Moab HPC Suite 9.1.2 Release Notes

In this topic:

- New Features
- Differences
- Installation and Upgrade Information
- Known Issues
- Resolved Issues

Document build: 03/21/2018 13:12 UTC-06

New Features

This topic contains a summary of key new features in Moab HPC Suite.

In this topic:

- General Suite
- Moab Workload Manager
- Moab Accounting Manager
- Moab Web Services
- Moab Insight
- Moab Viewpoint
- Torque Resource Manager
- Reporting Framework
- Automated Installer

General Suite

This section contains information applicable to more than one of the components in the Moab HPC Suite.

9.1.2

No new features.

9.1.1

Reporting Framework

The Reporting Framework enables administrators to collect, extract, and summarize data related to jobs, compute nodes, and users according to many different criteria. The data can be displayed in various report formats, such as tables, pie charts, line graphs, and bar graphs. Reports can be combined into dashboards that can be used to monitor system state, performance, etc. The Reporting Framework consists of three main components:

- Streaming Engine – a distributed Apache Spark-powered application that consumes the initial stream of data, calculates the pre-aggregated data and stores the data in a database.
- Reporting Framework Web Service – a standalone Java-based web service that provides an interface to the Streaming Engine
- Reporting User Interface – Viewpoint enhancements providing all the necessary screens to run reports and create dashboards and streaming pipelines.

See Reporting for more information about using the Reporting Framework in Viewpoint.

Moab Accounting Manager Integration With Slurm

Moab Accounting Manager now integrates with the Slurm resource manager without Moab. See Integrating With Slurm for more information about Moab Accounting Manager and Slurm integration.

Moab Passthrough Mode

Moab Passthrough mode enables job submission and monitoring with native resource managers, such as Slurm. In this mode, Moab scheduling is disabled and Moab acts as a mediator between Viewpoint and Slurm. Viewpoint add-ons, such as Remote Visualization, Nitro, and Reporting are supported in Passthrough mode. Currently, Slurm is the only resource manager supported by Moab Passthrough mode. Viewpoint enhancements support using Moab in Passthrough mode. See Moab Passthrough for more information about Moab Passthrough mode.

Elastic Computing Support

The Elastic Computing feature in Moab allows the Moab scheduler to take advantage of systems that can temporarily provide additional nodes (for example, to create new virtual machines or borrow physical nodes from another system) to fulfill a workload demand in a more timely manner. Enhancements to Moab and Viewpoint provide improved support for Elastic Computing features. See Elastic Computing Overview for more information about Elastic Computing.

Create NFS Automatically within Automated Installer

A shared file system is required when using the Automated Installer. A new checkbox has been added to enable the Automated Installer to set up NFS for you.

9.1.0

New Automated Installer Method

An automated installer is now available to install your Moab HPC Suite components. Using this method, the whole installation process takes roughly an hour.

> The automated installer method is not available for SUSE 11-based systems.

This method includes a user interface (web GUI) to let you easily plan your installation. For example, you can choose how many head nodes and which package to install and you will see the layout of your selected configuration. For example:

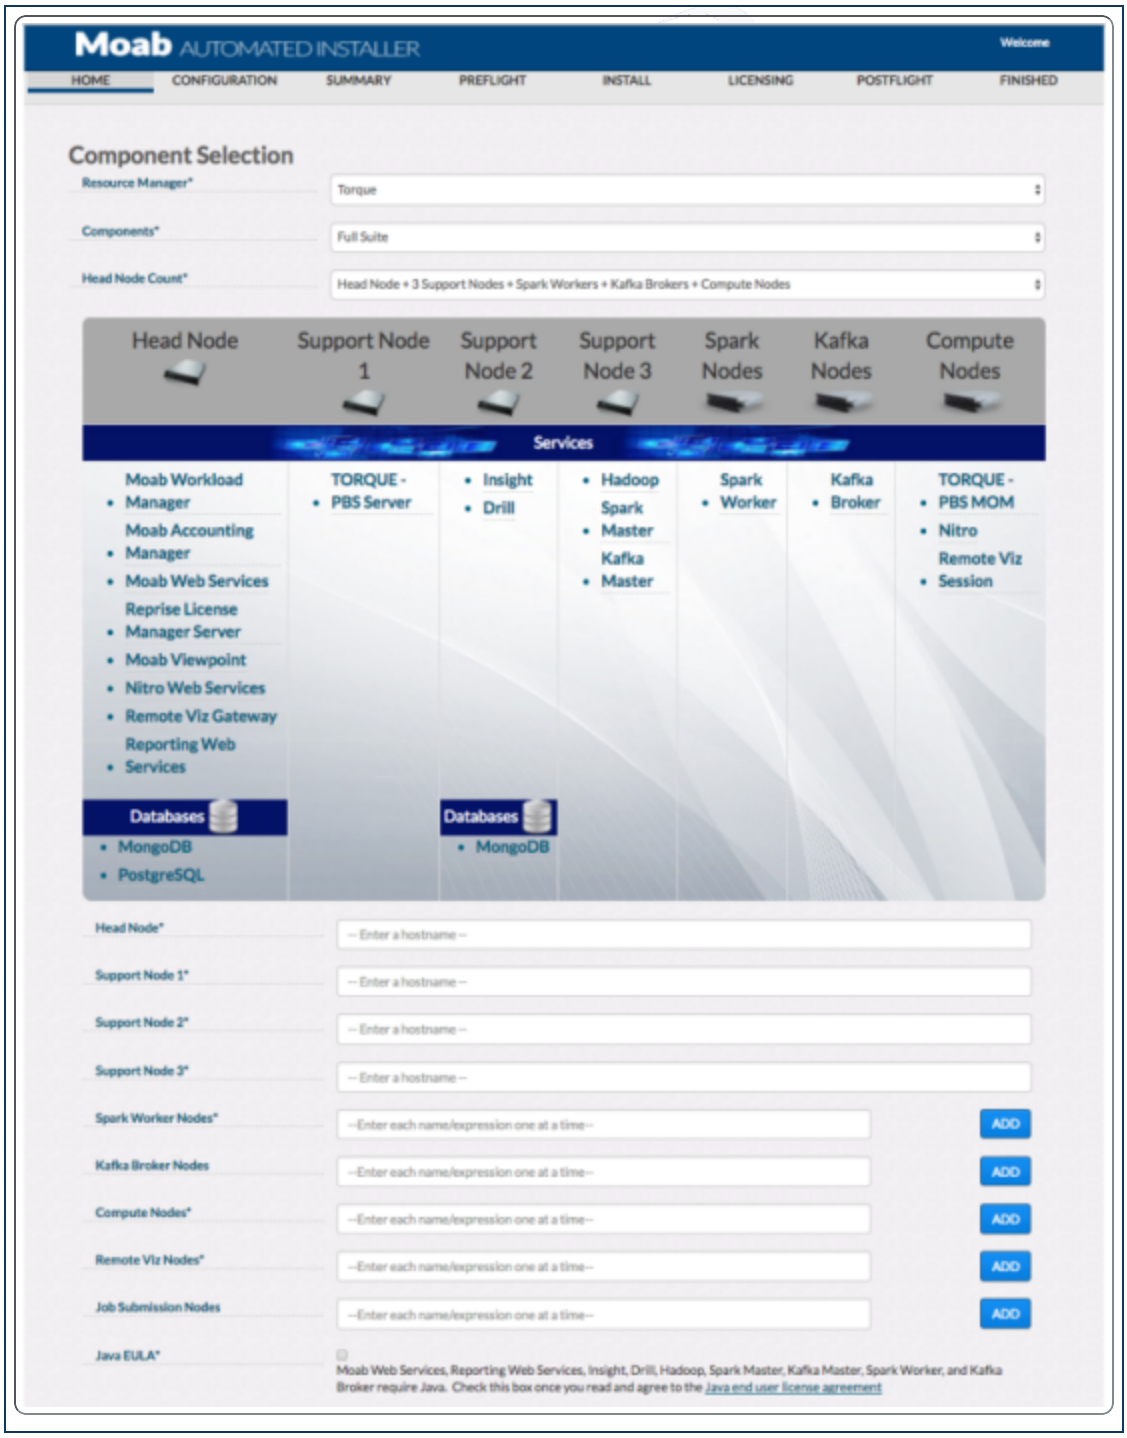

Once you have identified the configuration information (for example, the FQDN of the head node(s), the automated installer will then generate the necessary scripts and begin the installation process.

See Chapter 1 Automated installation Method for more information.

Client Authentication Now Supports SHA-512

Moab Workload Manager and Moab Web Services now support SHA-512 for message digests and message authentication codes. This support is in addition to the existing support for SHA-1.

- Moab Workload Manager has a new server checksum algorithm parameter named *SERVERCSALGO*. Valid values:
  - HMAC64: the default (SHA-1)
  - HMACSHA2: more secure (SHA-512)
- Moab Web Services has a new configuration parameter named *moab.messageDigestAlgorithm*. Valid values:
  - SHA-1: the default
  - SHA-512: more secure

If *SERVERCSALGO* is set to *HMAC64*, then *moab.messageDigestAlgorithm* must be set to *SHA-1*. Likewise, if *SERVERCSALGO* is set to *HMACSHA2*, then *moab.messageDigestAlgorithm* must be set to *SHA-512*.

NUMA Configuration for Cray KNL

Moab 9.0.2 and TORQUE 6.0.2 fully support re-provisioning Cray XC40 compute nodes based on the Intel(R) Xeon Phi(TM) Processor (formerly code-named Knights Landing or KNL) into different NUMA configurations of 1, 2, or 4 NUMA nodes and the high-bandwidth Multi-Channel DRAM (MCDRAM) memory configurations into all-cache, all-addressable, and hybrid (50/50 and 25/75) modes of operation.

Moab Workload Manager

9.1.2

Job Template Priority Can Be Set Per User

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

Ability to Bypass Memory Utilization Enforcement for Specific Users

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

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.

9.1.1

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.

Moab Pass-Through Support

Moab passthrough support gives customers the ability to use Viewpoint and other modules provided by Adaptive Computing on top of third-party schedulers like Slurm.

Script for Slurm Credentials

A Passthrough integration script is now available to gather Slurm credential information and report it back to Moab. This script is run under a IDCFG[] in the moab.cfg file.

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.

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.

9.1.0

Licensing by Sockets/Accelerators

Moab now uses RLM licenses to provide better support for licensing by sockets and acceleratores (NVIDIA GPUS or Intel MICs). Previously, licenses were generated based on the number procs. RLM licenses enable Adaptive Computing to generate licenses in better keeping with Moab's business model.

The new Moab "core" license will have a new name to reflect the RLM generation. Do *not* rename this license to moab.lic.

While the new Moab licenses use the RLM licenses technology, an RLM Server is not required to run most Moab functionality. Currently the only exception is if using the new Elastic Computing functionality, which will require a designated RLM to track usage.

The 9.1.0 licensing change does not affect legacy licenses; however, a new licensed may be required to use newer functionality.

Elastic Computing Management

Beginning with the 9.1.0 release, Elastic Computing now integrates with an RLM Server to track and report usage. This enables licenses by usage so that you are billed according to the total number of processor seconds used by the bursted (dynamic) nodes. Previously, Elastic Computing was licensed by number of nodes or sockets.

You can also limit over-bursting. For example, if your configuration allows 1000 processor seconds of use every day, then if a job needs to burst (and the used processor seconds reaches 1000 before the job can burst), the trigger to burst the job will not fire, and an error message is generated. You have the option to limit bursting by processor seconds or hours based on a maximum daily, monthly, quarterly, or yearly amount.

