Allocation and Limits

Storage Allocations

See our Storage page for:

- **More detailed information on storage allocations and policies**
- **Information on how to use xdisk**
- **Information on using UArizona Google Drive with HPC**

When you obtain a new HPC account, you will be provided with storage. The shared storage (/home, /groups, /xdisk) is accessible from any of the three production clusters: Puma, Ocelote and ElGato. The temporary (/tmp) space is unique to each compute node.

Location	Allocation	Usage
Permanent Storage		
/home/uxx/netid	50 GB	Individual allocations specific to each user.
/groups/PI	500 GB	Allocated as a communal space to each PI and their group members.
Temporary Storage		
/xdisk/PI	Up to 20 TB	Requested at the PI level. Available for up to 150 days with one 150 day extension possible for a total of 300 days.
/tmp	~1400GB NVMe (Puma) ~840GB spinning (Ocelote) ~840GB spinning (El Gato)	Local storage specific to each compute node. Usable as a scratch space for compute jobs. Not accessible once jobs end.

Job Allocations

All University of Arizona Principal Investigators (PIs; aka Faculty) that register for access to the UA High Performance Computing (HPC) receive these free allocations on the HPC machines which is shared among all members of their team. Currently all PIs receive:

HPC Machine	Standard Allocation Time per Month per PI	Windfall
Puma	100,000 CPU Hours per month	Unlimited but can be pre-empted
Ocelote	35,000 CPU Hours per month	Unlimited but can be pre-empted
El Gato	7,000 CPU Hours per month	Unlimited but can be pre-empted

Best practices

1. Use your standard allocation first! The standard allocation is guaranteed time on the HPC. It refreshes monthly and does not accrue (if a month's allocation isn't used it is lost).
2. Use the windfall queue when your standard allocation is exhausted. Windfall provides unlimited CPU-hours, but jobs in this queue can be stopped and restarted (pre-empted) by standard jobs.
3. If your group consistently needs more time than the free allocations, consider the HPC buy-in program.
4. Last resort for tight deadlines: PIs can request a special project allocation once per year (https://portal.hpc.arizona.edu/portal/; under the Support tab). Requesting a special project will provide qualified hours which are effectively the same as standard hours.
5. For several reasons we do not offer checkpointing. It may be desirable to have this capability in your code.

How to Find Your Remaining Allocation

To view your remaining allocation, use the command **va** in a terminal.
You can use this time on either the standard nodes which do not require special attributes in the scheduler script, or on the GPU nodes which do require special attributes.

Slurm and PBS Batch Queues

The batch queues on the different systems have the following memory, time and core limits.

Queue	Description
standard	Used to consume the monthly allocation of hours provided to each group
windfall	Used when standard is depleted but subject to preemption
high_priority	Used by 'buy-in' users for purchased nodes
qualified	Used by groups who have a temporary special project allocation

Job Limits

To check group, user, and job limitations on resource usage, use the command **`job-limits $YOUR_GROUP`** in the terminal.