

Moab HPC Suite 10.0.1 Release Notes

June 2024



Adaptive
C O M P U T I N G

Contents

1. New Features	3
Torque Resource Manager	3
Moab Workload Manager	5
Moab Accounting Manager	7
2. Differences	9
Torque Resource Manager	9
Moab Workload Manager	10
Moab Accounting Manager	12
3. Known Issues	13
Torque Resource Manager	13
Moab Workload Manager	13
Moab Accounting Manager	14
4. Resolved Issues	15
Torque Resource Manager	15
Moab Workload Manager	20
Moab Accounting Manager	24
5. Installation And Upgrade Information	25
RPM Distribution Tarball	25
Compatibility Requirements	25

1. New Features

In this topic:

- [Torque Resource Manager](#)
- [Moab Workload Manager](#)
- [Moab Accounting Manager](#)

Torque Resource Manager

7.0.1

RPMs Available for Red Hat 7, Red Hat 8, SUSE 12, and SUSE 15

In this release, RPMs are supported as an installation method with a Torque RPM distribution tarball bundle being provided for Red Hat 7, Red Hat 8, SUSE 12, and SUSE 15.

Support Added for hwloc Version 2

Torque is no longer restricted to using hwloc version 1, so it is now possible and recommended to use the hwloc packages provided by your OS vendor.

7.0.0

Support Added for Red Hat 8, SUSE 15, and Ubuntu

Support has been added for Red Hat 8 (e.g., RHEL 8, CentOS 8, Scientific Linux 8, Rocky Linux 8, Alma Linux 8), SUSE 15 (e.g. SLES 15, OpenSUSE Leap 15), and Ubuntu 18.04, 20.04, and 22.04.

Added Support for MIG Partitions on NVIDIA GPUs

Support has been added for Multi-Instance GPUs (MIGs). The partitioned MIG instances will be enumerated as real GPU instances. Due to current NVIDIA driver limitations, only one GPU should be requested per job.

Core Dumps Enabled by Default in systemd Unit Files

Core dumps for pbs_server, pbs_mom, and trqauthd crashes are now enabled by default when using systemd to start torque.

Start Count of Job Logged in Accounting Records

The number of times a job attempts to start before actually starting is now logged in server accounting records.

pbs_mom --about Displays commit hash

1. New Features

pbs_mom --about now displays the Git commit hash.

Specify Starting Privileged Port

The --with-reserved-port-start=PORT configuration option enables you to specify a port number for the lower bound of the reserved port range Torque will use when opening a reserved port.

Set Specific Memory or Swap Value per MOM

New pbs_mom configuration settings enable you to set specific memory or swap limits per MOM. \$max_physical_memory sets the amount of memory available to jobs on a node and \$max_swap_memory sets the swap memory limit.

Specify Port Range for Interactive Jobs

You can force interactive jobs to listen on a specified range of ports by setting INTERACTIVE_PORT_RANGE in torque.cfg.

Additional Information Shown in qstat -f Regarding Slot Limits

Slot limit for array jobs is now shown in 'qstat -f' output under the job_array_request attribute. Hold type due to a slot limit for an array job is now shown under the Hold_Types attribute ('l' hold type means the hold is due to a slot limit).

Added the -L request to the Accounting Log

Added the -L request, when made, to the accounting log.

Added -logfile Option to trqauthd

Added --logfile=<filepath> to trqauthd. If filepath is left blank it will default to /var/spool/torque/client_logs. Must set the file path each time trqauthd is run.

Additional Options in torque.setup

torque.setup script now permits specification of a config path (default is the path used when configuring Torque). New options have also been added to specify a prefix path (path to executables is in PREFIXPATH/bin and PREFIXPATH/sbin) and a new option has been added to log errors to syslog.

Improvements in torque.setup for Multi-Homed Systems

torque.setup now adds a queue operator and queue administrator entry for each hostname on a system with multiple network interfaces.

Log Start Count in the Accounting Records

Accounting log now reports the start count of a job.

Output and Error Paths Included in Abort Email

When walltime is hit on a job, it should now properly display the output and error paths if any.

momctl -u Option to Force Hardware Configuration Update

Add momctl -u to get the mom to force an update of its layout on pbs_server even if it wasn't started with -f.

Permit Specified Exit Code for All Canceled Jobs

Add the ability to set a parameter exit_code_canceled_job to force all canceled jobs to have the same exit code regardless of the state they were in when they were canceled.

Give jobstart and jobend Options of the \$node_check_interval Their Own Parameters

Adds the \$node_check_on_job_start and \$node_check_on_job_end variables to the pbs config as a replacement for the same behavior provided by the \$node_check_interval variable. The \$node_check_interval variable still behaves exactly the same, but the new variables should be used when wanting the check script to be run on job start or job end.

