

Moab Passthrough

Administrator Guide 9.1.3

September 2018



© 2018 Adaptive Computing Enterprises, Inc. All rights reserved.

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.

Adaptive Computing, Cluster Resources, Moab, Moab Workload Manager, Moab Viewpoint, Moab Cluster Manager, Moab Cluster Suite, Moab Grid Scheduler, Moab Grid Suite, Moab Access Portal, and other Adaptive Computing products are either registered trademarks or trademarks of Adaptive Computing Enterprises, Inc. The Adaptive Computing logo and the Cluster Resources logo are trademarks of Adaptive Computing Enterprises, Inc. All other company and product names may be trademarks of their respective companies.

Adaptive Computing Enterprises, Inc.
704 Goodlette Road North
Naples, FL 34102
+1 (239) 330-6083
www.adaptivecomputing.com



Scan to open online help

Contents

Welcome	5
Chapter 1: Moab Passthrough Overview	1
1.1 Considerations	2
1.2 Moab Passthrough And Viewpoint	3
Chapter 2: About Installing Moab Passthrough	5
2.1 Installing Moab Passthrough	6
2.2 Moab Passthrough Configuration	8
2.2.1 Configuring MinJobAge	8
2.2.2 Configuring Complete Job Purge Time	8
Chapter 3: Using Moab Passthrough	9
3.1 Moab Passthrough And Viewpoint	10
3.1.1 Passthrough Licensing	10
3.1.2 Viewpoint With Passthrough	10
3.2 Moab Passthrough Usage Notes	12
3.2.1 Determining Node State	12
3.2.2 Clearing Statistics/Failed Jobs	12

Welcome

Welcome to *Moab Passthrough 9.1.3 Administrator Guide*.

This guide is intended as a reference for system administrators.

Chapter 1: Moab Passthrough Overview

Moab passthrough mode enables limited job submission and monitoring with native resource managers, such as Slurm. Currently, Slurm is the only resource manager supported by Moab passthrough mode.

In this chapter:

1.1 Considerations	2
1.2 Moab Passthrough and Viewpoint	3

1.1 Considerations

i Passthrough mode requires a Moab passthrough license. Contact licenses@adaptivecomputing.com for information on obtaining licenses.

- Moab passthrough has been tested with Slurm versions 16.05.x and 17.02.x. Moab Passthrough may work with Slurm version 15.x, but this configuration has not been tested.
- Currently the only job modify actions supported are *job hold* and *job walltime*.
- Consider setting `MinJobAge=<#>` to 300 in `slurm.conf`. For systems that push through tens of thousands of jobs a day a minimum of 180 is suggested. If Moab does not have enough time to load the job from Slurm, consider increasing `MinJobAge` higher.

After adjusting the value of `MinJobAge` in `slurm.conf`, define the `JOBCPURGETIME` parameter in `moab.cfg`. `JOBCPURGETIME` should be a minimum of 60 seconds greater than the value of `MinJobAge`.

1.2 Moab Passthrough and Viewpoint

Some configuration is required to use Viewpoint with Moab passthrough mode. See the ViewpointReference Guide if you plan to use Viewpoint to submit jobs.

Job holds are all admin holds when submitting jobs through Viewpoint, so if a user wants to release a hold on a job, they will need to use Viewpoint. Due to how Slurm processes holds, you may not see the hold released until Slurm has processed the request. This can take a full Slurm scheduling iteration.

Chapter 2: About Installing Moab Passthrough

In this chapter:

2.1	Installing Moab Passthrough	6
2.2	Moab Passthrough Configuration	8
2.2.1	Configuring MinJobAge	8
2.2.2	Configuring Complete Job Purge Time	8

2.1 Installing Moab Passthrough

i The instructions in this section are for installing the Moab Passthrough feature, not for installing Slurm or any other resource manager.

Do the following:

1. You should have received a `moab.lic` file that enables Moab Passthrough. Copy the `moab.lic` file to `/opt/moab/etc/`.
2. If necessary, copy the passthrough scripts to `/opt/moab/tools`. The examples below assume they are in `/opt/moab/tools/passthrough/slurm`.
3. Change the `RM POLLINTERVAL` setting in `/opt/moab/etc/moab.cfg`:

```
RM POLLINTERVAL 0,10
```

4. Delete or comment out the following lines from `/opt/moab/etc/moab.cfg`:

```
RMCFG[pbs] TYPE=PBS SUBMITCMD=/usr/local/bin/qsub
MAXJOB 51200
JOBCFG[remote_visualization] FLAGS=usemoabjobid SELECT=TRUE
JOBCFG[nitro] FLAGS=usemoabjobid SELECT=TRUE
```

