Icahn School of Medicine at Mount Sinai

Minerva HPC and Data Ark Town Hall

Partnering with scientists to accelerate scientific discovery

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

The Minerva HPC team

Nov. 21, 2025

Agenda

1. Accomplishments and Updates
2. 2026 Roadmap
3. Appendix I – Minerva Usage

Accomplishments and Updates

Accomplishments and Updates

- Updates on NIST SP 800-171 Compliant Environment for Controlled Data
- Updates on Empire AI Cluster
- Minerva Expansion with $2M NIH Award for AI Mount Sinai (AIMS) Supercomputer
- Secure Internal Chatbot for Clinicians and Researchers
- Documentation and FAQs Updated On Minerva Website
- System Updates and Migration on NFS Home Directory, Proxy Server, LTO-5 Migration etc
- Updates on Data Ark Data Sets
- Reminder on Minerva Data Backup Policy
- Minerva Training Offered Fall 2025
- Staff Updates

Summary on NIST SP 800-171 Compliant Environment

► Executive Summary

- **Mount Sinai investigators can attest compliance** with NIH Security Best Practices for Users of Controlled-Access Data according to NIST800-171 for their project submission **while using Minerva**.
- Minerva HPC team is actively building the NIST 800-171 compliant environment **targeting June 2026.** – named as CUI (controlled unclassified information) cluster.
- The CUI cluster will be the only approved environment for NIH genomic data sets at Mount Sinai.

► Details: Why can we attest compliance now?

NIH confirmed via email that **an institution & their investigator can attest** they are meeting the security controls with **a self-assessment and plan** to address the gap.

Minerva developed **self-assessment** System Security Plan (SSP) and **Plans** of Action and Milestones (POAM). DTP security **approved** on 05/16/2025.

Our Plan and Actions to Build CUI Cluster by June 2026

- Purchased 10 petabytes raw of new encrypted storage
  - Hold all **raw and derived** data containing Controlled Unclassified Information (CUI)
  - Deliver expected in Dec. 2025 and in production by Feb, 2026

- We are creating a new locked down computational and data environment
  - Access restricted only to users with CUI or HIPAA/PHI data
  - There will be a dedicated login node and computational nodes with Linux hardening attached
  - These service will **not** be available:
    - X Open OnDemand; Database server; Web server; Posit Connect; Globus

- Data Movement
  - ✓ Data in: Allowed
  - ✓ Data out (Non-CUI): Allowed if the product doesn't contain CUI with the PI approval
  - X Data out (CUI): Policy requires users not to move CUI data out by default. Any requests to move **CUI data** will need the PI approval and the destination compliant to NIST SP 800-171

