

Minerva HPC Town Hall

Partnering with scientists to accelerate scientific discovery

Patricia Kovatch, Dean for Scientific Computing and Data

Lili Gai, PhD, Director for High Performance Computing and Data

The Minerva HPC Team

Scientific Computing and Data

May 29, 2026

Agenda

- 1 Accomplishments and Updates
- 2 2026 Roadmap
- 3 Appendix I – Minerva Usage
- 4 Appendix II – 2025 User Survey Results



Accomplishments and Updates (Oct 2025 - Current)

Updates on NIST SP 800-171 compliant environment to hold data from NIH-controlled genomic data repositories

The new Minerva Restricted Cluster (Minerva-Res) will be open targeting **July 2026**

- ✓ 2x dedicated login nodes accessible with security hardening
- ✓ GPFs filesystem (6.9 PB in total) with encryption at rest as */arionencrypt*
- ✓ LSF job scheduler including a set of high memory compute nodes with security hardening
- ✓ Open OnDemand for graphical interfaces with limited features

- These services will **not** be available:
 - Database server; Web server; Posit Connect; Globus

- Pending items:
 - System security plan (SSP) & policy approval, operation procedure & user documentation, friendly user testing

Minerva-Res is currently only approved for data obtained from NIH controlled genomic data repositories

New Procedures to Access the Minerva Restricted Cluster

▶ Project access

- PIs submit [Minerva Restricted Cluster: Project Resource Allocation Request Form](#)
- PIs review & agree to “[Annual Minerva Restricted Supercomputer Policies and Procedures](#)”

▶ User accounts

- Users review & agree to “[Annual Minerva Restricted Supercomputer Policies and Procedures](#)”
- PIs approve the user request
- ✗ Note: external accounts not supported currently due to the background screening security requirements.

▶ Network access: Deny all by default. No network access on all nodes

▶ Data movement will be strictly restricted

- ✓ Data in: Allowed with approval
- ✓ Data out (Non-controlled Data): [Allowed with the PI approval for every transfer](#)
- ✗ Data out (Controlled Data): Not permitted without explicit review; Contact hpchelp@hpc.mssm.edu [with PI approval](#)

Data Migration into Minerva-Res Cluster During Transition Period

Status:

- There are 131 projects currently handling controlled genomic data set on Minerva

Proposed Plan:

- For the data already on Minerva, users will move the data from Arion to Arionencrypt
 - Arion mounted as read-only on the Minerva-Res login nodes
 - A month transition period
- Users remove all controlled data from Arion then

Data Transfer on Minerva-Res Cluster

Status:

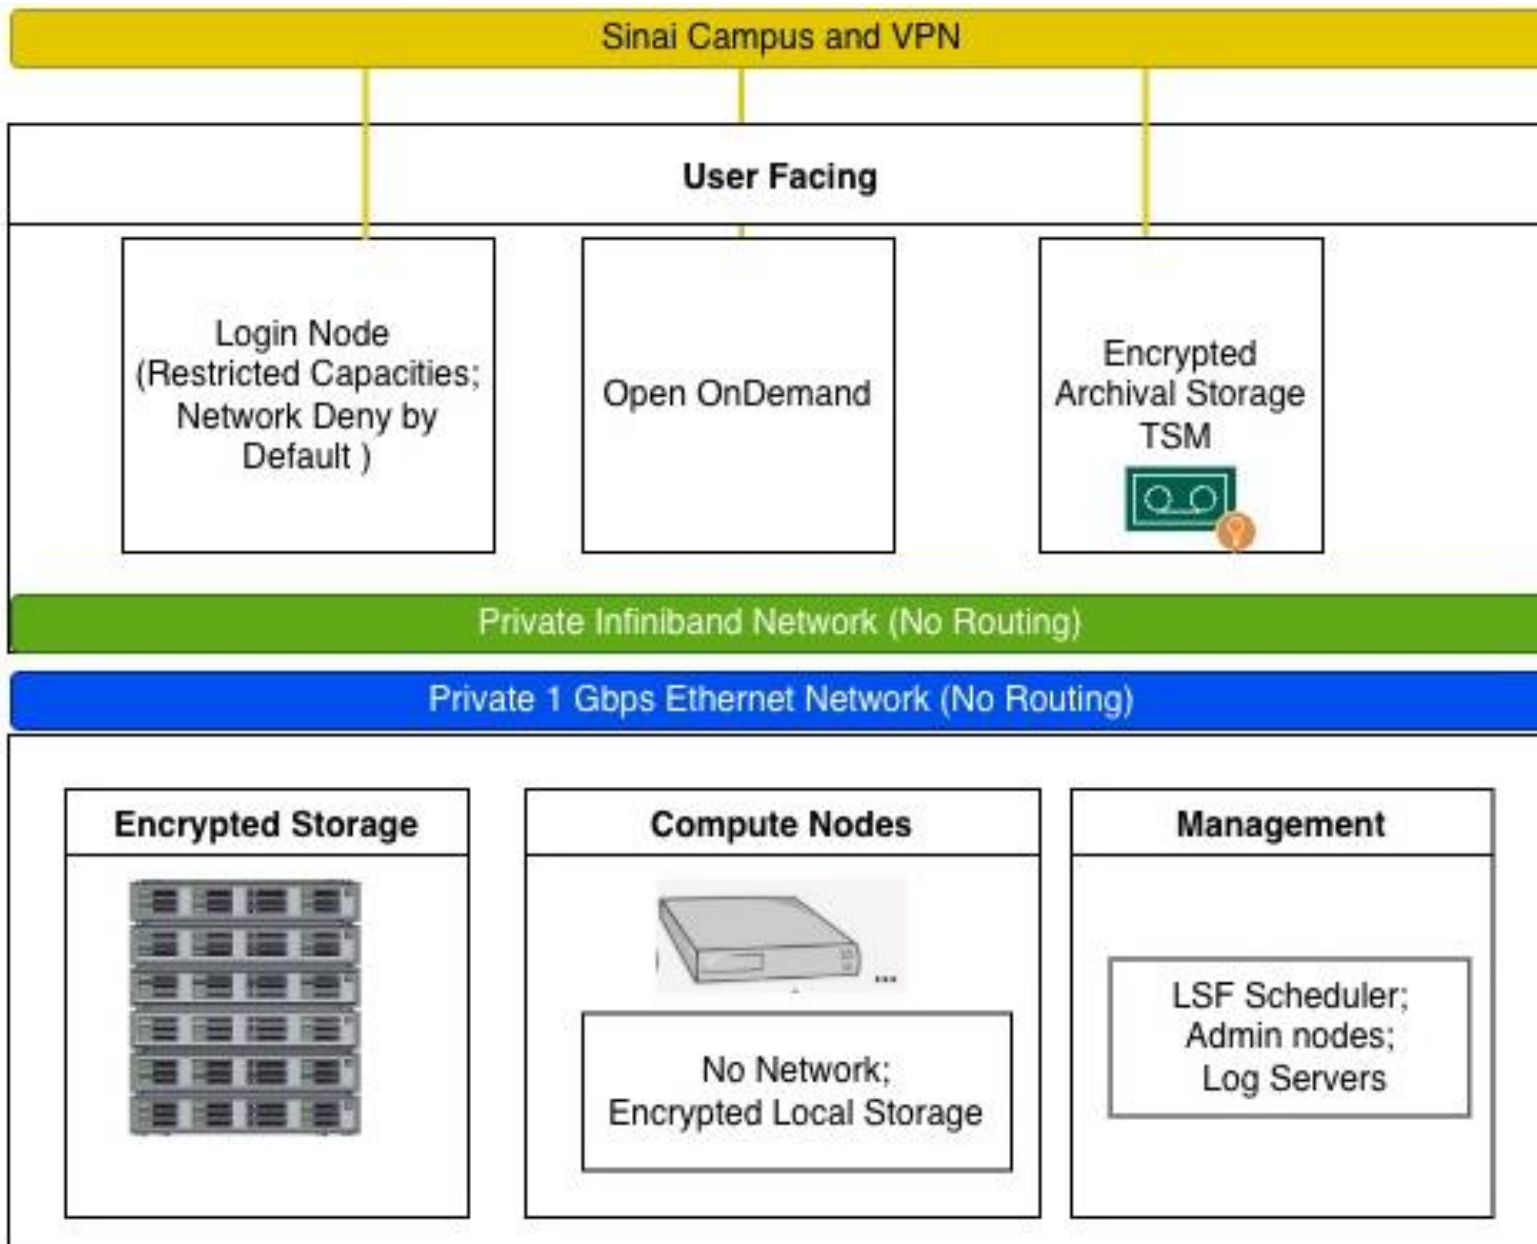
- Minerva-Res enforces both strict technical and policy controls to safeguard data aligning with NIST SP 800-171