New SCHEDCFG flag "ProvisionFirstReqOnly"

The *ProvisionFirstReqOnly* flag is added for NUMA-aware scheduling using the -L resource request syntax. This flag informs Moab that it should re-provision the compute nodes allocated only to the first resource request in a multiple-resource-request or "multi-req" job instead of *all* compute nodes.

With the 9.0.2 release, the *ProvisionFirstReqOnly* flag is only applicable on Cray KNL systems.

Ability to Choose which Format to Write the Event File

An *EVENTFILEFORMAT* configuration parameter has been added that lets you specify the format for the event file. Valid values are *JSON*, *WHITESPACE*, *WIKI*. The default is *WIKI*.

This change also deprecates the WIKIEVENTS configuration parameter.

totalDedicatedProcessors Field Reports Processor Usage

Moab now has includes totalDedicatedProcessors field that reports processor usage to MWS and Insight on a per-request basis.

Pass Job ID to Elastic Trigger Script

Elastic scripts have access to a *$JOBID* variable to obtain the job ID from Moab.

Disable Backfill on a Per-Class Basis

New Class Attribute credential "CLASSCFG[] FLAGS=NoBackfill" lets you disable backfill on a per-class basis.

Moab Handles Resources Reported by the Torque Resource Plug-In

Moab recognizes gres, gmetric, varattr, and feature resources reported by the Torque resource plug-in. If the resources show up in a ClusterQuery, their type will be preceded by a colon (for example, `gres:<gres_name>=<gres_value>`).

Remove Nodes from Classes

A node can be removed from a class using *mnodectl -m CFGCLASS-=<CLASS> <nodeexp>*.

New msub - n Option

msub -n is added to support node exclusivity.

createdDate Added to Block Messages

A *createdDate* has been added to the *job.blocks* array for MWS and Insight.

fullEnvironmentVariableList Job Attribute

Moab now provides a *fullEnvironmentVariableList* job attribute to MWS and Insight that includes environment variables not included in the environmentVariables field.

New USEMOABJOBID Job Flag

The *USEMOABJOBID* job flag is available for job submission and job templates. This flag specifies whether to return the Moab job ID when running "msub", or the resource manager's job ID if it is available.

The *USEMOABJOBID* flag setting is set a the job level. When set, it overrides the (global) setting in moabcfg (Appendix A Moab Parameters).

Administrator Policy for Max Preemption Count

The *JOBMAXPREEMPTCOUNT* policy can be set to specify the number of times a single job can be preempted.

showq Filter for Feature

Enable *"showq -w nodefeature="* to filter on node feature strings.

Job Blocked Processors and Processor Equivalents tracked in Accounting Manager

Moab now sends blocked processors and processor equivalents as job properties to the accounting manager so that they may be tracked and charged independent of the charge policy.

Reservation Reserved and Idle Processor Seconds tracked in Accounting Manager

Moab now sends reserved processor seconds and idle processor seconds as reservation properties to the accounting manager so that they may be tracked independently from the charge.

Requested Job Duration and Effective Queue Duration Reported to the Accounting Manager

Moab now sends the requested job wallclock limit *(RequestedDuration)* and the effective amount of time queued *(QueueDuration)* to the accounting manager for jobs. For these properties to be recorded in the MAM Usage Record, you must define custom usage record attributes in MAM for them.

Node Features Optionally Tracked in Accounting Manager

Moab can be configured to pass allocated node features as job properties to the accounting manager so they can be tracked and charged. The new *AMCFG[] INCLUDEFEATURES* parameter allows the option of sending feature counts or combinations on a per-node or per-task basis.

REQUIREDUSERLIST and REQUIREDACCTLIST for SRCFG

SRCFG[] has two new attributes *REQUIREDUSERLIST* and *REQUIREDACCTLIST*. Each of these takes a comma delimited list of users and accounts, respectively. When one of these is present any jobs that run in the reservation must match them. If both are present any job that runs inside the reservation must match both. It is recommended that any entries in *REQUIREDUSERLIST* and *REQUIREDACCTLIST* be present in the *USERLIST* and *ACCOUNTLIST* parameters to handle reservation affinities.

Command to Refresh LDAP Credentials Cache

`mcredctl -r uid user:<UID>` refreshes a user's UID and GID.

Reservation Profiles Can be Used for Standing Reservations

SRCFG[] PROFILE=<NAME> can be used to specify a reservation profile to be used for a standing reservation.

Moab Can Send Email from a Specified User

The *MAILFROMADDR* parameter can be used to specify an email address to go in the *From* field of emails Moab sends, so replies sent by recipients go to the specified address.

Commands to Display the Reason Jobs Are Blocked Or Deferred

showq -b -v and *showq -b -v --blocking* display the reason jobs are blocked or deferred.

Messages Showing a Resource Manager is Down Are Only Displayed to Admins

When a resource manager is down, the *showq*command only reports that the RM is down to admin users.

CLASSCFG[] MIN.TPN and MAX.TPN Used to Assign Jobs to Classes

If a job is submitted to a remap class, Moab attempts to determine the class to which the job belongs, based on the resources requested. *MIN.TPN* and *MAX.TPN* can be used to specify the minimum and maximum tasks per node for the job, respectively.

Moab Accounting Manager

9.1.2

No new features.

9.1.1

Slurm Integration

Moab Accounting Manager provides integration with Slurm by including a `contrib/slurm` directory that includes charge and reserve integration scripts to be used in the `Slurmctld` epilog and prolog as well as a Slurm patch that enables Slurm to enforce a configured failure action when unable to obtain a lien with MAM. The documentation has been extended to include a new section on integrating with Slurm.

9.1.0

Web Services

Moab Accounting Manager Web Services (MAMWS) enables full access to the Moab Accounting Manager API via a REST-like interface over HTTPS. MAMWS communicates with Moab Accounting Manager using the same wire protocol, message format and Perl API as the MAM command-line clients and GUI interfaces. MAMWS forwards the required username and password (collected via HTTP Basic Authentication) to the MAM server for authentication and authorization. Thus, each user that wants to be able to use MAM Web Services must first set a password in MAM (e.g with the mam-set-password client command). MAMWS runs under mod_perl from an Apache HTTP server.

New JSON DataType for Complex Object Properties

An attribute may have a DataType of JSON, which signals that it will store multiple complex values in the JSON object form: `"{\"<part_name>\":\"<part_value>\",...}"`. The part values can be extracted individually in a query with a --show option of the form: `<attribute_name>{<part_name>}` for certain client query commands having specifically added this support (e.g. mam-list-usagerecords, mam-shell). Added affiliated support for a new part attribute in `MAM::Selection` and `MAM::Condition` (as well the respective adders in `MAM::Request`). The use of partial selections and partial conditions may only be used by MAM sites using a PostgreSQL database. Partial queries are not supported for sites using a MySQL database.

New Usage Record Properties

There is new support for the following as default or custom attributes in the usage record object:

ExitCode. A new ExitCode usage record attribute has been added to store a job's exit code. The mam-charge, mam-create-usagerecord and mam-modify-usagerecord clients now provide a new `-x exit_code` option that will set the new exit code value in the usage record. This replaces the former custom CompletionCode usage record property. ExitCode must be added to the usagerecord.show string in the mam-client.conf file in order for this property to be seen in the default output in mam-list-usagerecords.

Features. Node features may now be tracked as a complex property. Features may be added as a custom usage record attribute to store the node features allocated by a job. The mam-charge, mam-reserve, mam-quote, mam-create-usage-record and mam-modify-usage-record clients now provide a new `-F "{\"<feature_name>\":<feature_count>,...}"` option that will set the node features allocated by a job. The mam-list-usagerecords client can use the `--show` option to display the entire features list by specifying the Features attribute or an individual feature value by using the `<Features>{<feature_part_name>}` syntax. Features must be added to the usagerecord.show string in the mam-client.conf file in order for this property to be seen in the default output in mam-list-usagerecords.

Licenses. Licenses may now be tracked and charged. A new complex Licenses usage record attribute has been added to store the licenses used by a job. The mam-charge, mam-reserve, mam-quote, mam-create-usage-record and mam-modify-usage-record clients now provide a new `-L "{\"<license_name>\":<license_count>,...}"` option that will set the licenses used by a job. The mam-list-usagerecords client can use the `--show` option to display the entire licenses list by specifying the

Licenses attribute or an individual license count by using the <Licenses>{<license_part_name>} syntax. Licenses must be added to the usagerecord.show string in the mam-client.conf file in order for this property to be seen in the default output in mam-list-usagerecords. It is possible to charge for licenses by creating a charge rate using Licenses as the charge rate name and the license name as the charge rate value.

Metrics. Generic metrics may now be tracked and charged. A new complex Metrics usage record attribute has been added to store the generic metrics for a job. The mam-charge, mam-reserve, mam-quote, mam-create-usage-record and mam-modify-usage-record clients now provide a new -Z "{\"<metric_name>\":<metric_count>,...}" option that will set the metrics for a job. The mam-list-usagerecords client can use the --show option to display the entire metrics list by specifying the Metrics attribute or an individual metric value by using the <Metrics>{<metric_part_name>} syntax. Metrics must be added to the usagerecord.show string in the mam-client.conf file in order for this property to be seen in the default output in mam-list-usagerecords. It is possible to charge for generic metrics by creating a charge rate using Metrics as the charge rate name and the metric name as the charge rate value.

Name. Name may be added as a custom usage record attribute to store the user-specified job name. The mam-charge, mam-quote, mam-reserve, mam-create-usagerecord and mam-modify-usagerecord clients now provide a new -n name option that will set the new name value in the usage record. Name must be added to the usagerecord.show string in the mam-client.conf file in order for this property to be seen in the default output in mam-list-usagerecords.

RequestedDuration. A new RequestedDuration usage record property has been added to record the requested wallclock limit for a job. The mam-charge, mam-reserve, mam-quote, mam-create-usage-record and mam-modify-usage-record clients now provide a new -W <requested_duration> option that will set the total estimated wallclock duration for a job. The mam-list-usagerecords client can use the --show option to display the requested duration by specifying the RequestedDuration attribute. RequestedDuration must be added to the usagerecord.show string in the mam-client.conf file in order for this property to be seen in the default output in mam-list-usagerecords.

Resources. Generic resources may now be tracked and charged. A new complex Resources usage record attribute has been added to store the generic resources consumed by a job. The mam-charge, mam-reserve, mam-quote, mam-create-usage-record and mam-modify-usage-record clients now provide a new -R "{\"<resource_name>\":<resource_count>,...}" option that will set the resources for a job. The mam-list-usagerecords client can use the --show option to display the entire resources list by specifying the Resources attribute or an individual resource count by using the <Resources>{<resource_part_name>} syntax. Resources must be added to the usagerecord.show string in the mam-client.conf file in order for this property to be seen in the default output in mam-list-usagerecords. It is possible to charge for generic resources by creating a charge rate using Resources as the charge rate name and the resource name as the charge rate value.

Variables. Job variables may now be tracked and charged. A new complex Variables usage record attribute has been added to store the job variables associated with a job. The mam-charge, mam-reserve, mam-quote, mam-create-usage-record and mam-modify-usage-record clients now provide a new -V "{\"<variable_name>\":\"<variable_value>\",...}" option that will set the variables for a job. The mam-list-usagerecords client can use the --show option to display the entire

variables list by specifying the Variables attribute or an individual variable value by using the `<Variables>{<variable_part_name>}` syntax. Variables must be added to the usagerecord.show string in the mam-client.conf file in order for this property to be seen in the default output in mam-list-usagerecords. It is possible to charge for job variables by creating a charge rate using Variables as the charge rate name and the complex variable (e.g. foo:bar) as the charge rate value.