Architecture for the New Locked Down Environment

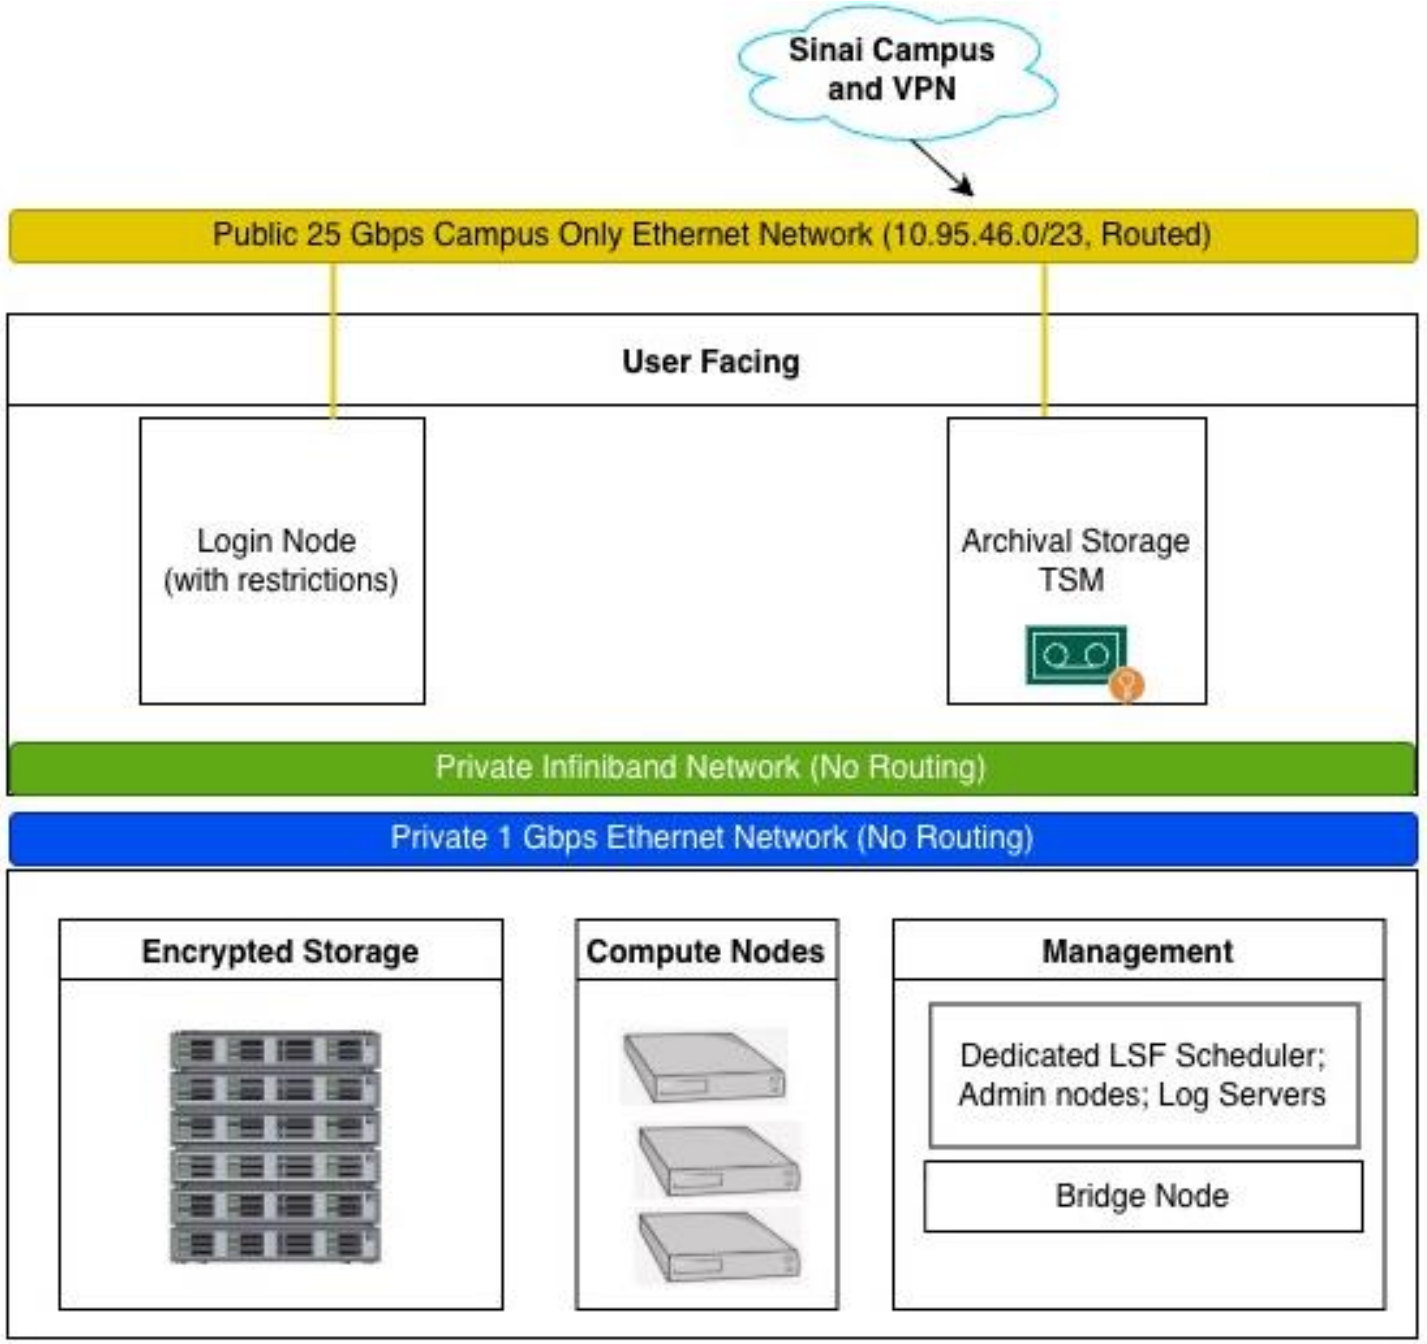

Good News! Mount Sinai Joined EMPIRE AI consortium

- **Thanks to Eric Nestler, Lisa Stump, Girish N. Nadkarni & Patricia Kovatch, we joined Empire AI,** enabling access to extra large-scale GPU resources and collaborations within New York State!

- Empire AI is a consortium of **ten** New York State institutions with support from New York State and private philanthropy, that oversees a shared computing facility to promote responsible research & development

Current Activities and Status			
Empire AI TownHall session with 94 attendees on 10/10/2025	Documentation and FAQ are available here	**6** PI groups onboarded to Alpha system with **19** users accounts	**4** project requests received for Beta allocation **- 30%** of our share (delayed production to Feb. 2026)

Minerva Expansion with $2M NIH S10 award for AI Mount Sinai (AIMS) Supercomputer

Timeline

- Awarded in late May
- The equipment will be delivered by Dec. 2025
- Aiming to get this into production in Feb. 2026

Operation Plan

- Will open to NIH-funded AI/GPU-related projects
- Will need NIH funding # and publications for report purpose as always. Thank you!

AIMS total system configuration	
