

# AI Cluster Policy Description

2. Also, soon, as we pass this feature testing, we soon need to talk about the "policy testing" (applying some weights, reasoning about how \$10K annual subscription converts to the SLURM weights).

Bringing this thread back. I apologize for the wall of text but it is necessary to for you to understand this before you can approve it.

After various iterations we (Bob, Har, and I) believe to have an implementation of the policy that meets the requirements discussed previously.

The relevant spreadsheet:

[AI Cluster Tracking and Priority calculation spreadsheet](#)

I'll be adding all in this chain as an editor to the sheet after I send this email.

Details:

By default the cluster uses a fair share algorithm to adjust job priorities.

More reading if you want, but not necessary to continue reading this

email.

[https://slurm.schedmd.com/fair\\_tree.html](https://slurm.schedmd.com/fair_tree.html)  
[https://slurm.schedmd.com/priority\\_multifactor.html](https://slurm.schedmd.com/priority_multifactor.html)

We adjust those priorities with partitions for those who have donated either monetarily or with hardware.

Here is a simplified version of the partition configuration as it stands now.

```
PartitionName=general Nodes=a[001-008]  
PartitionName=cdac-own Nodes=a[005-008] AllowGroups=cdac Priority=100  
PartitionName=cdac-contrib Nodes=a[001-008] AllowGroups=cdac Priority=5
```

general: For all users  $\${group}$ -own: Machines  $\$group$  has donated  $\${group}$ -contrib: A method to give slightly higher job priority to groups who have donated but do not own machines.

The key thing to notice before you continue reading is that nodes can be added to multiple partitions.

'general' and 'cdac-contrib' can submit to all nodes but with different priorities.

Priority is normalized in the sheet to be 0-100.

Understanding the partition configuration: 1. All users get access to partition 'general'.

It has a default priority of 0.

2. Group 'cdac' has donated nodes 'a005-a008'. They always get top priority on those machines (Priority=100).

This means that at most they would wait 4 hours for their job to be

submitted. 3. 'cdac-contrib' or "\${groupname}-contrib". Since 'cdac', in this example, has donated to the cluster they should get a higher priority on other machines as well.

We do the following calculation to determine the contrib partitions (cdac-contrib) usage over the past 30 days in comparison to total cluster usage.

$$\frac{\text{partition usage total time in seconds for 30 days}}{\text{all partition usage total time in seconds for 30 days}} = \text{percent used}$$

The percent will end up as an integer.

You'll see on the spreadsheet we take subtract this number from 100, ensure it's positive and call that "idleness".

Total amount of money contributed and "idleness" are the key factors in determining the priority of a groups 'contrib' partition.

This calculation will be run once a month and the relevant groups \${group}-contrib priority updated to reflect past months usage.

Note that the term "idleness" should not be taken literally. I don't know of way to actually calculate true idleness. I do believe that the current calculation reflects the intent of the term.

TODO:

- 1. There are multiple methods used to calculate priority reflected on

the spreadsheet.

- 1. Double check the math for correctness.
2. Does the math reflect our intent? I think it does.
3. Multiple methods in calculating priority are reflected (blue to

purple cells: har\_priority, normalized\_priority, weighted\_normalized\_priority

- 1. Choose one calculation to use. This decision doesn't prevent us

from changing it later if we find another calculation works better.

- 1. If you have a suggestion: Please show us the work in a new column

or by cloning the sheet.

Sheet usage:

- 1. All donations will get logged into the spreadsheet under the 'log'

sheet.

1. Hardware donation gets converted to USD by techstaff. A receipt of the purchase is good starting point.

1. donations get assigned a POSIX group. Group must have a primary contact, who then gets to set members for that group.

1. The group 'cs' is calculated on the spreadsheet but doesn't actually

get any priority set. 'general' is for all CS users. It needs to be part of the total sum but is treated as a special case with 0 priority.

## AI Cluster Admin

### TODO

Since I'm still working on it, I don't guarantee any uptime yet. Mainly I need to make sure TRES tracking is working like we want. This will involve restarting slurmd and slurmctld which will kill running jobs.

- ~~generate report of storage usage~~
- ~~groups (Slurm 'Accounts') created for PI's.~~
  - e.g. ericj\_group: ericj, user1, user1, etc
- ~~grab QOS data from somewhere (gsheet or some kind of DB)~~
- ~~Properly deploy sync script~~
  - ~~Systemd unit~~
  - ~~main loop~~
- ~~research on slurm plugin to force GRES selection on job submit. Might be able to use:~~
  - ~~SallocDefaultCommand~~
  - ~~Otherwise look for 'AccountingStorageTRES' and 'JobSubmitPlugins' and /etc/slurm-llnl/job\_submit.lua ← used to force user to specify '-gres'.~~
  - ~~jobs that do not specify a specific gpu type (e.g. gpu:rtx8000 or gpu:rtx2080ti) could be counted against either one but not specifically the you actually used.~~
  - ~~From 'AccountingStorageTRES' in slurm.conf: "Given a configuration of "AccountingStorageTRES=gres/gpu:tesla,gres/gpu:volta" Then "gres/gpu:tesla" and "gres/gpu:volta" will track jobs that explicitly request those GPU types. If a job requests GPUs, but does not explicitly specify the GPU type, then its resource allocation will be accounted for as either "gres/gpu:tesla" or "gres/gpu:volta", although the accounting may not match the actual GPU type allocated to the job and the GPUs allocated to the job could be heterogeneous. In an environment containing various GPU types, use of a job\_submit plugin may be desired in order to force jobs to explicitly specify some GPU type."~~
- ~~ganglia for Slurm: <http://ai-mgmt2.ai.cs.uchicago.edu>~~
  - ~~figure why summary view is no longer a thing.~~
- ~~update 'coolgpus'. Lose VTs when this is running.~~
  - ~~coolgpus: sets fan speeds of all gpus in system.~~
  - ~~Goal is to statically set fan speeds to 80%. The only way to do this is with fake Xservers...~~

but that means you lose all the VTs. Is this a compromise I'm willing to make? It is.

- home directory
  - ~~setup backups for home dirs~~
  - ~~default quota~~
  - ~~userland tool to check quota~~
  - ~~home directory usage report~~
- monitoring
  - ~~basic node monitor~~
  - ~~nfs or bandwidth monitoring~~
  - ~~gpu~~
- sync script
  - ~~fix bug to ensure accounts and users are created~~

From:

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

Permanent link:

<https://howto.cs.uchicago.edu/techstaff:aicluster-admin?rev=1612980230>

Last update: **2021/02/10 12:03**