Major Differences:

1. Data transfer is strictly restricted: Every data out must be approved by PI
2. No network access on compute and login nodes for your pipelines

Proposed Procedures for Non-Controlled Data Transfer:

1. Users request data transfer via submitting a request form
2. PIs will receive an email notification
3. Once PIs approved, data transfer access will be granted by HPC team



Security Controls

- MFA via RADIUS/NPS
- Centralized logging (Graylog/SIEM)
- Restricted network (deny-by-default)
- Ansible-based configuration management
- SELinux enforcing mode; FIPS mode
- Auditd enabled

Minerva Two-Factor Authentication Migrated to Azure MFA

- ▶ Minerva has switched to Microsoft's Azure multi-factor authentication (MFA) since **9:00 am Thursday, March 26, 2026**, to save cost on Symantec VIP
- ▶ Available authentication methods aligned with institutional MFA:
There are 3 methods enabled. You can choose either one of them:
 - Method 1: Phone Text passcode (Mount Sinai profile default),
 - Method 2: Microsoft Authenticator App – push notification
 - Method 3: Microsoft Authenticator App – one-time password code

By default, users who didn't setup Microsoft Authenticator will receive Phone Text SMS message codes (method 1).

Expanded Minerva with \$2M NIH S10 Award for an AI Mount Sinai (AIMS) Supercomputer: **in Production since Feb. 2026**

- ▶ Open to NIH-funded AI/GPU-related projects
- ▶ Submitted AIMS Final Research Performance Progress Report

2026/02/13 - 2026/05/05	
Num. of projects	19
Num. of users	40
Num. of jobs	2,839
Total % usage	40.2%

- ▶ Please acknowledge and associate the AIMS grant S10OD038231 in your publications. Thank you!

	AIMS
Performance (petaflops)	1.92 FP64 864 FP4 inference
Nodes	6 nodes
CPUs	672
# of GPUs	48
GPU Types	NVIDIA B200
System Memory (terabytes)	12
Storage	100 TB local NVME storage

Minerva System Migration, Patch and Monitoring

► Migration

- All service VMs were migrated to new Proxmox virtualization platform to save license cost
- Decommissioned end of life LTO5 tape library and LTO5 tapes; Secure offsite media destruction to reduce monthly cost on storage

► Security patching

- GPFs upgraded to version 5.3a, between February 9 and February 20
- OS kernel upgraded to Rocky 9.6 via rolling upgrade
- Emergency system-wide patch on critical CVEs on April 30 2026
- Implementing and testing weekly automatic patching on network-facing nodes
 - More Minerva downtime might be needed

► Monitoring jobs and cluster load actively via Grafana

- Direct user outreach and automated email notification to address misuse and resource waste

AIR.MS Data Sets

Radiology images

- All the DICOM data were copied to Minerva and prospective data synced daily
- 1.4 PB are transferred to Minerva
- **Next Step:** de-identify all DICOM data and build a research portal for user access
 - We will de-identify all images and make available to all
 - Radiology metadata will be available in SAP HANA in AIR.MS
 - Expected completion: CT and MRI for five years will be available in June

Kidney Precision Medicine Project (KPMP) open-access data

- Onboarded to AIR.MS data sets and open to researchers with DUA. Details [here](#).