Incremental Charging for Accumulating Usage Properties

Formerly, if you made an intermediate charge for an accumulating usage value such as cputime or energy, this value would be used in the charge as well as being stored in the usage record, but this could be mistaken for the cumulative value. Charge data now allows the specification of both the cumulative value and the incremental value, allowing you to make a charge for the incremental value while recording the cumulative value in the usage record.

Server Reconfig Option

After modifying configuration parameters used by the server, you can cause the server to reread its configuration by running `mam-server --reconfig` or by sending the HUP signal to the main server process.

Server Role Options

Two new options have been added to the server for improved support of high availability. The `--primary` option causes the server to start up in the primary server role. When running under the primary server role, events are enabled by default. The `--backup` option causes the server to start up in the backup server role. When running under the backup server role, events are disabled so they do not conflict with the primary server. If neither option are used, the server will attempt to determine its role by checking whether its hostname matches the server.host or backup.host entry in the site configuration file.

New --about Option on Clients

A new `--about` option has been added to the command line clients that displays product information including the product version, the revision number, the install directory, and the home directory.

New ExactMatch Filter Type for Fund Selection

A new ExactMatch filter type has been added to filter actions that involve selecting funds (Allocation Query, and Fund Deposit, Query, Reset, Stop and Withdraw). If the ExactMatch filter type is used, a fund will only be matched if the specified filters exactly match the fund constraints. This new filter type is now available in the following client commands: mam-balance, mam-deposit, mam-list-allocations, mam-list-funds, mam-modify-funds, mam-statement, mam-withdraw, and the contrib command: allocation-summary.

Data Types Returned With Responses

Responses that return data now also return metadata containing the data types for the data. This feature enables clients to better render the output appropriate to the data type.

Moab Web Services

9.1.2

No new features.

9.1.1

No new features.

9.1.0

Deleting a Job Now Accepts an Optional "Where" Condition

You can now specify a where condition to only delete the job if the where condition is satisfied. For example, to cancel the job only if it is idle, use
where={"state": "IDLE"}. The only supported where condition is state.

New dedicateAllProcessors Field

The job requirements array has a new field: dedicateAllProcessors

Within a requirement, if dedicateAllProcessors is true, then all processors on the node where the job runs will be dedicated to the job.

New MWS Health Checks

The affected calls are /rest/diag/health/summary and /rest/diag/health/detail.

- Insight's Kafka communication; this call preparatory for the upcoming Reporting Engine functionality and is disabled by default
- Plugins
- Insight Server; also added in 9.0.2

#PBS Requirements Passed to Job

#PBS directives included inline with job scripts are processed server-side. This means that #PBS requirements will be passed to a job, provided they do not conflict with other defined job requirements in the web service call.

Support for Accounting Resources' "complex" Properties

In support of the new complex usage record properties sent by Moab to the accounting manager, the accounting usage record queries have added support for partial selections and conditions (e.g. rest/accounting/usage-records?query={"variables":{"opsys":"linux"}}&fields=id,licenses{matlab}&api-version=3&pretty=true).

New Resource to Retrieve License Information

MWS contains a new resource, /rest/diag/licenses, that retrieves license information from Moab Workload Manager.

New RLM License Plugin

The RLM plugin for MWS is available beginning with version 9.0.1. This plugin polls a Reprise License Manager (RLM) for purchased and available licenses for a given independent software vendor (ISV) and product. It reports this license information to Moab Workload Manager as a resource on the GLOBAL node.

Accounting Queries Support Field Comparison and Timestamps

For accounting resources, the query parameter now permits simple field comparisons including equal, greater than, greater than or equal, less than, less than or equal, and not equal. Additionally, queries based on time stamps are now supported. As an example, the enhanced query parameter syntax will support queries such as `rest/accounting/allocations?query={"endTime": {$gt:"2016-03-01 00:00:00 UTC" }}`.

Moab Insight

9.1.2

No new features.

9.1.1

No new features.

9.1.0

Archived Node Reliability and Retention

Insight now stores archived nodes in the Insight MongoDB database. By default, nodes will remain in the MongoDB database two years past the last time they have been reported by Moab Workload Manager. However, a new parameter (archiver.nodes.expireAfterSeconds) is available that lets you specify when the nodes will be deleted from the MongoDB database (to save disk space).

Archived Reservation Reliability and Retention

Insight now stores archived reservations in the Insight MongoDB database. By default, reservations will remain in the MongoDB database two years past their end date. However, a new parameter (archiver-.reservations.expireAfterSeconds) is available that lets you specify when the ended reservations will be deleted from the MongoDB database (to save disk space).

Performance Improvement in Job Sample Processing

Job samples are now generated considerably faster.

Moab Viewpoint

9.1.2

Improved Create Report Page Editing

Autocompletion and syntax highlighting have been added to the SQL editor in the Create Report page.

Default Job Script for Remote Viz Application Templates

When you create a Remote Visualization application template, the template is now automatically given a default job script.

9.1.1

Elastic Node Support for Resource Job Timeline Page

Viewpoint's Resource Job Timeline page displays Elastic nodes. Filter settings support displaying only regular nodes, only Elastic nodes, or both node types.

Elastic Node Support for Node Details Page

Viewpoint's Node Details page display has been redesigned for Elastic nodes. Last Seen and TTL values are displayed. Interface elements for disabled nodes are deactivated.

Improvements to Pass-Through Support

When running in pass-through mode, Viewpoint hides user interface elements that are not supported by resource managers other than Torque. For example, Viewpoint does not display the CPU Statistics section of the Job Details page, because the necessary data is not provided by the resource manager. Likewise, Viewpoint Job and Node pages do not display generic resources that are not supported by the underlying resource manager.

Limit Length of Viewpoint User Session

You can new specify the length of a Viewpoint user session by setting the `VIEWPOINT_SESSION_AGE` variable in `viewpoint.cfg`.

9.1.0

Moab Pass-Through Support

Moab passthrough support gives customers the ability to use Viewpoint and other modules provided by Adaptive Computing on top of third-party schedulers like Slurm.

Set Min, Max, and Step Values for Resources

When creating application templates, you can now set minimum, maximum, and step values for setting the number of cores, nodes with core count, malleable core count, total memory, and memory per core.

Grid View Improvements

The sessions page grid view has been enhanced to provide a cleaner user experience. The session name is displayed by the session screen shot and icons next to the screen shot represent session operations that may be initiated from the grid view.

Colors Show Status of Nodes in Node List

The status of nodes in the node list is indicated by color.

Workload Tables Show Job Status By Color

The queue status of jobs in the workload tables is indicated by color.

Application Templates Can Request Malleable Cores

Application templates allow users to request malleable core jobs (relates to Moab's TRL parameter).

Allocated Node List Sorted and Role Indicated By Color

The allocated node list in job details page has been improved to reflect actual allocation order, as well as adding coloring so users can more easily determine the role of each of the nodes allocated to the user's job.

Job Messages Include Timestamp

Job messages on the Job Details page display a date/time stamp to help users determine if the message information is relevant.

Application Templates May Specify Total Memory Or Memory Per Core

Application templates allow users to request either total memory for the job or memory per core.

File Manager Configuration Page Includes A Connection Validation Test

The configuration page for File Manager includes Test button that verifies whether the server parameters the user entered are correct.

Search By Job Name On Workload List Page

On the Workload List page, users can search for jobs by job ID, job name, or submitter name.

New Default Roles and Changes to Role Permissions

Viewpoint's role management page has been greatly simplified. Admins no longer have to specify domain permissions and can now use a tree navigator to more easily manage their user's permissions.

Torque Resource Manager

6.1.2

Creation of Core Dumps Enabled by Default

For RHEL 7 and newer, Core dumps for `pbs_server`, `pbs_mom`, and `trqauthd` crashes are now enabled by default. For RHEL 6 and earlier, you can enable core dumping by adding `export DAEMON_COREFILE_LIMIT=unlimited` to the `/etc/init.d/pbs_mom` and `/etc/init.d/pbs_server` scripts.

Update pbs_server Hardware Configuration

momctl -u updates pbs_server with a hardware configuration from the MOM.

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

`pbs_mom --about` now displays the Git commit hash.

Specify Location of sendmail Executable

You can now specify the location of the sendmail executable (for sending notification emails) using the `sendmail_path` server parameter.

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 set the swap memory limit.

Specify Port Range for Interactive Jobs

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

6.1.1

Additional Options for Email Notifications

Additional % format specifiers are available for including information in email subject lines and messages. For example, %o inserts the name of the job owner and %u inserts a summary of the resources used by the job.

pbs_mom Switch to Force the Server to Accept a Hardware Update

pbs_mom -f forces the server to accept an update of the hardware on the node. You should use this switch the first time you start the MOM after an upgrade.

CUDA 8 Support

You can now set a `default_gpu_mode` parameter to control the default GPU mode for jobs that request GPUs.

6.1.0

Added Presetup Prologue for MOMs

The presetup_prologue executes before the MOMs attempt to set up jobs on the node. This prologue can be used, for example, to set up Kerberos.

Added default_gpu_mode Server Parameter

The default_gpu_mode server parameters sets the default value GPU mode. This should facilitate using CUDA 8.

Simplified Torque RPM Installation

The RPM installation no longer writes the server name into pbs_mom's config file, eliminating the need to configure the MOM config file in the default case.

Simplified Installation Process

"make install" and the package installer now install service/unit files, set up the path to the Torque libraries, and establish a shell path to binaries for sh and csh shells.

Resource Plug-In

There is now an API for creating resource plug-ins that report custom varattrs, generic resources, generic metrics, and features. Additionally, jobs can report custom resources through the plug-in.

-L Submission Recorded in the Accounting Log

Jobs submitted with the -L request syntax will now have the -L submission recorded in the accounting log.

NUMA-Aware cgroup Creation by Per Task or Per Job

A new Torque server parameter "cgroup_per_task" is available to let you specify whether cgroups are created per task or per job. The default is FALSE, meaning jobs submitted with the -L syntax will have *one* cgroup created per host; this behavior is similar to the pre-6.0 cpuset implementation.

The qsub/msub -L syntax is also modified to let you specify whether the cgroup is per task or per job at the job submission time.

Some MPI implementations are not compatible with using one cgroup per task.

Change hashname for a Torque Job to Match the Corresponding Moab Job ID

When jobs are submitted using qsub -J, the job's hashname is set to the job ID of the Moab job.

Job Arrays Now Only Partially Instantiated By Default

Job arrays are now only partially instantiated by default; additional idle jobs are added as the array is executed. This is controlled by the new idle_slot_limit server parameter.

A new qsub -i option has also be added to set the slot limit for the job array being submitted. If set for a non-array job, it will be rejected. If the user requests an idle slot limit that exceeds the server parameter's default, the job will be rejected.

Recover Array Subjobs

The new server parameter "ghost_array_recover" is added. pbs_server will now recover array subjobs even when the array (.AR file) couldn't be recovered. This parameter is set to TRUE by default.

Improved Job Cleanup Performance

The number of communications involved in job cleanup has been reduce, resulting in improved performance.

Epilogue Script Runs, Even if Output Files Cannot Be Appended

The epilogue script will now run when spool_as_final_name is configured, even if error and output files are not available and cannot be appended.

Cray-enabled Torque May Also Be Configured With cgroups

Cray-enabled Torque may also be configured with cgroups. On the login node, each job will have all of the CPUs and all of the memory controllers in its cgroup. For ALPS 1.7 (KNL), you can use the qsub-/msub -L syntax when submitting jobs.

