If campus network is fine, if you still experience a frequent drop of your ssh connections, you may use the following ssh configurations to keep your ssh session alive:
      On your own computer (not Minerva), edit the ~/.ssh/config file:
      ```
      Host minerva
      Hostname minerva.hpc.mssm.edu
      ```

```
ServerAliveInterval 240
ServerAliveCountMax 2
```

With this configuration, the SSH client sends a packet to the server every 240 seconds (4 minutes) to keep the connection alive. If the client does not receive a response after two tries (as specified by the ServerAliveCountMax setting), it closes the connection.

  - We also recommend you use a screen session so that you can return to your work after the drop of the connection.
  - If you have issue with setting up VPN account, VPN login, VIP token register, please send ticket to school IT at ASCIT@mssm.edu.
  - lease keep in mind the stability of your local internet can affect your connection speed.
  - We are talking with school IT to explore any possible improvements.

- CyberDuck requires entering my name/pwd 3-5 times before it lets me upload or download a single file.
  - While using VPN (the bandwidth is limited), data transfer from Minerva to your own computer will be extremely slow and it will also cause network congestion for other users. We recommend you use Globus for this situation which bypasses the VPN. Check our slides for Globus.
  - Please contact hpchelp@hpc.mssm.edu for specific CyberDuck upload/download issues.

- As I work from home at the moment, time limitation of the vpn connection is getting very inconvenient. It looks like the connection is open for 12 hours. After that time, it drops without any notifications or warnings. All unsaved work and terminals get lost. There is no option to continue the current session or change the connection time limit. Even a clock showing the rest of time would make things better. In addition to losing all unsaved data and terminals, I need to type in the security codes at least 6 times to reconnect properly.
  - The 12 hours connection limit is the policy set on school VPN. You can contact school IT at ASCIT@mssm.edu for the VPN alert settings.
  - Please refer to the above reply on “Remote connection issue; connection drop; broken pipe.” for additional information. .
  - Please use screen session for return to your work.

Job Scheduler & Queues

- The LSF queue structure works ok, and jobs run efficiently. However, I've run into issues with submission scripts, especially while passing variables. The same scripts work fine using Slurm, only to fail with LSF.
  - We are sorry to hear that. Please send us more details about your issues with the LSF submission script at hpchelp@hpc.mssm.edu. We will work with you on the issue.
- When it works, it's great! One small issue: I am on numerous private queues. This is fine, but I would like to be able to have a way to submit jobs to all but 1 specific private node.
  - We only have one private queue with different private nodes. If you submit your jobs to the private queue, LSF should be able to dispatch your jobs to the available private nodes you have access to in the queue. If it is not working, please send us a ticket at hpchelp@hpc.mssm.edu. We will work with you on the issue.

- However, it will be great if there is someway for users to monitor how "busy" minerva is so that we can know if our submitted jobs can run or not
  - You can use LSF commands such as bhosts, bqueues, bjobs to check the status of the cluster and your job. Our latest slides on LSF is at https://labs.icahn.mssm.edu/minervalab/wp-content/uploads/sites/342/2020/09/Minerva_LSF_2020-09-23.pdf.
  - We have installed a wrapper script authored by Harm van Bakel, which will make it easier to interact with the LSF job scheduler on Minerva, as suggested.
    To load them up, $ml LSFqueue
    To get more info on the module, $module help LSFqueue;
    And a detailed readme file at
    /hpc/packages/minerva-centos7/LSFqueue/1.0/README

  - We will include this in upcoming training session
  - You can get quick commands listed by running “minerva_help” on login nodes
- At times I need to wait for too long in the queue
  - The wait time depends on the status of the Minerva cluster and the resource request of your jobs. When the cluster hosts are all busy (check with bhosts) with many pending jobs, the wait time will be longer..
  - We added an extra 33 high memory nodes in Oct. 2020 to accommodate the memory-intensive jobs which helps to greatly reduce the wait time. In addition, we added 32 A100 GPUs in Nov. 2020 to reduce the wait time in GPU queues.

    - We monitor and set limit on running jobs for heavy users when the cluster is busy
    - We set maximum pending jobs for each user as 20,000 by default to prevent users from submitting several (tens of thousands) short jobs (0(minutes)) which result in many queries on LSF
    - We will adjust our queue structure and dedicated nodes to reduce the turnaround time and increase the throughput.
    - Some tips for users to reduce the wait time:
        - check whether your jobs run into maintenance : added default check for users with immediate screen message.
        - please check your LSF job logs for your successfully finished job, and give a close estimation of memory request for your LSF job. A larger memory request such as 20,000 MB may not be necessary for your own workflow. A smaller memory will help reduce your waiting time significantly.
        - Use LSF command “bqueues” to check the status of the queues. It will help to submit your jobs to the queue that is less occupied, if appropriate, such as express queue (which is up to 12h of walltime).
        - Use LSF command “bjobs -lp JOBID” to check the pending reason.
        - If you have urgent request, please contact us at hpchelp@hpc.mssm.edu for help
- Change to slurm
    - We plan to set up a test cluster with SLURM this year.