Add --xml Flag for Client Commands That Support -x

Add --xml synonym for -x to qstat and pbsnodes.

Add Node IP Address to momctl -d Output

Added IP address to momctl -d output.

Moab Workload Manager

10.0.1

RPMs Available for Red Hat 7, Red Hat 8, SUSE 12, and SUSE 15

RPMs are supported as an installation method in this release with a Moab HPC Basic RPM distribution tarball made available for Red Hat 7, Red Hat 8, SUSE 12, and SUSE 15.

MAXGREENSTANDBYPOOLSIZE Allows a Standby Pool Size of 0

A value of 0 for MAXGREENSTANDBYPOOLSIZE will now cause Moab to keep all nodes powered off until workload requires them to be powered on. A value of -1 disables the standby pool (leaving all nodes powered on by default).

10.0.0

Support for Red Hat 8, SUSE 15, and Ubuntu

Support has been added for Red Hat 8, SUSE 15, and Ubuntu 18.04, 20.04 and 22.04.

Allow Multiple Primary Objects (jobs, nodes) in checkjob, checknode

checkjob, checknode will now allow a comma-separated list of jobs, nodes to be specified.

Job Template Priority Can Be Set per User

The JOBCFG[] UPRIORITY attribute can be used to set a per-user job template priority.

New DEFAULTSTARTTIMEQUERY Parameter

A new DEFAULTSTARTTIMEQUERY parameter has been added that specifies the default start time estimation method to be used with the showstart command (and mjobctl -q starttime) when not otherwise specified with the command options. It can have values of HISTORICAL, PRIORITY or RESERVATION and defaults to PRIORITY.

Disable Memory Enforcement in RESOURCELIMITPOLICY

You can disable memory enforcement for a specific user (e.g., USERCFG[doug] FLAGS=DisableMemEnforcement).

Added FSDISABLEIFNEGATIVE Parameter

If set to TRUE, Moab will disable any fairshare components that are negative when determining a job's priority.

Optional Separate Client Configuration File

Clients may read client configuration parameters from moab-client.cfg, if present. Client configuration parameters in moab.cfg (if present) will take precedence over those in moab-client.cfg. The primary motivation for this feature is to allow RPMs to own separate configuration files in the client and server packages.

Ability to Bypass Memory Utilization Enforcement for Specific Users

RESOURCELIMITPOLICY memory utilization enforcement can be disabled on a per-user basis.

Reprovision of Nodes Regardless of OS

A new msub 'FORCEPROVISION' flag is added to override the default behavior when provisioning is enabled. When specified, a job will provision nodes whether or not they already have the requested OS.

Added Support for msub -W group_list=X

To submit a job as a secondary group to Torque, use the group_list=g_list option of the qsub -W command.

View Details for Running and Queued Jobs

showq -N shows node/task allocation for running and queued jobs.

Query Job Information for a Job in a Job Group

Information for a job in a job group can be obtained using `showq -w jobgroup=<jobgroup>` or `mdiag -j -w jobgroup=<jobgroup>`.

Users Can Set a Default Queue

Users can specify a default job queue. `msub -q` obtains the default queue from the `MOAB_DEFAULTQUEUE` environment variable.

Allow Asynchronous Requeue of a Job

Using `mjobctl -R --flags=force` forces an asynchronous requeue on Torque systems.

Options for Specifying Units for LOGFILEMAXSIZE

`LOGFILEMAXSIZE` can be specified in MB or GB.

Added Node Count per User Option

`showq -u <username>` now displays a node count for the user.

Check for Circular Dependencies No Longer Default

Moab formerly checked job dependencies for circular dependencies. This check is now only performed if the `SCHEDCFG CHECKCIRCULARDEPENDENCIES` flag is set.

Moab Accounting Manager

10.0.1

RPMs Available for Red Hat 7, Red Hat 8, SUSE 12, and SUSE 15

RPMs are supported as an installation method for Red Hat 7, Red Hat 8, SUSE 12, and SUSE 15 via a Moab HPC Suite RPM bundle.

10.0.0

Support Added for Red Hat 8, SUSE 15, and Ubuntu

Support has been added for Red Hat 8 (e.g., RHEL 8, CentOS 8, Scientific Linux 8, Rocky Linux 8, Alma Linux 8), SUSE 15 (e.g., SLES 15, OpenSUSE Leap 15), and Ubuntu 18.04, 20.04, and 22.04.

