

Moab HPC Suite

Installation and Configuration Guide 10.0.1 for Ubuntu-Based Systems

April 2024



Legal Notices

© 2015, 2024 Adaptive Computing Enterprises, Inc. All rights reserved.

Distribution of this document for commercial purposes in either hard or soft copy form is strictly prohibited without prior written consent from Adaptive Computing Enterprises, Inc.

This documentation and related software are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

This documentation and related software may provide access to or information about content, products, and services from third-parties. Adaptive Computing is not responsible for and expressly disclaims all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Adaptive Computing. Adaptive Computing will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Adaptive Computing.

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

Adaptive Computing Enterprises, Inc.
1100 5th Avenue South, Suite #201
Naples, FL 34102
+1 (239) 330-6093
www.adaptivecomputing.com

Contents

Introduction	5
Chapter 1: Planning Your Installation	6
1.1 Getting Started	6
1.1.1 Installation Terminology	7
1.1.2 Where to Start	7
1.2 Server Hardware Requirements	8
1.2.1 Topology	9
1.2.2 Hardware Requirements	12
1.2.3 Virtual Machines and Containers	17
1.3 Component Requirements	18
1.3.1 Torque	18
1.3.2 Moab Workload Manager	19
1.3.3 Moab Accounting Manager	20
1.3.4 RLM Server	20
1.4 RLM Licensing Model	21
1.4.1 Definitions	21
1.4.2 License File Names and Installation Locations	22
1.4.3 Licenses Issued by Adaptive Computing	23
Chapter 2: Manual Installation	25
2.1 Manual Installation Steps	25
2.1.1 Preparing for Manual Installation	26
2.1.2 Installing Torque Resource Manager	27
2.1.3 Installing Moab Workload Manager	32
2.1.4 Installing Moab Accounting Manager	37
2.1.5 Installing RLM Server	47
2.2 Additional Configuration	50
2.2.1 Opening Ports in a Firewall	50
2.2.2 Configuring SSL in Tomcat	50
2.2.3 Moab Workload Manager Configuration Options	50
2.2.4 Moab Accounting Manager Configuration Options	52
2.2.5 Using Multiple RLM Servers	54
2.2.6 Trusting Servers in Java	55
2.3 Manual Upgrade Steps	56
2.3.1 Upgrading Torque Resource Manager	57
2.3.2 Upgrading Moab Workload Manager	62

2.3.3 Upgrading Moab Accounting Manager	64
2.3.4 Upgrading RLM Server	67
2.3.5 Migrating the MAM Database from MySQL to PostgreSQL	68
Chapter 3: Troubleshooting	71
3.1 General Issues	71
3.1.1 Where do I set credentials and what are the default values?	71
3.2 Port Reference	77
3.3 Moab Workload Manager Issues	81
3.3.1 Moab error: cannot determine local hostname	81
3.3.2 Moab error: Moab will now exit due to license file not found	82

Introduction

Welcome to the *Moab HPC Suite 10.0.1 Installation and Configuration Guide* for Ubuntu-Based Systems

This guide includes detailed instructions for installing each component of the suite so that you can quickly get up and running.

This guide is intended for system administrators who are responsible for installing the Moab HPC Suite components.


 Depending on your system configuration and license, not all of the HPC Suite components may be available.

The 10.0.1 Moab HPC Suite contains the following components for Ubuntu-based systems:


- Torque Resource Manager 7.0.1
- Moab Workload Manager 10.0.1
- Moab Accounting Manager 10.0.1

Before commencing the installation or upgrade, see [Chapter 1: Planning Your Installation - page 6](#) to verify your system conforms to minimum prerequisites.

Chapter 1: Planning Your Installation

 We highly recommend that you *first* perform installations and upgrades in a *test environment*. Standard installation and upgrade procedures and use cases are tested prior to release. However, due to the wide range of possible configurations and customizations, it is important to exercise caution when deploying new versions of software into your production environments. This is especially true when the workload has vital bearing on your organization's day-to-day operations. We recommend that you test in an environment that mirrors your production environment's configuration, workflow and load as closely as possible. Contact your Adaptive Computing account manager for suggestions and options for installing/upgrading to newer versions.

There are many different ways to install and configure the Moab HPC Suite. Each environment has its own set of requirements and preferences. This chapter is intended to help an administrator understand how each of the Moab HPC Suite components interact, basic requirements, and configuration information to prepare for the installation.

 Code samples have been provided for convenience. Some code samples provide sample passwords (i.e., changeme!). We strongly recommend that you do not use these passwords during installation, as using the documented passwords could introduce unnecessary security vulnerabilities into your system.

In this chapter:

1.1 Getting Started	6
1.2 Server Hardware Requirements	8
1.3 Component Requirements	18
1.4 RLM Licensing Model	21

1.1 Getting Started

In this topic:

[1.1.1 Installation Terminology - page 7](#)

[1.1.2 Where to Start - page 7](#)

1.1.1 Installation Terminology

To aid in documentation clarity, Adaptive Computing uses the following terms in this Installation and Configuration Guide:

- **Components** – The different 'products' included in the Moab HPC Suite. For example, Moab Workload Manager, Moab Web Services.
- **Servers** – Also known as components, but specifically relating to the actual services. For example, the Moab Workload Manager component is referred to as the Moab Server for non-client services.
- **Host** – The actual box where a Moab HPC Suite component (server or client) is installed.

i Previous documentation typically used Head Node to designate a host or a Server.

1.1.2 Where to Start

You need to plan your environment and determine how many hosts you will need and for which components you will install using the Manual Installation method. The following are suggested steps to help you in your planning and installing process.

1. Determine whether you have a small, medium, High-Throughput or large environment; including an example, and required and recommended hardware requirements. See [1.2 Server Hardware Requirements - page 8](#).
2. Decide whether you will perform a Manual Installation for the various components. See [1.1 Identify the Manual Methods](#).

i The Manual Installation chapters have an 'Additional Configuration' section that provides additional information and instructions for optional, but recommended configurations (for example, Configuring SSL in Tomcat).

3. Review the software requirements for your components and set up your hosts accordingly. See [1.3 Component Requirements - page 18](#).
4. Install the individual components on their respective host(s). See [2.1.1 Preparing for Manual Installation - page 26](#) as applicable.

5. Refer to [Chapter 3: Troubleshooting - page 71](#) for assistance in addressing common problems during installation and configuration.

1.2 Server Hardware Requirements

The Moab HPC Suite is installed and configured differently for small, medium, or large environment types. This topic provides a general topology of the Moab HPC Suite and the server hardware requirements depending on your environment size.

In this topic:

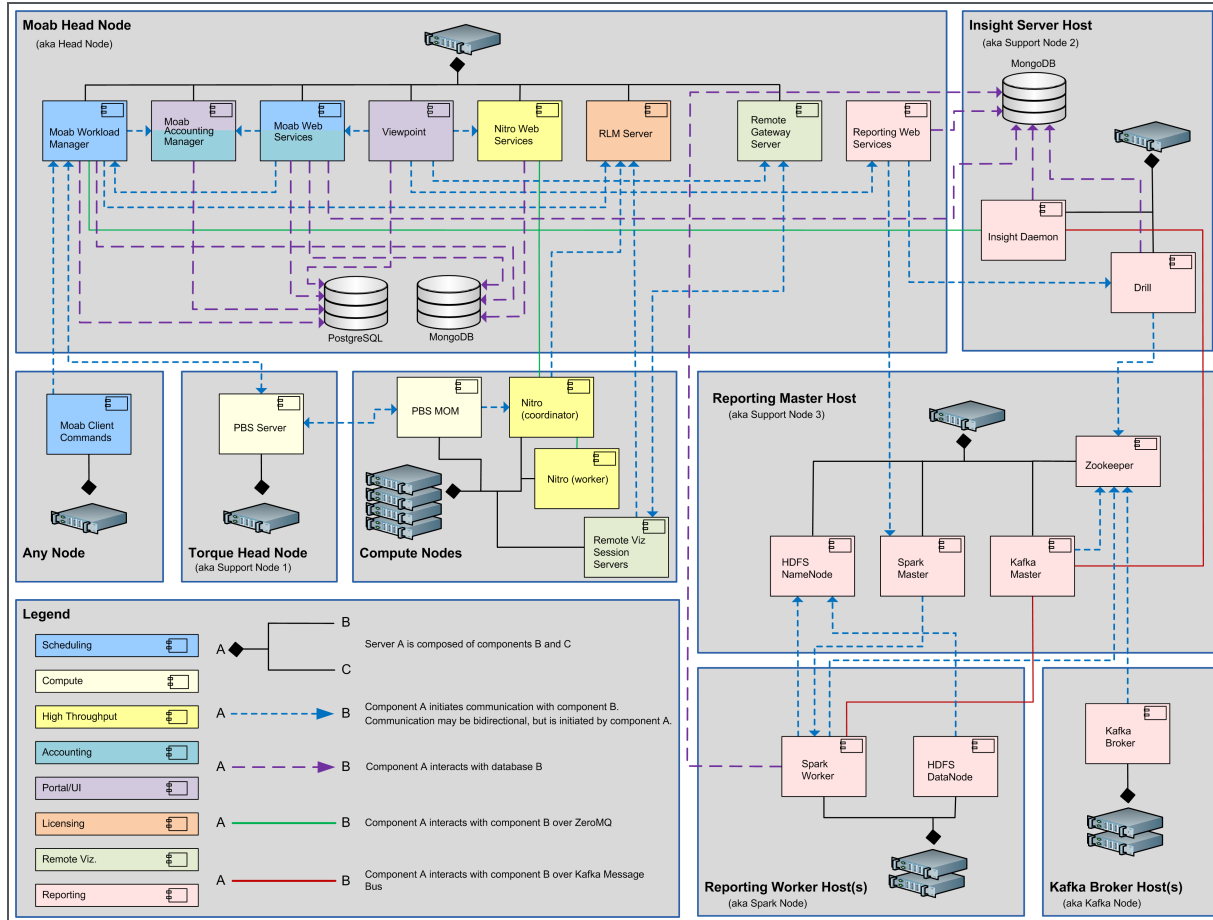
[1.2.1 Topology - page 9](#)

[1.2.2 Hardware Requirements - page 12](#)

[1.2.4 Virtual Machines and Containers](#)

1.2.1 Topology

The following diagram provides a general topology of the Moab HPC Suite for a medium (with high throughput) or a large environment:



Note the following:

- Smaller environments may elect to consolidate the Torque Server with the Moab Server on the same host, including PBS Server in the list of components installed on the same host.
- Although Moab Workload Manager and Moab Accounting Manager can share the same database instance, it is not a requirement. Two database instances can be used, one for each component.
- Larger systems will require more dedicated resources for each component, in which case it may be necessary to move individual components from the Moab Server Host (i.e., databases, Moab Accounting Manager, and/or Viewpoint) to their own respective servers.

Software components that may be included in a Moab HPC Suite installation are described in the table below:

Component	Description
Moab Workload Manager	A scheduling and management system designed for clusters and grids.
Torque Resource Manager - PBS Server	A resource manager for Moab HPC Suite. Torque provides the low-level functionality to discover and report cluster resources/features, and to start, hold, cancel, and monitor jobs. Required by Moab Workload Manager.
Torque Resource Manager - PBS MOM	Torque MOMs (Machine Oriented Mini-server) are agents installed on each compute node that complete tasks assigned to them by the Torque Server. When a multi-node job runs, one of the Torque MOMs is assigned the role of Mother Superior and all other nodes assigned to the job are sister nodes. Mother Superior manages the job across all the sister nodes by communicating with each of them and updating the Torque Server. Required by Torque.
Moab Accounting Manager	An accounting management system that allows for usage tracking, charge accounting, and allocation enforcements for resource usage in technical computing environments. Required by Moab Workload Manager and Moab Web Services.
Moab Web Services (MWS)	A component of the Moab HPC Suite that enables programmatic interaction with Moab Workload Manager via a RESTful interface. MWS lets you create and interact with Moab objects and properties such as jobs, nodes, virtual machines, and reservations. MWS is the preferred method for those wanting to create custom user interfaces for Moab and is the primary method by which Moab Viewpoint communicates with Moab HPC Suite. Required by Moab Viewpoint.
Reprise License Manager Server (RLM)	A flexible and easy-to-use license manager with the power to serve enterprise users. Required by Moab HPC Suite and Nitro.
Moab Insight	A component of the Moab HPC Suite that collects the data that Moab HPC Suite emits on its message queue and stores it in a database. The message queue is efficient, can be encrypted, and tolerates disconnections and restarts on either side. Required by Moab Viewpoint and Kafka Master.
Nitro	A highly powerful, yet simple task launching solution that operates as an independent product, but can also integrate seamlessly with any HPC scheduler. In the Moab HPC Suite, Nitro is fully integrated with Viewpoint for seamless