Minerva Training Offered Spring 2026

Five training sessions with training material (including slides & recording) online.

Session	Topic	Spring 2026 Date	Spring 2026 # of Attendees
1	Minerva Intro	Feb. 12	39
2	Essential Services Demo on Minerva	Feb. 19	27
3	Load Sharing Facility (LSF) Job Scheduler	Feb. 26	12
4	Introduction to GPU/AI resources on Minerva	Mar. 5	16
5	Leveraging Large Language Models in Biomedical Research	Mar. 12	17
Total			111

Staff Summary

The HPC team consists of **four** computational scientists/bioinformaticians

- Hyung Min Cho, PhD
- Jielin Yu, PhD
- Yiyuan Liu, PhD (part time Data Ark)
- Shamimul Hasan, PhD
(part time Empire AI cluster)



and **five** HPC architects/admins

- Wei Guo, PhD
- Tejas Rao
- Eric Rosenberg (part time Minerva TSM)
- Sumit Saluja
- **Joshua Baker-LePain joined on May 11th 2026**



All staff will be working part-time on supporting the new Minerva-Res cluster



What's Next?

Roadmap 2026

June	July	Aug	Sep	Oct	Nov
Launch and Support NIST 800-171 Compliant Minerva-Res Cluster					
Re-Assess HIPAA Compliance and Improve as Needed					
Security Patch and Compliance Review					
Continue to Support Empire AI Cluster					
Enable Minerva Authentication against Hospital AD			Fall Training		
Install New Storage to Replace End of Life Hardware					



Appendix I: Minerva Usage (Oct. 2025 – March. 2026)


Minerva Usage Summary

Accounts	Apr. 1 2025 - Sep. 30, 2025	Oct. 1 2025 - Mar. 31, 2026
Number of Active Users	1,213	1,082
Number of Total Users	5,020	5,267
Number of Project Groups	684 (484 active)	640 (521 active)
Storage		
High-Speed Storage Used (Arion)	23.3 PiB	25.5 PiB
Archival Storage Used	22.3 PiB	24.4 PiB
Compute		
Number of Jobs Run	119,666,873	51,033,400
Core-Hours Utilized	88,717,822 hrs	86,715,556 hrs
System		
Number of Maintenance Sessions	1 Full Maintenance 99.6% Uptime	1 Full Maintenance 99.9% Uptime

Jobs and Compute Cores By Partition

	Apr. 1 2025 - Sep. 30, 2025			Oct.1 2025 - Mar. 31, 2026		
Compute	# Jobs	CPU-hours	Utilization	# Jobs	CPU-hours	Utilization
Chimera	82,884,244	55,781,384	90.5%	28,297,609	47,579,846	77%
Old Hi-Memory Nodes	9,669,206	6,324,480	100%	6,259,515	6,289,920	100%
CATS	24,483,739	11,665,586	78.7%	8,921,330	10,828,640	73%
GPU nodes	2,629,684	14,946,372	58.9%	7,554,946	22,017,150	84%
Total	119,666,873	88,717,822	84%	51,033,400	86,715,556	80%

Job Mix



Top 10 Users Compute Core Hours

PI	Department	# Core-Hours	# Jobs
Raj, Towfique	Genetics and Genomic Sciences	9,979,255	621,454
Roussos, Panagiotis	Psychiatry	6,786,202	5,755,251
Goate, Alison	Genetics and Genomic Sciences	5,036,988	1,545,790
Charney, Alexander	Genetics and Genomic Sciences	4,180,408	1,263,197
Faith, Jeremiah	Genetics and Genomic Sciences	3,423,097	147,438
Filizola, Marta	Structural and Chemical Biology	3,397,768	1,346,851
Buxbaum, Joseph	Psychiatry	3,316,793	12,414,460
Campanella, Gabriele	AI and Human Health	3,172,636	67,920
Nadkarni, Girish	Medicine	2,172,273	709,370
Pejaver, Vikas	Institute for Genomic Health	2,072,598	3,441,649

