

Torque

Enabling and configuring support for cgroups in Torque

Question: What's involved in configuring the compute nodes to use cgroups?

Solution:

Before following these instructions, make sure to verify them against the latest "Installation and Configuration Guide". For more information, see the release notes and full product documentation for your OS & Torque versions at <http://docs.adaptivecomputing.com>

Torque version 6.0.0 added cgroups support. Prior to this, Torque supported cpusets, which is still available, but configuring with both options will cause `--enable-cpuset` to be ignored, since cgroups already includes them. Adaptive recommends configuring with cgroups on newer versions of Torque (and installing the most recent version). Be aware that not everything that worked with cpuset support is available or works the same way for cgroup support.

Summary: for `pbs_mom` to employ cgroups, you must do the following:

- 1) Have cgroups enabled and mounted on the compute nodes.
- 2) Have `hwloc` installed (on the nodes, and on the build host).
- 3) Configure with `"--enable-cgroups"` and `"--with-hwloc-path="`.

1) Enable cgroups:

* On RHEL/CentOS 7.x systems, run this command on the server and nodes:

```
# yum -y install libcgroup-tools
```

* On RHEL/CentOS 6.x systems, run these commands on the server and nodes:

```
# yum -y install libcgroup
```

```
# service cgconfig start
```

```
# chkconfig cgconfig on
```

* For other operating systems and full documentation, be sure to check the latest version of the Torque admin guide, in this section:

Installation and Configuration > Torque on NUMA Systems > Torque NUMA-Aware

Torque

Configuration

To confirm that cgroups have been mounted, you can run commands such as `lssubsys`, `lstopo`, or `cgsnapshot`.

2) Install hwloc:

The simplest way to do this is to run `"yum install hwloc-devel"`. If you think the hwloc libraries and tools may already be installed, you can run `"hwloc-info --version"` to check the version (if not found, try `"locate hwloc-info"` to make sure it's not already installed, but just missing from the path). You may also want to check your the system package management tool to see which versions have already been installed (for example, `"yum list installed | grep hwloc"` or `"rpm -qa | grep hwloc"`)

Torque requires hwloc 1.9.1, or version 1.11.1 if you have any nodes with the NVIDIA K80 GPU.

3) Configure with cgroups and hwloc:

When running `configure`, include `"---with-hwloc-path=<path to libhwloc.so> --enable-cgroups"`. Then build with `make` and install in the usual manner.

To confirm you have proper NUMA-aware reporting, run `pbsnodes`, which will display lines like this at the bottom of the output for each node:

...

```
total_sockets = 2
total_numa_nodes = 1
total_cores = 8
total_threads = 8
dedicated_sockets = 0
dedicated_numa_nodes = 0
dedicated_cores = 0
dedicated_threads = 8
Unique solution ID: #1171
Author: Rick McKay
Last update: 2017-06-02 03:51
```