TSM

- It seems that archive system is not stable. It crushes every 2 months
    - We currently have a total of 23.8 PB data on TSM which sits on more than 16,000 tapes. Unfortunately, the current TSM library is not big enough to accommodate all these tapes. When a retrieve command is evoked, the operators have to fetch the tape and load manually into the library. The response time depends on many factors but should be less than 1.5h. Once the tape is loaded, the data transfer starts at a constant rate of 150M/s. The data unavailable error is due to an error that is raised on the retrieve process and needs manual interference to fix. The cache full error is due to lack of scratch tapes. Those errors are inevitable but we will monitor the TSM more closely to avoid situations like this.

- There were a few tickets (mostly about archives) that have been without answering for weeks.
    - We had some resource constraints last year, which was challenging. This has caused slower than usual response time to help researchers on some specific topics. The reply to TSM related questions may be delayed. We have hired two new system admins last year and there is still one open system admin position. We will try our best to reduce the time taken for TSM ticket response this year.
    - Detailed documentation on TSM is available online which may help with some of the questions. Please see https://labs.icahn.mssm.edu/minervalab/archiving-data/ and please contact us for more assistance.

Software/packages

- Maybe available software should be tested better as sometimes errors during loading and/or during simple jobs occur (e.g. missing libraries)
    - Sorry for the inconvenience. We will review our testing processes and include more sanity checks before releasing the modules.
- Remove old, outdated packages.
    - We are trying our best to clean up the unnecessary packages, but we cannot really remove all the old software, since many users still rely on those for reproduction and consistency of their publications.
- Lately I have had some issues with remote visualization. Specifically, the VMD or the PyMOL modules produce error messages like the ones bellow when trying to load a file: - VMD: XRequest.150: GLXBadContext 0x4600105 XRequest.150: GLXBadContext 0x0 - PyMOL: QXcbConnection: XCB error: 8 (BadMatch), sequence: 646, resource id: 60817412, major code: 150 (Unknown), minor code: 11
    - Currently, for remote visualization over GUI, you need to use X11 forwarding. Make sure you have the appropriate XQuartz or Xming installed locally. Check page 7 and 8 at https://labs.icahn.mssm.edu/minervalab/wp-content/uploads/sites/342/2020/10/Minerva_Intro_-2020-09-16.pdf
    - If you still have a problem, please send us an email at hpchelp@hpc.mssm.edu. We will debug the specific issue with you.
- MATLAB license structure is not good for labs with multiple users.

    - The MATLAB license is administered by school ASCIT team. Please check with the ASCIT team (ASCIT@mssm.edu) to see if they can offer flexibility for labs with multiple users.
- I am not sure if singularity can do everything but I have a strong feeling that we will need to work with docker images more and more.
    - We can only support singularity on the shared HPC system due to security reasons. You can pull/run docker images with singularity. See our documents at Page 16 at https://labs.icahn.mssm.edu/minervalab/wp-content/uploads/sites/342/2020/10/Minerva_Intro_-2020-09-16.pdf

Tickets

- Ticket response is very fast but I could not use the link in the email to enter the ticket tracking/follow-up portal.
    - We are not able to reproduce this issue. Please send us a ticket at hpchelp@hpc.mssm.edu. We will work with you on the issue.

- Sometimes the response is slow, but overall I have been happy with the help. I've seen the ticket system was broken in a few instances, but overall, it's responsive and I am happy with it.
    - There were ~3 unscheduled outages on the ticket system in 2020 but otherwise it was available.
    - We will develop a strategy to make sure the ticket system is more highly available.

- Provide better troubleshooting help. I think that it is so difficult (as a student) to understand how these jobs can be submitted and how to deal with errors when they arise. There should be 1-1 sessions or better responses when questions arise.
    - You can always request 1-1 sessions by send us email at hpchelp@hpc.mssm.edu
    - There is documentation at https://labs.icahn.mssm.edu/minervalab/ to help guide you.
    - We have several training classes a year and announce these through the Minerva mailing list and post them on our hpc website.