Component	Description
	high-throughput job submission and monitoring.
Nitro Web Services	Enables programmatic interaction with Nitro for obtaining Nitro job status information within Viewpoint. Required by Moab Viewpoint.
Moab Viewpoint	A rich, easy-to-use portal for end-users and administrators, designed to increase productivity through its visual web-based interface, powerful job management features, and other workload functions. The portal provides greater self-sufficiency for end-users while reducing administrator overhead in High-Performance Computing. Nitro, and Reporting and Analytics features are also licensable for use with Viewpoint.
Reporting Web Services (RWS)	A component of Adaptive Computing Suites that enables programmatic interaction with Moab Reporting and Analytics via a RESTful interface. RWS is the preferred method for those wanting to create custom user interfaces for Moab Reporting and Analytics and is the primary method by which Moab Viewpoint communicates with Moab Reporting and Analytics.
Reporting and Analytics	Streams in massive amounts of workload and resource usage data from your High-Performance Computing (HPC), High-Throughput Computing (HTC) and Grid Computing environments, and then correlates that information against users, groups, and accounts, organizations so you can gain insights into exactly how your investment is being used and how well it aligns with your goals.
MongoDB	A free and open-source cross-platform document-oriented database program. Classified as a NoSQL database program, MongoDB uses JSON-like documents with schemas. Required by Moab Workload Manager, Moab Web Services, Nitro Web Services, Reporting Web Services, and Spark Worker.
PostgreSQL	An object-relational database (ORDBMS). That is, an RDBMS, with additional (optional use) object features – with an emphasis on extensibility and standards compliance. Required by Moab Workload Manager, Moab Accounting Manager, Moab Web Services, and Moab Viewpoint.
Drill	Apache Drill is an open-source software framework that supports data-intensive distributed applications for interactive analysis of large-scale datasets. Required by Reporting Web Services.
Hadoop	The Apache Hadoop software library is a framework that allows for the distributed processing of large data sets across clusters of computers using simple programming models. It is designed to scale up from single servers to thousands of machines, each offering local computation and storage. Rather than rely on hardware to deliver high availability, the library itself is designed to detect and handle failures at the application layer, so delivering a highly

Component	Description
	available service on top of a cluster of computers, each of which may be prone to failures. Required by Spark Worker.
Spark Master	Apache Spark is a fast and general engine for large-scale data processing. Spark Streaming is an extension of the core Spark API that enables scalable, high-throughput, fault-tolerant stream processing of live data streams. The Spark Master uses one or more Spark Workers when processing live data streams. Data can be ingested from many sources like Kafka, Flume, Kinesis, or TCP sockets, and can be processed using complex algorithms expressed with high-level functions like map, reduce, join and window. Finally, processed data can be pushed out to filesystems, databases, and live dashboards. Required by Reporting Web Services.
Spark Worker	The Spark Worker is used by a Spark Master when processing live data streams. Required by Spark Master.
Kafka Master	Apache Kafka is used for building real-time data pipelines and streaming apps. It is horizontally scalable, fault-tolerant, wicked fast, and runs in production in thousands of companies. Kafka Master uses one or more Kafka Brokers when pipelining and processing live data streams. Required by Spark Worker, and Insight.
Kafka Broker	Kafka Broker is used by a Kafka Master to pipeline and process live data streams. Apache Kafka is used for building real-time data pipelines and streaming apps. It is horizontally scalable, fault-tolerant, wicked fast, and runs in production in thousands of companies. Required by Kafka Master.

1.2.2 Hardware Requirements

The following tables show hardware requirements for Moab HPC Suite, Torque, and Reporting Framework environments of various deployment sizes.

1.2.2.A Moab HPC Suite and Torque Requirements

The following table identifies the minimum and recommended hardware requirements for the different environment types. Use this table as a guide when planning out your suite topology.

i Software requirements are listed per-component rather than suite-wide as the suite components reside on different hosts. See [1.3 Component Requirements - page 18](#)

Environment Type	# of Compute Nodes	Jobs/Week	Minimum Requirements (per Host Distribution)	Recommended Requirements (targeting minimum number of hosts)
Proof of Concept / Small Demo	50	<1k	<p>Moab Server+Torque Server Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 cores • At least 8 GB RAM • At least 100 GB dedicated disk space <p>Insight Server Host</p> <ul style="list-style-type: none"> • 8 Intel/AMD x86-64 cores • At least 16 GB RAM • At least 512 GB dedicated disk space 	Same as minimum
Medium	500	<100k	<p>Moab Server+Torque Server Host</p> <ul style="list-style-type: none"> • 8 Intel/AMD x86-64 cores • At least 16 GB RAM • At least 512 GB dedicated disk space 	<p>Moab Server+Torque Server Host</p> <ul style="list-style-type: none"> • 16 Intel/AMD x86-64 cores • At least 32 GB RAM • At least 1 TB dedicated disk space <p>Insight Server Host</p>

Environment Type	# of Compute Nodes	Jobs/Week	Minimum Requirements (per Host Distribution)	Recommended Requirements (targeting minimum number of hosts)
			Insight Server Host <ul style="list-style-type: none"> 8 Intel/AMD x86-64 cores At least 16 GB of RAM At least 1024 GB disk 	<ul style="list-style-type: none"> 8 Intel/AMD x86-64 cores At least 16 GB of RAM dedicated 1 Gbit channel between Insight and Moab 128 GB local SSD for swap At least 1024 GB disk
Medium with High Throughput or Larger	>500	>100k	Moab Server Host <ul style="list-style-type: none"> 8 Intel/AMD x86-64 cores At least 16 GB RAM At least 512 GB dedicated disk space Torque Server Host <ul style="list-style-type: none"> 8 Intel/AMD x86-64 cores At least 16 GB RAM At least 512 GB dedicated disk space 	<p>The Moab Server should <i>not</i> reside on the same host as the Torque Server.</p> <p>MWS Server <i>must</i> reside on the same host as the Moab Server (Moab Server Host).</p> <p>The MAM Server can reside on its own host, on the Moab Host (preferred), or another server's host (except for the Insight Host).</p> <p>The Viewpoint Server can reside on its own host, on the Moab Server Host (preferred), or another server's host (except for the Insight Server Host).</p> <p>Databases can also reside on the same or a different host from its server component.</p>

Environment Type	# of Compute Nodes	Jobs/Week	Minimum Requirements (per Host Distribution)	Recommended Requirements (targeting minimum number of hosts)
			Insight Server Host <ul style="list-style-type: none"> • 8 Intel/AMD x86-64 cores • At least 16 GB of RAM • At least 2048 GB disk 	

Note the following:

- All requirements above (minimum and recommended) target a minimum number of management servers. Administrators are encouraged to separate the Torque Server and the Moab Server onto different hosts where possible for better results; especially when High Throughput is enabled.
- Although many factors may have an impact on performance (network bandwidth, intended use and configuration, etc.), we consider High Throughput as something that makes a significant enough difference between minimum and recommended hardware requirements to merit mention in the table above.
- Moab and Torque are both multi-threaded and perform better with more processors.
- Due to the large amount of data Moab must send to Insight, Moab performs better without Insight enabled (for environments that do not require Viewpoint, or use Crystal Reporting).
- Regarding disk space, consideration should be given to requirements related to log files, log depth, number of jobs/nodes/reservations (more objects impact database journal size), average number of events generated (more events take more space), etc.

1.2.2.B Reporting Framework Requirements

The following table shows hardware requirements for the Reporting and Kafka hosts needed to support the addition of the Reporting Framework to a Moab HPC Suite environment. These requirements are *in addition* to the requirements shown in the table above.

Environment Type	Minimum Requirements (per Host Distribution)	Recommended Requirements (targeting minimum number of hosts)
Proof of Concept / Small Demo	<p>Reporting Master Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 cores • At least 8 GB RAM • At least 512 GB dedicated disk space <p>Reporting Worker Host</p> <ul style="list-style-type: none"> • 8 Intel/AMD x86-64 cores • At least 16 GB RAM • At least 512 GB dedicated disk space <p>Kafka Broker Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 cores • At least 6 GB RAM • At least 512 GB dedicated disk space 	Same as minimum
Medium	<p>Reporting Master Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 cores • At least 8 GB RAM • At least 1024 GB dedicated disk space <p>Reporting Worker Host</p> <ul style="list-style-type: none"> • 8 Intel/AMD x86-64 cores • At least 16 GB RAM • At least 512 GB dedicated disk space <p>Kafka Broker Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 	<p>Reporting Master Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 cores • At least 16 GB RAM • At least 1024 GB dedicated disk space <p>Reporting Worker Host</p> <ul style="list-style-type: none"> • 8 Intel/AMD x86-64 cores • At least 32 GB RAM • At least 512 GB dedicated disk space <p>Kafka Broker Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 cores • At least 6 GB RAM

Environment Type	Minimum Requirements (per Host Distribution)	Recommended Requirements (targeting minimum number of hosts)
	<ul style="list-style-type: none"> cores • At least 6 GB RAM • At least 1024 GB dedicated disk space 	<ul style="list-style-type: none"> • At least 1024 GB dedicated disk space
Medium with High Throughput or Larger	<p>Reporting Master Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 cores • At least 16 GB RAM • At least 2048 GB dedicated disk space <p>Reporting Worker Host</p> <ul style="list-style-type: none"> • 8 Intel/AMD x86-64 cores • At least 32 GB RAM • At least 512 GB dedicated disk space <p>Kafka Broker Host</p> <ul style="list-style-type: none"> • 4 Intel/AMD x86-64 cores • At least 6 GB RAM • At least 2048 GB dedicated disk space 	More than one Reporting Worker hosts is recommended.

1.2.3 Virtual Machines and Containers

The Moab HPC suite can be installed on virtual machines. There are both free and commercial virtual machine platforms available. However, installing the Moab HPC Suite on a container (e.g., Docker, Singularity, or LXC) has been known to be problematic and is not officially supported.

1.3 Component Requirements

In this topic:

[1.3.1 Torque - page 18](#)

[1.3.2 Moab Workload Manager - page 19](#)

[1.3.3 Moab Accounting Manager - page 20](#)

[1.3.4 RLM Server - page 20](#)

1.3.1 Torque



If you intend to use Torque 7.0.1 with Moab Workload Manager, you must run Moab version 8.0 or later. However, some Torque functionality may not be available. See the *Moab HPC Suite Release Notes* for more information.

1.3.1.A Supported Operating Systems

- Red Hat 7, 8
- SUSE 12, 15
- Ubuntu 18.04, 20.04, 22.04



Elastic Computing is not supported on Ubuntu.

1.3.1.B Software Requirements

- libxml2-devel package (package name may vary)
- openssl-devel package (package name may vary)
- Tcl/Tk version 8 or later if you plan to build the GUI portion of Torque, or use a Tcl-based scheduler
- cgroupv1:
 - cgroupv1 is recommended for the tarball install
 - cgroupv2 is not yet supported by Torque

- If you build Torque from source, the following additional software is required:
 - gcc
 - gcc-c++
 - posix-compatible version of make
 - libtool 1.5.22 or later
 - boost-devel 1.36.0 or later

1.3.2 Moab Workload Manager

1.3.2.A Supported Operating Systems

- Red Hat 7, 8
- SUSE 12, 15
- Ubuntu 18.04, 20.04, 22.04

1.3.2.B Software Requirements

- libcurl (<https://curl.haxx.se/libcurl/>)
- Perl 5.8.8 or later
- perl-CPAN (package name may vary)
- libxml2-devel (package name may vary)
- *(Optional)* Moab Accounting Manager
- *(Optional)* MySQL, PostgreSQL, or Oracle with ODBC driver (see 'Database Configuration' in the *Moab Workload Manager Administrator Guide* for details)

1.3.2.C Supported Resource Managers

- Torque

1.3.3 Moab Accounting Manager

i MAM is commonly installed on the same host as Moab Workload Manager; however, in some cases you might obtain better performance by installing them on different hosts.

1.3.3.A Supported Operating Systems

- Red Hat 7, 8
- SUSE 12, 15
- Ubuntu 18.04, 20.04, 22.04

1.3.3.B Software Requirements

- gcc
- perl-suidperl
- httpd
- mod_ssl
- rrdtool
- Moab Workload Manager 10.0.1
- Perl modules; see [2.1.4 Installing Moab Accounting Manager - page 37](#) (Manual Installation) [1.1 Installing Moab Accounting Manager](#) for more details

1.3.3.C Depends On (not necessarily on the same host)

MAM uses an RDBMS as a back end. We recommend that the database used by MAM does *not* reside on the same host as the database used by Insight.

- PostgreSQL 7.2 or later


1.3.4 RLM Server

Moab's feature Nitro requires access to a centralized Reprise License Manager (RLM) server.

Adaptive Computing *strongly* recommends that your RLM Server is the latest version.

This server is not load-extensive so it can be installed on any host within your Moab HPC Suite environment. It can also be installed on its own host.

 If your company already utilizes an RLM Server, you do not have to install another as long as the Moab HPC Suite components can access it.

 The host on which you install RLM Server must always be on and should have High Availability (uptime).

1.3.4.A Supported Operating Systems

- Red Hat 7, 8
- SUSE 12, 15
- Ubuntu 18.04, 20.04, 22.04

1.4 RLM Licensing Model

The Moab HPC Suite uses Reprise License Manager (RLM) to enforce nodelocked and floating licenses.