Top 10 PIs GPFS High Speed Storage

PI	Department	Storage Usage(TiB)
Roussos, Panagiotis	Psychiatry	1,297
Nadkarni, Girish	Medicine	1,203
Raj, Towfique	Genetics and Genomic Sciences	1,146
Charney, Alexander	Genetics and Genomic Sciences	1,104
Sealfon, Stuart	Neurology	769
Zhang, Bin	Genetics and Genomic Sciences	696
Goate, Alison	Genetics and Genomic Sciences	621
Sebra, Robert	Genetics and Genomic Sciences	591
Pejaver, Vikas	Institute for Genomic Health	417
Gelb, Bruce	Mindich Child Health and Development Institute	369

Top Compute and Storage Usage Department/Institute

Department/Institute	Compute Core Hours	Department/Institute	Storage (TiB)
Genetics and Genomic Sciences	27,904,341	Genetics and Genomic Sciences	5,859
Psychiatry	13,516,206	Psychiatry	2,269
Neuroscience	13,098,946	Medicine	1,733
AI and Human Health	5,274,033	Neuroscience	1,577
Medicine	4,752,040	Oncological Sciences	1,409
Structural and Chemical Biology	3,945,494	Neurology	892
Institute for Genomic Health	3,033,758	Institute for Genomic Health	622
Oncological Sciences	3,017,935	Microbiology	441
Pharmacological Sciences	1,993,969	Mindich Child Health and Development Institute	369
Health Science and Policy	1,279,425	Pathology	364

Top 10 PIs - GPU Usage Hours

PI	Department	GPU Hours	# Jobs
Filizola, Marta	Genetics and Genomic Sciences	122,707	34,886
Campanella, Gabriele	AI and Human Health	101,995	45,397
Nadkarni, Girish	Medicine	65,205	5,986
Raj, Towfique	Structural and Chemical Biology	64,047	89,073
Shi, Yi	Pharmacological Sciences	62,986	1,919,880
Shen, Li	Neuroscience	41,586	3,806
Beck, Erin	Neurology	28,627	49,018
Mei, Xueyan	BioMedical Engineering and Imaging Institute	26,094	28,067
Chowell, Diego	Precision Immunology Institute	18,934	6,513
Lin, Baihan	AI and Human Health	16,782	5,848

Total TSM archival storage usage

Current archive storage usage	
Archived data	24.4 PB
Total data with offsite copy	48.8 PB
Number of tapes used	3,318

Statistics of Oct 2025 - Mar 2026			
Amount of archived data	2,803 TB	Amount of retrieved data	593 TB
# of users who have issued archive commands	76	# of users who have issued retrieve operations	43

LTO5 to LTO9 Migration			
Amount of data migrated so far	14 PB	Est data remaining max	0 TB
		Est migration completion	Complete



Minerva HPCHelp Tickets (hpchelp@hpc.mssm.edu)

Jan - May 2026	
Number of tickets created	2,145
Average Response Time	0.80 days

of HPC Tickets opened is increasing largely over years.

Year	2026	2025	2024	2023	2022	2021
# of tickets	2145	5,228	5,040	3,754	3,240	2,915

Minerva Publications > 2,000 Since 2012

We collect publications twice a year (Jan & May). Thank you!!!