Peak performance (petaFLOPS) [NVIDIA]	1.92 FP64 3.84 FP32 216 FP16 432 FP8 864 FP4
# of GPUs	48 NVIDIA B200
GPU memory size (TB) and type	9 HBM3e
GPU memory bandwidth (TB/s)	384
NVLink bandwidth (TB/s)	87
# of CPU cores	672
System memory (TB)	12
# of nodes & type	6 Lenovo ThinkSystem SR780a V3 8-way GPU
Node interconnect	1x IB NDR400 per GPU
Node interconnect Bandwidth (Tb/s)	19.2

Minerva supercomputer configuration in total:
13 petaflops of peak 64-bit compute power
25k CPUs + >408 NVIDIA GPUs
450 TB memory
50 PB raw parallel file system

Secure Internal Chatbot for Clinicians and Researchers

Goal

- Develop a **cost-effective and secure internal chatbot** with selectable large language models (LLMs) on Minerva, enabling use of AI models and institutional data for research and clinical innovation safely

Status

- Open for friendly user testing
- Hosted within a HIPAA-approved secure enclave, validated by Mount Sinai's Cybersecurity team

Joint Effort by AIR.MS and Minerva

- **Frontend:** Powered by Open WebUI, managed by AIR.MS team
- **Backend:** All LLM and Ollama runtime operated on Minerva H100 GPUs, ensuring high performance
- This design allows computationally expensive LLM tasks to run on Minerva, while users interact seamless via web browser or API

ATTENTION: NO DEEPSEEK MODEL. NO NON-US MODELS.