In this topic:

[1.4.1 Definitions - page 21](#)

[1.4.2 License File Names and Installation Locations - page 22](#)

[1.4.3 Licenses Issued by Adaptive Computing - page 23](#)

1.4.1 Definitions

Term	Definition
Nodelocked License	A nodelocked license is a license grant that allows software to be used on a particular computer, and on that computer only. A license server is not required to enforce a nodelocked license. However, a unique unmodifiable ID unique to the computer is required.
Floating License	Floating licensing is a software licensing approach in which a limited number of licenses for a software application are shared among a larger number of

Term	Definition
	users/clients over time. When an authorized user wants to run the application, they request a license from a central license server. A license server is used for floating licenses and logging of usage data.
RLM HostID	<p>The host ID of the RLM server. Adaptive Computing requests that you provide the 'ether' RLM HostID when the RLM HostID is required.</p> <p>RLM's <i>rlmhostid</i> command can be used to obtain the 'ether' RLM HostID. However, the <i>rlmhostid</i> command may not be available on every host (typically only available where RLM server is installed). The <i>rlmhostid</i> command takes an optional parameter <i>hostid type</i>, which defaults to <i>ether</i> when not given.</p> <p>When the <i>rlmhostid</i> command is unavailable, provide the Ethernet MAC address of one of the host's network interface cards (NICs).</p>

1.4.2 License File Names and Installation Locations

You will be provided license files with file names depicted below. Install each file in the designated Destination Directory and rename the files to the Installed File Name. When issued RLM-signed licenses, technically, the only Installed File Name requirement is that the file name end in `.lic`. However, we suggest you use the Installed File Name.

i Do not install more than one of each of the following licenses in the respective Destination Directory. Having more than one `moab-rlm--xxxxx.lic` file in `/opt/moab/etc`, for example, will cause problems. Therefore, the suggested Installed File Names.

File Description	File Name	Destination Directory	Installed File Name	Host
Moab Workload Manager (version >= 9.1.0)	<code>moab-rlm--xxxxx.lic</code>	<code>/opt/moab/etc</code>	<code>moab-rlm.lic</code>	Moab Server Host
Moab Workload Manager (version < 9.1.0)	<code>moab_xxxx.lic</code> <code>moab_l_xxxx.lic</code> <code>moab_c_xxxx.lic</code>	<code>/opt/moab/etc</code>	<code>moab.lic</code>	Moab Server Host

File Description	File Name	Destination Directory	Installed File Name	Host
Viewpoint	moab-vp--xxxxx.lic	/opt/viewpoint /etc/license.d	license.lic	Viewpoint Host
Nitro	nitro-rlm--xxxxx.lic	/opt/rlm/	nitro.lic	RLM Host
Nitro	nitro-rlm--xxxxx.lic	/opt/nitro/etc	nitro.lic	Nitro Coordinator Host - All Nitro Compute Nodes
Adaptive Computing ISV RLM	adaptiveco.set adaptiveco.opt	/opt/rlm	adaptiveco.set adaptiveco.opt	RLM Host
StarNet ISV RLM	starnet.set	/opt/rlm	starnet.set	RLM Host

1.4.3 Licenses Issued by Adaptive Computing

Instructions for obtaining licenses from Adaptive Computing are given below.

1.4.3.A Nodelocked/Uncounted Licenses

Moab HPC Suite License

This license is limited by the number of sockets the scheduler will schedule. The socket count includes physical sockets, GPUs, and Xeon Phis in your cluster.

Moab does not delegate license enforcement by socket to an external licensing server. Doing so allows Moab to give you visibility into all resources reported by your resource managers, while scheduling within the limits of your licensing agreement.

Send the following to licenses@adaptivecomputing.com to obtain a Moab License:

- RLM Hostid of the Moab Server host
- Total number of physical sockets, GPUs, and Xeon Phis in your cluster

Install this license as `/opt/moab/etc/moab-rlm.lic` on the Moab Server host.

Viewpoint License

Send the following to licenses@adaptivecomputing.com to obtain a Viewpoint License:

- RLM Hostid of the Viewpoint Server host

Install this license as `/opt/viewpoint/etc/license.d/license.lic` on the Viewpoint Server host.

1.4.3.B Floating/Counted Licenses

Nitro License

Nitro can be licensed based on number of concurrent Nitro Workers or number of cores.

Send the following to licenses@adaptivecomputing.com to obtain a Nitro License:

- RLM Server Hostname
- RLM Hostid of the RLM server
- Number of Nitro Workers or cores

Chapter 2: Manual Installation

This chapter provides installation, configuration, and upgrading information using the Manual Installation method.

Be aware of the following:

- Manual Installation is not available for Insight or Viewpoint.
- Because many system-level files and directories are accessed during the installation, the instructions in this guide should be executed with root privileges. You will see that the instructions execute commands as the root user. Also be aware that the same commands will work for a non-root user with the `sudo` command.

In this chapter:

2.1 Manual Installation Steps	25
2.2 Additional Configuration	50
2.3 Manual Upgrade Steps	56

Related Topics

- [Chapter 1: Planning Your Installation](#) - page 6
- [2.1.1 Preparing for Manual Installation](#) - page 26

2.1 Manual Installation Steps

This section provides instructions and other information for installing your Moab HPC Suite components using the Manual installation method.

In this section:

- [2.1.1 Preparing for Manual Installation](#) - page 26
- [2.1.2 Installing Torque Resource Manager](#) - page 27
- [2.1.3 Installing Moab Workload Manager](#) - page 32
- [2.1.4 Installing Moab Accounting Manager](#) - page 37
- [2.1.5 Installing RLM Server](#) - page 47

2.1.1 Preparing for Manual Installation

The manual installation process of the Moab HPC Suite includes installing the different components in the suite.

i Many individual components have dependencies on other components (see [Chapter 1: Planning Your Installation - page 6](#)). However, if you do not require a certain component, you do not have to install it.

The install instructions for each component include information about system requirements and dependencies. Some include prerequisite instructions that you will need to complete before you begin the install. Read this information carefully, and make sure you have installed all the dependencies and packages that are necessary in order to avoid errors during the Moab HPC Suite install process.

i Because many system-level files and directories are accessed during the installation, the instructions in this guide should be executed with root privileges.

You will see that the instructions execute commands as the root user. Note that the same commands will work for a non-root user with the `sudo` command.

This topic contains prerequisite instructions that you will need to complete before you begin the installations.

In this topic:

[2.1.1.A Set Up Proxies - page 26](#)

[2.1.1.B Update Your System Software to the Latest Version - page 27](#)

[2.1.1.C Ensure Hostname Resolution for All Hosts - page 27](#)

[2.1.1.D Install the Moab HPC Suite Software Components - page 27](#)

2.1.1.A Set Up Proxies

If your site uses an external repository to install python dependencies (for example, the host where you install Viewpoint might need to download extra packages), set up pip to use a proxy. Do the following:

```
export https_proxy=https://<proxy_server_id>:<port>
```

2.1.1.B Update Your System Software to the Latest Version

We recommend that you update your system software to the latest version before installing Moab HPC Suite components.

On *each* host where you will install the Moab HPC Suite components, run the following command:

```
[root]# apt upgrade
```

2.1.1.C Ensure Hostname Resolution for All Hosts


Each host should be resolvable from all other hosts in the cluster. Usually this is implemented by having all hosts in DNS. Alternatively, each host can include all other hosts (with the correct IP address) in its `/etc/hosts` file.

2.1.1.D Install the Moab HPC Suite Software Components

To install the Moab HPC Suite, install the packages in the following order:

1. Torque. See [2.1.2 Installing Torque Resource Manager - page 27](#).
2. Moab Workload Manager. See [2.1.3 Installing Moab Workload Manager - page 32](#).
3. Moab Accounting Manager. See [2.1.4 Installing Moab Accounting Manager - page 37](#).

2.1.2 Installing Torque Resource Manager

 If you intend to use Torque Resource Manager 7.0.1 with Moab Workload Manager, you must run Moab HPC Suite version 8.0 or later. However, some Torque functionality may not be available. See the *Moab HPC Suite Release Notes* for more information.

This topic contains instructions on how to install and start Torque Resource Manager (Torque).

In this topic:

[2.1.2.A Open Necessary Ports - page 28](#)

[2.1.2.B Install Dependencies, Packages, or Clients - page 28](#)

- [2.1.2.C Install Torque Server - page 29](#)
- [2.1.2.D Install Torque MOMs - page 30](#)
- [2.1.2.E Install Torque Clients - page 31](#)
- [2.1.2.F Configure Data Management - page 32](#)

2.1.2.A Open Necessary Ports

Torque requires certain ports to be open for essential communication.

If your site is running firewall software on its hosts, you will need to configure the firewall to allow connections to the necessary ports:

Location	Ports	Functions	When Needed
Torque Server Host	15001	Torque Client and MOM communication to Torque Server	Always
Torque MOM Host (Compute Nodes)	15002	Torque Server communication to Torque MOMs	Always
Torque MOM Host (Compute Nodes)	15003	Torque MOM communication to other Torque MOMs	Always

If using the MOM hierarchy (documented in 'Setting Up the MOM Hierarchy' in the *Torque Resource Manager Administrator Guide*), you must also open port 15003 from the server to the nodes.

See also:

- [2.2.1 Opening Ports in a Firewall - page 50](#) for general instructions and an example of how to open ports in the firewall.
- 'Configuring Ports' in the *Torque Resource Manager Administrator Guide* for more information on how to configure the ports that Torque uses for communication.

2.1.2.B Install Dependencies, Packages, or Clients

Install Packages

On the Torque Server Host, run the following command to install the `libxml2-devel`, `openssl-devel`, and `boost-devel` packages:

```
[root]# apt install build-essential libxml2-dev libssl-dev libboost-dev libtool pkg-
config zlib1g-dev
```

Install hwloc

On the Torque Server Host, each Torque MOM Host, and each Torque Client Host, install the hwloc development package:

```
[root]# apt install libhwloc-dev
```

2.1.2.C Install Torque Server

i You must complete the tasks to install the dependencies, packages, or clients before installing Torque Server. See [2.1.2.B Install Dependencies, Packages, or Clients - page 28](#).

If your configuration uses firewalls, you *must also* open the necessary ports before installing the Torque Server. See [2.1.2.A Open Necessary Ports - page 28](#).

On the Torque Server Host, do the following.

1. Download the latest Torque tarball from the [Adaptive Computing](#) website:

```
[root]# tar -xzvf torque-7.0.1.tar.gz
[root]# cd torque-7.0.1/
```

2. Determine which `./configure` command options you need to add, based on your system configuration. At a minimum, you add: `--enable-cgroups`

i These instructions assume you are using `cgroups`. When `cgroups` are supported, `cpusets` are handled by the `cgroup cpuset` subsystem. If you are not using `cgroups`, use `--enable-cpusets` instead.

i If `--enable-gui` is part of your configuration, run the following commands:

```
$ cd /usr/lib64
$ ln -s libXext.so.6.4.0 libXext.so
$ ln -s libXss.so.1 libXss.so
```

When finished, `cd` back to your install directory.

See 'Customizing the Install' in the *Torque Resource Manager Administrator Guide* for more information on which options are available to customize the `./configure` command.

3. Run each of the following commands:

```
[root]# ./configure --enable-cgroups # add any other specified options
[root]# make
[root]# make install
```

4. Source the appropriate profile file to add `/usr/local/bin` and `/usr/local/sbin` to your path:

```
[root]# . /etc/profile.d/torque.sh
```

5. Initialize `serverdb` by executing the `torque.setup` script:

```
[root]# ./torque.setup root
```

6. Add nodes to the `/var/spool/torque/server_priv/nodes` file. See 'Specifying Compute Nodes' in the *Torque Resource Manager Administrator Guide* for information on syntax and options for specifying compute nodes.7. Configure `pbs_server` to start automatically at system boot, and then start the daemon:

```
[root]# qterm
[root]# systemctl enable pbs_server.service
[root]# systemctl start pbs_server.service
```

2.1.2.D Install Torque MOMs

In most installations, you will install a Torque MOM on each of your compute nodes.

i See 'Specifying Compute Nodes' or 'Configuring Torque on Compute Nodes' in the *Torque Resource Manager Administrator Guide* for more information.

Do the following.

1. On the Torque Server Host, do the following:

a. Create the self-extracting packages that are copied and executed on your nodes:

```
[root]# make packages
```

b. Copy the self-extracting MOM packages to each Torque MOM Host. We recommend that you use a remote shell, such as SSH, to install packages on remote systems. Set up shared SSH keys if you do not want to supply a password for each Torque MOM Host.

```
[root]# scp torque-package-mom-linux-x86_64.sh <mom-node>:
```

- c. Copy the `pbs_mom` startup script to each Torque MOM Host:

```
[root]# scp contrib/systemd/pbs_mom.service <mom-node>:/lib/systemd/system/
```

2. On *each* Torque MOM Host, do the following:

- a. If installing Torque with `cgroups` support enabled (recommended), install `cgroup-tools`:

```
[root]# apt install cgroup-tools
```

i Torque currently only supports `cgroupv1`. If you are using a version of Ubuntu that uses `cgroupv2` by default (e.g., 22.04), you will need to switch to use `cgroupv1`:

```
[root]# mount | grep cgroup # If you see cgroup2 you will need to switch
to cgroupv1
[root]# vim /etc/default/grub
GRUB_CMDLINE_LINUX_DEFAULT="systemd.unified_cgroup_hierarchy=0"
[root]# update-grub
[root]# reboot
```

- b. Install the self-extracting MOM package:

```
[root]# ./torque-package-mom-linux-x86_64.sh --install
```

- c. (Optional) If you expect your jobs to require more than the default 12 MB of stack space, increase the stack limit by editing the `LimitSTACK` setting in `/usr/lib/systemd/system/pbs_mom.service`:

```
LimitSTACK=infinity
```

- d. Configure `pbs_mom` to start at system boot, and then start the daemon:

```
[root]# systemctl enable pbs_mom.service
[root]# systemctl start pbs_mom.service
```

2.1.2.E Install Torque Clients

If you want to have the Torque client commands installed on hosts other than the Torque Server Host (such as the compute nodes or separate login nodes), do the following.