Tutorial/Training

- Containerized computing; I would be interested to learn more about containers and docker images. From practical point of view, not just list of commands.; I'd be interested in a session about jupyter notebook on minerva.
    - We will schedule a session this semester. We will send out announcement when it is scheduled
- Training videos would be helpful.
    - Good suggestion. We will work on developing training videos.
- As new people join the lab it is always good to have those sessions so they can attend and for current people to get more advanced.
    - We have put our training slides on our website and the latest slide is at https://labs.icahn.mssm.edu/minervalab/the-minerva-user-group-and-training-classes/
    - We offered four tutorial training in 2020, including both "Introduction to Minerva environment" and "Submitting jobs via LSF". The 2021 training sessions will be announced once they are scheduled.
    - If your research group would like to schedule an additional tutorial or tutorials focused on special HPC topics, you can reach out to us at hpchelp@hpc.mssm.edu

- I had some trouble making A100 working for my exiting tensorflow workflow. I think A100 needs a more recent version of TensorFlow and cuda/cudnn library. It would be nice to have a session on this topic.
    - We are investigating and updating the software as needed. It will be helpful if you can send us details on this for debug at hpchelp@hpc.mssm.edu
- I've been to a Minerva basics session, and it was not helpful. The bar for figuring out how to do work on Minerva is too high, and keeps researchers away.
    - You can always request 1-1 sessions on Minerva environment by sending us email at hpchelp@hpc.mssm.edu
    - Other resources available include
        - For basic linux command you can visit tutorials from Google such as https://www.guru99.com/linux-commands-cheat-sheet.html
        - There is also course on Introduction to Scientific Computing BSR1015 (a two credit course with an expanded lab)
    - We are also investigating other portals to provide more convenient access of our HPC resources such as open onDemand or TurvoVNC. We will send out announcement when it is ready

Documentation on Website

- The documentation is horrible; Documentation online could be improved. Inclusions such as using anaconda environments and whatnot would be appreciated. Much of the documentation for this stuff exists within the package's own sites, but pointers there would be great.
    - We are continuously working to improve our documentation. We have been updating the changes on the system and adding new services. As suggested, we add more pointers to the manuals on Minerva's website.
    - For the most recent updates and guides, please read/follow our announcement and Newsletter (subscribe to hpcusers@mssm.edu). We will update the documentation on our website in a timely manner to reflect the changes in our system.

Others

- Can we get ondemand on minerva (https://openondemand.org/)? It would allow users to kick off jupyter/Rstudio session from the browser and explore files.
    - Evaluation of OnDemand and similar tools is in progress. . We plan to have it available by April ?.
    -
- I think that much can be done to improve the user-ability of the server; The server is not very user friendly. I am not very proficient at coding, but have used other servers for large data sets at my prior institute that was easier to navigate.
    - Please send us a ticket at hpchelp@hpc.mssm.edu to discuss the specific questions. We welcome any enhancement requests and will follow with you on feasibility and inclusion on future updates to the system.
- It would be nice if you add Hadoop env
    - MapReduce workloads represent only a small fraction of our overall workload so far, but require their own standalone environment. Our HPC cluster uses a robust parallel file system, GPFS, that is sufficient, and, in some cases, superior for many MapReduce workloads. Thus, you can run MapReduce workloads as regular parallel jobs. Please contact hpchelp@hpc.mssm.edu if you need assistance in optimizing your workload.
- Lack of efficient local-server connections. E.g., jupyter-lab notebooks.
    - We are working on the on-the-fly jupyter notebook script. Currently, you can access it by minerva-jupyter-web.sh (for help message minerva-jupyter-web.sh -h), and some details here. We are still working to improve it and writing documentation on this.
- Change server 'apache' like settings to be compatible with WashU epigenome browser

    - Please provide us (hpchelp@hpc.mssm.edu) the setting that you want to change for making it compatible with WashU epigenome browser. We will review the settings, and provide suggestions accordingly.
- Being able to setup user-owned environment (e.g., using linuxbrew, etc.) in user directory, w/o relying on packages being available at the system level.
    - The linuxbrew should be working on our system. There may be something hard-coded in the version of linuxbrew you use, and failed to detect your user-owned brew environment. Please send us a ticket at hpchelp@hpc.mssm.edu to discuss your specific issue.

Thank you for your nice words!

- Most of the time it works very well.
- There is always a place for improvement but overall, excellent!!
- I'm sure there are improvements to be made, but even right now, can be considered excellent!
- The server are usually great with great supports.
- The Minerva staff is amazing. Very responsive and supportive. Thank you for that.
- HPC staff is fantastic. Super conscientious and always helpful.
- HPC helpdesk operators are always very professional and willing to help.
- Missing packages are quickly installed when requested.
- Team does a great job updating software packages!
- The team is very responsive and helpful with issues.
- Responsiveness of HCP is usually great!
- Generally respond quickly.
- User support is outstanding!!!
- No complaint. Reasonable quick response and very helpful
- I really appreciate the transformation from a very unreliable to a very solid computational environment that I witnessed last year. Thank you for making it possible.
- Staff is responsive.
- You people are the best. Very kind and helpful
- Great job.
- Everything is perfect. The HPC staff is very kind and helpful.
- The service is great and provides us with a platform to perform data analysis. The people are great to work with