Documentation and FAQs Updated On Minerva Website

- We reviewed and updated all the documentation on HPC website this Fall
- To make FAQs more useful and current, we developed an AI-powered script using a Large Language Model (Llama 3.3 latest) to analyze over 5,000 user tickets (Apr 2024 – Apr 2025)
  - **Knowledge Extraction:** It identifies the core question asked and the solution provided by the HPC team
  - **FAQ Identification:** The LLM analyzes all interactions to surface the most common and impactful user questions
  - **Outcome:** Based on this analysis, the AI-powered script identified **the top 30 common questions** and we updated the FAQs on our website

Service Upgraded and Migrated to New Servers with Rocky OS 9.5

Improve stability on Minerva NFS home directory

- NFS home directory server OS updated to licensed RedHat with support
- NFS server controller firmware update (to prevent another unexpected outage)
- NFS kernel patch to fix a NFS hang issue

Improve security of the Minerva mail and proxy server

- Migrated to new server with Rocky9 new hardware for improved robustness and security

Improve security of the TSM key manager server

- Migration of key manager to new hardware, and software updates for security and support

Almost finished with the migration from the old LTO5 tape system to the new LTO9

- Only 12TB left on LTO5 tapes! (Two-year long project)

Jira service desk to enable status tracking and history tickets etc (on hold)

- Central Jira service doesn’t support AD authentication for “portal-only” customers who are not licensed users, causing security concern

Data Ark Data Sets

Radiology images

- Work with a new vendor to transfer all the DICOM data to Minerva via direct copy
- 1.0 PB out of 1.4 PB are transferred to Minerva (~ 50 days)
- **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: June 2026

Kidney Precision Medicine Project (KPMP) open-access data

- Onboarding to Data Ark; Pending on Data Use Agreement approval.

Reminder on Minerva Data Backup Policy

We are sending emails to reminder users about data backup policy on Minerva

This is a routine reminder on Minerva data backup policy.

1. **We do not backup any user files on any nodes including the private nodes. Please archive/backup your important files by yourselves.**
    a. We have included this in the message of the day (MOTD) after you log into Minerva, User Responsibilities and Acceptable Usage Policy, and Annual HIPAA Policy.
    b. We will keep reminding you every quarter via email.
    c. Please archive/backup your files following the guide at https://labs.icahn.mssm.edu/minervalab/documentation/access-tsm-with-command-line/

2. **Please don't set the permission of your Minerva files as rwx (read, write and execute) for everyone/others.**
    a. This can result in file deletion by others. Please double check your file permission on Minerva especially for your project directory.
    b. Limiting file permissions is the user's responsibility according to the annual HIPAA compliance requirement/training.

Minerva Training Offered Fall 2025

Nine training sessions with training material (including slides & recording) online

Session	Topic	Spring 2025 # of Attendees	Fall 2025 # of Attendees
1	Minerva Intro	71	49
2	Essential Services Demo on Minerva	-	19
3	Load Sharing Facility (LSF) Job Scheduler	47	20
4	Introduction to GPU/AI resources on Minerva	50	14
5	Accelerating Biomedical Data Science with GPUs: Practical Approaches And Tools	35	11
6	Leveraging Large Language Models in Biomedical Research	84	24
7	Access Minerva via web browser Open OnDemand	30	16
8	How to Accelerate Genome Analysis Toolkit (GATK) by using Parabricks	24	9
9	Introduction to Data Ark	23	3
10	Empire AI Town Hall	-	94
	Total	**364**	**259**

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




and **four** HPC architects/admins

- Wei Guo, PhD
- Tejas Rao
- Eric Rosenberg (part time Minerva TSM)
- Sumit Saluja




Three open positions (two pending approval)

- Sr. HPC Admin
- *Lead HPC Architect – cybersecurity*
- *Sr. System Admin (part time on Minerva)*

What’s Next?