1. On the Torque Server Host, do the following:
 - a. Copy the self-extracting client package to *each* Torque Client Host.

i We recommend that you use a remote shell, such as SSH, to install packages on remote systems. Set up shared SSH keys if you do not want to supply a password for each Torque Client Host.

```
[root]# scp torque-package-clients-linux-x86_64.sh <torque-client-host>:
```

- b. Copy the trqauthd startup script to *each* Torque Client Host:

```
[root]# scp contrib/systemd/trqauthd.service <torque-client-host>:/lib/systemd/system/
```

2. On *each* Torque Client Host, do the following:

- a. Install the self-extracting client package:

```
[root]# ./torque-package-clients-linux-x86_64.sh --install
```

- b. Enable and start the trqauthd service:

```
[root]# systemctl enable trqauthd.service
[root]# systemctl start trqauthd.service
```

2.1.2.F Configure Data Management

When a batch job completes, stdout and stderr files are generated and placed in the spool directory on the master Torque MOM Host for the job instead of the submit host. You can configure the Torque batch environment to copy the stdout and stderr files back to the submit host. See 'Configuring Data Management' in the *Torque Resource Manager Administrator Guide* for more information.

Related Topics

- [2.1.1 Preparing for Manual Installation - page 26](#)

2.1.3 Installing Moab Workload Manager

This topic contains instructions on how to install and start Moab Workload Manager (Moab HPC Suite).

In this topic:

- [2.1.3.A Open Necessary Ports - page 33](#)
- [2.1.3.B Install Dependencies, Packages, or Clients - page 33](#)
- [2.1.3.C Install Moab Server - page 34](#)
- [2.1.3.D Configure Torque to Trust Moab - page 36](#)
- [2.1.3.E Verify the Installation - page 36](#)
- [2.1.3.F \(Optional\) Install Moab Client - page 36](#)

2.1.3.A Open Necessary Ports

If your site is running firewall software on its hosts, you will need to configure the firewall to allow connections to the necessary ports:

Location	Ports	Functions	When Needed
Moab Server Host	42559	Moab Server Port	If you intend to run client commands on a host different from the Moab Server Host <i>or</i> if you will be using Moab in a grid

See [2.2.1 Opening Ports in a Firewall - page 50](#) for general instructions and an example of how to open ports in the firewall.

2.1.3.B Install Dependencies, Packages, or Clients

Dependencies and Packages

On the Moab Server Host, run the following command to install the required Moab HPC Suite dependencies and packages:

```
[root]# apt install build-essential libxml2-dev pkg-config
```

Torque Client

If you are using Torque and are installing the Torque Server on a different host (Torque Server Host) from the Moab Server (Moab Server Host), install the Torque client on the Moab Server Host in order for Moab to interact with Torque.

Follow the instructions in [Install hwloc - page 29](#) and [2.1.2.E Install Torque Clients - page 31](#) using the Moab Server Host as the Torque Client Host; with the exception that you must copy and install the `torque-package-devel-linux-<arch>.sh` self-extracting package in addition to the `torque-package-clients-linux-<arch>.sh` package:

```
[root]# scp torque-package-devel-linux-x86_64.sh <torque-client-host>:
```

```
[root]# ./torque-package-devel-linux-x86_64.sh --install
```

2.1.3.C Install Moab Server

i You must complete the tasks to install the dependencies, packages, or clients before installing Moab Server. See [2.1.3.B Install Dependencies, Packages, or Clients - page 33](#).

If your configuration uses firewalls, you *must also* open the necessary ports before installing the Moab Server. See [2.1.3.A Open Necessary Ports - page 33](#).

On the Moab Server Host, do the following.

1. Download the latest Moab Workload Manager tarball from the [Adaptive Computing website](#).
2. As the root user, run the following commands:

```
[root]# tar -xzvf moab-10.0.1-<OS>.tar.gz
[root]# cd moab-10.0.1-<OS>
```

i The variable marked `<OS>` indicates the OS for which the build was designed.

4. Configure Moab. If you are installing Moab Accounting Manager, configure Moab with the `--with-am` option.

```
[root]# ./configure <options>
```

i See [2.2.3 Moab Workload Manager Configuration Options - page 50](#) for a list of commonly used options or use `./configure --help` for a complete list of available options.

5. *ONLY* if you are using green computing *or* if you are using a resource manager other than Torque, run the `make perldeps` command to install the necessary perl modules using CPAN. When first running CPAN, you will be asked for configuration information. We recommend that you choose an automatic configuration. You will be prompted to provide input during module installation; running the `make perldeps` command with a script is not recommended.

```
[root]# make perldeps
```

6. Install Moab:

```
[root]# make install
```

7. Modify the Moab configuration file:

```
[root]# vi /opt/moab/etc/moab.cfg
```

Do one of the following:

- If using Torque Resource Manager:
 - i. Verify that `SUBMITCMD` is set up for your Torque resource manager and that it points to a valid `qsub` executable, for example:


```
RMCFG[torque] SUBMITCMD=/usr/local/bin/qsub
```
 - ii. If you installed the Torque Server on a different host (Torque Server Host), configure the `RMCFG HOST` parameter to tell Moab the host on which Torque Server is running:


```
RMCFG[torque] HOST=<torque_server_hostname>
```
- If using a NATIVE resource manager, see 'Managing Resources Directly with the Native Interface' in the *Moab Workload Manager Administrator Guide* for configuration information.

8. Source the appropriate profile script to add the Moab HPC Suite executable directories to your current shell `$PATH` environment:

```
[root]# . /etc/profile.d/moab.sh
```

9. Copy your license file into the same directory as `moab.cfg` (`/opt/moab/etc/` by default):

```
[root]# cp moab.lic $MOABHOMEDIR/moab.lic
```

a. To verify the current status of your license, run the following command:

```
[root]# moab --about 2>&1 | grep License
```

You should get something similar to the following in the response:

```
Moab Workload Manager Version '10.0.1' License Information:
Current License: Max Procs = 10000
Current License: Valid Until - Jul 13 19:42:10 2024
```

i A license is required for Moab. A trial license may be included in your Moab installation enabling you to run Moab for a limited time and with limited features. Email licenses@adaptivecomputing.com for information on obtaining licenses.

10. Start Moab:

```
[root]# systemctl start moab.service
```

2.1.3.D Configure Torque to Trust Moab

If you are using Torque as a resource manager and you installed the Torque Server on a different host (Torque Server Host), which is recommended, do the following:

On the Torque Server Host, add the name of the Moab Server Host (where Moab Server is installed) as a manager and as a submit host:

```
[root]# qmgr
Qmgr: set server managers += root@<moab_server_hostname>
Qmgr: set server submit_hosts += <moab_server_hostname>
Qmgr: exit
```

2.1.3.E Verify the Installation

If you have a resource manager configured, verify that the scheduler is able to schedule a job. Do the following.

Submit a sleep job as a non-root user (adaptive is used in this example) and verify the job is running:

```
[root]# su - adaptive
[adaptive]$ echo sleep 150 | msub
[adaptive]$ showq
[adaptive]$ exit
```

2.1.3.F (Optional) Install Moab Client

After you have installed Moab Server, you can create a client tarball to install just the Moab client commands on a login/client host. This tarball uses a single `tar` command to install the binary Moab client command files and their man pages. The tarball also contains a `moab.cfg` file configured with the Moab Server host name and port number so you do not have to manually configure this information on the login/client node. To create the Moab client package, the Moab service must first be stopped. Also, you need to `cd` into the `moab<build-info>` directory from the original Moab install.

i If your site needs secure communication and authentication between Moab Client Host and the Moab Server Host, create a site-specific key and place it in the same directory as your `moab.cfg` file. By default, this would be `$MOABHOMEDIR/etc/.moab.key`. When the Moab server and client commands detect the presence of those two files they will use the key in those files to authenticate and communicate, instead of the default key. If `.moab.key` is used, this protected file needs to be on each host that is authorized to run Moab client commands. See 'Mauth Authentication' in the *Moab Workload Manager Administrator Guide* for more information.

Do the following.

1. On the Moab Server Host, create the client tarball:

```
[root]# make client-pkg
```

2. Copy the tarball to the root directory of the Moab Client Host.
3. Copy the `/etc/profile.d/moab.sh` to the Moab Client Host.
4. On the Moab Client Host, run the tarball to install the Moab client commands:

```
[root]# tar xvf client.tgz
```

This creates an `opt/moab/` directory in the CWD; it does not create `/opt/moab/`. To use the current `client.tgz`, you must `cd` to `/`, then `untar` the package.

5. Copy the `/opt/moab/etc/.moab.key` file to the same location on the Moab Client Host.

Related Topics

- [2.1.1 Preparing for Manual Installation - page 26](#)

2.1.4 Installing Moab Accounting Manager

This topic contains instructions on how to install and start Moab Accounting Manager (MAM).

In this topic:

- [2.1.4.A Plan Your Installation - page 38](#)
- [2.1.4.B Open Necessary Ports - page 38](#)
- [2.1.4.C Install and Initialize PostgreSQL Server - page 39](#)
- [2.1.4.D Install Dependencies, Packages, or Clients - page 40](#)
- [2.1.4.E Install MAM Server - page 40](#)
- [2.1.4.F Configure the MAM GUI - page 43](#)
- [2.1.4.G Configure MAM Web Services - page 44](#)
- [2.1.4.H Access the MAM GUI - page 45](#)
- [2.1.4.I Access MAM Web Services - page 45](#)

[2.1.4.J Configure Moab Workload Manager to Use MAM - page 46](#)

[2.1.4.K Initialize Moab Accounting Manager - page 47](#)

2.1.4.A Plan Your Installation

The first step is determining the number of different hosts (physical machines) required for your MAM installation.

Your MAM installation includes:

- MAM Server
- MAM Database
- MAM Clients (possibly several hosts)
- MAM Web Server (optional, for the MAM GUI and/or MAM Web Services)

Each of these components can be installed on their own hosts (meaning the actual physical machine) or can be combined on the same hosts. For example, the MAM Database can be installed on the same *host* as the MAM Server. Or the MAM Server can be installed on the same host on which you installed the Moab Server.

Once you have determined which components are installed on which hosts, complete the rest of the instructions for the MAM installation.

i The instructions that follow in this topic will use the term *host* after each component to indicate the physical machine on which the component is installed (for example, MAM Server Host and MAM Database Host). Depending on your configuration, the *host* may refer to the component installed on its own machine or installed on the same machine as another component.

2.1.4.B Open Necessary Ports

If your site is running firewall software on its hosts, configure the firewall to allow connections to the necessary ports:

Location	Ports	Functions	When Needed
MAM Server Host	7112	MAM Server Port	If you will be installing the MAM Server on a different host from where you installed the Moab Server <i>or</i> you will be installing the MAM Clients on other hosts

Location	Ports	Functions	When Needed
MAM Web Server Host	443	HTTPS Port	If using the MAM GUI or MAM Web Services
MAM Database Host	5432	MAM PostgreSQL Server Port	If you will be installing the MAM Database on a different host from the MAM Server

See [2.2.1 Opening Ports in a Firewall - page 50](#) for general instructions and an example of how to open ports in the firewall.

2.1.4.C Install and Initialize PostgreSQL Server

Moab Accounting Manager uses a database for transactions and data persistence. The PostgreSQL database can be installed on a different host from the MAM Server; however, it is often convenient to install them on the same host. For example, the PostgreSQL database can be installed on:

- the same host as the MAM Server.
- a separate PostgreSQL database host.
- a separate *shared* PostgreSQL database host.

On the host where the MAM PostgreSQL database will reside, do the following.

1. Install and initialize the PostgreSQL Server:

```
[root]# apt install postgresql
```

2. If the MAM Database Host is installed on a *different* host from where you will install the MAM Server, configure PostgreSQL to accept connections from the MAM Server Host:

```
[root]# vi /etc/postgresql/*/main/postgresql.conf

# Replace <mam-server-host> with the TCP/IP address on which the database server is
# to listen for connections
# from the MAM server. This will normally be the hostname or ip address of the MAM
# Database Host.
listen_addresses = '<mam-database-host>'
```

3. Start or restart the database:

```
[root]# systemctl restart postgresql.service
```

2.1.4.D Install Dependencies, Packages, or Clients

Use the following instructions to install the required Moab Accounting Manager dependencies, packages, or clients.

i Depending on your configuration, the MAM Server Host and the MAM Web Server Host may be installed on the same host. The MAM Client Host is automatically installed on the same host as the MAM Server Host; however, you can also install the MAM Client Host on any other hosts on which you want to have the MAM client commands available to users or administrators.

i If any of the Perl module packages fail to install or are unavailable for your system, you can install it from CPAN by running `cpan MODULENAME` where `MODULENAME` is the respective perl module name.

