

Torque

How does Torque calculate and report memory usage?

Torque pulls information about memory usage for jobs from `/proc/<pid>/stat`. It pulls information on the Resident Set Size (RSS) to populate the `mem` field and it pulls the Virtual Size to populate `vmem`. The RSS is primarily the physical memory and Virtual Size is physical memory and swap combined. These numbers can get skewed slightly because they also include shared libraries and other types of memory.

It's possible to see memory utilization be higher than the amount of memory available on a node. Torque collects the high water mark for memory usage for each process it tracks and sums them together. For example, if a job spawned 10 processes and they each used 7gb of memory at their peak usage, it would report 70gb of usage for the job.

Unique solution ID: #1067

Author: Ben Roberts

Last update: 2015-09-04 15:59