Roadmap 2026 (Q1-Q2)

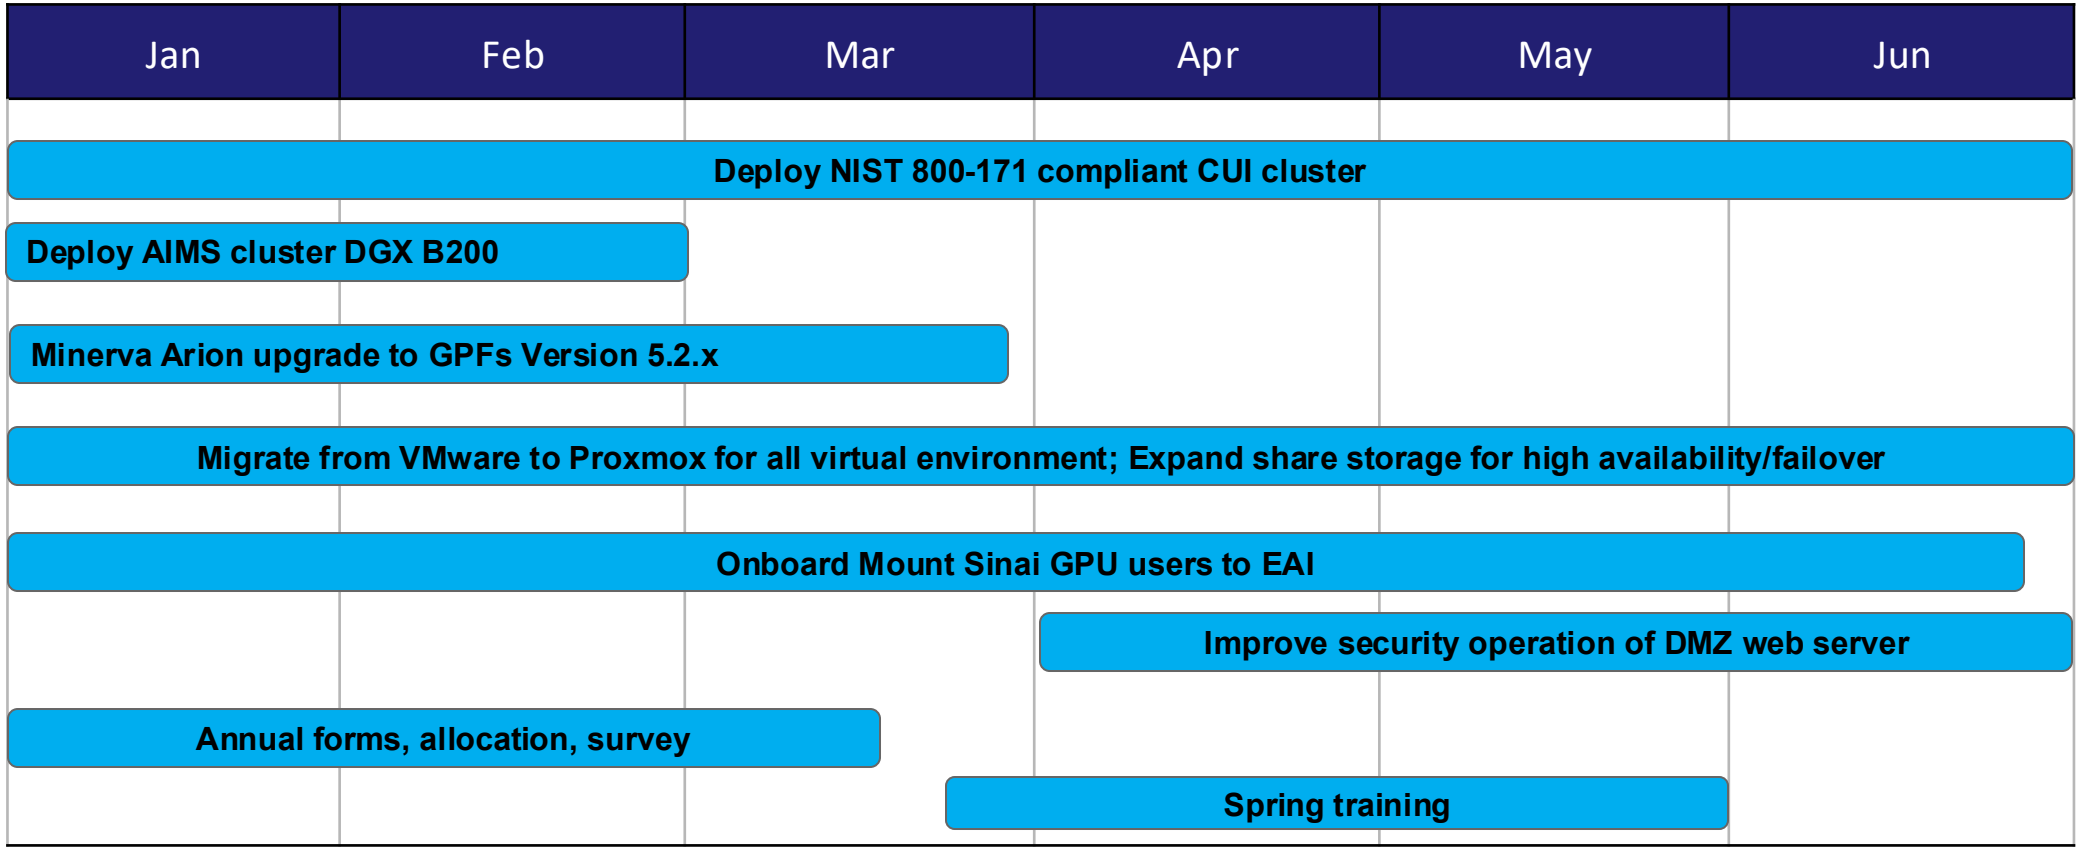

Appendix I: Minerva Usage (April 1 2025 – Sep. 30 2025)

Minerva Usage Summary (April 1 2025- Sep. 30 2025)

Accounts	Oct.1 2024- Mar. 31, 2025	April.1 2025- Sep. 30, 2025
Number of Active Users	977	1,213
Number of Total Users	4,675	5,020
Number of Project Groups	617 (445 active)	684 (484 active)
Storage		
High-Speed Storage Used (Arion)	21.6 PiB (68% utilization)	23.3 PiB (72.8%)
Archival Storage Used	20.5 PB	22.3 PB
Compute		
Number of Jobs Run	52,719,740	119,666,873
Core-Hours Utilized	78,230,911 hrs*	89,562,728 hrs
System		
Number of Maintenance Sessions	1 unplanned outage and 2 Full Maintenance 98% Uptime	1 Full Maintenance 99.6% Uptime

* Didn’t count all the retiring node during the transition for a month

Jobs and Compute Cores By Partition

	Oct.1 2024- Mar. 31, 2025			April.1 2025- Sep. 30, 2025		
Compute	# Jobs	CPU-hours	Utilization	# Jobs	CPU-hours	Utilization
Chimera	38,493,981	49,664,845	82.4%	82,884,244	55,781,384	90.5%
Old Hi-Memory Nodes	4,553,928	6,267,942	100%	9,669,206	7,169,386	100%
CATS	7,916,931	11,706,802	80.8%	24483739	11,665,586	78.7%
GPU nodes	1,580,180	9,342,661	45.8%	2629684	14,946,372	58.9%*
BODE2 (Retired in Stages)	174,720	1,248,661	-	-	-	-
Total	52,719,740	78,230,911	76.0%	119,666,873	89,562,728	84%

* This is for all GPU models. If breakdown by GPU models, L40S and V100 usage is well below average, while the new H100 is about 70%.