1. On the MAM Server Host, the MAM Web Server Host, and the MAM Client Hosts, run the following commands:

```
[root]# apt install gcc libauthen-pam-perl libconfig-tiny-perl libcrypt-cbc-perl
libcrypt-des-perl libcrypt-des-ede3-perl libdigest-hmac-perl libdigest-sha-perl
liberror-perl libjson-perl liblog-dispatch-filerotate-perl liblog-log4perl-perl
libparams-validate-perl libxml-libxml-perl make
```

2. On the MAM Server Host, run the following command:

```
[root]# apt-get install postgresql-common postgresql-client libdate-manip-perl
libdbd-pg-perl libdbi-perl
```

3. If you plan to use the MAM GUI, on the MAM Web Server Host, run the following:

```
[root]# apt install apache2 libcgi-session-perl
```

4. If you plan to use MAM Web Services, on the MAM Web Server Host, run the following command:

```
[root]# apt install apache2 libapache2-mod-perl2
```

5. On each of the MAM Client Hosts (including the MAM Server Host), run the following command:

```
[root]# apt install libssl-dev libterm-readline-gnu-perl libterm-readkey-perl
```

2.1.4.E Install MAM Server

If you will be installing the MAM GUI or MAM Web Services and you intend to use PAM authentication (see 'Integrating with PAM' in the *Moab Accounting Manager Administrator Guide*), the MAM server will need to be installed and run using the root user. If this is not

required in your setup, we recommend installing and running the MAM server using a non-root user. These instructions will show examples for installing MAM using a non-root user, `mam`. If using root, or a different user, you will need to make appropriate substitutions for your selected user in the remaining steps and sections below.

On the MAM Server Host, do the following.

1. Create a user called `mam` and switch to that user:

```
[root]# useradd -m mam -s $(which bash)
[root]# su - mam
[mam]$ mkdir src
[mam]$ cd src
```

2. Download the latest MAM tarball from the [Adaptive Computing](#) website.
3. As the `mam` user, run the following commands:

```
[mam]$ tar -zxvf mam-10.0.1.tar.gz
[mam]$ cd mam-10.0.1
```

4. Configure the software. For a list of all the configuration options, see [2.2.4 Moab Accounting Manager Configuration Options - page 52](#).

```
[mam]$ ./configure
```

i If you are planning to use the GUI or web services and you want to use PAM for Unix password authentication, use the `--with-pam` option. This will configure MAM to run as root and configure the GUI and web services to use PAM for user password authentication.

5. Compile the software:

```
[mam]$ make
```

i Depending on your configuration, you may need to replace 'make' with a `make` command that includes additional functionality. Specifically:

- If you only need to install the clients on a particular system, use `make clients-only`.
- If you only need to install the web GUI on a particular system, use `make gui-only`.
- If you only need to install the web services on a particular system, use `make ws-only`.

6. Install the software:

```
[mam]$ exit
[root]# cd ~/mam/src/mam-10.0.1
[root]# make install
```

i Depending on your configuration, you may need to replace 'make install' with a make command that includes additional functionality. Specifically:

- If you only need to install the clients on a particular system, use `make install-clients-only`.
- If you only need to install the web GUI on a particular system, use `make install-gui-only`.
- If you only need to install the web services on a particular system, use `make install-ws-only`.

7. As the database user, create a mam user with the ability to create a database.

i PostgreSQL should have previously been installed using the instructions in [2.1.1 Preparing for Manual Installation](#) - page 26.

```
[root]# su - postgres
[postgres]$ psql
CREATE USER mam WITH CREATEDB PASSWORD 'changeme!';
\q
[postgres]$ exit
```

The password you define must be synchronized with the `database.password` value in `/opt/mam/etc/mam-server.conf`:

```
[root]# vi /opt/mam/etc/mam-server.conf
database.password = changeme!
```

For systems with a separate PostgreSQL host, add `database.datasource` to `/opt/mam/etc/mam-server.conf`:

```
database.datasource=DBI:Pg:dbname=mam;host=remote-host
```

8. Create the mam database. Run the `hpc.sql` script to populate the database with objects, actions, and attributes necessary to function as an Accounting Manager:

```
[root]# su - mam
[mam]$ cd src/mam-10.0.1
[mam]$ createdb mam
[mam]$ psql mam < hpc.sql
[mam]$ exit
```

9. Configure MAM to automatically start up at system boot; start the `mam` service:

```
[root]# systemctl enable mam.service
[root]# systemctl start mam.service
```

2.1.4.F Configure the MAM GUI

If you plan to use the web GUI, then on the MAM Web Server Host, do the following.

1. As `root`, add or edit the SSL virtual host definition as appropriate for your environment. To do so, configure the `cgi-bin` directory in `ssl.conf`. Below the `cgi-bin` directory element, create an alias for `/cgi-bin` pointing to your `cgi-bin` directory. If you chose to install to a `cgi-bin` sub-directory, you might want to create an alias for that as well. Also, add `index.cgi` to the `DirectoryIndex` so you can use the shorter sub-directory name.

```
[root]# a2enmod cgi
[root]# a2enmod ssl
[root]# a2ensite default-ssl

[root]# vi /etc/apache2/sites-enabled/default-ssl.conf

# Configure your cgi-bin directory
<Directory "/usr/lib/cgi-bin">
    Options ExecCGI
    AddHandler cgi-script .cgi
    AllowOverride All
    Require all granted
</Directory>

# Aliases for /cgi-bin
Alias /cgi-bin/ /usr/lib/cgi-bin/
Alias /mam /usr/lib/cgi-bin/mam/

# Make shorter sub-dir name available
DirectoryIndex index.cgi
```

2. For the highest security, we recommend that you install a public key certificate that has been signed by a certificate authority. The exact steps to do this are specific to your distribution and the chosen certificate authority. An overview of this process for CentOS 7 is documented at <https://access.redhat.com/products>.

Alternatively, if your network domain can be secured from man-in-the-middle attacks, you could use a self-signed certificate. Often this does not require any additional steps because in many distributions, such as Ubuntu, the Apache SSL configuration provides self-signed certificates by default.

If your configuration uses self-signed certificates, no action is required; Ubuntu ships with ready-made certificates.

3. Configure the GUI to use the desired authentication method. Edit the GUI configuration file and ensure that the `authentication.method` parameter is set to the desired value.

This parameter can be set to a value of `mam-password` to use the built-in MAM database Password table, or to a value of `pam` to authenticate the user using Linux PAM. If you want to use PAM for Unix password authentication, you should have run `configure` with the `--with-pam` option (or `--with-user=root`) in order to configure MAM to run as root. See 'Integrating with PAM' in the *Moab Accounting Manager Administrator Guide* for more information about the steps required to configure PAM authentication.

```
[mam]$ vi /opt/mam/etc/mam-gui.conf
authentication.method = mam-password
```

4. Start or restart the HTTP server daemon:

```
[root]# systemctl restart apache2.service
```

2.1.4.G Configure MAM Web Services

If you plan to use MAM Web Services, then on the MAM Web Server Host, do the following.

1. Edit the SSL virtual host definition in `ssl.conf` to include the `mamws` location, for example:

```
[root]# a2enmod SSL
[root]# a2ensite default-ssl

[root]# vi /etc/apache2/sites-enabled/default-ssl.conf

# Place the following within the 443 VirtualHost definition
PerlOptions +Parent
PerlSwitches -Mlib=/opt/mam/lib
PerlModule MAM::WSResponseHandler
PerlModule MAM::WSAuthenHandler
<Location /mamws>
    SetHandler perl-script
    PerlResponseHandler MAM::WSResponseHandler
    Options +ExecCGI

    AuthName MAM
    PerlAuthenHandler MAM::WSAuthenHandler
    Require valid-user

    Order allow,deny
    Allow from all
</Location>
```

2. For the highest security, we recommend that you install a public key certificate that has been signed by a certificate authority. The exact steps to do this are specific to your distribution and the chosen certificate authority. An overview of this process for CentOS 7 is documented at <https://access.redhat.com/products>.

Alternatively, if your network domain can be secured from man-in-the-middle attacks, you could use a self-signed certificate. Often this does not require any additional steps

because in many distributions, such as Ubuntu, the Apache SSL configuration provides self-signed certificates by default.

If your configuration uses self-signed certificates, no action is required; Ubuntu ships with ready-made certificates.

3. Configure MAM Web Services to use the desired authentication method. Edit the web services configuration file and ensure that the `authentication.method` parameter is set to the desired value. This parameter can be set to a value of `mam-password` to use the built-in MAM database Password table, or to a value of `pam` to authenticate the user using Linux PAM. If you want to use PAM for Unix password authentication, you should have run `configure` with the `--with-pam` option (or `--with-user=root`) in order to configure MAM to run as root. See 'Integrating with PAM' in the *Moab Accounting Manager Administrator Guide* for more information about the steps required to configure PAM authentication.

```
[mam]$ vi /opt/mam/etc/mam-ws.conf
authentication.method = mam-password
```

4. Start or restart the HTTP server daemon:

```
[root]# systemctl restart apache2.service
```

2.1.4.H Access the MAM GUI

If you plan to use the web GUI, then on the MAM Server Host, do the following.

1. If your GUI authentication method is `mam-password`, create a password for the `mam` user that you want to access the MAM GUI:

```
[root]# su - mam
[mam]$ mam-set-password
[mam]$ exit
```

2. Verify the connection:
 - a. Open a web browser and navigate to `https://<mam_web_server_host>/cgi-bin/mam`.
 - b. Log in as the `mam` user with the password you set in step 1.

2.1.4.I Access MAM Web Services

If you plan to use MAM web services, then on a MAM Client Host, do the following.

1. If your web services authentication method is `mam-password`, create a password for the `mam` user that you wish to access the MAM Web Services:

```
[root]# su - mam
[mam]$ mam-set-password
[mam]$ exit
```

2. Make a call to web services:

```
[root]# curl -k -X GET --basic -u mam:changeme! 'https://<mam_web_server_
host>/mamws/system'
```

Alternatively, for queries, you can use the browser to access the URL, for example: 'https://<mam_web_server_host>/mamws/system'.

2.1.4.J Configure Moab Workload Manager to Use MAM

If integrating with Moab Workload Manager, do the following.

1. Configure Moab to talk to MAM. Do *one* of the following:

- **MAM Option.** If you will be using the MAM (direct network) accounting manager interface with Moab Workload Manager (this is the default), do the following:
 - a. On the Moab Server Host, edit the Moab configuration file, uncomment the AMCFG lines and set the TYPE to MAM and set the HOST. If the Moab Server and the MAM Server are on the same host, set HOST to 'localhost'; otherwise, set HOST to the host name for the MAM Server (MAM Server Host).

```
[root]# vi /opt/moab/etc/moab.cfg
AMCFG[mam] TYPE=MAM HOST=<mam_server_host>
```

Customize additionally as needed. See 'Accounting, Charging, and Allocation Management' in the *Moab Workload Manager Administrator Guide*

- b. Configure Moab to authenticate with MAM using the MAM secret key:
 - i. On the MAM Server Host, copy the auto-generated secret key from the token.value value in the /opt/mam/etc/mam-site.conf file.
 - ii. On the Moab Server Host, add the secret key to the moab-private.cfg file as the value of the CLIENTCFG KEY attribute:

```
[root]# vi /opt/moab/etc/moab-private.cfg
CLIENTCFG[AM:mam] KEY=<MAMSecretKey>
```

- **Native Option.** If you will be using the Native (custom script) accounting manager interface with Moab Workload Manager, do the following:
 - a. On the Moab Server Host, edit the Moab configuration file, uncomment the AMCFG lines and set the TYPE to NATIVE:

```
[root]# vi /opt/moab/etc/moab.cfg
AMCFG[mam] TYPE=NATIVE
```

- b. If you are installing Moab Accounting Manager on a different host (MAM Server Host) from the Moab Server (Moab Server Host), install the Moab Accounting Manager client on the Moab Server Host in order for the custom scripts to use the MAM API.

On the Moab HPC Suite Server Host, follow the instructions in [2.1.4.D Install Dependencies, Packages, or Clients - page 40](#) and [2.1.4.E Install MAM Server - page 40](#); with the following exceptions:

- Install only the dependent packages applicable to MAM Client Hosts.
- Use the configure option `--without-init`.
- Instead of running `make`, use `make clients-only`.
- Instead of running `make install`, use `make install-clients-only`.
- Omit the step to create the database and all of the steps thereafter.

2. On the Moab Server Host, restart Moab:

```
[root]# systemctl restart moab.service
```

2.1.4.K Initialize Moab Accounting Manager

You will need to initialize Moab Accounting Manager to function in the way that is most applicable to the needs of your site. See 'Initial Setup' in the *Moab Accounting Manager Administrator Guide* to set up Moab Accounting Manager for your desired accounting mode.

Related Topics

- [2.1.1 Preparing for Manual Installation - page 26](#)

2.1.5 Installing RLM Server


This topic contains instructions on how to install an RLM Server.

The RLM Server can run multiple licenses.

Access to a Reprise License Manager (RLM) server is required when using Moab's feature Nitro.

As the RLM Server can run multiple licenses, we recommend that you install *one* RLM Server for your configuration. If your company already uses an RLM Server, you do not need to install a new one for Adaptive Computing products. However, Adaptive Computing

strongly recommends that your RLM Server is the latest version, and the Adaptive Computing products may use a different port than the default RLM Server port (5053).

 If your system configuration requires more than one RLM Server, additional configuration may be needed. See [2.2.5 Using Multiple RLM Servers - page 54](#) for more information.

 Expired RLM license files must be removed and the RLM service restarted whenever a new license is installed.


In this topic:

[2.1.5.A Open Necessary Ports - page 48](#)

[2.1.5.B Install the RLM Server - page 49](#)

[2.1.5.C Change the Default Passwords - page 49](#)

2.1.5.A Open Necessary Ports

 These instructions assume you are using the default ports. If your configuration will use other ports, then substitute your port numbers when opening the ports.