Added Support for PAM Authentication for Web Services and the Web UI

Enabled support for the SSSRMAP Cleartext token type, which will be used with PAM authentication to pass the username and password to the server. This mechanism should only be used under SSL. A new `-c (cleartext)` option was added to `mam-auth` to extract the server connection attributes without creating a signature or requiring the request body as input. Added a new `authentication.method` parameter in the gui and web services configuration files. This parameter specifies which authentication mechanism to use when

basic auth is used to prompt for a user name and password. It can assume a value of 'pam', which uses the Cleartext token type to send the username and password to PAM (Pluggable Authentication Module) for authentication, or a value of 'mam-password', which uses the Password token type to compare a hashed value of the password with the value in the MAM database. The default value of this parameter depends on the value of the --with-pam configure option described in the next section.

New --with-pam Configuration Option

A new --with-pam configure option has been added to indicate whether to enable PAM authentication or not. When PAM is asserted (e.g., by specifying --with-pam), the PAM config file will be installed and PAM becomes the default authentication method for the GUI and web services. When PAM is disasserted (e.g., by specifying --without-pam), the PAM config file is not installed and the mam-password authentication method is used for the GUI and web services. If the --with-pam option is not specified, the PAM config file is installed but the authentication method for the GUI and web services defaults to mam-password. PAM authentication will not work without the PAM config file being installed. When PAM authentication is configured to use Unix password authentication, MAM will need to be running as root in order to perform authentication for the users. Thus, when --with-pam is specified, the accounting admin user (normally mam) will default to root, unless overridden by the --with-user configure option.

Extended mam-read-configuration to Parse the Web Services Configuration File

The mam-read-configuration client command now parses and displays configuration from mam-ws.conf.

Migration Scripts Also Now Adjust the Configuration Files

The migration scripts have been augmented to transform the configuration files as well as the database schema. Formerly, when upgrading to recent release versions, the make install step would copy the newly distributed configuration files alongside the current configuration files. The customer was responsible for merging changes into the configuration files. As a result of this change, when upgrading, it should no longer be necessary to tweak the config files manually by following pointers in the release notes.

2. Differences

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

In this topic:

- [Torque Resource Manager](#)
- [Moab Workload Manager](#)
- [Moab Accounting Manager](#)

Torque Resource Manager

7.0.1

None.

7.0.0

Support for Red Hat 8, SUSE 15, and Ubuntu

Support has been added for Red Hat 8, SUSE 15, and Ubuntu 18.04, 20.04 and 22.04. Red Hat 6 and SUSE 11 are not supported in version 7.0.

Outbound Sockets Bound to Daemon Host Address

trqauthd, pbs_server and pbs_mom now bind outbound sockets to the host address where the daemon is running. This ensures that, if a host contains multiple network interfaces, only the one associated with the host name is used. If this behavior is not desired, it may be disabled by rebuilding Torque using the `--disable-bind-outbound-sockets` switch with `configure`. See also [3. Known Issues](#).

pbsnodes, qstat, and qsub Output to stdout

The `--version` options of `pbsnodes`, `qstat`, and `qsub` now display to `stdout`, rather than `stderr`.

make rpm Uses the Configured Prefix

When creating an RPM using `make rpm`, the RPM will now default to using the configured prefix. Since `configure` defaults to using a prefix of `/usr/local`, `make rpm` will now likewise default to using a prefix of `/usr/local` rather than `/usr` for the resulting RPM. When creating an RPM, the configured prefix can be overridden by using the `RPMOPT` `make` variable (e.g., `make RPMOPTS="--define '_prefix /opt/torque-latest'" rpm`).

Changed Pbsnodes Output to Use a '+' Symbol Instead of a ',' Symbol to Delimit Jobs in the Node Job List

This change overcomes a conflict occurring when ranges are not strictly sequential (i.e., include commas such as 0-1,6-7).

Moab Workload Manager

10.0.1

Default (disabled) Value for MAXGREENSTANDBYPOOLSIZE Changed to -1

The default value for MAXGREENSTANDBYPOOLSIZE (causing it to be disabled) has changed from 0 to -1. A value of 0 is now used to maintain a standby pool size of zero (all nodes will be kept powered off until needed by the current workload).

10.0.0

Slurm is no Longer Supported

Slurm is no longer supported as a resource manager for Moab.

Support for Red Hat 8, SUSE 15, and Ubuntu

Support has been added for Red Hat 8, SUSE 15, and Ubuntu 18.04, 20.04 and 22.04.

Red Hat 6 and SUSE 11 are no longer supported in version 10.0.0.

