

About Machines

Last updated: May 2026

AI Cluster - Slurm

Please send in a ticket requesting to be added if it is your first time using the AI cluster. You will need to be involved in research with a CS faculty member.

Feedback is requested. Find us in the Slack ai-cluster channel (channel ID: C02KW3M0BDK).

Infrastructure

Summary of nodes installed on the cluster:

AI Cluster Specs

CPU Cores: 2960 System Mem: 34389 GB GPU Memory: 8032 GB GPUs:

92	A40	48GB
52	L40S	48GB
14	H100	80GB

Storage: 483 TB

Computer/GPU Nodes

We like the alphabet, so we have compute node groups for just about every letter in it.

- "a" series: 3 nodes, each with 64 CPU threads, 192GB RAM, four RTX2080ti GPUs
- "aa" series: 2 nodes, each with 32 CPU threads, 32GB RAM, four RTX2080 GPUs
- "b", "d", "e", "k", "r" series: 15 nodes, each with 64 CPU threads, 512GB RAM, four A40's
- "c" series: 1 node with 48 CPU threads, 64GB RAM, two A30's
- "f" & "j" series: 6 nodes, each with 32 CPU threads, 128GB RAM, four A40's
- "g" & "q" series: 4 nodes, each with 96 CPU threads, 1TB RAM, eight L40S GPUs
- "h" series: 1 node with 96 CPU threads, 1TB RAM, four H100 SXM GPUs
- "l" series: 1 node with 256 CPU threads, 1.5TB RAM, six H100 PCI GPUs
- "m" series: 3 nodes with 128 CPU threads, 1.5TB RAM, no GPU's
- "n" series: 1 node with 96 CPU threads, 1.5TB RAM, four H100 SXM GPUs
- "t" series: 5 nodes with 48 CPU threads, 512GB RAM, four L40S GPUs

- all compute nodes:
 - Each node has a /local space for times when it's beneficial to not write over NFS. Space in /local varies from node to node. Please try to clean up when you're done.
 - Home directories and project space are mounted over NFS. Default quota for home directories is 50GB, but it may be increased as needed with permission.
 - Research groups may additionally be allocated project space that exists outside the home directory quota on different storage servers, for collaboration and shared storage.

Storage

- ai-storage1:
 - 63T total storage
 - uplink to cluster network: 100G
 - /home/<username>
 - 50G quota per user.
- ai-storage2:
 - 63T total storage
 - uplink to cluster network: 100G
 - /net/scratch: Create yourself a directory /net/scratch/\$USER. Use it for whatever you want.
 - Eventually data will be auto deleted after X amount of time. Maybe 90 days or whatever we determine makes sense.
- ai-storage3:
 - zfs mirror with previous snapshots of ai-storage1 and ai-storage4.
- ai-storage4:
 - 70TB total storage
 - uplink to cluster network: 100G
 - /net/projects:
 - Idea would be to create a dataset with a quota for people in a collaboration group to use.
 - Normal LDAP groups that you are used to and available everywhere else would control access to these directories. e.g. jonaslab, sandlab
- peanut-storage1:
 - 273TB total storage
 - uplink to cluster network: 100G fiber
 - /net/bulk:
 - A nice place for large datasets that either don't change much, or are being used and re-used a lot.
 - /net/archive:
 - A place to keep data from projects2 that isn't actively being worked on.
- peanut-storage2:
 - 546TB total storage
 - uplink to cluster network: 100G fiber
 - backups from peanut-storage3 and peanut-storage4
- peanut-storage3:
 - 224TB total storage
 - uplink to cluster network: 100G fiber

- /net/projects2:
 - Even more project space for your projects that you can put your projects in.
- peanut-storage4:
 - 28TB total storage
 - uplink to cluster 100G fiber
 - instructional storage for Peanut Cluster (NOT research)

From:

<https://howto.cs.uchicago.edu/> - **How do I?**

Permanent link:

<https://howto.cs.uchicago.edu/aicluster:machines>

Last update: **2026/05/04 09:50**