If your site is running firewall software on its hosts, you will need to configure the firewall to allow connections to the necessary ports:

Location	Ports	Functions	When Needed
RLMServer Host	5053	RLM Server Port	Always
RLM Server Host	5054	RLM Web Interface Port	Always
RLM Server Host	5135	ISV adaptiveco Port (for the Adaptive license-enabled products)	For Moab Workload Manager <i>and</i> if Nitro is part of your configuration

See [2.2.1 Opening Ports in a Firewall - page 50](#) for general instructions and an example of how to open ports in the firewall.

2.1.5.B Install the RLM Server

On the host where the RLM Server will reside, do the following.

1. Download the latest RLM build from the [Adaptive Computing](#) website.
2. Create a non-root user and group (rlm is used in the example):

```
[root]# groupadd -r rlm
[root]# useradd -r -g rlm -d /opt/rlm -c "A non-root user under which to run
Reprise License Manager" rlm
```

3. Create a directory and install the tarball files in that location (we are using /opt/rlm as the install location in the example):

```
[root]# mkdir -p -m 0744 /opt/rlm
[root]# cd /opt/rlm
[root]# tar -xzvf /tmp/ac-rlm-12.1.2.tar.gz --strip-components=1
[root]# chown -R rlm:rlm /opt/rlm
```

i The `--strip-components=1` removes the "ac-rlm-12.1/" from the relative path so that they are extracted into the current directory.

4. Install the startup scripts.

i If you are using a user:group other than rlm:rlm or a location other than /opt/rlm, then edit the following files to reflect those changes after copying them.

5. Start the services and configure the RLM Server to start automatically at system reboot:

```
[root]# systemctl enable rlm.service
[root]# systemctl start rlm.service
```

2.1.5.C Change the Default Passwords

The RLM Web interface includes two usernames (admin and user) by default. These usernames have the default password 'changeme!'.

! If you do not change this password, RLM will not be secure. For tips on choosing a good password, see <https://www.us-cert.gov/ncas/tips/ST04-002>.

Do the following for *both* the user and the admin usernames.

1. Using a web browser, navigate to your RLM instance (`https://<RLM_host>:5054`; where <RLM_host> is the IP address or name of the RLM Server Host).
2. Log in.

3. Select `Change Password` and change the password according to your password security process.

2.2 Additional Configuration

In this section:

[2.2.1 Opening Ports in a Firewall - page 50](#)

[2.2.2 Configuring SSL in Tomcat - page 50](#)

[2.2.3 Moab Workload Manager Configuration Options - page 50](#)

[2.2.4 Moab Accounting Manager Configuration Options - page 52](#)

[2.2.5 Using Multiple RLM Servers - page 54](#)

[2.2.6 Trusting Servers in Java - page 55](#)

2.2.1 Opening Ports in a Firewall

If your site is running firewall software on its hosts, configure the firewall to allow connections to the products in your installation.

This topic provides an example and general instructions for how to open ports in your firewall. See [3.2 Port Reference - page 77](#) for the actual port numbers for the various products.

Ubuntu-based systems use `ufw` as the default firewall software. If you use different firewall software, refer to your firewall documentation for opening ports in your firewall.

The following is an example of adding port 1234 when using `ufw`:

```
[root]# ufw allow 1234/tcp
```

2.2.2 Configuring SSL in Tomcat

To configure SSL in Tomcat, refer to the Apache Tomcat [documentation](#).

2.2.3 Moab Workload Manager Configuration Options

The following is a list of commonly used configure options. For a complete list, use `./configure --help` when configuring Moab HPC Suite.

Option	Description	Example
--prefix	Specifies the location of the binaries and libraries of the Moab install. The default location is <code>/opt/moab</code> .	<pre>[root]# ./configure --prefix=/usr/local</pre>
--with-am	Specifies that you want to configure Moab with Moab Accounting Manager.	<pre>[root]# ./configure --with-am</pre>
--with-am-dir	Uses the specified prefix directory for the accounting manager if installed in a non-default location.	<pre>[root]# ./configure --with-am-dir=/opt/mam-10.0.1</pre>
--with-flexlm	Causes Moab to install the <code>license.mon.flexLM.pl</code> script in the <code>/opt/moab/tools</code> directory. For more information about this script, see the section 'Interfacing with FlexNet (formerly FLEXlm)' in the <i>Moab Workload Manager Administrator Guide</i> .	<pre>[root]# ./configure --with-flexlm</pre>
--with-homedir	Specifies the location of the Moab configuration directory and the MOABHOMEDIR environment variable. The default location is <code>/opt/moab</code> . i By default, MOABHOMEDIR is automatically set during installation. Use the <code>--without-profile</code> option to disable installed scripts.	<pre>[root]# ./configure --with-homedir=/var/moab</pre> <i>The Moab HPC Suite home directory will be /var/moab instead of the default /opt/moab.</i>
--without-init	Disables the installation of a distribution-specific, Moab service startup file. By default,	<pre>[root]# ./configure --without-init</pre>

Option	Description	Example
	'make install' will install an init.d or systemd service startup file as appropriate for your distribution. The installed file (/etc/init.d/moab or /usr/lib/systemd/system/moab.service) can be customized to your needs. If you do not want this file to be installed, use this option to exclude it.	
--without-profile	Disables the installation of a distribution-specific shell profile for bash and C shell. By default, 'make install' will install the Moab shell initialization scripts as appropriate for your operating system. These scripts help to establish the MOABHOMEDIR, PERL5LIB, PATH and MANPATH environment variables to specify where the new moab configuration, scripts, binaries and man pages reside. The installed scripts (/etc/profile.d/moab.{csh,sh}) can be customized to your needs. If you do not want these scripts to be installed, use this option to exclude them.	<pre>[root]# ./configure --without-profile</pre>

2.2.4 Moab Accounting Manager Configuration Options

The following table comprises commonly-used configure options:

Option	Description
-h,--help	Run <code>./configure --help</code> to see the list of configure options.

Option	Description
--localstatedir=DIR	Home directory where per-configuration subdirectories (such as <code>etc</code> , <code>log</code> , <code>data</code>) will be installed (defaults to <code>PREFIX</code>).
--prefix=PREFIX	Base installation directory where all subdirectories will be installed unless otherwise designated (defaults to <code>/opt/mam</code>).
--with-cgi-bin=DIR	If you intend to use the web GUI, use <code>--with-cgi-bin</code> to specify the directory where you want the Moab Accounting Manager CGI files to reside (defaults to <code>/var/www/cgi-bin/mam</code>).
--with-db-name=NAME	Name of the SQL database that the server will sync with (defaults to <code>mam</code>).
--with-legacy-links	Creates symbolic links allowing the use of the old client and server command names (for example, <code>mam-list-users</code> would be created as symbolic link to <code>mam-list-users</code>). When running a command under its old name, the command will issue a deprecation warning. This warning can be disabled by setting <code>client.deprecationwarning = false</code> in the <code>mam-client.conf</code> file. The default is not to install the legacy links.
--with-mam-libs=local site	Use <code>--with-mam-libs</code> to indicate whether you want to install the Perl MAM modules in a local directory (<code>\${exec_prefix}/lib</code>) or in the default system site-perl directory (defaults to <code>local</code>).
--with-promotion=mamauth suidperl	Command-line clients and scripts using the API need to use a privilege promotion method to authenticate and encrypt the communication using the symmetric key. The default is <code>suidperl</code> if it is installed on the system, otherwise the default is <code>mamauth</code> . See the description for the <code>promotion.method</code> configuration parameter in the Client Configuration section for more information about the two security promotion methods.
--with-user=USER	Use <code>--with-user</code> to specify the accounting admin userid that the server will run under and who will have full administrative privileges (defaults to <code>mam</code>). We recommend that this be a non-privileged user for the

Option	Description
	highest security.
--without-gui	Specifies whether to install the CGI web GUI. If you do not intend to use the CGI web GUI, you can specify <code>--without-gui</code> to not install the CGI scripts. Otherwise, the default is to install the GUI CGI scripts.
--without-init	If you do not intend to use the <code>mam init.d</code> service, you can use <code>--without-init</code> to specify that Moab HPC Suite should not install the <code>mam init.d</code> script. Otherwise, the script is installed by default.
--with[out]-pam	Indicates whether to use PAM authentication for the GUI and web services. If <code>--with-pam</code> is specified, the PAM configuration file is installed and the GUI and web services will default to using <code>pam</code> as the authentication method. If <code>--without-pam</code> is specified, the PAM configuration file is not installed and the GUI and web services will default to using <code>mam-password</code> as the authentication method. If this option is not specified, the PAM configuration file is installed but the GUI and web services will default to using <code>mam-password</code> as the authentication method. When <code>--with-pam</code> option is specified, the accounting admin user will default to <code>root</code> , unless overridden with the <code>--with-user</code> configuration option.
--without-profile	If you do not intend to use the <code>mam profile.d</code> environment scripts, you can use <code>--without-profile</code> to specify that Moab HPC Suite should not install the <code>mam profile.d</code> scripts. Otherwise, the scripts are installed by default.

2.2.5 Using Multiple RLM Servers

As the RLM Server can run multiple licenses, we recommend that you install *one* RLM Server for your configuration.

However, if your configuration requires more than one RLM Server, you will *need* to configure the Adaptive Computing products to connect to a specific RLM Server. If not configured to connect to a specific RLM Server, the Adaptive Computing product will scan the network and connect to the first RLM Server it finds listening to request the license. If the first RLM Server does *not* have the product's license, the RLM connection will fail.

If you are using multiple RLM Servers, do the following to configure an Adaptive Computing product to connect to a specific RLM Server.

1. Modify the RLM Server not to accept the network search connections.
 - a. Edit the init script in `/opt/rlm/` to add `-noudp`:

```
start() {
  su -l $rlmuser -s /bin/bash -c "$rlmdir/rlm -l -dlog $debuglog -noudp &"
}
```

2. Enable the Adaptive Computing product to connect to a specific RLM. On the host where the Adaptive Computing product resides, do the following:
 - a. Create a new text file and name it with the `.lic` extension (typically, `remote.lic`) and save it in the same location as the other Adaptive Computing licenses. Be careful not to override an existing license.
 - b. Edit the new `remote.lic` file to point to the specific RLM Server hostname and port. Port 5053 is the default. If you use a different port number for the RLM Server, specify that port number in the `remote.lic` file:

```
HOST <hostname> ANY 5053
```

Repeat as needed for each Adaptive Computing product that you want to connect to a specific RLM Server.

2.2.6 Trusting Servers in Java

In this topic:

[2.2.6.A Prerequisites - page 55](#)

[2.2.6.B Retrieve the Server's X.509 Public Certificate - page 56](#)

[2.2.6.C Add the Server's Certificate to Java's Keystore - page 56](#)

2.2.6.A Prerequisites

Some of these instructions refer to `JAVA_HOME`, which must point to the same directory that Tomcat uses. To set `JAVA_HOME`, do this:

```
[root]# source /etc/tomcat/tomcat.conf
```

Your system administrator might have defined Tomcat's `JAVA_HOME` in a different file.

2.2.6.B Retrieve the Server's X.509 Public Certificate

To retrieve the server's certificate, use the following command:

```
[root]# $JAVA_HOME/bin/keytool -printcert -rfc -sslserver <servername>:<port>
/tmp/public.cert.pem
```

Replace `<servername>` with the server's host name and `<port>` with the secure port number. The default port for https is 443. The default port for LDAP is 636. If successful, `/tmp/public.cert.pem` contains the server's public certificate. Otherwise, `/tmp/public.cert.pem` contains an error message. This message is typical:
`keytool error: java.lang.Exception: No certificate from the SSL server.` This message suggests that the server name or port is incorrect. Consult your IT department to determine the correct server name and port.

2.2.6.C Add the Server's Certificate to Java's Keystore

Java stores trusted certificates in a database known as the keystore. Because each new version of Java has its own keystore, you need to add the server certificate to the Java keystore (using the steps below) every time you install a new version of Java.

Java's keystore is located at `$JAVA_HOME/lib/security/cacerts`. If Tomcat's `JAVA_HOME` points to a JDK, then the keystore is located at `$JAVA_HOME/jre/lib/security/cacerts`. To add the server certificate to the keystore, run the following command:

```
[root]# $JAVA_HOME/bin/keytool -import -trustcacerts -file /tmp/public.cert.pem -alias
<servername> -keystore $JAVA_HOME/lib/security/cacerts
```

You will be prompted for the keystore password, which is 'changeit' by default.

i Your system administrator might have changed this password.


After you've entered the keystore password, you'll see the description of the server's certificate. At the end of the description, it prompts you to trust the certificate:


```
Trust this certificate? [no]:
```

Type `yes` and press `Enter` to add the certificate to the keystore.

2.3 Manual Upgrade Steps

This section provides instructions and other information when upgrading your Moab HPC Suite components using the Manual upgrade method.

 We highly recommend that you *first* perform upgrades in a *test environment*. Installation and upgrade procedures are tested prior to release; however, due to customizable variations that may be utilized by your configuration, it is not recommended to drop new versions of software directly into production environments. This is especially true when the workload has vital bearing. Contact Adaptive Computing Professional Services for more information.

