

# Moab HPC Suite 7.2.7 – Enterprise Edition Release Notes

The release notes file contains the following sections:

- [New features on page 1](#)
- [Differences on page 2](#)
- [Installation and upgrade information on page 2](#)
- [Known issues on page 2](#)
- [Resolved issues on page 3](#)
- [Product documentation on page 6](#)

## New features

The following is a summary of key new features in Moab HPC Suite 7.2.7 – Enterprise Edition.

### Moab Workload Manager

#### *EnableSlurmMemPerCPU scheduler flag*

A new scheduler flag called *EnableSlurmMemPerCPU* allows you to request pmem in an msub command when you use SLURM as your resource manager.

```
SCHEDCFG[moab] FLAGS=ENABLESLURMMEMPERCPU
```

For more information, see "[SCHEDCFG flags](#)" in the Moab Workload Manager Administrator Guide.

#### *UsePhysicalMemory resource manager flag*

A new resource manager flag called *UsePhysicalMemory* causes Moab to adjust the reported Available Memory when a job takes from either physical or virtual memory. Previously, if a node had 12 GB of RAM and 12 GB of swap space (24 GB of virtual memory), the "Available Memory" would not change when jobs were scheduled on the node.

```
RMCFG[torque] FLAGS=UsePhysicalMemory
```

For more information, see "[Resource Manager Configuration](#)" in the Moab Workload Manager Administrator Guide.

#### *FreeCompletedJobSubmitString scheduler flag*

A new scheduler flag called *FREECOMPLETEDJOBSUBMITSTRING* frees up the submit string for completed jobs to decrease the memory needed for operation. For more information, see "[SCHEDCFG flags](#)" in the Moab Workload Manager Administrator Guide.

```
SCHEDCFG[] FLAGS=FREECOMPLETEDJOBSUBMITSTRING
```

### Default features for nodes in a partition

You can now specify default node features for nodes in a given partition. For more information, see "[Partitions](#)" in the Moab Workload Manager Administrator Guide.

```
PARCFG[] NODECFG[DEFAULT] FEATURES+=feature
```

## Differences

This section contains differences in previously existing features that require a change in configuration or routine.

### TORQUE Resource Manager

#### *trqauthd can only load once*

You cannot load `trqauthd` more than once. It opens a UNIX domain name file to do its communication with client commands. If the UNIX domain name file exists, `trqauthd` will not load. By default, this file is `/tmp/trqauthd-unix`, but you can configure it to point to a different directory. If `trqauthd` will not start and you know there are no other instances of `trqauthd` running, delete the UNIX domain file and try again.

## Installation and upgrade information

### Installing Moab HPC Suite 7.2.7 – Enterprise Edition

Complete Moab HPC Suite 7.2.7 installation instructions and requirements can be found in HTML or PDF format.

- 7.2.7 RPM Installation Guide – [HTML](#) – [PDF](#)
- 7.2.7 Tarball Installation Guide – [HTML](#) – [PDF](#)

### Upgrading to Moab HPC Suite 7.2.7 – Enterprise Edition

Moab Accounting Manager 7.2.7 is compatible with Moab Workload Manager versions 7.2.x, 7.3.x, and 7.5.x. It is only compatible with a Moab Accounting Manager database schema version of 7.5. To determine your current database schema version, run `goldsh System Query`. If your current database schema version is less than 7.5, you need to migrate your database to 7.5. For more information, see "[Upgrading Moab Accounting Manager](#)" in the *Moab HPC Suite – Enterprise Edition Installation Guide*.

## Known issues

The following are known issues in the Moab HPC Suite 7.2.7 – Enterprise Edition. Following each issue description is an associated issue number in parentheses.