Job Script Path is an Argument to Prologue and Epilogue Scripts

A new positional parameter contains the full path of a job's job script to the job's prologue and epilogue scripts when the Torque pbs_mom "mother superior" launches the scripts. The prologue script is the new 8th positional parameter. The epilogue script is the new 11th positional parameter.

RPM Install for SLES 12

You can now use RPM to install and upgrade Torque on SLES 12.

qrun Can Interpret the -L Syntax

You can run jobs that request resources using *qrun -L* without a hostlist.

User May Request No Emails be Sent, Even on Failure

Specifying -m p at qsub time makes it so no emails are sent for the job, even on failure.

User settable kill_delay Through qsub -K Option

Added a user settable, per-job kill delay, called kill_delay. It is settable via the new qsub -K option.

Torque Commands Appear in User's Login Shell Path

RPMs created via build-torque now create /etc/profile.d/torque.sh (and torque.csh) files so that Torque commands appear in a user's login shell path.

Option to Email User on All Non-zero Exit Codes

The qsub -m f option sends an email if a job has a non-zero exit code. The qsub -m f option can be used with the a, b, and e options, but not with n or p.

Added an Option for When to Send Mail Notifications

Added a -m f option to qsub to send mail when a job terminates with a non-zero exit code.

Option to Disable Reading of RUR Information

Added $cray_check_rur configure option to disable reading of Resource Utilization Reporting (RUR) energy usage for Cray login nodes. If set to false, login MOMs will not look at the energy resource information used for each job. Disabling this may improve performance.

Improved Job Submission and Start Time Performance

Transactions per job submitted have been reduced from 11-13 to 3, improving job submission time by an average of 11-13%.

Added ghost_queue Queue Attribute

If pbs_server restarts and recovers a job but cannot find that job's queue, a new queue with the 'ghost_queue' attribute is created. That job is added and then run on the new queue. Jobs cannot be added to ghost queues. Once an administrator reviews and corrects the queue's settings, the 'ghost_queue' attribute can be removed, allowing the new queue to function normally.

New "email_batch_seconds" Server Parameter

The new server parameter "email_batch_seconds" lets you control at what frequency a batch of emails are sent to each user.

MOM Parameter to Overwrite Output Files

The $force_overwrite MOM parameter, if set to true, forces the output files to be overwritten each time a job is started.

Option to Turn Off trqauthd Logging

Added a -n option to trqauthd that disables trqauthd from logging anything.

Reporting Framework

9.1.2

X-Value Display Strategy Added to Report Design

An option was added to specify a display strategy to be used when the labels on a line or column chart overlap.

Automated Installer

9.1.2

No new features.

Differences

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

In this topic:

- General Suite
- Moab Workload Manager
- Moab Accounting Manager
- Moab Web Services
- Moab Insight
- Moab Viewpoint
- Torque Resource Manager
- Reporting Framework
- Automated Installer

General Suite

This section contains information applicable to more than one of the components in the Moab HPC Suite.

9.1.2

No known differences.

9.1.1

No known differences.

9.1.0

No known differences.

Moab Workload Manager

9.1.2

SLES 11 RPM Spec File Removed from Binary Tarball

The SLES 11 RPM spec file has been removed from the Moab binary tarball, because Adaptive Computing does not support RPMs for SLES 11.

9.1.1

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.

9.1.0

EPEL Installation Changes for RHEL 6 and 7

Installation of the EPEL RPM repository for Red Hat Enterprise Linux (RHEL) is no longer required. The RPMs that were being pulled form EPEL are now included in the Moab RPM download. Previously, EPEL was automatically installed by the install-rpm-repos.sh script.

Nodes Listed in Numerical Order in checkjob Rejection Summary

When listing the reason for rejection for a job, the *checkjob* command now sorts the list of nodes for a reservation in numerical order.

Cray KNL Management Scripts Now Require CAPMC 1.1 or Later

KNL reference scripts now align with the CAPMC 1.1 specification, and will allow the Moab-Cray-KNL integration to scale better.

Turn Off Idle Node Purging on a Per Node Basis

For elastic computing, you can turn off the purging of *individual* dynamic nodes output using the varattr attribute in Torque or through the wiki interface.

CHECKSUSPENDEDJOBPRIORITY Parameter Default Change

The *CHECKSUSPENDEDJOBPRIORITY* parameter default is now FALSE.

> This change may affect additional parameters or flags that assumed CHECKSUSPENDEDJOBPRIORITY is set to TRUE.

Changes to Base PE Calculations

Base PE calculations are now on online resources, rather than total resources.

showres Does Not Report Down or Idle Nodes

showres -f does not include down or idle nodes in its output.

Enhancement to the job.blocks Array (MOAB-8556)

Added *createdDate* to the *job.blocks* array for MWS and Insight.

Unset the MAXLOAD on a Specific Node with nodecfg

Added ability to set MAXLOAD to -1 on a node to unset the default.

Moved OutputFormat Tag into Job Tag

The OutputFormat tag has been moved into the job tag so that usage of OutputFormat is XML compliant.

Moab Handles qsub Job Arrays the Same as msub

Moab now understands Torque job arrays the same as its own job arrays. Job arrays submitted to Torque will show up as job arrays in Moab.

Documented Account Balance Cache Query (MOAB-8398)

The command "mrmctl -q AccountBalanceCache AM" has been documented let you query Moab's internal account balance cache.

This command is only applicable when the fast-allocation accounting mode is configured.

Changes to Job Accounting Properties Reported to the Accounting Manager

While Moab previously sent generic resources (including licenses), generic metrics and job variables as direct job properties to the accounting manager, Moab now sends generic resources, licenses, generic metrics and job variables in a consolidated form to the accounting manager under the Resources, Licenses, Metrics and Variables aggregate properties respectively. This change lets you distinguish the source of a property from among the disparate property types (including those having identical names) and limits the number of usage record properties needed to be configured within the accounting manager to track the desired aspects of the jobs.

Moab now additionally includes BlockedProcessors and ProcessorEquivalents as direct Job properties sent to the accounting manager instead of scaling the Processors charge by a consumptionRate attribute. This change allows these properties to be logged in the usage record, allows for better accuracy, and moves the control of what is being charged for into MAM. Moab now sends the exit code as ExitCode to the accounting manager (previously sent as CompletionCode). Moab now sends the hardware accelerators as GPUs and MICs (previously sent as Gpus and Mics). Moab now also sends RequestedDuration (requested wallclock limit) and QueueDuration (effective duration queued) to the accounting manager for jobs.

Configuration Changes Required to use Alternate Processor Charge Policies

Rather than specifying in Moab whether you want to charge for allocated processors, blocked processors, processor equivalents or cpu time, Moab now sends all of these resource usage properties to the accounting manager so that they may be tracked and charged independently.

- If you were using one of AMCFG[] CHARGEPOLICY=Debit{All,Successful}Blocked, comparable charging can be obtained by setting a charge rate for BlockedProcessors in Moab Accounting Manager.
- If you were using one of AMCFG[] CHARGEPOLICY=Debit{All,Successful}CPU, comparable charging can be obtained by setting a charge rate for CPUTime in Moab Accounting Manager.
- If you were using one of AMCFG[] CHARGEPOLICY=Debit{All,Successful}PE, comparable charging can be obtained by setting a charge rate for ProcessorEquivalents in Moab Accounting Manager.
- If you were using one of AMCFG[] CHARGEPOLICY=Debit{All,Successful}WC, comparable charging can be obtained by setting a charge rate for Processors in Moab Accounting Manager (which is the default).

See Managing Charge Rates in the *Moab Accounting ManagerAdministrator Guide* for examples on configuring charge rates for these alternate processor charge policies.

Job Charge Policy

A new AMCFG[] JOBCHARGEPOLICY parameter has been added that dictates the behavior of which jobs to track or charge with the accounting manager. This parameter permits values of All, None or Successful to enable charging of resources used for all, none or successful jobs respectively.

The related AMCFG[] CHARGEPOLICY parameter is obsolete.

- If you were using one of AMCFG[] CHARGEPOLICY=DebitAll{Blocked,CPU,PE,WC} to track and charge for all jobs, comparable filtering can be obtained by using the new AMCFG[] JOBCHARGEPOLICY=All (which is the default).
- If you were using one of AMCFG[] CHARGEPOLICY=DebitSuccessful {Blocked,CPU,PE,WC} to track and charge for only successful jobs, comparable filtering can be obtained by using the new AMCFG[] JOBCHARGEPOLICY=Successful. Tracking and charging for jobs can also be disabled entirely by setting AMCFG[] JOBCHARGEPOLICY=None.

See Charging for Jobs in the *Moab Workload ManagerAdministrator Guide*.

Changes to Reservation Accounting Properties Reported to the Accounting Manager

Moab now includes IdleProcessorSeconds and ReservedProcessorSeconds as direct Reservation properties sent to the accounting manager; instead of scaling the Processors charge by a consumptionRate attribute.

These properties additionally include an increment attribute to differentiate between the charged processor seconds and the cumulative processor seconds (which is recorded in the usage record). These changes permit more accurate charging for reservation idle cycles, particularly when periodic charging is enabled. It similarly moves the control of the charging algorithm into MAM and allows for different charge rates to be used for jobs and reservations.

Configuration Changes Required to Continue to Charge for Reservations

The mechanism and configuration steps for charging for the unused cycles in reservations has changed.

If you were previously charging for reservations in Moab Accounting Manager, you will need to add the IdleProcessorSeconds and ReservedProcessorSeconds properties to the usage record object in MAM *and* adjust your charge rates to include a reservation-based charge rate for IdleProcessorSeconds in order to continue charging for reservations. See the example for Enabling Reservation Statistics (in Customizing the Usage Record Object) and see the example for Charging for the unused cycles in reservations (in Creating Charge Rates) in the *Moab Accounting ManagerAdministrator Guide* for how to do this.

Reservation Charge Policy

A new AMCFG[] RESERVATIONCHARGEPOLICY parameter has been added that dictates the behavior of tracking or charging for the total and unused cycles in a reservation with the accounting manager. This parameter permits values of All, None or Select to enable charging of unused cycles for all, none or select reservations respectively.

The related AMCFG[] ALWAYSCHARGERESERVATIONS parameter is now `deprecated.`

- If you were using AMCFG[] ALWAYSCHARGERESERVATIONS=True, comparable functionality can be obtained by using the new AMCFG[] RESERVATIONCHARGEPOLICY=All.
- If you were using AMCFG[] ALWAYSCHARGERESERVATIONS=False, comparable functionality can be obtained by using the new AMCFG[] RESERVATIONCHARGEPOLICY-Y=Select.

See Charging for Reservations in the *Moab Workload ManagerAdministrator Guide*.

Changes to Reservation Statistics

Current blocked processor seconds (cbps, STATCBPS), current reserved processor seconds (crps, STATCRPS), total blocked processor seconds (tbps, STATTBPS) and total reserved processor seconds (trps, STATTRPS) have been added as new reservation statistics and have been added to the ODBC database schema.

Current active processor seconds (caps, STATCAPS), current idle processor seconds (cips, STATCIPS), total active processor seconds (taps, STATTAPS) and total idle processor seconds (tips, STATTIPS) `are no longer reported in the reservation statistics and have been removed from the ODBC database schema.` The former active processor seconds statistics are roughly equivalent to the new blocked processor seconds statistics and the former idle processor seconds statistics are roughly equivalent to the new reserved processor seconds statistics minus the new blocked processor seconds statistics.

Added UseRMCompletedJobStats SCHEDCFG Flag

Use RM reported CPUTime when calculating job efficiency with SCHEDCFG[] FLAGS-S=UseRMCompletedJobStats.