 Because many system-level files and directories are accessed during the upgrade, the upgrade instructions in this guide should be executed with root privileges. You will see that the instructions execute commands as the root user. Note that the same commands will work for a non-root user with the `sudo` command.

In this section:

- [2.3.1 Upgrading Torque Resource Manager - page 57](#)
- [2.3.2 Upgrading Moab Workload Manager - page 62](#)
- [2.3.3 Upgrading Moab Accounting Manager - page 64](#)
- [2.3.4 Upgrading RLM Server - page 67](#)
- [2.3.5 Migrating the MAM Database from MySQL to PostgreSQL - page 68](#)

2.3.1 Upgrading Torque Resource Manager

Torque 7.0 binaries are backward compatible with Torque 5.0 or later. However they are not backward compatible with Torque versions prior to 5.0. When you upgrade to Torque 7.0.1 from versions prior to 5.0, all MOM and server daemons must be upgraded at the same time.

The job format is compatible between 7.0 and previous versions of Torque and any queued jobs will upgrade to the new version. We do *not* recommend upgrading Torque while jobs are in a running state.

This topic contains instructions on how to upgrade and start Torque Resource Manager (Torque).

 If you need to upgrade a Torque version prior to 4.0, contact [Adaptive Computing](#).

i See 'Considerations Before Upgrading' in the *Torque Resource Manager Administrator Guide* for additional important information, including about how to handle running jobs during an upgrade, mixed server/MOM versions, and the possibility of upgrading the MOMs without having to take compute nodes offline.

In this topic:

- [2.3.1.A Before You Upgrade - page 58](#)
- [2.3.1.B Stop Torque Services - page 59](#)
- [2.3.1.C Upgrade the Torque Server - page 59](#)
- [2.3.1.D Update the Torque MOMs - page 60](#)
- [2.3.1.E Update the Torque Clients - page 61](#)
- [2.3.1.F Start Torque Services - page 61](#)
- [2.3.1.G Perform Status and Error Checks - page 62](#)

2.3.1.A Before You Upgrade

This section contains information you should be aware of before upgrading.

Running Jobs

Before upgrading the system, all running jobs must complete. To prevent queued jobs from starting, nodes can be set to offline or all queues can be disabled (using the `started` queue attribute). See 'pbsnodes' or 'Queue Attributes' in the *Torque Resource Manager Administrator Guide* for more information.

hwloc

Torque version 7.0.1 and later permit and recommend the use of the OS vendor provided version of hwloc.

1. If you are upgrading from a version of Torque earlier than 7.0.1, on the Torque Server Host, each Torque MOM Host, and each Torque Client Host, run the following commands:

```
[root]# apt install libhwloc-dev pkg-config
[root]# cd hwloc-1.9.1
[root]# make uninstall
```

2. On the Torque Server Host, run the following commands:

```
[root]# rm /etc/ld.so.conf.d/hwloc.conf
[root]# ldconfig
```

GPU Support

Because Torque GPU support has evolved over time, upgrading may require a re-examination of the cluster's GPU setup, especially if the upgrade will include configuration changes to take advantage of cgroups and/or NVIDIA/NVML support. See 'Scheduling GPUs' in the Accelerators chapter of the *Moab Workload Manager Administrator Guide* for an overview of currently-available options.


2.3.1.B Stop Torque Services

Do the following.

1. On the Torque Server Host, shut down the Torque server:

```
[root]# systemctl stop pbs_server.service
```

2. On *each* host where the Torque MOM Host resides (regardless of whether it resides on the Torque Server Host), shut down the Torque MOM service.

 Confirm all jobs have completed before stopping `pbs_mom`. You can do this by typing `momctl -d3`. If there are no jobs running, you will see the message 'NOTE: no local jobs detected' towards the bottom of the output. If jobs are still running and the MOM is shutdown, you will only be able to track when the job completes and you will not be able to get completion codes or statistics.

```
[root]# systemctl stop pbs_mom.service
```

3. On *each* host where the Torque Client Host resides (regardless of whether it resides on the Moab Server Host, the Torque Server Host, or the Torque MOM Hosts), shut down the `trqauthd` service:

```
[root]# systemctl stop trqauthd.service
```

2.3.1.C Upgrade the Torque Server

 You *must* complete all the previous upgrade steps in this topic before upgrading Torque server. See the list of steps at the beginning of this topic.

On the Torque Server Host, do the following.

1. Back up your `server_priv` directory:

```
[root]# tar -cvf backup.tar.gz TORQUE_HOME/server_priv
```

2. Download the latest Torque tarball from the [Adaptive Computing](#) website.

3. Depending on your system configuration, you will need to add `./configure` command options. At a minimum, you add: `--enable-cgroups`

i These instructions assume you are using `cgroups`. When `cgroups` are supported, `cpusets` are handled by the `cgroup` `cpuset` subsystem. If you are not using `cgroups`, use `--enable-cpusets` instead.

i If `--enable-gui` is part of your configuration, run the following commands:

```
$ cd /usr/lib64
$ ln -s libXext.so.6.4.0 libXext.so
$ ln -s libXss.so.1 libXss.so
```

When finished, `cd` back to your `install` directory.

See 'Customizing the Install' in the *Torque Resource Manager Administrator Guide* for more information on which options are available to customize the `./configure` command.

4. Install the latest Torque tarball:

```
[root]# tar xzvf torque-7.0.1.tar.gz
[root]# cd torque-7.0.1
[root]# ./configure --enable-cgroups # add any other required options
[root]# make
[root]# make install
```

2.3.1.D Update the Torque MOMs

Do the following.

1. On the Torque Server Host, do the following:
 - a. Create the self-extracting packages that are copied and executed on your nodes:

```
[root]# make packages
```

- b. Copy the self-extracting mom package to *each* Torque MOM Host. We recommend that you use a remote shell, such as SSH, to install packages on remote systems. Set up shared SSH keys if you do not want to supply a password for each Torque MOM Host.

```
[root]# scp torque-package-mom-linux-x86_64.sh <torque-mom-host>:
```

- c. Copy the `pbs_mom` startup script to each Torque MOM Host:

```
[root]# scp contrib/systemd/pbs_mom.service <mom-node>:/lib/systemd/system/
```

2. On *each* Torque MOM Host, install the self-extracting MOM package:

```
[root]# ./torque-package-mom-linux-x86_64.sh --install
```

2.3.1.E Update the Torque Clients

This section contains instructions on updating the Torque clients on the Torque Client Hosts (including the Moab Server Host and Torque MOM Hosts, if applicable).

1. On the Torque Server Host, do the following:
 - a. Copy the self-extracting client package to *each* Torque Client Host. We recommend that you use a remote shell, such as SSH, to install packages on remote systems. Set up shared SSH keys if you do not want to supply a password for each Torque MOM Host.

```
[root]# scp torque-package-clients-linux-x86_64.sh <torque-client-host>:
```

- b. If Moab Workload Manager is part of your configuration, copy the self-extracting devel package to the Moab Server Host:

```
[root]# scp torque-package-devel-linux-x86_64.sh <moab-server-host>:
```

- c. Copy the trqauthd startup script to each Torque Client Host:

```
[root]# scp contrib/systemd/trqauthd.service <torque-client-host>:/lib/systemd/system/
```

2. On *each* Torque Client Host, do the following.

i This step can be done from the Torque server from a remote shell, such as SSH. Set up shared SSH keys if you do not want to supply a password for each Torque Client Host.

```
[root]# ./torque-package-clients-linux-x86_64.sh --install
```

3. If Moab Workload Manager is part of your configuration, run the following command on the Moab Server Host:

```
[root]# ./torque-package-devel-linux-x86_64.sh --install
```

2.3.1.F Start Torque Services

Do the following.

1. On the Torque Server Host, start up the Torque server:

```
[root]# systemctl daemon-reload
[root]# systemctl start pbs_server.service
```

2. On *each* Torque MOM Host, start up the Torque MOM service:

```
[root]# systemctl daemon-reload
[root]# systemctl start pbs_mom.service
```

3. On *each* Torque Client Host (including the Moab Server Host, Torque Server Host, and Torque MOM Hosts, if applicable), start up the `trqauthd` service:

```
[root]# systemctl daemon-reload
[root]# systemctl start trqauthd.service
```

2.3.1.G Perform Status and Error Checks

On the Torque Server Host, verify that the status of the nodes and jobs are as expected:

```
[root]# pbsnodes
[root]# qstat
```

2.3.2 Upgrading Moab Workload Manager

This topic provides instructions to upgrade Moab Workload Manager to the latest release version. Depending on which version of Moab you are presently running, upgrade instructions may vary.


Moab Workload Manager uses the standard `configure`, `make`, and `make install` steps for upgrades. This topic provides a number of sample steps referenced to a particular installation on a Linux platform using the bash shell. These steps indicate the user ID in brackets performing the step. The exact commands to be performed and the user that issues them will vary based on the platform, shell, installation preferences, and other factors.

In this topic:

[2.3.2.A Recommendations - page 63](#)

[2.3.2.B Upgrade Moab Workload Manager - page 63](#)

2.3.2.A Recommendations

 We highly recommend that you *first* perform upgrades in a *test environment*. See the warning in [1.1 Preparing for Upgrade](#). We also recommend that you verify the policies, scripts, and queues work the way you want them to in this test environment. See 'Testing New Releases and Policies' in the *Moab Workload Manager Administrator Guide* for more information.

2.3.2.B Upgrade Moab Workload Manager

On the Moab Server Host, do the following.

1. Download the latest Moab Workload Manager tarball from the [Adaptive Computing](#) website.
2. Untar the distribution file, for example:

```
[root]# tar -xzvf moab-10.0.1-<OS>.tar.gz
```

 The variable marked <OS> indicates the OS for which the build was designed.

3. Change directory into the extracted directory:

```
[root]# cd moab-10.0.1-<OS>
```

4. Configure the installation package. Use the same configure options as when Moab was installed previously. If you cannot remember which options were used previously, check the `config.log` file in the directory where the previous version of Moab was installed from. For a complete list of configure options, use `./configure --help`.
5. Stop Moab:

```
[root]# systemctl stop moab.service
```

 While Moab is down, all currently running jobs continue to run on the nodes, the job queue remains intact, and new jobs cannot be submitted to Moab.

6. Back up your Moab Workload Manager home directory (`/opt/moab/` by default) before continuing.
7. Install Moab:

```
[root]# make install
```

i Default configuration files are installed during `make install`. Existing configuration files are not overwritten and the new files are given a `.dist` extension.

8. If you use ODBC, you must confirm the database schema compatibility. See 'Migrating Your Database to Newer Versions of Moab' in the *Moab Workload Manager Administrator Guide* for more information.
9. Verify the version number is correct before starting the new server version:

```
[root]# moab --about
```

You should get something similar to the following in the response:

```
Defaults:  server=:42559  cfgdir=/opt/moab (env)  vardir=/opt/moab
Build dir:  /tmp/jenkins/workspace/MWM-10.0.1/label/build-<OS>
Build host: us-devops-build10
Build date: Oct 09 13:00:00 MST 2023
Build args: NA
Compiler Flags:  -D_M64 -D_BUILDDATETIME="2023100913" -DMUSEZEROMQ -
DMUSEWEBSERVICES -DMUSEMONGODB -DMMAX_GRES=512 -DMMAX_RANGE=2048 -DMMAX_TASK=32768
-fPIC -gdwarf-3 -Wall -Wextra -DVALGRIND -Og -x c++ -std=c++11 -DDMAX_PJOB=512 -D_
GNU_SOURCE
Compiled as little endian.
Version: moab server 10.0.1 (revision 2023100913, changeset
14dee972ebcee919207e48054e9f285db9f6a555)
```

10. Start Moab:

```
[root]# systemctl daemon-reload
[root]# systemctl start moab.service
```

2.3.3 Upgrading Moab Accounting Manager

This topic provides instructions to upgrade MAM to the latest release version. It includes instructions for migrating your database schema to a new version if necessary.

Moab Accounting Manager uses the standard `configure`, `make`, and `make install` steps for upgrades. This document provides a number of sample steps referenced to a particular installation on a Linux platform using the bash shell. These steps indicate the user ID in brackets performing the step. The exact commands to be performed and the user that issues them will vary based on the platform, shell, installation preferences, and other factors.

On the MAM Server Host, do the following.

1. Determine the MAM Accounting admin user and change to that user:

```
[root]# mam-list-users | grep 'Accounting Admin'
mam      True                               Accounting
Admin
[root]# su - mam
```

2. Determine whether a migration is necessary.

- a. Determine your database version:

```
[mam]$ mam-shell System Query
```

- b. If the current version is lower than 10.0, then your database and configuration files will need to be migrated. The steps required to do so are incorporated in the remaining steps for this topic.

3. Stop the server daemon:

```
[mam]$ su -c "systemctl stop mam.service"
```

4. If a migration is required, create a database backup:

```
[mam]$ pg_dump -U <mam_database_user> -W <old_database_name> > /tmp/<old_database_name>.sql
```

i MySQL is no longer a supported database for MAM. If you are using MySQL for your MAM database, follow the instructions in [2.3.5 Migrating the MAM Database from MySQL to PostgreSQL - page 68](#) to convert your database.