- A default job template is not applied to job array sub-jobs until after a Moab recycle. Restart Moab to apply the job template (MOAB-5121).
- When the resource manager reports a wiki attribute that Moab does not recognize and it contains "OS," Moab considers it the OS (MOAB-5120).
- Specifications in the Moab configuration files overwrite conflicting specifications in the `moab.dat` file. This means that if you dynamically change Moab configurations found in `moab.cfg` or its included configuration files, the changes may be lost upon restart (MOAB-4246).
- Node flags cannot be removed via the `moab.cfg`. They should be set and removed dynamically by running `mschedctl -m config` or using a resource manager (MOAB-4123).
- When you upgrade Moab and TORQUE, depending on the versions, you could encounter a problem where the core files are created frequently in `/opt/moab`. You can resolve this problem by removing the old library files from `/usr/local/lib` (TRQ-1082).

## Resolved issues

The following is a list of some key bugs fixed in Moab HPC Suite 7.2.7 – Enterprise Edition. Following each issue description is an associated issue number in parentheses.

- **The `mstat_converter` tool did not read `$MOABHOMEDIR/etc/moab.cfg` as documented.** The Moab Workload Manager Administrator Guide and the `mstat_converter` README file have been modified to correctly state that it reads `$MOABHOMEDIR/moab.cfg` (MOAB-6627).
- **When nodes were unavailable for a while and reintroduced to Moab after reboot, Moab would put them in the wrong partition.** Moab now returns nodes to the correct partition after reboot (MOAB-6539).
- **`msub -I -V` failed when an environment variable contained a space; in addition, environment variables were not being inserted into an interactive SLURM job with `msub -I -E`.** Environment variables are no longer truncated and accept spaces in a SLURM environment (MOAB-6594 and MOAB-5839).
- **Moab incorrectly parsed the node state for `checknode`.** Spaces and semicolons are now supported (MOAB-6612)
- **Moab could not find the job ID of dependencies when they were submitted.** Moab is again able to find the job ID of dependencies (MOAB-6614).
- **The `showstate` command would fail with "exceeded stack limit" errors with a stack size within limits.** The `showstate` command no longer crashes with the default stack size of 8192 KB (MOAB-5078).
- **Moab ignored host list requirements for multi-req jobs.** Moab again honors the host list requirements given to jobs (MOAB-6475).
- **Moab Accounting Manager charged long-running jobs twice.** The final charge is now incremental if `flushinterval` is configured (MOAB-6622).

- **When you tried to modify a job that had already migrated, Moab would report a success even though this cannot be done.** Moab now prints an error when you try to modify a job that has migrated (MOAB-6600).
- **In a SLURM environment, Moab did not properly handle an `msub` node or `ppn` specification.** Moab now handles node and `ppn` requests correctly in SLURM environments (MOAB-6361 and MOAB-6580).
- **Submitting CRAY jobs with `mpp*` attributes did not allow you to filter based on nodes.** Moab will filter by node for certain `mpp*` attributes (MOAB-6509).
- **The `mppnodes` attributes did not work correctly when a job requested to use a specific node list exclusively.** Moab now honors an exclusive node list requested at job submission with `mppnodes` (MOAB-6245).
- **Moab did not update existing reservations when only the credential list changed for a standing reservation.** Moab now updates the reservation with the updated settings if the credential list is different (MOAB-6451).
- **In job XML, Moab would ignore the `JobName`.** Moab now sets the job name from the `<JobName>` XML tag (MOAB-6535).
- **`PARCFG[] DEFAULTNODEFEATURES` would overwrite SLURM-defined features.** Moab now sets the SLURM-reported features each iteration so that they are not overwritten (MOAB-6489).
- **When TORQUE or Moab took a long time to schedule at the beginning of an iteration, the reservations would back up and Moab would refrain from scheduling the reservations.** Moab now schedules the reservations in this scenario (MOAB-6480).
- **Submitting a job with a `-l partition=ALL` would allow a credential to run the job on a partition to which it should not have access.** Moab no longer makes this error (MOAB-6379).
- **Moab would indicate a connection to MongoDB even when authentication failed.** Moab now reports the Mongo connection as down when authentication fails (MOAB-6391).
- **Moab would not run a job as root in SLURM with `ALLOWROOTJOBS` set to `TRUE`.** Moab now successfully runs these jobs (MOAB-6449).
- **Reservations with a description spanning multiple lines would disappear from the system after a Moab restart.** Moab preserves these reservations when it restarts (MOAB-6231).
- **The Moab green `ipmi.mon.py` file did indicate when the `node-bmc.txt` file was missing.** The IPMI scripts now handle a missing file error case (MOAB-5821).
- **In some cases, jobs were running on the wrong partition.** Moab will only run jobs on the assigned partition (MOAB-6428).
- **The IPMI scripts would generate an error if a non-IPMI node was in a node list.** The scripts no longer generate this error (MOAB-6377).
- **When a job was submitted in MWS by root with a proxy user specified, the job was assigned to root instead of the user's primary group.** MWS now looks up the group if it is not specified at job submission (MOAB-6487).