Configuration Changes Required to Distinguish Licenses from Generic Resources when sent to the Accounting Manager

It is now possible to distinguish licenses from generic resources when they are reported to the accounting manager. Any generic resources seen on nodes reported via a native resource manager having the server configuration parameter RMCFG[] RESOURCETYPE=License will be consolidated under the new Licenses property of the job. License generic resources defined in Moab may be identified as license types by using the Moab configuration parameter GRESCFG[] LICENSE=TRUE.

LIMITEDJOBCP Now Has a Default Value of TRUE

Checkpointing on a per job basis can negatively affect Moab on systems with many jobs. To avoid this performance impact, LIMITEDJOBCP now has a default value of TRUE, which means Moab will only checkpoint a job when the state changes.

Ability to the Specify Path with the Log Level

The mschedctl -L command has a new format of [<LOGLEVEL>[:<*LOG_FILE*>]]. In addition to specifying the level for which you want to save log messages, you can specify the path where you want to save those log files.

Moab Accounting Manager

9.1.2

No known differences.

9.1.1

No known differences.

9.1.0

MySQL Use Deprecated

The use of a MySQL (or MariaDB) database as the backend database for MAM is deprecated. Beginning with this release, some features (such as partial queries) may not be available when using a MySQL database. Support for MySQL may be removed in a future release. If you are using a MySQL database for your MAM backend, it is highly recommended that you migrate the MAM database from MySQL to PostgreSQL and reconfigure Moab Accounting Manager to use the PostgreSQL database.

Configure --with-user Defaults to mam

Configure has changed to use mam as the default accounting admin userid if not specified via the `--with-user` option. Previously, configure would use the user running the configure command as the default accounting admin userid.

Re-enabled Support for the Raw Format

Support for the raw format in query client commands has been re-enabled. When using the `--format=raw` option, output fields are delimited with a pipe (|) symbol.

Sample Allocation Summary Report Provided

A sample allocation summary report that summarizes the status of current and expired allocations has been added to the contrib directory. By default, it will display fund constraints and allocation and usage details for the currently active allocations. If a timeframe is given, both active and expired allocations overlapping the specified time period will be displayed. Filters can be specified to narrow the set of allocations displayed. Specific fields can be specified for display. A totals line is displayed unless you are using the csv or raw format. Additional fields include the number of jobs and reservations charged, processor-seconds used, and percent of idle reservation cycles used.

Custom CompletionCode Usage Record Attribute Changed to ExitCode

Previously, Moab would send the exit code via the CompletionCode usage record property, which would have optionally been added as a custom property. Moab has been modified to send the exit code by default in the new ExitCode usage record property. If the database migration script detects a CompletionCode usage record property, it will rename the usage record property to be ExitCode, otherwise a new ExitCode usage record property will be created.

Disk No Longer Included as Default usagerecord Property

The Disk property has been removed from the default display list for usage records (usagerecord.show) and is no longer added as a usage record attribute for new installations. It can still be added as a custom attribute, and support still exists for the -D <disk> option in the various commands to manipulate it. Disk is not removed as a usage record attribute from existing installations.

CPUTime No Longer Included in Default Display List for Usage Records

The CPUTime property has been removed from the default display list for usage records (usagerecord.show). However, it is still created as a default usage record attribute for new installations. It simply needs to be added to the default display list for mam-list-usagerecords by editing the the usagerecord.show parameter in the client configuration file (mam-client.conf).

Incremental Values Can Be Specified in Charge-Related Actions

It is now possible to specify both the incremental as well as the cumulative values for an accumulating usage property in charges, liens and quotes, making it possible for you to charge for the incremental value while recording the cumulative value in the usage record. The usage record property XML data permits a new increment attribute in support of this feature. Additionally, the mam-charge, mam-reserve, and mam-quote clients allow a new syntax for accumulating properties of the form: `[<cumulative_value>][(<incremental_value>)]`.

Enhancements to the MAM Perl API

The new `MAM::Datum->addProperty` method has been added as an alternative to `MAM::Datum->setProperty` which replaces similarly named properties. `MAM::Request->set{Object,Selection,Assignment,Condition}` methods have been renamed to `MAM::Request->add{Object,Selection,Assignment,Condition}` as a more consistent description of what the functions do.

Improved Handling for Server Termination

When the server is shutdown via `mam-server --shutdown,-k`, it now logs the termination in the server log. Additionally, if the TERM signal is sent to the main server process, the pid file will be removed, the termination will be logged, and the server process will be ended.

-R Charge Rate Option Renamed to --rate charge_rate in Client Commands

All clients using the former `-R charge_rate_name[{charge_rate_value}]=charge_rate_amount,...` option have been changed to use the new `--rate charge_rate_name[{charge_rate_value}]=Icharge_rate_amount,...` option in order to support the new `-R {"resource_name":resource_count,...}` option. This change affects the mam-charge, mam-quote, and mam-reserve clients.

-E email_address Option Renamed to --email email_address in Client Commands

The former `-E email_address` option has been deprecated in favor of the new `--email email_address` option in order to facilitate the new -E energy option. This change affects the mam-create-user and mam-modify-user clients.

-F phone_number Option Renamed to --phone phone_number in Client Commands

The former `-F phone_number` option has been deprecated in favor of the new `--phone phone_number` option in order to facilitate the new `-F feature_name:feature_count,...` option. This change affects the mam-create-user and mam-modify-user clients.

The Effect of the -t Option in mam-charge, mam-reserve, and mam-quote Has Changed

The -t option was previously used to express the incremental charge, lien, or quote duration. However, in the mam-create-usagerecord and mam-modify-usagerecord, -t was used to express the total actual duration. Additionally, there was no direct way in mam-charge, mam-reserve, and mam-quote to set the total actual duration (one had to use `-X Duration=<actual_duration>`). A new -zt option has been added to mam-charge, mam-reserve, and mam-quote to specify the incremental charge, lien, or quote duration. The -t option in mam-charge has been changed to indicate the total actual duration. As it is not applicable to set the total actual duration for a quote or lien, the -t option has been removed from mam-quote and mam-reserve. This change is expected to have low impact because most calls of the

charge, lien, and quote actions occur through the wire protocol or via the API, and in the typical use case for calling mam-charge directly to make up for missed charges, the previous method for setting the total actual duration (`-X Duration=<actual_duration>`) still works as previously. Using the new -t option in the same way as previously has the same overall effect when making non-incremental charges.

The Effect of the -s Option in mam-charge, mam-reserve, and mam-quote Has Changed

The -s option was previously used to express the charge, lien, or quote start time. However, there was no direct way to express the original start time for the usage (`one had to use -X StartTime=<start_time>`). A new -zs option has been added to mam-charge, mam-reserve, and mam-quote for incremental charging. The -s option in mam-charge and mam-reserve has been changed to set the original start time for the usage. As it is not applicable to set the original start time for a quote, the -s option has been removed from mam-quote. This change is expected to have low impact since the former way of setting the original start time (`-X StartTime=<start_time>`) will still work as previously. Using the new -s option as previously has the same overall effect when making non-incremental charges.

The Effect of the -e Option in mam-charge, mam-reserve, and mam-quote Has Changed

The -e option was previously used to express the charge, lien, or quote end time. However, there was no direct way to express the final end time for the usage (one had to use `-X EndTime=<end_time>`). The -e option in mam-charge has been changed to set the final end time for the usage. As it is not applicable to set the final end time for a lien or quote, the -e option has been removed from mam-reserve and mam-quote. This change is expected to have low impact since the former way of setting the final end time (`-X EndTime=<end_time>`) will still work as previously. The new use of the -e option will have a similar overall effect as previously when making non-incremental charges.

The Effect of the -d Option in mam-charge, mam-reserve, and mam-quote Has Changed

The -d option was previously used to annotate a description in the charge, lien, or quote transactions. The -d option now annotates a description in the usage record. This change is deemed to be more useful and is better aligned with the use of the -d option in the mam-create-usagerecord and mam-modify-usagerecord clients.

Events Enabled by Default

The event scheduler is now enabled by default.

Allocations and Notifications Refreshed Via Events

New events have been created to explicitly refresh allocations and notifications instead of doing this as a side-effect of queries. This enables queries to avoid obtaining a request ID and allows them to be run on a read-only database replica.

Events Disabled for Backup Server

Since both the primary and backup servers may run simultaneously, events are disabled for the backup server so they do not conflict with events triggered by the primary server.

The super.user Server Configuration Parameter Has Been Removed

Actions requiring escalated (super user) privileges should always use the configured primary accounting admin user, which defaults to mam but may be overridden with `configure --with-user`.

Moab Web Services

9.1.2

MWS Runs Under SELinux

The Automated Installer now installs a policy that will allow MWS to run under SELinux on CentOS/Red Hat hosts

9.1.1

No known differences.

9.1.0

The Insight Database Query API is No Longer Available

The /rest/insight resource is no longer available. Use the ViewpointQueryHelper plugin instead.

/mws/rest/diag/licenses Resource Now Returns Elastic Usage Information

In addition to the existing license information, the /mws/rest/diag/licenses resource now returns Elastic usage information (such as daily, monthly, and yearly processor seconds).

New Fields for Job Resource

These fields were added to the jobs resource:

- blocks.partition
- dependencies.count
- durationMinimum
- partitionAccessListScheduler
- rank

Changes to Reservation Statistics

Blocked processor seconds (blockedProcessorSeconds) and reserved processor seconds (reservedProcessorSeconds) have been added as new reservation statistics. The current active processor seconds (caps), current idle processor seconds (cips), total active processor seconds (taps), and total idle processor seconds (tips) are no longer reported in the reservation statistics.

The new blocked processor seconds value is roughly equivalent to the sum of what was previously reported as the current active processor seconds and total active processor seconds. The new reserved processor seconds value is roughly equivalent to the sum of what was previously reported as current active processor seconds, previous total active processor seconds, current idle processor seconds, and total idle processor seconds.

Deprecation of Accounting Resources

With the advent of the new MAM Web Services component in Moab Accounting Manager, the services available through Moab Web Services for interacting with objects in Moab Accounting Manager are deprecated. The MWS Accounting resources and services will not be enhanced and may be removed in a future release.

Moab Insight

9.1.2

No known differences.

9.1.1

No known differences.

9.1.0

Finer Granularity for Job Sample Pruning

Job sample records in Insight's MongoDB will be pruned such that granularity is high for recent samples but low (once per day) for samples older than one month.

Insight Can Now Run on Secure Sites

Insight no longer requires execute privileges in the /tmp directory, which is not allowed in some secure sites.

Insight Now Uses MongoDB Instead of PostgreSQL

Insight no longer requires access to both PostgreSQL and MongoDB. This makes the installation and management process of Insight a lot simpler. Performance improvements are expected from this and other changess.

Moab Viewpoint

9.1.2

Dedicated Processors Calculation on CPU Statistics Graph

The CPU statistics graph now uses the instantaneous value of dedicated processors at sample time to calculate the number of dedicated processors reported, rather than the accumulated dedicated processor seconds over the sample interval divided by the sample interval length.

Limit User Access to Configuration Tab

Only the viewpoint-admin and moab-admin users are now permitted to configure Viewpoint's connection to MWS.

9.1.1.1

User's Moab Credentials stored in the User's Session

Viewpoint now stores the user's Moab credentials (accounts, classes, groups, and QoSes) in the user's Viewpoint session. If there are any changes to your group, account, class, or QoS assignments, log out and back in again to refresh your credentials.

