

# Job-submission queues and charges

Most **Cheyenne** batch queues are for *exclusive* use, so jobs submitted to those queues are charged for all 36 cores on each node that is used. Jobs in the *shared use* "share" queue are charged only for the cores that are used.

The "regular" queue, which has a 12-hour wall-clock limit, meets most users' needs for running batch jobs on Cheyenne.

## Page contents

- [Cheyenne queue details](#)
- [Calculating charges](#)
  - [Exclusive nodes](#)
  - [Shared nodes \(Cheyenne and Casper\)](#)
- [Checking and managing charges](#)

## Cheyenne queue details

Queue name	Priority order	Wall clock limit (hours)	Queue factor	Cheyenne queue description
premium	1	12	1.5	Jobs are charged at 150% of the regular rate.
regular	2	12	1	Most production batch jobs run in this queue; also accepts interactive jobs.
economy	3	12	0.7	Production batch jobs are charged at 70% of the regular rate.
share	NA	6	1	Interactive and serial batch use for debugging and other tasks on a single, shared, 128-GB node. An individual job can use up to 18 cores. A user can run multiple jobs in the share queue concurrently if the total number of cores they require is no more than 18. Additional jobs that the user submits remain in the queue to run later.

Some additional queues on the system are for dedicated purposes and accessible only to authorized users.

**Casper:** See [this page](#) for information on running jobs on Casper nodes.

## Calculating charges

### Exclusive nodes

Charges for use of Cheyenne are calculated in terms of core-hours. Jobs run in Cheyenne queues other than "share" are charged for exclusive use of the nodes by this formula:

**wall-clock hours × nodes used × cores per node × queue factor**

### Number of nodes used

Your batch script indicates how many Cheyenne nodes your job will use. In this example, you have selected 2 nodes, each of which has 36 cores. Your job will be charged for the use of 72 cores.

```
#PBS -l select=2:ncpus=36:mpiprocs=36
```

### Shared nodes (Cheyenne and Casper)

Charges for jobs that you run on a shared node, including Casper nodes, are calculated by this formula:

**core-seconds/3600 (core-hours)**

## Checking and managing charges

Users can check computing and storage charges through the CISL Systems Accounting Manager. (Go to [documentation](#) or to [SAM app](#).)

If you have concerns about using your allocation most efficiently, [contact the NCAR Research Computing help desk](#) for guidance. Sometimes jobs can be configured to make better use of the processors, and you may be able to save by using a less expensive queue.

CISL can refund core-hours if system failures cause jobs to fail and the failed jobs are reported promptly. Use this [core-hours refund request form](#) (login required) if you think a refund is warranted. Technical limitations prevent us from verifying refund eligibility for jobs that are more than seven days old.