Please cite Minerva platform and the technologies adopted in all your publications using:

- *Kovatch P, Gai L, Cho HM, Fluder E, Jiang D. Optimizing High-Performance Computing Systems for Biomedical Workloads. IEEE Int Symp Parallel Distrib Process Workshops Phd Forum. 2020 May;2020:183-192. doi: 10.1109/ipdpsw50202.2020.00040. Epub 2020 Jul 28. PMID: 33088611; PMCID: PMC7575271*

All publications utilizing Minerva resources must include one of the following acknowledgements depending on your funding status:

- **If your NIH-funded projects only use the AIMS B200 GPUs, you must include the following in all your publications:** “*This work was supported in part through the Minerva computational and data resources and staff expertise provided by Scientific Computing and Data at the Icahn School of Medicine at Mount Sinai and supported by the Clinical and Translational Science Awards (CTSA) grant UL1TR004419 from the National Center for Advancing Translational Sciences. Research reported in this publication was also supported by the Office of Research Infrastructure of the National Institutes of Health under award number S10OD038231. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.*”

To associate the CTSA grant UL1TR004419, the CATS grant S10OD030463, or the AIMS grant S10OD038231 to an existing publication, please follow [these instructions](#) from the NIH (see the section “Associating Funding to your Publications”).

Year	# Pubs
2012	55
2013	59
2014	62
2015	115
2016	149
2017	165
2018	133
2019	178
2020	146
2021	234
2022	174
2023	239
2024	214
2025	162


Data Ark - Usage

Dataset	Size (GB)	# of Times Data Accessed	# of unique users
Digital Pathology Slides	2,220,236	23,620,295	5
gnomAD	8,627	1,867,427	2
Genebass	902	581,998	2
1000 Genome	245	9,103	9
MSDW OMOP	3,500	1,571	7
Reference Genome	142	306	5
eQTLGen	38	237	1
GTEX	1,887	104	3
GWAS summary statistics	6,825	42	1
Living Brain Project	< 1	12	1
UKBiobank LD	2,866	6	2

of unique active users: **31**

of support tickets (including tickets automatically generated from user DUA submissions): **98**

Appendix II: 2025 User Survey Results



2025 Minerva Survey Results

We asked 5 questions:

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).

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

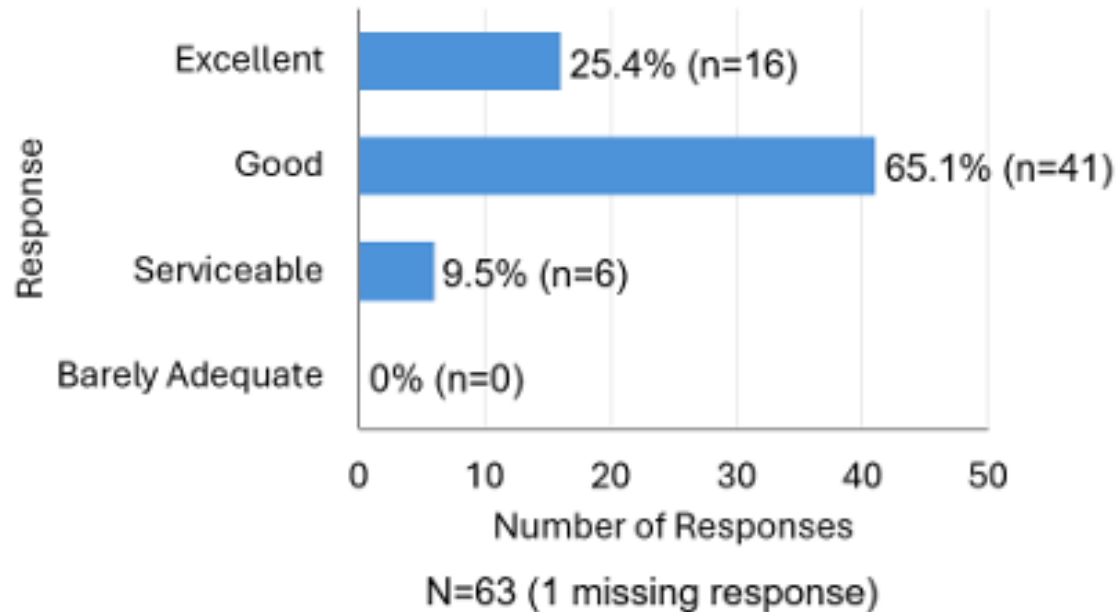
Q4: Is there any other new training session you would prefer adding to 2025 Minerva HPC training serials?