Job Mix

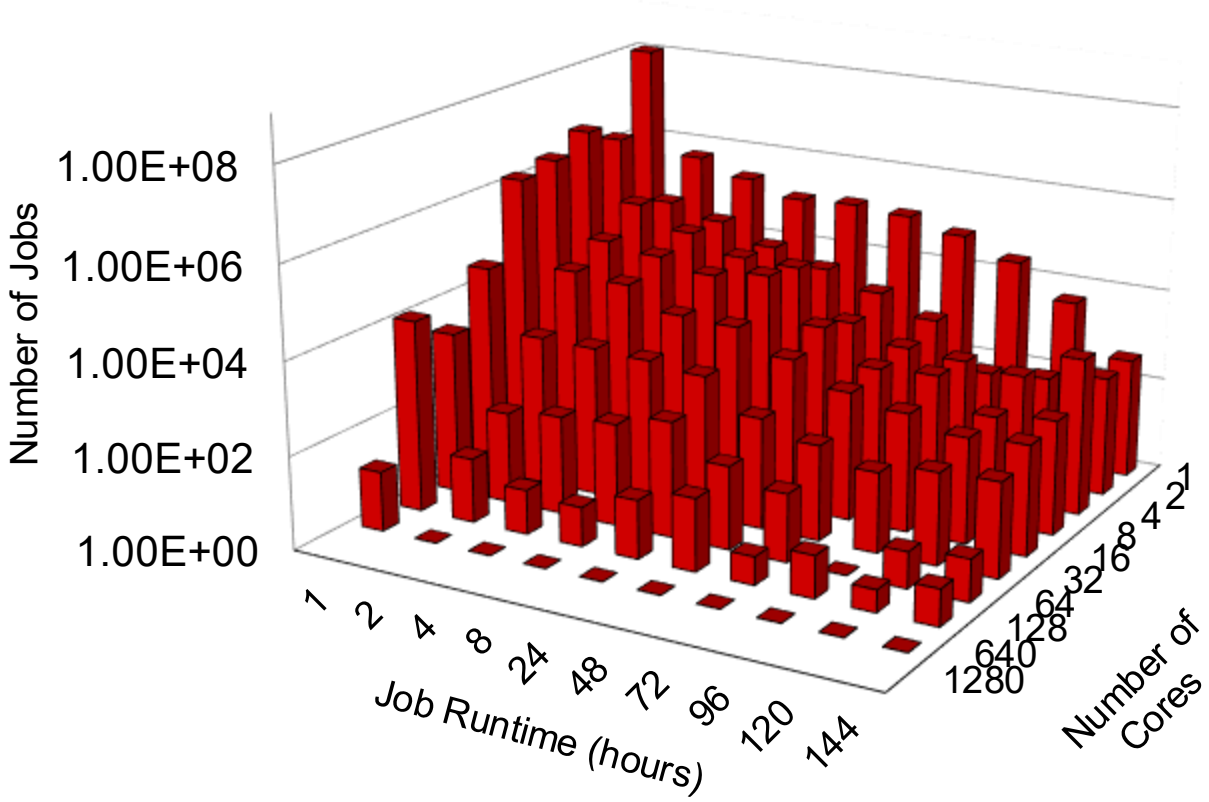
Number of Jobs
1.00E+08
1.00E+06
1.00E+04
1.00E+02
1.00E+00
1
2
4
8
24
48
72
96
120
144
Job Runtime (hours)
1
2
4
8
16
32
64
128
640
1280
Number of Cores

Top 10 Users Compute Core Hours

PI	Department	# Core-Hours	# Jobs
Zhang, Bin	Genetics and Genomic Sciences	9,465,878	563,493
Raj, Towfique	Genetics and Genomic Sciences	6,544,962	3,140,963
Pejaver, Vikas	Institute for Genomic Health	5,510,633	1,496,165
Roussos, Panagiotis	Psychiatry	4,793,963	4,815,743
Buxbaum, Joseph	Psychiatry	4,668,463	81,973,771
Reva, Boris	Genetics and Genomic Sciences	3,839,010	1,461,570
Goate, Alison	Genetics and Genomic Sciences	3,457,513	600,156
Clemente, Jose	Genetics and Genomic Sciences	3,186,115	3,759,100
Nadkarni, Girish	Medicine	2,652,871	904,902
Kenny, Eimear	Institute for Genomic Health	2,481,218	12,046,274