Removed Obsolete VM-related Clients and Documentation

The mvmctl client command is obsolete and has been removed. This includes the removal of obsolete VM-related parameters and features such as generic events for VM detection, VM-related node attributes (e.g., FLAGS, OVERCOMMIT, VMOCTHRESHOLD), mshow -a vmusage attribute, VM triggers, per-partition settings (VMCREATEDURATION, VMDELETEDURATION, VMMIGRATEDURATION), NOVMMIGRATE Job flag, Logging events (VMCREATE, VMDESTROY, VMMIGRATE, VMPOWEROFF, VMPOWERON).

Credential Diagnostic Commands Improved to Show More Information in Verbose Mode and xml Format

Many credential attributes and limits that were not previously displayed in mdiag - {a,c,g,q,u} can now be displayed in both human-readable and xml formats. Extended credential information (previously displayed on separate additional lines) are now only displayed when using the verbose (-v) option.

Changed Disk-Related Values to Use long Instead of int

All disk resource values, metrics and statistics have been promoted to use a data type of long (8 bytes) rather than int (4 bytes) in order to support systems with several petabytes or more of total disk storage.

Updating Checkpoint File Using LIMITEDJOBSCP Parameter

The behavior of the LIMITEDJOBSCP parameter has been changed such that when it is set to a value of FALSE, Moab will update the checkpoint file and cache for all jobs each iteration.

Removed Querying of Events via Command-Line Client

The mdiag -e client command is obsolete and has been removed.

Log Level Displayed in Logging Statement

Moab log statements have been enhanced such that when there is no relevant event code to display, the fifth field will instead show the hex value of the log level required to display that log statement.

Moab-Slurm Integration Via the sched/wiki2 Interface is Deprecated

As of Slurm version 17.02, the wiki scheduling interface used by Moab to schedule Slurm as a resource manager has been removed from Slurm. The Moab-Slurm integration via the sched/wiki2 interface is deprecated and may be removed in a future release.

LOGFILEMAXSIZE may be specified in MB or GB

The default unit for the LOGFILEMAXSIZE parameter is bytes, but it can also be specified in mb, MB, gb, or GB.

Estimated Start Time Query Now Defaults to Use the Priority Estimation Method

When querying the estimated start time for a job (e.g., using showstart or mjobctl -q starttime), the estimation method employed now defaults to priority (formerly reservation) unless specified by using the new DEFAULTSTARTTIMEQUERY parameter.

DATARM Documentation

Removed the Resource Manager Attribute DATARM.

Grid Data Management Documentation

Removed the section Grid Data Management.

mdiag -P Documentation

Removed references of the obsolete mdiag -P command.

Moab Accounting Manager

10.0.1

RPMs Changed to Install MAM as root

The RPMs were changed to install MAM as root so that PAM authentication may be used.

10.0.0

Support for Red Hat 8, SUSE 15, and Ubuntu

Support has been added for Red Hat 8, SUSE 15, and Ubuntu 18.04, 20.04, and 22.04. Red Hat 6 and SUSE 11 are not supported in version 10.0.

Renamed the security.* Configuration Parameters

The security.* configuration parameters have been renamed to allow better consistency when configuring new aspects of authentication. The security.promotion parameter has been renamed to promotion.method. The security.authentication parameter has been renamed to authentication.enabled. The security.encryption parameter has been renamed to encryption.enabled.

3. Known Issues

This topic lists known issues in Moab. Following each issue description is an associated issue number in parentheses. Known issues are aggregated and grouped by the release version for which they were first reported.

In this topic:

- [Torque Resource Manager](#)
- [Moab Workload Manager](#)
- [Moab Accounting Manager](#)

Torque Resource Manager

7.0.1

- No known issues.

7.0.0

- If you are using a multi-homed host (one that has multiple active network interfaces) either as a submit host or pbs_mom host, communication to pbs_server on a different host may not work as expected and it may be necessary to disable binding outbound sockets. For example, your submit host is on both a public network and an internal private network. Your pbs_server host is on the internal private network only. trqauthd on the submit host will not be able to connect to the server host. In this case you should rebuild Torque with the binding of outbound sockets disabled (see [2. Differences](#)). (TRQ-4072)
- Hwloc may not detect GPUs. When that occurs, they are reported on socket 0 numa node 0. (TRQ-3978)

Moab Workload Manager

10.0.1

- No known issues.

10.0.0

- When the NODEAVAILABILITYPOLICY is configured with UTILIZED:PROC, it currently behaves as though configured with COMBINED:PROC. Moab will not schedule jobs that would result in the requested processors exceeding the configured processors on the nodes. (MOAB-9122)

Moab Accounting Manager

10.0.1

No known issues.

10.0.0

No known issues.

4. Resolved Issues

Resolved issues are aggregated and grouped by the release version in which they were resolved. When applicable, each resolved issue has the associated issue number in parentheses.

