

Moab-SGE Integration Notes

Moab-SGE Integration Notes

Copyright © 2011 Adaptive Computing Enterprises, Inc.

This document provides information on the steps to integrate Moab with an existing functional installation of SGE.

Notice

Distribution of this document for commercial purposes in either hard or soft copy form is strictly prohibited without prior written consent from Adaptive Computing Enterprises, Inc.

Overview

Moab's native resource manager interface can be used to manage an SGE resource manager. The integration steps simply involve the creation of a complex variable and a default request definition. The Moab tools directory contains a collection of customizable scripts which are used to interact with sge. This directory also contains a configuration file for the sge tools.

Moab Integration Steps

You should follow the regular steps for installing Moab with the following exceptions:

Run Configure with the `--with-sge` option

When running the configure command, use the `--with-sge` option to specify the use of the native resource manager interface with the sge resource manager subtype. This will place a line similar to the following in the Moab configuration file (`moab.cfg`):

```
RMCFG[clustername]    TYPE=NATIVE:sge
```

Example 1. Running configure

```
$ ./configure --prefix=/opt/moab --with-homedir=/var/moab --with-sge
```

Customize the moab configuration file

In order to allow the specification of a parallel environment (`-l pe`) via `msub`, you will need to tell Moab to pass through arbitrary resource types.

Example 2. Edit `moab.cfg`

```
# vi /var/moab/moab.cfg

# Transmit arbitrary resource types (ie. pe) from msub into the job-start script
CLIENTCFG[Moab] FLAGS=AllowUnknownResource

# Allow regular users to awaken the scheduler for responsive msubs
ADMINCFG[5] USERS=ALL SERVICES=mschedctl:resume
```

Customize the sge tools configuration file

You may need to customize the `$MOABHOMEDIR/etc/config.sge.pl` file to include the correct `SGE_ROOT` and `PATH`, and set other configuration parameters.

Example 3. Edit `config.sge.pl`

```
# vi /var/moab/etc/config.sge.pl

# Set the SGE_ROOT environment variable
$ENV{SGE_ROOT} = "/opt/sge-root";

# Set the PATH to include directories for sge commands -- qhost, etc.
$ENV{PATH} = "$ENV{SGE_ROOT}/bin/lx24-x86:$ENV{PATH}";
```

SGE Integration Steps

After installing SGE on your cluster and verifying that it is running serial and parallel jobs satisfactorily, you should perform the following steps:

Define a new complex variable named `nodelist`

Use the `qconf -mc` command to edit the complex variable list and add a new requestable variable of the name `nodelist` and the type `RESTRING`.

```
# qconf -mc

nodelist          nodelist          RESTRING          ==          YES          NO          NONE          0
```

Add a default `nodelist` request definition

This step will set the `nodelist` complex variable for all jobs to the unassigned state until they are ready to run, at which time the job will be assigned a `nodelist` directing which nodes it can run on.

Example 4. Edit `sge_request`

```
# vi /opt/sge-root/default/common/sge_request

# Set the job's nodelist variable to the unassigned state until it is ready to
# start at which time it will be reset to the list of nodes it is designated to
# run on
-l nodelist=unassigned
```

Populate the node's `nodelist` variable

This step will set the `nodelist` complex variable for all exec hosts to their own short hostnames. This will allow jobs to start when their `nodelist` value matches up with a set of nodes.

Example 5. `qconf -rattr exechost complex_values nodelist=$hostname $hostname`

```
# for i in `qconf -sel | sed 's/\..*//'; do echo $i; qconf -rattr exechost
complex_values nodelist=$i $i; done
```

Shorten the scheduler interval

Use the `qconf -msconf` command to edit the `schedule_interval` setting to be less than or equal to one half the time of the Moab `RMPOLLINTERVAL` (seen with `showconfig | grep RMPOLLINTERVAL`).

```
# qconf -msconf
schedule_interval                0:0:15
```

Add the sge ports to the services file

In order for the sge client commands to know what port to use when communicating with the sge qmaster, the ports should be listed in the `/etc/services` file. (Alternatively, the `SGE_QMASTER_PORT` environment variable must be set in the `config.sge.pl` file).

Example 6. Edit `/etc/services`

```
# vi /etc/services
sge_qmaster    536/tcp          # SGE QMaster
sge_execd     537/tcp          # SGE Execd
```