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

We received 64 responses and 31 comments from 1,200 active users in Jan 2026 (5.0% response rate).

2025 Survey Results Question 1

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



User satisfaction(\geq Good)

2025: 91%

2024: 83%

2023: 81%

2022: 84%

2021: 85%

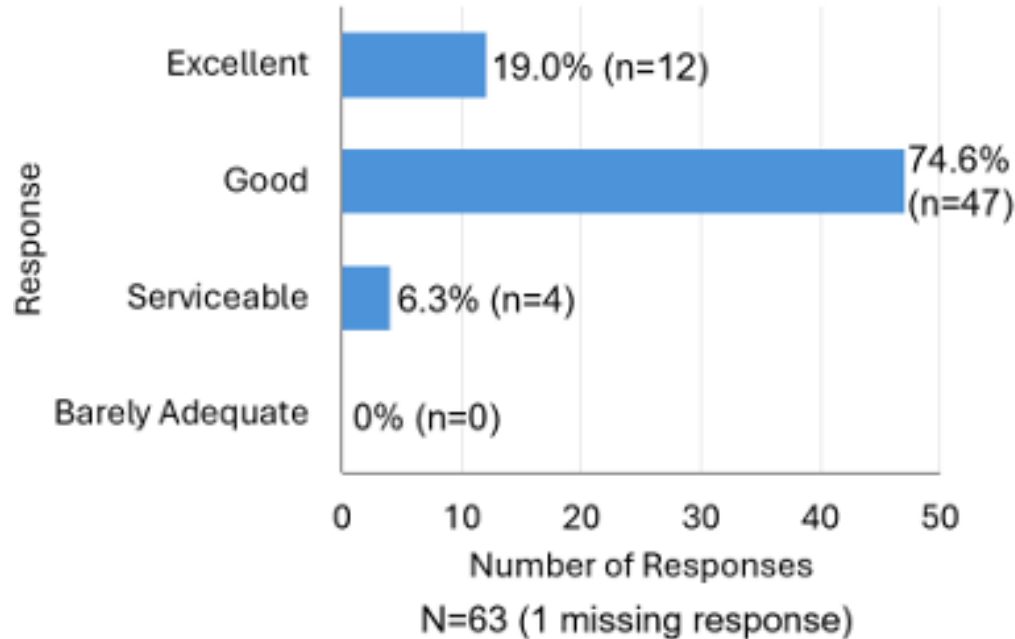
2020: 81%

2019: 65%

2018: 54%

2025 Survey Results Question 2

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



User satisfaction(\geq Good)

2025: 94%

2024: 77%

2023: 85%

2022: 93%

2021: 81%


2020: 80%

2019: 80%

2018: 67%

2025 Survey Results Question 3

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



User satisfaction(>=Good)