Top 10 PIs GPFS High Speed Storage

PI	Department	Storage Usage(TiB)
Nadkarni, Girish	Medicine	1,320
Roussos, Panagiotis	Psychiatry	1,313
Raj, Towfique	Genetics and Genomic Sciences	1,257
Sebra, Robert	Genetics and Genomic Sciences	1,041
Charney, Alexander	Genetics and Genomic Sciences	1,029
Sealfon, Stuart	Neurology	824
Zhang, Bin	Genetics and Genomic Sciences	753
Goate, Alison	Genetics and Genomic Sciences	676
Hasson, Dan	Oncological Sciences	435
Pejaver, Vikas	Institute for Genomic Health	409

Top Compute and Storage Usage Department/Institute

Department/Institute	Compute Core Hours
Genetics and Genomic Sciences	37,057,815
Psychiatry	12,361,280
Neuroscience	84,504,501
Institute for Genomic Health	8,242,803
Medicine	5,374,697
Oncological Sciences	3,701,238
AI and Human Health	3,483,593
Structural and Chemical Biology	2,947,907
Pharmacology	1,949,430
Neurology	1,743,003

Department/Institute	Storage (TiB)
Genetics and Genomic Sciences	6,301
Psychiatry	2,336
Medicine	1,962
Neuroscience	1,643
Oncological Sciences	1,392
Neurology	930
Institute for Genomic Health	771
Microbiology	470
Pathology	249
Mindich Child Health and Development Institute	248

Top 10 PIs - GPU Usage Hours

PI	Department	GPU Hours	# Jobs
Campanella, Gabriele	AI and Human Health	97,769	420,788
Nadkarni, Girish	Medicine	79,769	7,269
Raj, Towfique	Structural and Chemical Biology	59,466	559,381
Filizola, Marta	Genetics and Genomic Sciences	57,216	26,184
Shen, Li	Neuroscience	41,955	3,464
Shi, Yi	Pharmacological Sciences	30,193	471,114
Roussos, Panagiotis	Psychiatry	20,892	45,239
Osman, Roman	Structural and Chemical Biology	19,373	2,125
Tsankov, Alexander	Genetics and Genomic Sciences	16,427	6,437
Peter, Inga	Genetics and Genomic Sciences	16,288	29,124

Total TSM archival storage usage (Apr 2025- Sep 2025)

Current archive storage usage	
Archived data	22.3 PB
Total data with offsite copy	44.6 PB
Number of tapes used	3,022

Statistics of Oct 2024 - Mar 2025			
Amount of archived data	2,128 TB	Amount of retrieved data	798 TB
# of users who have issued archive commands	67	# of users who have issued retrieve operations	42

LTO5 to LTO9 Migration			
Amount of data migrated so far	13.9 PB	Est data remaining max	12 TB
		Est migration completion	Q4 2025

Minerva Help desk tickets (hpchelp@hpc.mssm.edu)

Jan - Oct 2025	
Number of tickets created	5,228
Average Response Time	0.8 days

# of HPC Tickets opened is increasing largely over years.

Year	2025	2024	2023	2022	2021
# of tickets	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 you *have NIH-funded projects* on Minerva, 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 S10OD026880 and S10OD030463. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.*"

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	**80**

To associate the CTSA grant UL1TR004419 or the CATS grant *S10OD030463* to an existing publication, please follow these instructions from the NIH (see the section "Associating Funding to your Publications").

Data Ark - Usage

Dataset	Size (GB)	# of Times Data Accessed	# of unique users
UK Biobank LD	2,866	31,346	7
1000 Genome	143	5,285	14
GWAS Summary Statistics	6,826	4,680	4
CBIPM-BioMe	4,608	3,828	12
gnomAD	8,628	2,223	5
GTEX	1,888	1,976	7
Reference Genome	142	838	6
TCGA	154	121	2
LD Score Regression	173	99	4
eQTLGen	39	2	1
Genebass	903	2	1

# of unique active users: **37**
# of support tickets (including tickets automatically generated from user DUA submissions): **101**

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).

Thank You

Icahn School of Medicine at Mount Sinai