9.1.0

Additional Base Roles

Viewpoint now configures with six default (base) roles. The HPCAdmin and HPCUser are the same as they were in previous Viewpoint versions. The other four roles are:

- NitroAdmin – Administrative user, with permission to create Nitro application templates and manage other user's Nitro jobs.
- NitroUser – Basic user, with permission to create and manage their own Nitro jobs.
- RemoteVizAdmin – Administrative user, with permission to create remote visualization application templates and manage other user's remote visualization jobs.
- RemoteVizUser – Basic user, with permission to create and manage their own remote visualization jobs.

All six default roles are present regardless of whether your configuration includes Nitro or Remote Visualization.

Torque Resource Manager

6.1.2

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 the Known Issues section.

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`.

6.1.1

$CUDA_VISIBLE_DEVICES is not set by default

$CUDA_VISIBLE_DEVICES is not set by default if you're using cgroups. Also, this makes its contents compatible with varying CUDA versions.

6.1.0

Use --enable-cgroups Recommended; --enable-cpusets is Deprecated

--enable-cgroups is the recommended configuration option. When cgroups are supported, cpusets are handled by the cgroup cpuset subsystem.

--enable-cpusets is deprecated; meaning that the functionality is still supported; however, no new features will be added.

pbs_server Does Not Shut Down When trqauthd is Stopped

The *trqauthd* service is no longer a requirement of *pbs_server*, so *pbs_server* no longer shuts down when *trqauthd* is stopped.

TorqueConfigure Script Does Not Enable the GUI Component

When the devel-tk and devel-tcl libraries have been installed, the Torqueconfiguration script does not enable the GUI component. To build the GUI component, specify --enable-gui when building from source.

Default Value for keep_completed Parameter Changed

The default value for the Torque `keep_completed` parameter has been changed to 300 seconds. The `keep_competed` parameter specifies the number of seconds jobs are kept in the job execution queue after they are completed. Keeping completed jobs in the queue enables Torque to report on the status of the jobs.

legacy_vmem Server Parameter Affects Behavior of the -l vmem Option

legacy_vmem is a new server parameter that affects the behavior of the -l vmem option. When set to true, the vmem request will be the amount of memory requested for each node of the job. When it is unset or false, vmem will be the amount of memory for the entire job and will be divided accordingly.

Queue Support for Both resources_default.* and req_information_default.* Settings

When queues have both resources_default.* and req_information_default.* set then they are applied according to their resource request type. resources_default.* settings are applied to jobs that do not explicitly use the -L syntax, while req_information_default.* settings are applied only to jobs that explicitly use the -L resource syntax.

Prohibited Mode Not Allowed for User Jobs

Setting the compute mode of an NVIDIA GPU to prohibited makes it so the GPU cannot be used at all. In previous versions of Torque users were allowed to set a GPU to prohibited mode. But then it could not set the mode to anything else since the GPU was now prohibited. This change went into effect for version 5.1.3, 6.0.2, and later.

Cray-Enabled Torque May Be Configured with cgroups

Support is added for Cray-enabled Torque configured with cgroups.

- On the login node, each job will have all of the cpus and all of the memory controllers in it's cgroup.
- For ALPS 1.7 (KNL), you can use the qsub/msub -L syntax when submitting jobs.

Support for Single Job Dependencies and Array Dependencies at the Same Time

Jobs can depend on single job dependencies and array dependencies at the same time.

Added a Way to Exclude Compute Nodes From allow_node_submit.

Set 'allow_node_submit=true' in qmgr, and then add a remote mom to node_submit_exceptions with 'set server node_submit_exceptions = <nodename>' in qmgr. This setting will disable job submissions on the <nodename> specified by the user. Previously, the process of disallowing compute nodes from submitting jobs would require making a separate entry in qmgr for each individual compute node.

Reduced the Number of Logging Statements

Reduced the number of logging statements when a node isn't up and therefore can't receive the mom hierarchy.

Added tcp_incoming_timeout Server Parameter

tcp_incoming_timeout specifies the number of seconds before incoming connections timeout. tcp_ timeout now specifies the timeout for outgoing connections or connections initiated by pbs_server. tcp_ incoming_timeout functions exactly the same as tcp_timeout, but governs incoming connections while tcp_timeout governs only outgoing connections (or connections initiated by pbs_server).

pbs_server Enhancement for Very Large Number of Jobs.

pbs_server has been enhanced to better handle a very large number of jobs (several hundred thousand or more) by enabling an alternate way for it to store job-related files in the directories $PBS_ HOME/server_priv/jobs and $PBS_HOME/server_priv/arrays.

A new boolean server attribute, use_jobs_subdirs, lets an administrator direct the way pbs_server will store its job-related files. When use_jobs_subdirs is unset (or set to false), job and job array files will be stored directly under $PBS_HOME/server_priv/jobs and $PBS_HOME/server_priv/arrays. This is the default behavior and the way the server has stored these files in the past. When use_job_subdirs is set to true, job and job array files will be distributed over 10 subdirectories under their respective parent directories. This method helps to keep a smaller number of files in a given directory.

If an administrator wishes to change the use_jobs_subdirs attribute from its previous value (or when setting it to true when it has not previously been set), it is highly recommended that Torque be drained of all jobs. Failing to take this action may result in the loss of existing jobs.

Reporting Framework

9.1.2

Upgrade to Spark 2.1

The Reporting Framework now uses Spark version 2.1.2.

Automated Installer

9.1.2

No known differences.

Installation and Upgrade Information

This section identifies information useful when installing and upgrading.

When installing or upgrading, it is *strongly* recommended that administrators configure Moab with mauth authentication with a complex key value. See "Mauth Authentication" in the *Moab Workload Manager Administrator Guide* for more information.

In this topic:

- Compatibility Requirements
- Installing Moab 9.1.2
- Upgrading to Moab 9.1.2

Compatibility Requirements

For the 9.1.x release, all Moab HPC Components are compatible with other 9.1.x versions. Compatibility with prior versions (such as 9.0.x) is not supported; however exceptions can occur between Moab Workload Manager and Torque Resource Manager.

The following identifies additional information regarding compatibility:

- Torque Resource Manager and Moab Workload Manager
  - If you are using Torque 6.1, you must run Moab 8.0 or later. However, if running Moab less than 9.1, the following Torque functionality will not be available.
    - NUMA-aware scheduling and cgroup support; requires Moab 9.0 or later.
    - Torque plugin, requires Moab 9.1 or later.
  - Moab supports Torque version 4.29 or later.
- Moab Web Services and Moab Workload Manager
  - Moab Web Services does not support SUSE 11-based systems. The Moab Server Host on which you will also install MWS needs to have a MWS-compatible OS.

Installing Moab 9.1.2

Please see the *Moab Installation and Configuration Guide* for manual or RPM-based installation instructions.

Upgrading to Moab 9.1.2

Please see the *Moab Installation and Configuration Guide* for manual or RPM-based installation instructions.

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:

- General Suite
- Moab Workload Manager
- Moab Web Services
- Moab Accounting Manager
- Moab Insight
- Moab Viewpoint
- Torque Resource Manager
- Reporting Framework
- Automated Installer

General Suite

9.1.2

No known issues.

9.1.1

No known issues.

9.1.0

- If the new, automated installer falls mid-install, Apache may not have been configured correctly. As a result, when you run the automated installer again, you may see an error that the Apache restart failed. If you see this error, uninstall Apache and then run the automated installer again. (DOC-3022)

Moab Workload Manager

9.1.2

- When the `NODEAVAILABILITYPOLICY` is configured with `UTILIZED:PROC`, it currently behaves as though configured with `COMBINED:PROC`. Moab will not schedule jobs, which

would result in the requested processors exceeding the configured processors on the nodes. (MOAB-9122)

9.1.2

No known issues.

9.1.1

- Moab fails to start; error "Could not find a valid RLM license (base license keyword 'moab_license'): System clock has been set back (-40)". A script is available to correct this problem. Contact your Adaptive Computing account manager to resolve this issue. (DOC-2977)

9.1.0

- Certain multi-node/multi-task jobs submitted using the new -L syntax will start correctly but on subsequent iterations the tasks per node will revert to 1. (MOAB-8718)

Moab Web Services

9.1.2

No known issues.

9.1.1

No known issues.

9.1.0

No known issues.

Moab Accounting Manager

9.1.2

No known issues.

9.1.1

No known issues.

9.1.0

No known issues.

Moab Insight

9.1.2

No known issues.

9.1.1

No known issues.

9.1.0

No known issues.

Moab Viewpoint

9.1.2

No known issues.

9.1.1

- An internal server error message is displayed when the user attempts to save an application template with the same name as an existing template. (IRIS-3294)
- Vertical lines in chart reports disappear when the window is resized. (Occurs only in Firefox 52 version and only when the page has a vertical scrollbar.) (RF-100)
- New Spark applications sometimes start unexpectedly when the network connection between Spark master and worker nodes is unstable. Possible workarounds include increasing the spark.-worker.timeout setting in the spark-env.sh script, updating Spark to v2.1.0, and the steps below. (RF-81)

  If this occurs:

  1. Kill all Spark applications.
  2. Check the number of Spark processes on each node.

     ```
     # ps ax | grep /opt/spark*
     ```

     There should be one spark-master process and no more than one spark-worker process (possibly zero).
  3. If there are more than two running processes, kill unneeded processes.

     ```
     # ps ax | grep /opt/spark* | grep -v grep | grep -v
     worker.Worker | grep -v master.Master | awk '{print $1}' |
     ```

```
xargs kill -9
```

4. Restart Tomcat

9.1.0

- When migrating jobs to Torque from Viewpoint, Moab will translate the request into the equivalent *qsub* command with the proper *-l procs* syntax. In some situations, Torque's queues may have been configured with a *default_resources.nodes* setting that is incompatible with the job's *-l procs* request. In this situation, the *default_resources.nodes* setting should be removed from the queue or the job should be submitted to a queue that does not have a *default_resources.nodes* setting.
- There is a known issue in RHEL/CentOS 7.2.1511 and 7.3.1611 (https://bugs.centos.org/view.php?id=12722) where the pip generates a Traceback in the pyparsing package. Check the version of pyparsing that is installed. Install pyparsing 2.0.3 or later. Either run 'yum update', 'yum update pyparsing', or locate and install pyparsing 2.0.3 or later.
- Job arrays sometimes do not expand to show subjobs correctly on the Workload and Node Details pages. (IRIS-2949)

Torque Resource Manager

6.1.2

- If you are using a multi-homed host (one that has multiple active network interfaces) as a submit 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 Differences). (TRQ-4072)
- Hwloc may not detect GPUs. When that occurs, they are reported on socket 0 numa node 0. (TRQ-3978)

6.1.1

- Updating pbs_server when GPU jobs are running can result in job issues. Make sure no GPU jobs are running before updating pbs_server. (DOC-3132)

6.1.0

- qstat -f accounting logs sometimes report "resources_used.vmem=0". This is due to how Linux reports memory usage. (TRQ-3760)

- Torque will not compile when the tk-devel and tcl-devel packages are installed on the host. To work around this, disable building of the gui component by using --disable-gui when executing configure. (TRQ-3723)

Reporting Framework

9.1.2

No known issues.

Automated Installer

9.1.2

No known issues.