5. If your `/opt/moab/etc/moab.cfg` file contains the following lines, delete them also:

```
RMCFG[mws] TYPE=MWS
RMCFG[mws] BASEURL=http://slurm-install.vapor2.ac:8080/mws
```

6. Add the following lines to your `/opt/moab/etc/moab.cfg` file:

i Change the path to `sbatch` to the appropriate path for your environment.

```
#####
# Moab.cfg configuration lines
#####
RMCFG[slurm] SUBMITCMD=/usr/bin/sbatch
RMCFG[slurm] AUTHTYPE=CHECKSUM
RMCFG[slurm] FLAGS=PROXYJOBSUBMISSION
RMCFG[slurm] JOBIDFORMAT=integer
RMCFG[slurm] CLUSTERQUERYURL=exec:///opt/moab/tools/passthrough/slurm/node_query.py
RMCFG[slurm] JOBCANCELURL=exec:///opt/moab/tools/passthrough/slurm/job_cancel.py
RMCFG[slurm] JOBMODIFYURL=exec:///opt/moab/tools/passthrough/slurm/job_modify.py
RMCFG[slurm] WORKLOADQUERYURL=exec:///opt/moab/tools/passthrough/slurm/workload_
query.py
RMCFG[slurm] SYNCJOBID=FALSE
RMCFG[slurm] TYPE=WIKI:SLURM
IDCFG[slurm] SERVER=exec:///opt/moab/tools/passthrough/slurm/slurm_credentials.py
REFRESHPERIOD=hour CREATECRED=TRUE
```

7. Edit the `/opt/moab/tools/passthrough/slurm/passthrough.cfg` file to change the `scontrol_path` setting to the location of Slurm commands. If desired, change the debug

level of the script and the logging path.

```
[MASTER]
version = 1.0

# Must end with a /
scontrol_path = /usr/local/bin/

# INFO|ERROR|DEBUG
log_level = ERROR

log_file = /opt/moab/log/passthrough.log
```

8. Edit the `/opt/moab/tools/passthrough/slurm/RmSlurm.py` file. Set `self.config.read` to match the location to your `passthrough.conf` file (`/opt/moab/tools/passthrough/slurm/passthrough.cfg`).
9. Do one of the following:
 - If the Moab process is not running, start Moab (`service moab start`).
 - If Moab is running, restart Moab (`mschedctl -R`).
10. Check the output of `mdiag -R -v` and `mdiag -n` to validate your setup is working. The output of `mdiag -R -v` should look similar to the following:

```
[root]# mdiag -R -v
diagnosing resource managers

RM[internal] State: --- Type: SSS
Max Fail/Iteration: 0
JobID Format: integer
JobCounter: 2683
DefaultClass: gpu
RM Performance: AvgTime=0.56s MaxTime=0.71s (5 samples)
RM Languages: -
RM Sub-Languages: -

RM[slurm] State: Active Type: WIKI:SLURM ResourceType: COMPUTE
Server: localhost:7321
Timeout: 30000.00 ms
Version: 'slurm 16.05.4'
Password Set: Yes
Cluster Query URL: exec://$HOME/tools/passthrough/slurm/node_query.py
Workload Query URL: exec://$HOME/tools/passthrough/slurm/workload_query.py
Job Cancel URL: exec://$HOME/tools/passthrough/slurm/job_cancel.py
Job Modify URL: exec://$HOME/tools/passthrough/slurm/job_modify.py
Objects Reported: Nodes=2 (2 procs) Jobs=4851
Nodes Reported: 2 (N/A)
Flags: executionServer,proxyjobsubmission
Partition: slurm
NOTE: checksum authorization set (secret key: 1...)
Event Management: EPORT=10777 (no events received)
NOTE: delta node query enabled
NOTE: delta job query enabled
RM Performance: AvgTime=2.34s MaxTime=3.51s (16 samples)
RM Languages: WIKI
RM Sub-Languages: WIKI
```

2.2 Moab Passthrough Configuration

This topic provides notes for configuring Moab HPC Suite suite components for Moab Passthrough.

In this topic:

[2.2.1 Configuring MinJobAge - page 8](#)

[2.2.2 Configuring Complete Job Purge Time - page 8](#)

2.2.1 Configuring MinJobAge

Consider setting `slurm.conf` `MinJobAge=<#>` to 300. For systems that push through tens of thousands of jobs a day a minimum of 180 is suggested. If Moab does not have enough time to load the job from Slurm, consider increasing the Slurm `MinJobAge` higher.