- **In a grid environment, Moab would send jobs to clusters that did not have the requested class.** Moab now sends the jobs to a cluster with the requested class (MOAB-6419).
- **Moab commands would sometimes return a "could not authenticate client using .moab.key" error.** The key authentication procedure causing the errors has been fixed (MOAB-6152).
- **The database init scripts shipped with Moab did not correctly populate the schema.** These schema problems have been resolved (MOAB-6117).
- **An error would occur when service containers contained service containers.** MWS once again supports service containers within service containers (WS-2090).
- **SLES did not handle http authentication, causing MWS to return a 401 error.** Viewpoint now sends the authorization header every time it accesses MWS (VEW-5389).
- **You could not update the Automatic Virtual Machine Migration policy for generic metrics.** Viewpoint now updates the policy for generic metrics (VEW-5671).
- **MAM could not be configured in Viewpoint.** An Accounting section has been added to `navigation.xml` to make MAM configuration possible (VEW-5700).
- **Queries did not always return aggregation information.** Queries now return aggregation information even when you specify options (GOLD-213).
- **The status bar that confirmed an account creation did not return any feedback in newer browsers.** It now provides feedback when actions are taken in all supported browsers (GOLD-222).
- **Newly-created transfer allocations had their Active field set to NULL.** Newly-created allocations from transfers now have an Active field value of True or False. Additionally, Allocation Refresh now repairs allocations with Active=NULL (GOLD-231).
- **The offset arrows in the MAM GUI List pages were large and did not function.** The size of the arrows has been reduced and their functionality restored (GOLD-238).
- The source directory of legacy GUI scripts and images were renamed from standalone to legacy so the GUI installation will not fail.
- **Nodes could be allocated to array sub jobs even after the job was deleted.** TORQUE no longer allocates to array sub jobs when the job is deleted (TRQ-2329).
- **MOMs at a version earlier than 4.2.6 could not run jobs if the server was at version 4.2.6.** The 4.2.7 server is now compatible to earlier MOMs (TRQ-2351).
- **Cray `np=X` did not restrict job counts for login nodes.** Login nodes restrict the number of jobs to the number specified by `np=X` (TRQ-2373).
- **There was a potential overflow in user job counts; in addition, a user was considered a different user if he submitted a job from a different submit host.** Both problems have been resolved (TRQ-2354).
- **Jobs were sometimes assigned to the same CPU sets after restarting `pbs_mom`.** CPU sets are no longer duplicated (TRQ-2369).
- **If a job removed with `qdel` were kept around long enough with `keep_completed` and submitted with a job wait time (`-a`), they would run.** Deleted jobs no longer run (TRQ-2377).

- **TORQUE would segfault when re-sending batch requests at restart time.** This error no longer occurs (TRQ-2347).
- **TORQUE continued to report nodes in a free state when the host had not reported more than node\_check\_state seconds (the timeout was failing).** Nodes are now marked as down when the host is down (TRQ-2270).
- **Non-Cray nodes reporting to a Cray pbs\_server received incorrect task counts assigned for running jobs.** pbs\_server now correctly distributes the task counts (TRQ-2395).
- **Reporter mom would fork to send its update, sometimes causing defunct pbs\_mom processes.** This no longer occurs (TRQ-2299).

# Product documentation

## Technical documentation

The online help for Moab HPC Suite 7.2.7 – Enterprise Edition is available in HTML and PDF format on the [Adaptive Computing Documentation page](#).