In this topic:

- [Torque Resource Manager](#)
- [Moab Workload Manager](#)
- [Moab Accounting Manager](#)

Torque Resource Manager

7.0.1

- No known issues.

7.0.0

- qstat with many job arguments resulted in communication failures. (TRQ-4451)
- Fixed server segfault when processing mom update. (TRQ-4358)
- Variables specified with qsub -v are not passed to the script. (TRQ-4349)
- Fixed request memory calculation when using the ncpus resource which now factors in execution slots. (TRQ-4300)
- Close syslog before forking in pbs_mom to avoid potential syslog deadlock in child process. (TRQ-4264)
- Fixed job placement when requesting mics. (TRQ-4237)
- Correct array index problem in pbs_server when deleting array jobs. (TRQ-4223)
- Fixed a problem where pbs_server would not properly read in assigned processors following a pbs_server restart while jobs were running. (TRQ-4213)
- When looking up NVIDIA devices from the NVML library, convert the PCI bus id string returned by the library to lower case before referencing the device in sysfs. (TRQ-4212)
- Fixed a bug in pbs_server where cores were not correctly selected when free cores were not in a contiguous ordering. (TRQ-4208)

4. Resolved Issues

- Revert former behavior when using `qsub -t`. Use of duplicate array ids is again permitted at submission (for example, `'qsub -t 1-1'`). Internally, the array ids are not duplicated. In the example listed only 1 array job would be created. (TRQ-4189)
- Resolve issue where per-user server and queue queued job counts were not being properly tabulated for array jobs. (TRQ-4186)
- Properly allocate space for long property names to avoid a buffer overrun in `pbs_server`. (TRQ-4184)
- Fixed a bug that prevented a job submission when using a secondary group (`-W group_list=<grp>`) and when `VALIDATEGROUP=true` in `torque.cfg`. (TRQ-4183)
- Fixed a bug that prevented `'qsub -L tasks=N'` from working correctly. (TRQ-4179)
- Users were able to delete jobs without proper permissions. (TRQ-4161)
- Fixed an issue where server was failing to properly maintain the per user job count for a queue and for the server when a job was submitted with a failed dependency. (TRQ-4138)
- Avoid deadlock in server related to orphaned reservation handling. (TRQ-4128)
- Avoid `pbs_server` deadlock when checking whether or not job is finished. (TRQ-4115)
- Permit non-integer memory values when using `qsub -L` and the code now conforms to the documentation where it says that petabytes and exabytes can be specified. (TRQ-4097)
- When multiple `qsub -l` options were specified for the same resource, the last one given was the one sent to the server, as in previous versions of Torque. (TRQ-4085)
- Bind client-side outbound sockets to public-facing IP address so that one common interface is used for both incoming and outgoing network requests. (TRQ-4072)
- Allow launching of more than one GUI application in an interactive X-forwarding session. (TRQ-4070)
- Add logic to array processing to avoid a `pbs_server` crash. (TRQ-4065)
- `pbs_server` crashed when searching the environment (TRQ-4036)
- `pbs_server` crashed when a KNL node booted using `cgroups` and had memory-only nodes. (TRQ-4035)
- Submitting a job that specified a GPU mode caused the GPU mode of a GPU not assigned to the job to have its mode changed. (TRQ-4017)
- Fixed an issue with job names when `hostname:port` was specified in the `server_name` file. (TRQ-4015)
- Fixed a client-side segfault when using Torque's `drmaa` library. (TRQ-4003)

4. Resolved Issues

- Configure build script now prevents incompatible `--enable-cgroups` and `--enable-geometry-requests` from being specified concurrently. (TRQ-4000)
- Fixed a `pbs_server` deadlock issue when array task ids are not supplied in increasing order. (TRQ-3994)
- Properly handle spaces in job output files when delivering them at job end. (TRQ-3992)
- Made it so that the node health checker does not inherit `pbs_mom`'s file descriptors. (TRQ-3991)
- `pbs_server` did not preserve GPU counts across restarts. (TRQ-3989)
- `make rpm` failed with installed (but unpackaged) file(s) error. (TRQ-3988)
- Work around an `hwloc` bug (`hwloc` fails to detect all NVIDIA devices on some systems) by instead using the NVML API to look up the relevant information about the devices. (TRQ-3978)
- Make `pbs_server` automatically take a new layout when the number of GPUs has increased. (TRQ-3965)
- Prevent a MOM from running the epilogue unless it has run/tried to run the prologue. (TRQ-3957)
- Job dependencies were not evaluated correctly for jobs in a routing queue. (TRQ-3951)
- Torque could potentially hang when a job obtained a dependency on itself. (TRQ-3950)
- Reject jobs with unknown resources even when they are submitted to routing queues. (TRQ-3944)
- Fixed a crash around a mistyped array submission. (TRQ-3941)
- Fixed a potential infinite loop for moms running multi-node jobs. (TRQ-3940)
- Add protections against a segfault around the resource plugin. (TRQ-3937)
- `qsub` did not correctly translate memory requirements between `-l` and `-L`. (TRQ-3934)
- Fixed a crash around a mistyped array submission. (TRQ-3933)
- Queueing two arrays and applying `max_user_queueable` caused the arrays to drop elements and not complete. (TRQ-3932)
- When `VALIDATEPATH` unset (or set to `FALSE`) in `torque.cfg`, `qsub` now properly does not validate the `-w` (working directory) path. (TRQ-3929)