2.2.2 Configuring Complete Job Purge Time

After adjusting the value of `MinJobAge` in `slurm.conf`, define the `JOBCPURGETIME` parameter in `moab.cfg`. `JOBCPURGETIME` should be a minimum of 60 seconds greater than the value of `MinJobAge`.

By default, Moab maintains detailed information on completed jobs for 5 minutes. You may want to set the completed job purge time to preserve completed job data for a longer period. See "`JOBCPURGETIME`" in the *Moab Workload Manager Administrator Guide* for more information.

Chapter 3: Using Moab Passthrough

In this chapter:

3.1	Moab Passthrough and Viewpoint	10
3.1.1	Passthrough Licensing	10
3.1.2	Viewpoint With Passthrough	10
3.2	Moab Passthrough Usage Notes	12
3.2.1	Determining Node State	12
3.2.2	Clearing Statistics/Failed Jobs	12

3.1 Moab Passthrough and Viewpoint

Moab passthrough mode enables limited job submission and monitoring with native resource managers, such as Slurm. Currently, Slurm is the only resource manager supported by Moab passthrough mode. This topic describes changes in how Viewpoint operates in passthrough mode.

In this topic:

- [3.1.1 Passthrough Licensing - page 10](#)
- [3.1.2 Viewpoint With Passthrough - page 10](#)

3.1.1 Passthrough Licensing

Passthrough mode is active when both a Moab Passthrough license and a Viewpoint Passthrough Support license are enabled. See 1.1 Licensed Features Page in the Moab Viewpoint Reference Guide for more information about verifying which license feature are active. See [Chapter 2: About Installing Moab Passthrough - page 5](#) for more information about installing and configuring passthrough mode.

3.1.2 Viewpoint With Passthrough

When Moab passthrough mode is active, Viewpoint hides features that are not available. The disabled features on each of Viewpoint's pages are listed below.

3.1.2.A Node Details Page

The following features are not available on the Node Details Page when running in passthrough mode:

- Adding features to a node. The Features box is present and shows configured features for the selected node, but you cannot add additional node features.
- Reservations
- The node's processor architecture

3.1.2.B Template Details Page

The following features are not available on the Template Details Page when running in passthrough mode:

- The Delay Start By widget in Time Management settings
- The Architecture widget in Resources settings

- The Merge Streams widget in Advanced Settings
- The Moab Environment Variables widget in Advanced Settings
- The Excluded Features widget in Advanced Settings
- The Operating System widget in Advanced Features
- The Email Input widget in Advanced Features
- The Email Options widget in Advanced Features
- The Default Moab Template widget in Advanced Features
- All Node Policies settings except Nodes Requested
- Moab inputs and environment variables in Script Builder

3.1.2.C Update or Create Role Pages

The following features are not available on the Update or Create Roles Pages when running in passthrough mode:

- Job Details - Priority Analysis permission.

3.1.2.D Job Details Page

The following features are not available on the Job Details Page when running in passthrough mode:

- Job Priority
- Scheduler Analysis cannot be performed on a blocked job.
- QoS setting cannot be changed for a job.
- Reservation names.
- Processor architecture.
- Operating system.

Related Topics

- [Chapter 3: Using Moab Passthrough - page 9](#)
- [1.2 Moab Passthrough and Viewpoint - page 3](#)

3.2 Moab Passthrough Usage Notes

In this topic:

[3.2.1 Determining Node State - page 12](#)

[3.2.2 Clearing Statistics/Failed Jobs - page 12](#)

3.2.1 Determining Node State

`mdiag -n` provides information about the state of nodes Moab is tracking.

```
[root]# mdiag -n -v
compute node summary
Name          State    Procs   Memory   Disk    Swap    Speed   Opsys   Arch  Par
Load Classes
support-sn1   ---     0:1     74:74    0:0     0:0     1.00    Linux  x86_64  slu
0.00 [normal][intr][medium][share][extra][gpu] a,b,c,d,ghi
GRES=socket:1,gpu:4,lic:2 NODEINDEX=1
support-sn2   ---     0:1     76:76    0:0     0:0     1.00    Linux  x86_64  slu
0.00 [normal][intr][medium][share][extra] a,b,c,d,ghi
GRES=socket:1
NODEINDEX=2
-----     ---     0:2     150:150  0:0     0:0

Total Nodes: 2 (Active: 0 Idle: 0 Down: 2)
```

3.2.2 Clearing Statistics/Failed Jobs

You can use `mrmctl -f` to clear reported failures, performance data, and messages for a specified resource manager.

```
[root]# mrmctl -f messages slurm
```