

Hyak Job Scheduler Changes

Beginning Tuesday, January 11 at 09:00am, the Hyak job scheduler will undergo a number of changes to improve cluster reliability, performance, and utilization. Below you will find a summary of changes in behavior that will affect users.

MAJOR CHANGES

1. Node local scratch disks (`n####:/tmp`, `n####:/scr`, `n####:/var/tmp`) will be cleaned up after the completion of each job. Groups that would like to keep data on the disks in their nodes should save any data into a directory with their group name. If your UNIX group is `hyak-mygroup`, your directory should be named `mygroup`.
2. Jobs must use all the processors on a node. For many users, this will require no changes. If your jobs use fewer than 8 processors, you will need to create job scripts that start multiple jobs. Sample scripts are available on the Hyak User Wiki: ([Hyak_Serial_Job_Scripts](#)).
3. Users will now be able to submit to a special queue, `bf`, and those jobs will run on the 56 nodes owned by the eScience Institute. Users can specify an alternate queue using the `-q msub` option, `msub -q bf myjob.sh`. These jobs can be canceled and requeued by eScience job submissions, so short jobs or jobs that use checkpointing are best suited to this new queue. Interactive jobs cannot be run in the `bf` queue.

OTHER CHANGES

1. When a job is submitted, there may be a delay of up to thirty seconds before the job is scheduled. This delay will be visible when running `showq` or starting an interactive session.
2. Users who submit hundreds of jobs at once should use `qsub` in place of `msub`.
3. A resource called `file` will now be available for each node indicating how much space is available on the local scratch disk. `msub -l file=50GB myjob.sh` would only select nodes that have at least 50GB free on their local scratch disk.
4. Users with access to multiple groups' nodes will now have to specify their group, rather than the partition name when they are submitting jobs to their non-default group/partition. Most users are only in one group and will not have to change their behavior. If your default UNIX group is `hyak-mygroup` and your secondary group is `hyak-othergroup` you'll have to run `qsub -W group_list=hyak-othergroup myjob.sh` if you'd like to use the nodes owned by `othergroup`. You must also use `qsub` instead of `msub`.