4. Resolved Issues

- Changed pbs_mom systemd unit file shutdown action to attempt a graceful shutdown before killing the daemon. (TRQ-3927)
- Fixed a regression which disallowed job scripts larger than 64k. (TRQ-3923)
- A deprecated GPU mode was not handled properly when configuring NVIDIA support and building with CUDA 8 libraries. (TRQ-3913)
- qsub could crash if the port on which an interactive job was listening was reset. (TRQ-3911)
- Fixed potential for pbs_server deadlock when many successive qmgr commands operating on the same queue are executed. (TRQ-3910)
- pbs_mom ignored time specified in \$varattr and updated dynamic values every 45 seconds. (TRQ-3901)
- Fixed an issue with jobs not moving through the routing queue in a FIFO fashion. (TRQ-3890)
- Fixed a rare crash around array subjobs terminating. (TRQ-3886)
- The default_gpu_mode did not take effect after a restart. (TRQ-3885)
- pbs_server reported job aborts when jobs finished normally. (TRQ-3882)
- MOM daemons were preventing GPUs from being reset. (TRQ-3881)
- Fixed a problem with the req_information about running the job not getting cleared on job failures. (TRQ-3873)
- Fixed a precedence problem between the submit filter and the command line for interactive jobs. (TRQ-3862)
- The /etc/profile.d/torque.csh script caused errors when installing from RPM files. (TRQ-3852)
- select() system calls replaced with poll() to be more adaptable to systems with a larger open file descriptor limit. (TRQ-3847)
- Fixed an error in detecting the amount of disk space on a system. (TRQ-3832)
- Made queues not apply conflicting defaults, such as when a job requests nodes and the queue has a procs default value. (TRQ-3830)
- Make a job's cgroups limit RSS memory. (TRQ-3827)
- Verify that the time string set for a resource is correct. (TRQ-3825)
- Added new sendmail_path server attribute that can be used to point to a sendmail executable. Default path to sendmail for RPM installations is now /usr/sbin. (TRQ-3807)

4. Resolved Issues

- Kill all processes in a cpuset before removing a cgroup. (TRQ-3753)
- Set GPU modes for -l procs jobs as well as -l nodes jobs. (TRQ-3213)
- Output and error paths were not displayed properly when a job was canceled for exceeding its walltime. (TRQ-3793)
- Fixed a problem with stdout not being properly delivered on an NFS file system when Torque was configured with --disable-spool. (TRQ-3792)
- NVIDIA device feature query failures were logged as errors if the feature was not present on the device. (TRQ-3766)
- Fixed an error with job dependencies and job_suffix_alias. (TRQ-3735)
- make rpm did not use the configured prefix. (TRQ-3716)
- Fixed the code where unsatisfied array dependencies leave held jobs forever. (TRQ-3685)
- qstat did not correctly list the job output path when qsub -k was used. (TRQ-3634)
- Enforce resources_max.procs correctly. (TRQ-3627)
- Make qsub fail if the job's working directory is invalid at submission time. (TRQ-3609)
- Sister nodes did not clean up job files when the jobs were no longer recognized by the mother superior. (TRQ-3524)
- The \$pbsclient setting previously did nothing. Now, a client can be specified using that variable which will be allowed to run momctl remotely. (TRQ-3512)
- Configure script did not search /usr/bin when searching for xauth when a path was not explicitly provided. (TRQ-3489)
- Configure script showed pbs_iff as the authentication type, rather than trqauthd. (TRQ-3413)
- Interactive jobs were not notified when a job was canceled while waiting to start. (TRQ-3031)
- Fixed a condition that was incorrectly producing a "Bad UID for job execution" message in the pbs_mom logs. (TRQ-2978)

Moab Workload Manager

10.0.1.1

- Moab is not authorized to update the Moab MongoDB database on Red Hat 7-based systems. (MOAB-9337)

10.0.1

- Processor id is mistakenly used as the task count when used with Torque 7.0 and higher.
- mdiag -q diag can crash Moab under some circumstances.
- Moab does not create multiple reservations with ReservationDepth > 1 when a reservation was specified for the job to run in.
- mschedctl -m config can fail intermittently for RMCFG parameter modifications.

