

Moab HPC Suite 7.2.9 – Basic Edition Release Notes

The release notes file contains the following sections:

- [New features on page 1](#)
- [Differences on page 1](#)
- [Installation and upgrade information on page 3](#)
- [Known issues on page 3](#)
- [Resolved issues on page 4](#)
- [Product documentation on page 9](#)

New features

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

Moab Workload Manager

7.2.9

No new features.

7.2.8

checkjob, showbf, and showstart support --blocking

The `checkjob`, `showbf`, and `showstart` commands now support the `--blocking` option, which produces real time information rather than cached information.

JOBFREETIME component of node allocation policy PRIORITY

A new component of the node allocation policy *PRIORITY* called `JOBFREETIME` specifies the number of seconds that the node will be idle between now and when the job is scheduled to start. For more information, see "[Node Allocation Policies](#)" in the Moab Workload Manager Administrator Guide.

Differences

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

- [Moab Workload Manager on page 2](#)
- [TORQUE Resource Manager on page 2](#)

Moab Workload Manager

7.2.9

MAXPROCPERNODE can be also be configured per node

Added capability to configure MAXPROCPERNODE per class and per node. Before, MAXPROCPERNODE could be configured per class but it applied to all nodes. Now you can configure something like the following:

```
CLASSCFG [cpu] MAXPROCPERNODE [n1, n2, n3]=20 MAXPROCPERNODE [n4, n5, n6]=10
```

ALERT log message

Added a missing newline character to a log statement.

7.2.8

TORQUE job ID added to event records

The TORQUE job ID now appears in the JOBSTART and JOBEND event records.

Failed erasing job alerts escalated

The Moab log alert "Failed erasing job" has a higher log level.

TORQUE Resource Manager

4.2.9

qmgr server parameter "copy_on_rerun"

A new qmgr option: set server copy_on_rerun=[True|False] is available. When set to True, Torque will copy the OU, ER files over to the user-specified directory when the qrerun command is executed (i.e a job preemption). This setting requires a pbs_server restart for the new value to take in effect. Note that the MOMs and the pbs_server must be updated to this version before setting copy_on_rerun=True will behave as expected.

qmgr server parameter "job_exclusive_on_use"

A new qmgr option: job_exclusive_on_use=[True|False] is available. When set to True, pbsnodes will report job-exclusive anytime 1 or more processors are in use. This resolves discrepancies between Moab and TORQUE node reports in cases where Moab is configured with a SINGLEJOB policy.

scan_for_terminated

Improved performance by moving scan_for_terminated to its own thread.

TORQUE accounting improvements

Two new fields were added to the accounting file for completed jobs: total_execution_slots and unique_node_count. total_execution_slots should be 20 for a job that requests nodes=2:ppn=10. unique_node_count should be the number of unique hosts the job occupied.

4.2.8

Thread pool manager and thread stack sizes modified

The thread pool manager now frees idle nodes. Additionally, the default thread stack sizes changed to a maximum of 8 MB and a minimum of 1 MB.

Exit job improvements

The way TORQUE exits jobs on moms has been improved to reduce the number of stray jobs, make epilogue run only once, and stop the occurrence of single jobs showing up on nodes they should not.

Installation and upgrade information



When installing or upgrading, it is *strongly* recommended that administrators configure Moab with mauth authentication with a complex key value. See [Mauth Authentication](#) for more information.

Installing Moab HPC Suite 7.2.9 – Basic Edition

Complete Moab HPC Suite 7.2.9 installation instructions and requirements can be found in HTML or PDF format.

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

Known issues

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

Known issues are aggregated and grouped by the release version for which they first occurred or where reported.