9.1.1

- When installing Remote Visualization on a Red Hat-6 based system where you have configured separate Remote Visualization Session Servers, users cannot connect to the Remote Visualization jobs. This problem is remedied by a one-time restart of the FastX service on the Remote Visualization Gateway Host. (INST-121)

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:

- Moab Workload Manager
- Moab Accounting Manager
- Moab Web Services
- Moab Insight
- Moab Viewpoint
- Torque Resource Manager
- Reporting Framework
- Automated Installer
- Remote Visualization

Moab Workload Manager

9.1.2

- 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)
- mrsvctl -G gave incorrect nodes for certain shapes and origins. (MOAB-9017)
- Canceled jobs were no longer receiving JOBEND events. (MOAB-8975)
- 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)

9.1.1

- 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)
- Processes in Z (defunct) state were not being closed. (MOAB-8933)
- SCHEDCFG[] USEDATABASE was not enabled for HA installations. (MOAB-8926)
- NODESET definition in a class overrode reservation affinity. (MOAB-8903)
- Moab was not creating standing reservations on all partitions by default. (MOAB-8888)
- mdiag -S --xml did not display 5-minute statistics, as it did in non-xml mode. (MOAB-6896)

9.1.0

- Jobs with an invalid minwclimit were causing Moab to crash. (MOAB-8863)
- Moab was not reporting the correct task count for jobs that allocate ALLPROCS . (MOAB-8847)
- The systemd file generated with the --prefix option was not using the specified path. (MOAB-8727)
- systemd was terminating the active Moab server instead of the local Moab server. (MOAB-8667)
- Scheduler Analysis Block Reasons section displayed duplicate messages. (MOAB-8657)
- Task count was incorrect when requesting generic resource scheduling (GRES) on the global node. (MOAB-8627)
- Moab block messages were not always informative. (MOAB-8612)
- If jobs with a custom name were canceled, subsequent jobs with a similar custom name failed. (MOAB-8661)
- Jobs were rejected when requesting NUMA resources and NODEACCESSPOLICY-Y=SINGLEJOB. (MOAB-8565)
- msub did not allow you to specify a working directory, as with qsub -w <path>. (MOAB-8551)
- The class of GPU jobs that request more than one task was not being mapped correctly. (MOAB-8538)
- DPROC was not being calculated correctly on GPU jobs. (MOAB-8536)
- mcredctl was not reporting variables in credentials. (MOAB-8510)
- Jobs place on hold or batchhold were not being correctly recognized as being on hold. (MOAB-8433)
- Requeuing a job that was restarting caused it to be canceled. (MOAB-8422)
- Requesting a node's name as a feature prevents other jobs from getting that node. Jobs that request node names as features will now be placed on hold. (MOAB-8420)

- Issues reported for groups, users, and accounts with spaces in the name. Enabled SCHEDCFG[] FLAGS=AllowCredentialsWithSpaces to let Moab ignore POSIX standards and allow groups, users, and accounts with spaces in the name. (MOAB-8399)
- Jobs with the "AFTER" dependency were not auto-cancelling. (MOAB-8387)
- Job ID collisions were occurring in a grid environment. (MOAB-8385)
- Fixed issue with multi-req jobs and "EnforceReservedNodes". (MOAB-8105)
- In a Master/Slave grid configuration, the Master Moab did not honor the TTL attribute of jobs running on the slave node. (MOAB-7871)
- mdiag -C gave an error when specifying an SRCFG GRES resource. (MOAB-7547)

Moab Accounting Manager

9.1.2

- 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)

9.1.1

- NodeHours was not reported correctly (was reported in seconds instead of hours). (MAM-399)
- Web services did not permit condition operator expressions and other special characters in the query string. (MAM-395)
- Web service queries did not permit selection of aliases containing dashes. (MAM-393)
- Query offsets were not expressed independently from query limits(MAM-392)
- Future (DEPTH>1) liens did not charge the correct allocation. (MAM-390)
- mam-charge-retry.pl script failed for array jobs. (MAM-388)
- 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)

9.1.0

- Events did not log notifications to the Notification table by default, as documented.
- MAM GUI sometimes failed after an RPM upgrade due to changed ownership of the GUI log file. (MAM-373)
- Clients using the mamauth promotion method were intermittently returning an empty digest and signature causing the failure message "Incoming digest does not match calculated digest." (MAM-372, MAM-379)

- Multi-chunk queries sometimes caused server to time out and client to hang. (MAM-361)
- An event alarm scoping problem resulted in client "Connection reset by peer" failures, and in rare occasions by server failure. (MAM-340)
- The systemd unit file did not use the correct sbin directory for ExecStart. (MAM-338)

Moab Web Services

9.1.2

- The viewpoint-query-helper plugin was very slow. (WS-2569)
- MWS did not pass the number of GPUs requested to Moab correctly. (WS-2568)

9.1.1

- MWS could not handle a large volume of messages from ZMQ. (WS-2554)

9.1.0

- During a POST to /mws/rest/jobs, MWS now returns a 400 error if tasksPerNode field is defined but taskCount is not defined. (WS-2504)
- CPUTime in accounting usage record was being rendered as a date-time. (WS-2469)
- Field does not appear in the MWS job field documentation. Added javadoc for job field "submitCommandFile". (WS-2454)
- The MWS health check page did not correctly indicate LDAP DNs errors. (WS-2442)
- Performance when loading large LDAP trees with many attributes needed improvement. (WS-2254)

Moab Insight

9.1.2

- Insight Kafka `Linger Time in Milliseconds` and `Maximum Block Time in Milliseconds` fields were always blank on the MWS health page (IN-585)
- Insight did not report the status of the connection between Insight and Kafka to MWS. (IN-584)
- JobSyncService incorrectly converted database objects during sync with Moab Mongo. (IN-579)
- `sync.py` did not set the `completedDate` for jobs known to be terminated. (IN-578)
- Insight submitted multiple messages to the reporting framework for a completed job. (IN-570)
- The state of completed jobs was not set to completed in Insight's database. (IN-569)

9.1.1

- Destroyed Elastic nodes were still being counted when cluster/system samples were taken. (IN-563)

9.1.0

- Insight filed to process messages from Moab that had unexpected fields. (IN-507)

Moab Viewpoint

9.1.2

- Viewpoint showed completed jobs as active. (IN-562)
- File manager would not create a folder with a name containing a "%" character. (IRIS-3355)
- File manager did not validate the value entered in the field specifying the number of lines of a file to be displayed in a file preview. (IRIS-3354)
- File manager would not display or set permissions on files in folders with a "%" character in the name. (IRIS-3347)
- The incorrect page was displayed after the Application Templates page was refreshed. (IRIS-3343)
- The progress wheel was not displayed while loading the version history for an application template. (IRIS-3340)
- The template type filter for application templates did not work. (IRIS-3339)
- File manager would continue to display server validation messages when the Server Verify SSL was turned off. (IRIS-3335)
- Menu items on the Templates page would not open in a new tab. (IRIS-3324)
- The Create Template page did not validate the length of job names. (IRIS-3322)
- Long job names caused other page components to be obscured. (IRIS-3321)
- Trying to create a job template with the same name as an existing template caused an internal server error message to be displayed. (IRIS-3294)
- The Update Role page did not display a Reset Permissions button for all provided roles. (IRIS-3292)
- When the user name of the default Viewpoint user (e.g. viewpoint-admin) matched the name of the default MWS user (e.g. moab-admin), Viewpoint only displayed the Configuration and Licensed Features options in the menu. Recovery required manual modification of the database. (IRIS-3235)
- A warning message was not displayed when closing the Aggregated View designer with unsaved changes. (IRIS-3232)

- When jobs were created from an application template with total cores, nodes with core count, and malleable core count activated, malleable core count was the default selection rather than total cores. (IRIS-3160)
- The Node Details page displayed incorrect information after an elastic node was removed. (IRIS-3142)
- Job arrays sometimes did not expand to show subjobs correctly on the Workload and Node Details pages. (IRIS-2949)
- Viewpoint did not display an error message when attempting to create a principal with a blank name. (IRIS-2522)
- Viewpoint migrate displayed a misleading error message when `/proc/sys/fs/inotify/max_user_instances` was set too low. (IRIS-2374)

9.1.1.1

- The Job Details page loaded very slowly when there was a large number of jobs. (WS-2569)

9.1.1

- A Hold button was displayed for jobs with holds, rather than the expected Release button. (IRIS-3295)
- The Application Templates tab in the Configuration page showed templates for products that were not licensed. (IRIS-3287)
- On the Remote Visualization application template, it was not possible to make both the memory per core and memory per job options selectable. (IRIS-3285)
- The job script in the Remote Visualization application template included a Geometry request that had no effect. (IRIS-3283)
- Pagination on the Configuration Role Management page was sometimes incorrect when the number of entries per page setting was changed. (IRIS-3280)
- When creating a job, users could select an invalid generic resource. (IRIS-3275)
- The Remote Visualization application template was not handling generic resources correctly. (IRIS-3270)
- Roles and permissions for unlicensed features were displayed on the Configuration Page. (IRIS-3269)
- The Nodes Requested field of a job template did not work in pass-through mode. (IRIS-3264)
- The RemoteVizAdmin role was able to view other users' jobs. (IRIS-3256)
- Users were able to cancel other users' jobs. (IRIS-3255)
- Fixed an infinite loop on the Resource Job Timeline page. (IRIS-3251)
- The initial node status on the Resource Job Timeline page was not displayed correctly. (IRIS-3225)
- Closing a Remote Visualization session did not remove the session from the Sessions tab. (IRIS-3191)

- File Manager now considers all groups to which a user belongs when making decisions on access rights. (Previously, it only considered the user's primary group.) (IRIS-3161)

9.1.0

- File Manager displayed no feedback message after a copy/move operation. (IRIS-2593)
- When uploading files, File Manger overwrote files with the same name. (IRIS-2591)
- File Manager displayed an internal server error when a user without read permission attempted to read a file or its details. (IRIS-2542)
- File Manager did not handle Unicode file and directory names correctly. (IRIS-2511)
- Custom job submission script disappeared when Esc was pressed to exit the script editor. (IRIS-2450)
- Graphs on Home page were not displaying correctly when internet access was limited. (IRIS-2449)
- Create Job returned unclear error message when the ID of a new job matched the ID of an existing job. (IRIS-2355)

Torque Resource Manager

6.1.2

- Users were able to delete jobs without proper permissions. (TRQ-4161)
- 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)
- `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)
- `pbs_server` did not preserve GPU counts across restarts. (TRQ-3989)
- Job dependencies were not evaluated correctly for jobs in a routing queue. (TRQ-3951)
- `make rpm` failed with installed (but unpackaged) file(s) error (TRQ-3988)
- Torque could potentially hang when a job obtained a dependency on itself. (TRQ-3950)
- `qsub` did not correctly translate memory requirements between -l and -L. (TRQ-3934)
- Queueing two arrays and applying `max_user_queueable` caused the arrays to drop elements and not complete. (TRQ-3932)
- 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)

- pbs_mom ignored time specified in $varattr and updated dynamic values every 45 seconds. (TRQ-3901)
- pbs_server reported job aborts when jobs finished normally. (TRQ-3882)
- MOM daemons were preventing GPUs from being reset. (TRQ-3881)
- The /etc/profile.d/torque.csh script caused errors when installing from RPM files. (TRQ-3852)
- Output and error paths were not displayed properly when a job was canceled for exceeding its walltime. (TRQ-3793)
- NVIDIA device feature query failures were logged as errors if the feature was not present on the device. (TRQ-3766)
- make rpm did not use the configured prefix (TRQ-3716)
- qstat did not correctly list the job output path when qsub -k was used. (TRQ-3634)
- qsub did not enforce resources_max.procs correctly. (TRQ-3627)
- Notes on Cray compute nodes were not tracked after pbs_server restarts. (TRQ-3577)
- Sister nodes did not clean up job files when the jobs were no longer recognized by the mother superior. (TRQ-3524)
- 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)