10.0.0.1

- Processor id is mistakenly used as the task count when used with Torque 7.0 and higher.

10.0.0

- showbf -m does not reflect correct resource availability windows. (MOAB-9300)
- mdiag -t --xml portrayed up resources as utilized resources (*.utl). (MOAB-5727)
- When a job includes a period in its name (such as when a template dependency is automatically created), its short name can mistakenly be resolved to be the same as the base job causing unanticipated results (such as the base job completing without running). (QA-1188)
- When multiple end events are recorded for a reservation (e.g., when allocated resources change), incorrect end times and [CT][BR]PS values can be reported. (MOAB-9292)
- Fixed mdiag -t to permit the specification of a partition name and correctly filter the output by it. (QA-1180)
- Memory leaks could occur when a database and node profiling are configured. (MOAB-9297)
- Using mshow -a to request a transaction in the future provides near-term transactions instead when intervening reservations are present. (MOAB-9287)
- Many credential attributes and limits were not displayed via mdiag -\{a,c,g,q,u\}.

(MOAB-9290)

- `mrsvctl --format=xml` does not always display successful responses in xml. (MOAB-9281)
- `msub -F` fails if argument starts with a hyphen. (MOAB-9276)
- Moab will fail to start when using systemd unit files with a customized Moab home directory. (MOAB-9260)
- When submitting a job using `msub` with the job script specified as a command line argument, lines longer than 1024 characters could be corrupted. (MOAB-9258)
- When a job ends that has required a single use reservation by reservation group with multiple members, the overlapping reservation may not be the one removed. (MOAB-9254)
- Specifying the `procspecified` flag as the `jobflags` resource manager extension attribute value via XML in a job submission (e.g., by MWS/Viewpoint) resulted in the job's requested processor count being set as the node count if the job node match policy was set to `ExactNode`. (MOAB-9241)
- Non-blocking `showq` showed held vacated jobs in the blocked jobs section. (MOAB-9229)
- A deadlock could occur if Moab was processing a logging statement when a `SIGUSR1` or `SIGUSR2` signal was received. (MOAB-9228)
- The estimated node count was not being calculated correctly when applying the `MAXNODE` throttling policy. (MOAB-9226)
- Moab would sometimes miscalculate the `pmem` value sent to Torque (passing it a value of per-task memory rather than per-process memory). (MOAB-9224)
- When applying a job template that changes the per-task processors for the job, Moab was not sending the updated per-task processor information to Torque when the job is migrated. (MOAB-9224)
- Some job dependencies were not being marked as satisfied if others had already been determined to be unsatisfiable in a previous iteration. (MOAB-9216)
- Job create triggers do not fire if at least one resource manager extension value is not specified in the job submission arguments. (MOAB-9206)
- The job efficiency and processor seconds utilized statistics were overcalculated. (MOAB-9191)
- An innocuous parser error could have been emitted from `msub` if the job script contained a less-than sign. (MOAB-9186)
- Available memory could be mistakenly set to configured memory if the `UsePhysicalMemory` flag is being used and no memory is available. (MOAB-9173)

4. Resolved Issues

- Moab could create duplicate liens in MAM when using a non-blocking start stage, only one of which was removed at the end of the job. (MOAB-9159)
- Moab crashed on startup on some job array checkpoint files. (MOAB-9144)
- Jobs sometimes failed to run in a grid, citing 'partition <partition_name> does not exist' due to the total disk value for the partition overflowing its int (4-byte) value limit. (MOAB-9121)
- The first node in an job's allocated node list was sometimes remapped to a required hostlist after Moab was restarted. (MOAB-9092)
- Submitting a job exceeding the maximum permitted number of feature requirements caused Moab to crash. (MOAB-9076)
- showhist command sometimes reported incorrect value for Node Count (MOAB-9074)
- Segmentation fault in VerifyTorqueArrayStatus. (MOAB-9073)
- A job submission timeout can lead to a defer loop if SYNCJOBID and a JUSTINTIME migration policy are enabled. (MOAB-9072)
- When both per-task and per-job requested memory were specified in a job, Moab did not enforce the least restrictive request, as prescribed by the Torque documentation. (MOAB-9051)
- Jobs were assigned to nodes with insufficient memory. (MOAB-9028)
- Job gets reservation for fewer processors than requested and won't run. (MOAB-9025)
- Infinite loop on startup with mrsvctl -c in init script logging many 'cannot remove node' lines. (MOAB-9015)
- Make msub -b <retry_count> work. (MOAB-9012)
- mrsvctl should reject reservation created with invalid flag/s. (MOAB-8977)
- Canceled jobs were no longer receiving JOBEND events. (MOAB-8975)
- Fixed crash in MjobAddToNL. (MOAB-8990)
- Zombie moab processes not cleaned up when SPOOLDIRKEEPTIME configured. (MOAB-8933)
- Client RPM provided moab-client.cfg but clients did not read from it. (MOAB-8931)
- mam-charge-retry.pl script failed for array jobs. (MAM-388)
- Moab handled GLOBAL node resources differently with a moab-rlm.lic license than with a legacy moab.lic license. (MOAB-8960)
- ADMINCFG[] SERVICES=setspri was not configuring correctly. (MOAB-8957)