7.2.9/4.2.9

- Resources requested through Job submissions from MWS are doubled by Moab. To avoid this, when the **JOBNODEMATCHPOLICY** is set to **EXACTNODE** in Moab, avoid using `requirements.tasksPerNode`. Use `requirements.resourcesPerTask.processors.dedicated` instead. (MOAB-7424)
- **BACKFILLPOLICY BESTFIT** does not support multi-req jobs. Only **FIRSTFIT** supports multi-req jobs. (MOAB-6824)
- Some limitations exist in the way that pbsdsh can be used. Please note the following situations are not currently supported:

- Running multiple instances of pbsdsh concurrently within a single job. (TRQ-2851)
- Using the -o and -s options concurrently; although requesting these options together is permitted, only the output from the first node is displayed rather than output from every node in the chain. (TRQ-2690)
- Moab does not currently consider JOBNODEMATCHPOLICY when it handles preemption and results in unexpected behavior. This issue has been known for earlier versions (was not introduced with 7.2.9). Updated "Canceling jobs with preemption" and "Testing and troubleshooting preemption" in the *Moab Workload Manager Administrator Guide* to specify that JOBNODEMATCHPOLICY should *not* be set when enabling preemption. (DOC-2161)
- Jobs submitted with invalid credentials are put in a held state, instead of rejected, until the administrator can respond. The checkjob command gives administrators further information regarding why the job is held. Blindly assuming that all held jobs should in fact be running RIGHT NOW is not only unsafe, but circumvents intentional Moab policies and workflow. An administrator should exercise care when resolving held jobs. (CVE-2014-5375, MOAB-7478, MOAB-7526)
- When installing or upgrading, it is *strongly* recommended that administrators configure Moab with mauth authentication with a complex key value. See [Mauth Authentication](#) for more information. (CVE-2014-5376, MOAB-7525, MOAB-7480)

7.2.8/4.2.8

- 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 – Basic Edition. Following each issue description is an associated issue number in parentheses.

Resolved issues are aggregated and grouped by the release version in which they were resolved.

7.2.9/4.2.9

- **Moab was still blocking on *checkjob* and *showres*.** Added *showres* to the list of commands available when UIMANAGEMENTPOLICY is set to FORK. (MOAB-7233)
- **Moab Authentication Bypass issue.** This has been fixed. (CVE-2014-5300, MOAB-7100, MOAB-7524)
- **ALERT log message did not end in a new line.** Added a missing new line character to a log message. (MOAB-7329)
- **XML job log error.** Corrected mismatched <Job_Id> XML tags in the job log. (TRQ-2692)
- **Moab was not decrementing GRES correctly within a Cray environment.** Fixed a GRES bug found in Cray environments. (MOAB-7213)
- **Moab changes task count on jobs on restart.** Fixed bug around incorrect task counts on multi-req jobs across a Moab restart in grid environment. (MOAB-7104)
- **TORQUE accounting problems - jobs without accounting records.** Fixed bug related to accounting records on large systems. (TRQ-2367)
- **GRESTOJOBATTR entry in moab.cfg breaks MAXGRES set in IDCFG.** Fixed bug where MAXGRES ignored when reading IDCFG information from file and GRESTOJOBATTR set in moab.cfg. (MOAB-7078)
- **FLAGS=SharedMem and MEMWEIGHT did not work together.** Fixed bug where MEMWEIGHT not applied when FLAGS=SharedMem was configured. (MOAB-7033)
- **TORQUE was leaving behind error and out files when a job was preempted or requeued.** Fixed bug where OU files were being left in spool when job was preempted or requeued. (TRQ-2732)
- **Viewpoint was not unsetting Node Allocation Priority Policy.** Fixed bug where Viewpoint wasn't correctly unsetting Node Allocation Priority Policy. (VEW-5721)
- **momctl -q clearmsg didn't seem to clear error messages permanently.** Fixed bug where `momctl -q clearmsg` didn't properly clear error messages. (TRQ-2828)
- **Moab XML was adding incorrect information.** Fixed bug where incorrect data was being added to job XML. (MOAB-7108)
- **max_user_queuable limit reached, however, there were no jobs in the queue.** Fixed bug where jobs rejected due to max_user_queuable limit reached, yet no jobs in the queue. (TRQ-2795)
- **Jobs were going to wrong partition when partition(s) down.** Fixed bug where jobs were being placed in wrong partition when one or more requested partitions were down. (MOAB-7103)
- **Preemptor Jobs were not getting started correctly.** Fixed bug where preemptor jobs were not getting started correctly. (MOAB-7205)
- **qsub did not process arguments correctly when a submitfilter is used.** Fixed bug where qsub did not process args correctly when using a submit filter. (TRQ-2646)
- **Reported cput was incorrect.** Fixed bug where reported cput was incorrect. (TRQ-2759)
- **Jobs were not being rejected even when the jobs violate fairshare tree policy.** Fixed bug where some jobs were not rejected even though they violated fairshare policy. (MOAB-7077)