6.1.1.1

- Jobs would not run after upgrading the server from 6.1.0 to 6.1.1 when using cgroups. (TRQ-3997)
- Torque installation failed on some multi-homed systems. (TRQ-3993)

6.1.1

- qsub did not properly validate the working directory (-w) path when VALIDATEPATH was unset or set to FALSE. (TRQ-3929)
- Jobs were sometimes lost during pbs_mom systemd unit file shutdown. (TRQ-3927)
- The shutdown function for pbs_server in the init.d script was sometimes failing. (TRQ-3924)
- A large number of acl_users could deadlock pbs_server. (TRQ-3910)
- The default_gpu_mode did not take effect after a restart. (TRQ-3885)
- Command-line directives did not override submit filter directives for interactive jobs. (TRQ-3862)
- GPU models k40s and k80s were not detected correctly when using GPUs and cgroups. (TRQ-3861)
- pbs_server crashed when loading jobs after a server shutdown. (TRQ-3839)

- stdout was not properly delivered on NFS file systems when Torque was configured with --disable-spool.(TRQ-3792)
- Torque was not allocating resources correctly for cgroup jobs. (TRQ-3790)
- Torque had communication issues when IPv6 was configured but not used. (TRQ-3773)
- Torque was not terminating all processes in a cgroup before removing it.. (TRQ-3753)
- Job holds were not updated when a slot limit was changed for a job array. (TRQ-2360)

6.1.0

- Systemd unit files do not correctly start/stop Torque daemons. A new -F (don't fork) switch has been added to pbs_server, pbs_mom, and trqauthd. (TRQ-3795)
- pbs_server was not starting when an array dependency was cleared before pbs_server started. (TRQ-3774)
- Torque was not waiting for a compute node to do a copy out process. (TRQ-3762)
- Fixed node locking issues that were causing pbs_server to abort when built with newer versions of the threading library. (TRQ-3755)
- Server build was failing with --disable-spool configure option. (TRQ-3751)
- Forced requeue command (qrerun -f) was not obeying kill_delay settings. (TRQ-3737)
- When adding a dynamic node, the server was initially putting the node in free state, but not reporting this state to the server. (TRQ-3733)
- The memory cgroup was not being set correctly for jobs requesting pmem and procs. (TRQ-3728)
- Jobs were not being removed from node_usage files after ending. (TRQ-3727)
- Jobs were becoming stuck after not starting correctly. (TRQ-3720)
- pbs_server aborted under ghost array recovery when job array files were missing. (TRQ-3719)
- momctl was not displaying the $varattr script or the correct time and date. (TRQ-3712) (TRQ-3617)
- pbs_mom crashed when started on the reporter node without $apbasil_protocol set in the config. (TRQ-3715)
- Torque was not allocating enough memory controllers to satisfy memory requests. (TRQ-3681)
- pbs_server was not being properly shut down when in HA mode. (TRQ-3670)
- pbs_server was not detecting and updating total_threads when a node's hyperthreading was enabled. (TRQ-3662)
- pbs_server was not properly restarted when running "service pbs_server restart" during installation. (TRQ-3657)
- Memory and swap limits were not set in cgroup. For information on how memory and swap options are used, see 1.1 -L NUMA Resource Request in the *Torque Resource Manager Administrator Guide*. (TRQ-3656)
- Jobs submitted with -l option with exclusive access to the node were not receiving all CPUs and memory controllers in the cgroup. (TRQ-3649)

- Corrected logging to only log that a signal is sent to a process when it is actually issued. (TRQ-3638)
- pmem was not getting set correctly. With cgroups enabled, pmem is the amount of resident memory allocated per process where are process is given by the value of ppn. For example: `qsub -l nodes=1:ppn=2,pmem=250mb` will allocate a total of 500 MB on the node where the job is run, 250 MB per ppn. (TRQ-3628)
- NUMA -L syntax defaulted to override user-specified parameters. (TRQ-3623)
- qstat -x returned nothing (instead of an empty XML document) when there are not jobs queued. (TRQ-3622)
- Jobs in which a task required more than one socket could not be started using NUMA -L syntax. (TRQ-3618)
- $usecp parameter was ignored when specifying which directories should be staged. (TRQ-3613)
- Server deadlocked when job_save() failed. (TRQ-3605)
- Tasks' memory usage was sometimes not reported. (TRQ-3601)
- Interactive jobs skipped submit filter directives if the first line was not #PBS. (TRQ-3585)
- Client commands were not attempting to contact the fallback server when the primary server was down. (TRQ-3582)
- Crash/infinite loop when loading certain node usage files. (TRQ-3576)
- Issue reported with alps login nodes. Updated cpusets for alps login nodes so that all of the cpus are in the job's cpuset. (TRQ-3568)
- Torque crashed intermittently when using the -L syntax. (TRQ-3566)
- Torque was not reporting nodes in the correct order. (TRQ-3559)
- Torque returned non-specific network failure messages to Moab. (TRQ-3539)
- Completed jobs were still reported in pbsnodes. (TRQ-3525)
- A deadlock occurred when handling job dependencies. (TRQ-3519)
- cgroup directories were not removed when jobs were completed. (TRQ-3515)
- drmaa unable to link with Torque. (TRQ-3511)
- Memory calculation issues reported when cgroups enabled and -l vmem|pmem|mem are used. (TRQ-3499)
- Epilogue not showing up in momctl -d3 output. (TRQ-3495)
- Job dependencies were not being cleared with High Availability server. (TRQ-3477)
- Array subjobs did not have a queued entry in the accounting log. (TRQ-3470)
- Segfault in create_alps_subnode with node_note populated. (TRQ-3445)
- Problems with clearing a node note. Removed length restriction on a node note. (TRQ-3439)
- Jobs that never ran were receiving end records. (TRQ-3432)
- Array templates were being reported as jobs. (TRQ-3405)
- A shell escape in pbs_mom's config file when specifying GRES did not show up in pbsnodes or Moab. (TRQ-3393)

- Resources_used.walltime changed to seconds from HH:MM:SS in accounting logs. (TRQ-3385)
- libtorque.so was not being created. (TRQ-3374)
- pbs_server timed out connection to pbs_mom. Added load balancing to login nodes when they start to get busy. (TRQ-3367)
- pbs_mom would hang when sending status from a child. Added a timeout for node health check scripts so that they cannot make the mom daemon hang. (TRQ-3364)
- pbs_mom hangs on restart with init script. Ensured that necessary services have been brought up before starting the Torque daemons and that the Torque daemons are shutdown before their required services are shutdown. (TRQ-3345)
- Numbered directories in server_priv/jobs (and arrays) were missing when the server attribute use_jobs_subdirs was set to TRUE. (TRQ-3185/)
- qrls gave no response and logged no problem when a failure occurred due to a slot limit restriction. (TRQ-3328)
- Fixed a memory leak when jobs were being started asynchronously. (TRQ-3326)
- qsub -W stage-in was not working. Fixed failures where the group name showed up in the log as the problem but the user did not belong to the group name given in the error. (TRQ-3312)
- Multiple moms sent invalid destroy_alps_reservation/req_delete_reservation. Only allows one kill orphaned reservation request per reservation at one time. (TRQ-3299)
- Problems building RPMs on Red Hat 6/CentOS 6 systems. (TRQ-3283)
- Jobs with square brackets in the name were aborted on restart if they weren't array subjobs. An issue was fixed with jobs getting aborted if they are named with "[]" in the name but aren't Torque array jobs. (TRQ-3214)
- Logs filled with messages about not sending hierarchy to mom. Failures are only logged the first time it can't send the hierarchy to a mom. (TRQ-3156)
- Jobs started even if mother superior could not resolve the hostname for a sister node. (TRQ-3134)
- qsig was not working correctly when display_job_server_suffix = false. (TRQ-3102)
- Error condition where the mom's port would be inserted into the .JB file name. (TRQ-3090)
- Torque was not able to release holds on job arrays. Running qrls on an array subjob allows pbs_server to correct slot limit holds for the array to which it belongs. (TRQ-3088)
- Down/offline nodes caused TORQUE to not online elastic nodes. pbs_server is now able to bring up new nodes even when there are nodes in the system that are down or offline. (TRQ-3066)
- Completed jobs were not getting cleaned up. Fixed various issues when restarting dependency jobs, including them not getting removed even after completion. (TRQ-3175)
- Several log messages were unclear. (TRQ-2860)
- qsub -v was not processing environment variables with no value correctly. (TRQ-2699)
- Node recovers when behind processing requests. pbs_server now detects when a node is failing too frequently and mark it down temporarily if this happens. Once a node is marked down, it will be marked up again if either two consecutive communications from pbs_server to the node receive successful replies, or after five minutes of staying offline (whichever comes first). A node is

considered to be failing too frequently if it has three failures to reply to a server request without have two consecutive successes in between. (TRQ-2517)

- qstat -f was not displaying init_work_dir. (TRQ-2459)
- Job holds were not updated when the slot limit was changed for a job array. (TRQ-2360)
- Torque enabled unused scheduler_iteration parameter by default. (TRQ-2161)

Reporting Framework

9.1.2

- The report creation page would duplicate the SQL table prefix when the user entered it manually. (RF-142)
- Spark stderr file size increased without limit. (RF-121)
- Default values of control panel variables were incorrect when reports added to dashboards. (RF-120)
- The page of data shown in a report in a dashboard did not match the page label. (RF-119)
- `Save and Close` and `Save and Run` buttons were disabled when returning to the page by clicking the Back button. (RF-111)
- Unauthorized users could see data for other users' aggregated views by editing the URL. (RF-94)
- New Spark applications sometimes started unexpectedly when the network connection between Spark master and worker nodes was unstable. (RF-81)
- In advanced mode, users were able to access aggregated views not shared with them. (RF-48)
- The Dashboard report control panel showed controls corresponding to unused variables. (IRIS-3310)
- In Firefox, the custom date range selector on the Edit Dashboard page would not allow the user to manually enter a date. (IRIS-3226)
- The Dashboards page did not correctly display dashboard descriptions containing special characters. (IRIS-3183)
- The drop-down list in the Dashboard Designer showed reports that had already been added to the dashboard. (IRIS-3091)
- After all aggregated views were deleted, the Spark reporting application repeatedly restarted. (IRIS-3013)
- When previewing a report in the Aggregated View Explorer, long titles or values did not display correctly. (IRIS-2743)

Automated Installer

9.1.2

- The installer failed when it was unable to download a pyparsing library on CentOS 7. (INST-156)

9.1.1.1

- The installer did not complete correctly when installing on Red Hat 7 on a physical machine. (INST-135)
- The installer did not count processors correctly on physical machines. (INST-133)

9.1.1

- CentOS 7.3.1611 introduced the following bug: https://bugs.centos.org/view.php?id=12722. The Automated Install installs a newer version of pyparsing as a workaround. (INST-59)
- Ansible install preflight check failed on RHEL-6 saying "No package libffi-devel available." (INST-54)
- mdiag -l showed no license because installer incorrectly names the Moab license file. (INST-21)
- Moab was not configured to use usage-tracking mode when accounting was enabled. (INST-18)
- Jobs were not seen in Viewpoint Workload – Insight was configured incorrectly when installed on a separate host. (INST-15)

Remote Visualization

9.1.2

- The Fastx service did not restart because the installer did not receive an LSB-compatible return code from the service status command. (RVIZ-9)
- The Remote Sessions page did not show new sessions until the page was manually refreshed. (IRIS-3296)

9.1.1

- Return code from service fastx status command was not LSB-compatible. (RVIZ-9)