4. Resolved Issues

- When provisioning is enabled (on KNL systems for example) by default Moab does not provision a node if the current OS is the same as what is being requested. However, a job can request Moab to provision all its nodes with the 'ForceProvision' flag. (AC-9367)
- Processes in Z (defunct) state were not being closed. (MOAB-8933)
- SCHEDCFG[] USEDATABASE was not enabled for HA installations. (MOAB-8926)
- Crash if you have a dependency based on a jobname and you have multiple jobs with same job name. (MOAB-8928)
- NODESET definition in a class overrode reservation affinity. (MOAB-8903)
- Invalid value message in mdiag -C for QOSREJECTPOLICY. (MOAB-8923)
- Abort around NomaNode Gres. (MOAB-8916)
- Unable to remove a project from a user's access list via mschedctl. (MOAB-8890)
- Moab was not creating standing reservations on all partitions by default. (MOAB-8888)
- If the node specified by the hostlist does not currently have the resources specified by the standing reservation, Moab could sometimes ignore the hostlist and apply the reservation to another node. (MOAB-8792)
- Future (DEPTH>1) liens do not charge the correct allocation. (MAM-390)
- Fixed a crash when dumping node stats. (MOAB-8897)
- Canceling a job with an invalid WHERE argument runs the base command. (MOAB-8841)
- mdiag -S --xml did not display 5-minute statistics, as it did in non-xml mode. (MOAB-6896)
- If an HA restart fails you could get two masters. (MOAB-8609)
- Effective queue duration could be miscalculated slightly at submit time infrequently affecting job priority calculations and job start order. (QA-1219)
- Jobs that need their nodes to be provisioned to a different operating system did not always run immediately after the provision job resulting in cyclic provisioning. (MOAB-9209)
- Moab was returning elastic limits for QOSs that have none in them. (MOAB-8884)

Moab Accounting Manager

10.0.1

- No known issues.

10.0.0

- mybalance fails with error message 'Global symbol "\$user" requires explicit package name' (MAM-436).
- Too many simultaneous calls to the contributed Slurm-MAM integration scripts can deplete memory and cause failures in MAM and Slurm (MAM-432).
- The lien failure message in the job was truncated when MAM was used with Slurm in certain environments. (MAM-420)
- mam-list-usagerecords did not handle aliased aggregations correctly. (MAM-416)
- The migrate_9.0-9.1.pl script did not properly handle MySQL schema upgrades. (MAM-407)
- Removed the taint switch from clients on systems that don't have suidperl installed. This is in part to circumvent a bug in some versions of Date::Manip. (MAM-387)
- mam-list-usagerecords NodeHours field was not dividing by 3600. (MAM-399)
- Decode web service url-encoding from the browser to permit condition operator expressions and other special characters in the query string. (MAM-395)
- Permit selection aliases having dashes in web services queries. (MAM-393)
- Allow a query offset to be expressed independently of a query limit. (MAM-392)

5. Installation and Upgrade Information

This section identifies information useful when installing and upgrading.

In this topic:

- [RPM Distribution Tarball](#)
- [Compatibility Requirements](#)

RPM Distribution Tarball

A Moab HPC Basic Suite RPM distribution tarball is available at version 10.0.1 for Red Hat 7, Red Hat 8, SUSE 12 and SUSE 15 that provides RPMs for Torque 7.0.1 and Moab Workload Manager 10.0.1.

Compatibility Requirements

Elastic Computing is not supported on Ubuntu.

Torque 7.0 or later is not backward compatible with Torque 6.1 or earlier. When upgrading to Torque 7.0, the server, moms, and clients must be upgraded at the same time. The job format is compatible between Torque 7.0 and previous versions of Torque. After upgrading, any queued jobs should continue to work with the new version.