- **Large environment variables (>1024 chars) crashed interactive jobs.** Fixed bug where very large environment variables caused failures in interactive jobs. (MOAB-7021)
- **Moab was losing limit set with changeparam.** Fixed bug with changeparam and multidimensional policies. (MOAB-7234)
- **mschedctl -m config "IDCFG[accounts] REFRESHPERIOD=minute|hour|infinity" did not change setting in Moab.** Fixed bug with mschedctl -m config related to IDCFG and REFRESHPERIOD. (MOAB-6887)
- **Moab reservation corruption was blocking jobs.** Fixed bug with reservations across moab restarts. (MOAB-7140)
- **qsub was passing arguments to submit filter all together with interactive jobs.** Fixed bug with submit filter arguments not being parsed during interactive jobs. (TRQ-2707)
- **TORQUE appeared to be truncating environment variables when the node list is larger (such as procs=1952).** Fixed bug with truncated job environment variables in cases where node lists are very long. (MOAB-7136)
- **Build bug reported with MIC libraries.** Fixed build bug related to newer Intel MIC libraries installing in different locations. (TRQ-2653)
- **TORQUE was not dividing GPUS amongst NUMA nodes.** Fixed problem where GPUs were not split between NUMA nodes. You now need to specify which gpus belong to each node board in the mom.layout file. (TRQ-2730)

A sample mom.layout file might look like

```
nodes=0 gpu=0
nodes=1 gpu=1
```

 This only works if you use nvml. The nvidia-smi command is not supported.

- **Performance bug found in ReportJob.groovy.** Fixed scalability issue with large sample sets (e.g., node utilization). (WS-2273)
- **Crash on start up when reading empty array file.** Fixed start up bug related to empty job array (.AR) files. (TRQ-2787)
- **Long startup time reported with 7.2.8.** Improved performance during startup. (MOAB-7243)
- **qstat wouldn't parse anything after a bad job ID.** Improved qstat behavior in cases where bad job IDs were referenced in the command. (TRQ-2410)
- **pbsdsh required FQDN even if other elements didn't.** pbsdsh no longer requires FQDN. (TRQ-2632)
- **multireq jobs take hours to start.** Fix bug where multi-req jobs were slow to start in certain cases. (MOAB-6824).
- **SRCFG with ROLLBACKOFFSET not creating reservations.** Fixed bug where SRCFG with ROLLBACKOFFSET was not creating expected reservations (MOAB-7148)

- **Node going down briefly caused node conflict with reservations.** Fixed bug related to standing reservation conflicts over a node reported as Down. (MOAB-7378)
- **Reservation was not being created when some resources are available.** Fixed an issue with standing reservations that occasionally weren't created when they should be. (MOAB-7384)
- **Separate headers for multiple jobs IDs provided to qstat.** Fixed output format bug in cases where multiple job IDs are passed into qstat. (TRQ-2411)
- **MWS should allow job script inside the POST body.** Job scripts can now be included in the POST body when submitting jobs via MWS. (WS-2112)

