

# Chapter 10: Job Priority and Queue

## Management

Users submit read and write jobs to **encp** (often through an interface, e.g., dCache). **Encp** sends a request for each job to an Enstore library manager for processing. The library manager receives these requests, stores them in a queue, assigns a priority to each one, and passes them to a mover for actual data transfer. A request's priority determines when it will get processed relative to others in the queue, and its priority may change as circumstances change while it waits in the queue. There are four items that factor into determining priority: category, numerical value within the category, "Fair Share", and ownership of resources by a group/experiment.

### 10.1 Job Priority Categories

---

There are two categories of job priority: normal and DAQ/Admin. The default priority is "normal". DAQ/Admin priority, as its name implies, is reserved for high-priority jobs.

DAQ/Admin priority is granted only to job requests that satisfy certain preconfigured conditions. Conditions, if set, filter on the request's originating username, group, node, and so on. For example, conditions could be set to grant DAQ/Admin priority to all jobs submitted by user *joe* from node *mynode1*, by users *george* or *kim* from any node, and by any user from node *myspecialnode*; all other jobs get normal priority.



Setting conditions for DAQ/Admin priority is an administrative function.

Requests that get submitted and receive DAQ/Admin priority get moved to the head of the LM's request queue. If there is more than one at a time, the other priority-related factors determine the order in which these requests get processed. Any normal priority transfer that is in progress is allowed to complete, but the system does not then process other normal priority requests in the queue that are waiting for the same volume. At the completion of the current transfer, the tape is replaced in the drive if necessary, then the (first) DAQ/Admin request gets processed.

## 10.2 Numerical Priority Values

---

Within its priority category, each request is also assigned an initial numeric priority value. The numeric priority is set by default according to preconfigured conditions. It can be overridden on the command line using the **encp** options described in section 5.5 *Encp Command Options*, although we strongly recommend against doing this. The numeric priority may change while the request waits in the queue depending on (a) the **encp** options used, (b) the elapsed time in the queue.

## 10.3 Fair Share Resource Allotment

---

Enstore queue management has an algorithm called “Fair Share”, that helps to keep any one storage group (experiment or group) from monopolizing tape drives.

When a mover which does not have a mounted tape asks the Library Manager for the next request, the Fair Share determines which storage groups currently have requests in progress (at a mover) and which ones don't, then gives preference to requests associated with those that don't. This helps ensure that there is a fair distribution of resources available to all groups currently using the system.

## 10.4 Resource Ownership

---

When an experiment or group purchases one or more tape drives, the drives go into the pool of Enstore resources, accessible by all users, with the recognition that the purchasing group has preferential access to this number of drives. The Fair Share configuration (section 10.3) gets modified to guarantee this access. The priorities of other requests in the queue may be perturbed when Enstore needs to free up one or more tape drives for jobs submitted by the purchasing group.

In contrast to DAQ/Admin requests, these (normal priority) requests, identified by their storage group, must wait until a tape has dismounted in the normal way before being processed. A tape normally dismounts after all requests in the queue requiring that tape have completed.