2025: 81%

2024: 75%

2023: 82%

2022: 88%

2021: 86%

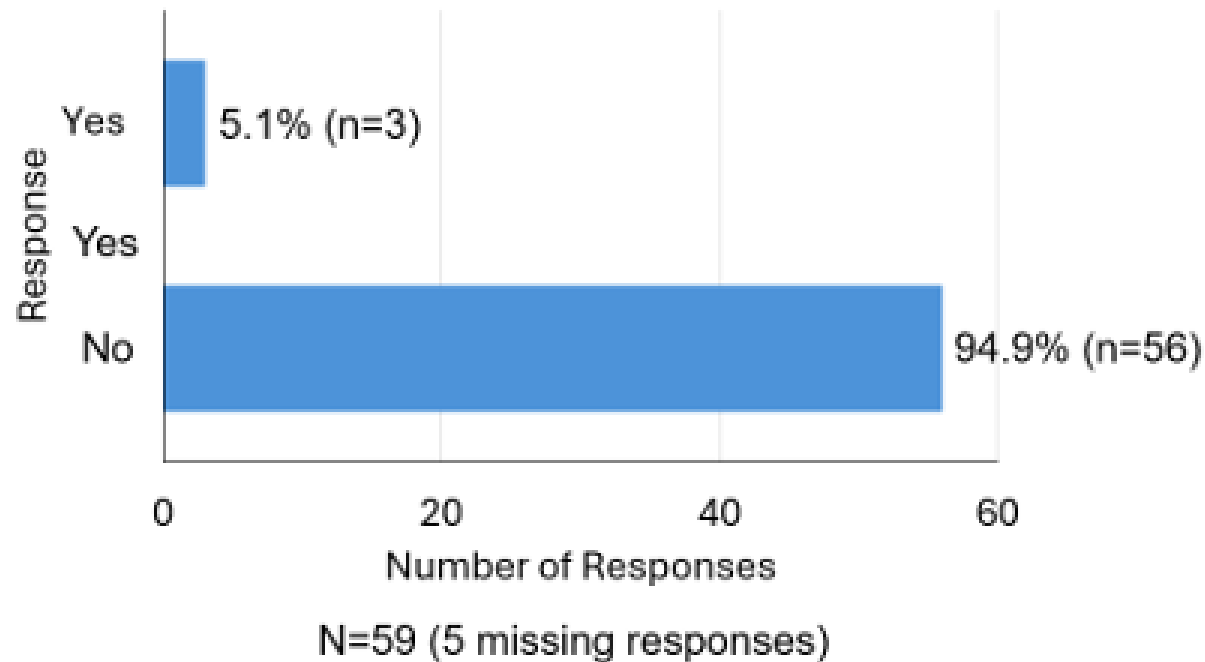
2020: 91%

2019: 73%

2018: 80%

2025 Survey Results Question 4

Q4: Is there any other new training session you would prefer adding to 2025 Minerva HPC training serials?



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.*
- GPU/AI use for sequencing data.
 - Fine tuning large language models on Minerva.
 - Building web interfaces that can access data securely from the Minerva file system.

2025 Data Ark Survey Results

We asked 4 questions:

Q1: How would you rate your satisfaction with the Data Ark: Data Commons quality and availability at Mount Sinai?

Q2: How would you rate your satisfaction with Data Ark support?


Q3: What barriers exist preventing your usage of Data Ark?

Q4: Please enter any other feedback on Mount Sinai's Data Ark: Data Commons including suggestions for data sets, availability, accessibility, support, new features and development.

We received 20 responses and 2 comments from 74 active users in Jan 2026 (27% response rate).

2025 Survey Results Question 1


Q1: How would you rate your satisfaction with the Data Ark: Data Commons quality and availability at Mount Sinai?



User satisfaction(\geq Satisfied)
2025: 80%
2024: 78%
2023: 80%
2022: 72%

2025 Survey Results Question 2

Q2: How would you rate your satisfaction with Data Ark support?



User satisfaction(\geq Satisfied)
2025: 82%
2024: 78%
2023: 84%
2022: 72%

2025 Survey Results Question 3

Q3: What barriers exist preventing your usage of Data Ark?

Comment

- I find there are not clear enough "get started" guides. Much of what's online is very rudimentary and not user-friendly (instead of screenshots and plain text on a webpage, maybe consider making polished powerpoints or similar). I think this is a major barrier for uptake across the health system broadly and is one of the things that has made getting started difficult for me.
- Getting access to the restricted part is cumbersome. There are no real guidelines when requesting large scale access, for big data computational biology studies.

Acknowledgements

- Supported by the Clinical and Translational Science Awards (CTSA) grant UL1TR004419 from the National Center for Advancing Translational Sciences (NCATS), National Institutes of Health (NIH).

CTSA Clinical & Translational[®]
Science Awards



Thank You