7.2.8/4.2.8

- **Moab did not recognize users' default groups configured in /etc/groups.** Moab now recognizes default groups. (MOAB-6522)
- **Moab had several memory leaks.** Moab no longer has these memory leaks, use of uninit variable, and double free. (MOAB-6811)
- **In a peer-to-peer grid, an issue with hop counts caused a node to disappear.** The cluster query race condition on hop counts has been resolved. (MOAB-7034)
- **In an HA environment where the secondary server had taken over, Moab started writing to an old log file.** Moab now writes to the correct log file in this scenario.
- **In a Cray environment, submitting a job would fail when you specified a range in mppnodes.** Moab again supports specifying a range in mppnodes. (MOAB-6808)
- **In a Cray environment, specifying mpp* parameters in qsub caused TORQUE and Moab to ignore feature requests.** Feature requests now work correctly when you use mpp* parameters with qsub. (MOAB-7004)
- **Modifying a reservation start time did not cause relevant jobs to change start times and the mdiag -r output did not immediately reflect the update.** Moab now modifies job start time when its reservation has been modified and the mdiag -r output displays it immediately. (MOAB-6861, MOAB-6933)
- **Moab would override the partition specified in a job template with the default partition.** Moab now only sets the default partition on a job template if you do not specify one. (MOAB-6898)
- **Moab did not consistently evaluate secondary groups in environments where the scheduler flag EXTENDEDGROUPSUPPORT was set.** Moab now correctly evaluates the secondary groups. (MOAB-6784)
- **Moab would crash when you wrote invalid configuration parameters in mom_priv.config.** Moab no longer crashes when the file contains an invalid option. (MOAB-6905)
- **The job taskcount reverted to its old value after modification.** The taskcount no longer reverts after you run an mjobctl -m command. (MOAB-6926)
- **MAXPS accepted invalid integers.** It no longer accepts integers larger than 64 bit. (MOAB-6925)
- **Moab incorrectly allocated reservations with a superset hostlist.** Moab now correctly allocates these reservations. (MOAB-6714)

- **The mstat_converter README file incorrectly documented the moab.cfg location.** The mstat_converter README has been updated to reflect the correct location of `$MOABHOMEDIR/etc/moab.cfg`. (MOAB-6627, MOAB-6135)
- **Moab took progressively more time to process and load jobs that contained dependencies.** The speed of loading dependency jobs has increased substantially. (MOAB-6900)
- **Moab iterations extended when it processed node features.** The Moab process for checking features has been optimized to improve iteration processing time. (MOAB-6889)
- **Running mschedctl-R caused array jobs to break.** Recycling Moab no longer breaks non-running sub jobs. (MOAB-6836)
- **Moab would create standing reservation instances in the wrong order.** Moab now creates standing reservations in order according to start time. (MOAB-6104)
- **When you added new nodes to Moab after creating a reservation with the -t ALL flag, Moab did not automatically add them to the reservation.** Moab automatically adds new nodes to reservations configured with `-t ALL`. (MOAB-6556)
- **Moab did not close database connections on restart.** Database connections close when you run `mschedctl -R`. (MOAB-6160)
- **Moab would not honor the class hostlist when acl_host_enable was set.** Moab applies the class hostlist under these conditions. (MOAB-6492)
- **When Moab scheduled multi-req jobs with EXACTNODE set, an incorrect node count caused job starvation.** The node count has been corrected and job starvation no longer occurs. (MOAB-6670)
- **In a SLURM environment, Moab did not consider the ppn setting when the NODEACCESSPOLICY was SINGLEJOB.** Moab now considers ppn as it allocates jobs in a `SINGLEJOB` environment. (MOAB-6580)
- **Moab crashed when it restarted with a job running and the MAXNODE limit had been exceeded.** Moab no longer crashes under these conditions. (MOAB-6607)
- **Moab crashed due to an undefined buffer in a vsnprintf call.** This crash no longer occurs. (MOAB-6726)
- **TORQUE would track the wrong number of processors.** The issue where the total number of execution slots having a count off by one for every Cray compute node has been fixed. (TRQ-2501)
- **A memory leak occurred on asynchronous qrun -a commands.** The memory leak no longer occurs, and a write after free error that could lead to memory corruption was also resolved (TRQ-2498).
- **For newer versions of nvidia drivers, the GPU status was only displayed at pbs_mom startup.** The GPUs now appear in the `pbsnodes` output. (TRQ-2647)

Product documentation

Technical documentation

The online help for Moab HPC Suite 7.2.9 – Basic Edition is available in HTML and PDF format on the [Adaptive Computing Documentation page](#).