5. Verify that each of the prerequisites listed in [2.1.4 Installing Moab Accounting Manager - page 37](#) have been satisfied.
6. Download the latest MAM tarball from the [Adaptive Computing](#) website.
7. Unpack the tar archive and change directory into the top directory of the distribution:

```
[mam]$ tar -zxvf mam-10.0.1.tar.gz
[mam]$ cd mam-10.0.1
```

8. Configure Moab Accounting Manager by running `configure` with the desired options.

We recommend that you use the same configure options that were used in the previous installation. You can examine the `config.log` file where you unpacked your previous distribution to help determine the configuration options that were used to install the prior version of MAM.

```
[mam]$ ./configure
```

9. Run `make` to compile the program:

```
[mam]$ make
```

i Depending on your configuration, you may need to replace 'make' with a make command that includes additional functionality. Specifically:

- If you only need to install the clients on a particular system, use `clients-only`.
- If you only need to install the web GUI on a particular system, use `make gui-only`.
- If you only need to install the web services on a particular system, use `make ws-only`.

10. Run `make install` as root to install Moab Accounting Manager:

```
[mam]$ su -c "make install"
```

i Depending on your configuration, you may need to replace 'make install' with a make command that includes additional functionality. Specifically:

- If you only need to install the clients on a particular system, use `make install-clients-only`.
- If you only need to install the web GUI on a particular system, use `make install-gui-only`.
- If you only need to install the web services on a particular system, use `make install-ws-only`.

11. Start the server daemon:

```
[mam]$ su -c "systemctl daemon-reload"
[mam]$ su -c "systemctl start mam.service"
```

12. If you need to migrate, you will do so by running one or more migration scripts. You must run every incremental migration script between the version you are currently using and the new version (10.0). These scripts are designed to be rerunnable, so if you encounter a failure, resolve the failure and rerun the migration script. If you are unable to resolve the failure and complete the migration, contact [Support](#).

For example, if you are migrating from Moab Accounting Manager version 8.0, you must run four migration scripts: the first to migrate the database schema from 8.0 to 8.1, the second to migrate from 8.1 to 9.0, the third to migrate the database schema from 9.0 to 9.1 and the fourth to migrate the database schema from 9.1 to 10.0:

```
[mam]$ sbin/migrate_8.0-8.1.pl
[mam]$ sbin/migrate_8.1-9.0.pl
[mam]$ sbin/migrate_9.0-9.1.pl
[mam]$ sbin/migrate_9.1-10.0.pl
```

13. Verify that the resulting database schema version is 10.0:

```
[mam]$ mam-shell System Query
```

Name	Version	Description
Moab Accounting Manager	10.0	Commercial Release

14. Verify that the executables have been upgraded to 10.0.1:

```
[mam]$ mam-server -v
```

```
Moab Accounting Manager version 10.0.1
```

15. If you are upgrading MAM from a version prior to 9.1.0, and you want to use MAM Web Services, perform the following procedures (provided in the Installing Moab Accounting Manager topic):

- [2.1.4.G Configure MAM Web Services - page 44](#)
- [2.1.4.I Access MAM Web Services - page 45](#)

2.3.4 Upgrading RLM Server

Adaptive Computing *strongly* recommends that your RLM Server is the latest version.

In this topic:

- [2.3.4.A Confirm if an Upgrade is Needed - page 67](#)
- [2.3.4.B Upgrade the RLM Server - page 67](#)

2.3.4.A Confirm if an Upgrade is Needed

Run the following command to determine your current version of RLM Server:

```
[root]# /opt/rlm/rlm -v
```

If the version reported is less than 12.1BL2, continue with the section to Upgrade the RLM Server below.

2.3.4.B Upgrade the RLM Server

i These instructions assume you used /opt/rlm as the install location.

On the RLM Server Host, do the following.

1. Download the latest RLM build (`ac-rlm-12.1.2.tar.gz`) from the [Adaptive Computing](#) website.

2. Stop the RLM service:

```
[root]# systemctl stop rlm.service
```

3. Archive the existing RLM installation, including the license file(s):

```
[root]# mv /opt/rlm/ /opt/rlm-<archive_version>/
```

4. Install the new tarball files:

```
[root]# mkdir -p -m 0744 /opt/rlm
[root]# cd /opt/rlm
[root]# tar -xzvf /<unpack-directory>/ac-rlm-12.1.2.tar.gz --strip-components=1
[root]# chown -R rlm:rlm /opt/rlm
```

i The `--strip-components=1` removes the `'ac-rlm-12.1.2/'` from the relative path so that they are extracted into the current directory.

5. Install the startup scripts.

i If you are using a user:group other than `rlm:rlm` or a location other than `/opt/rlm`, then edit the following files to reflect those changes after copying them.

```
[root]# cp systemd/rlm.service /etc/systemd/system
```

6. Restore the license file(s):

```
[root]# cp /opt/rlm-<archive_version>/*.lic /opt/rlm/
```

7. Restart the RLM service:

```
[root]# systemctl daemon-reload
[root]# systemctl restart rlm.service
```

2.3.5 Migrating the MAM Database from MySQL to PostgreSQL

PostgreSQL is the preferred DBMS for MAM. If you have already installed MySQL as the DBMS for MAM, you are not required to migrate their database to use PostgreSQL at this time. However, MySQL is considered deprecated and new installations will only use PostgreSQL.

i PostgreSQL does not provide a standard procedure for migrating an existing database from MySQL to PostgreSQL. Adaptive Computing has had success using the `py-mysql2pgsql` tools for migrating/converting/exporting data from MySQL to PostgreSQL. See <https://github.com/philipsoutham/py-mysql2pgsql> for additional details.

To Migrate the MAM Database

This procedure was successfully tested on an actual customer MySQL database with millions of transactions. It completed in less than an hour.

1. Make a backup copy of your MySQL mam database:

```
[root]# mysqldump mam > /archive/mam.mysql
```

2. Follow the instructions to Install PostgreSQL:

- **Manual Install - 1.1 Installing Moab Web Services**

3. Install the prerequisite packages:

4. Install `pg-mysql2pgsql` (from source):

```
[root]# cd /software
[root]# git clone git://github.com/philipsoutham/py-mysql2pgsql.git
[root]# cd py-mysql2pgsql
[root]# python setup.py install
```

5. Run `pg-mysql2pgsql` once to create a template yaml config file:

```
[root]# py-mysql2pgsql -v
```

6. Edit the config file to specify the MySQL database connection information and a file to output the result:

```
[root]# vi mysql2pgsql.yml
```

```
mysql:
  hostname: localhost
  port: 3306
  socket:
  username: mam
  password: changeme
  database: mam
  compress: false
  destination:
  # if file is given, output goes to file, else postgres
  file: /archive/mam.pgsql
  postgres:
  hostname: localhost
  port: 5432
  username:
  password:
```

```
database:
```

7. Run the `pg-mysql2pgsql` program again to convert the database:

```
[root]# py-mysql2pgsql -v
```

8. Create the `mam` database in PostgreSQL:

```
[root]# su - postgres
[postgres]$ psql
postgres=# create database "mam";
postgres=# create user mam with password 'changeme!';
postgres=# \q
[postgres]$ exit
```

9. Import the converted data into the PostgreSQL database:

```
[root]# su - mam
[mam]$ psql mam < /archive/mam.pgsql
```

10. Point MAM to use the new postgresql database:

```
[mam]$ cd /software/mam-latest
[mam]$ ./configure # This will generate an etc/mam-
server.conf.dist file
[mam]$ vi /opt/mam/etc/mam-server.conf # Merge in the database.datasource from
etc/mam-server.conf.dist
```

11. Restart Moab Accounting Manager:

```
[mam]$ mam-server -r
```

Chapter 3: Troubleshooting

This chapter details some common problems and general solutions. Additional troubleshooting can be found in the individual Moab HPC Suite component documentation.

Note: If you currently have a support services contract and encounter an installation problem that you can't resolve, please submit an [online support case](#), and a technical support specialist will contact you.

In this chapter:

3.1 General Issues	71
3.2 Port Reference	77
3.3 Moab Workload Manager Issues	81

3.1 General Issues

This topic details some common problems and general solutions.

In this topic:

3.1.1 Where do I set credentials and what are the default values? - page 71

3.1.1 Where do I set credentials and what are the default values?

Communication and cooperation between various components of the Moab HPC Suite requires credentials to be properly configured. For ease of use, the credential information, including where credentials are set, default values, and where they are used are grouped by database and product.

3.1.1.A Database Credentials

MongoDB

Database	User	Default Password	Used By	Parameters
admin	admin_user	secret1	system admins	NA
moab	moab_user	secret2	/opt/moab/etc/moab-private.cfg	MONGouser, MONGOPASSWORD
moab	mws_user	secret3	/opt/mws/etc/mws-config.groovy	grails.mongo.username, grails.mongo.password
moab	insight_user	secret4	/opt/insight/etc/config.groovy	moab.mongo.username, moab.mongo.password
mws	mws_user	secret3	/opt/mws/etc/mws-config.groovy	grails.mongo.username, grails.mongo.password
insight	insight_user	secret4	/opt/insight/etc/config.groovy	mongo.username, mongo.password
insight	mws_user	secret3	https://<mws_server>:8080/mws/admin/plugins/edit/viewpoint-query-helper	user, password
nitro-db	nitro_user	secret5	/opt/nitro-web-services/etc/nitro.cfg	db_username, db_password
reporting	reporting_user	secret6	/opt/reporting/application.conf	database.username, database.password

i The following characters must be escaped in strings in the `/opt/insight/etc/config.groovy` and `/opt/mws/etc/mws-config.groovy` files (such as when used in a password): `\` (backslash), `"` (double quote), `'` (single quote), `$` (dollar sign). Example: `mongo.password="my\$cool\$password"`. We recommend that you avoid using these characters.

PostgreSQL

Database	User	Default Password	Used By	Parameters
moab_viewpoint	moab_viewpoint	changeme!	<code>/opt/viewpoint/etc/viewpoint.cfg</code>	VIEWPOINT_DATABASE_USER, VIEWPOINT_DATABASE_PASSWORD
mam	mam	changeme!	<code>/opt/mam/etc/mam-server.cfg</code>	database.user, database.password

Apache Drill

The Drill host should have a user that Reporting Web Services can use to authenticate to Drill.

Host	User	File	Parameter Name	Default Value
Drill host	drilluser	<code>/opt/reporting-web-services/etc/application.properties</code>	reporting.rest.drill.username	changeme!

3.1.1.B Product Credentials

Moab Workload Manager

Declared Parameter		Used By		Default Value
File	Parameter Name	File	Parameter Name	
/opt/moab/etc/moab-private.cfg	MESSAGEQUEUES SECRETKEY	/opt/mws/etc/mws-config.groovy	moab.messageQueue.secretKey	NA
		/opt/insight/etc/config.groovy	messageQueue.secretKey	
/opt/moab/etc/moab.key	NA	/opt/mws/etc/mws-config.groovy	moab.secretKey	NA

Moab Accounting Manager

Declared Parameter		Used By		Default Value
File	Parameter Name	File	Parameter Name	
/opt/mam/etc/mam-site.conf	token.value	/opt/moab/etc/moab-private.cfg	CLIENTCFG [AM:mam] KEY	NA

Moab Web Services

Declared Parameter		Used By		Default Value
File	Parameter Name	File	Parameter Name	
/opt/mws/etc/mws-config.groovy	auth.defaultUser.username	https://<viewpoint_server>:8081/configuration/	Username	moab-admin

Declared Parameter		Used By		Default Value
File	Parameter Name	File	Parameter Name	
		/opt/moab/etc/moab-private.cfg	CLIENTCFG [RM:mws] USERNAME	
/opt/mws/etc/mws-config.groovy	auth.defaultUser.password	https://<viewpoint_server>:8081/configuration/	Password	change!
		/opt/moab/etc/moab-private.cfg	CLIENTCFG [RM:mws] PASSWORD	
/opt/mws/etc/mws-config.groovy	grails.plugin.springsecurity.oauth.Provider.clients[0].clientSecret	https://<viewpoint_server>:8081/configuration/	Client Secret	NA

i The following characters must be escaped in strings in the `/opt/insight/etc/config.groovy` and `/opt/mws/etc/mws-config.groovy` files (such as when used in a password): `\` (backslash), `"` (double quote), `'` (single quote), `$` (dollar sign). Example: `mongo.password="my\$cool$password"`. We recommend that you avoid using these characters.

Nitro Web Services

Declared Parameter		Used By	Default Value
File	Parameter Name		
/opt/nitro-web-services/etc/nitro.cfg	ws_admin_password	Installation - default NWS API user creation	ChangeMe2!
/opt/nitro-web-services/etc/nitro.cfg	ws_readonly_username	Installation - default NWS API user creation https://<viewpoint_server>:8081/configuration/ -> Nitro Services -> Username	nitro-readonly-user
/opt/nitro-web-services/etc/nitro.cfg	ws_readonly_password	Installation - default NWS API user creation https://<viewpoint_server>:8081/configuration/ -> Nitro Services -> Password	ChangeMe3!
/opt/nitro-web-services/etc/nitro.cfg	ws_writeonly_username	Installation - default NWS API user creation /opt/nitro-web-services/etc/zmq_job_status_adapter.cfg -> username	nitro-writeonly-user
/opt/nitro-web-services/etc/nitro.cfg	ws_writeonly_password	Installation - default NWS API user creation /opt/nitro-web-services/etc/zmq_job_status_adapter.cfg -> password	ChangeMe4!

Viewpoint

Declared Parameter		Used By	Default Value
File	Parameter Name		
/opt/viewpoint/etc/viewpoint.cfg	username	https://<viewpoint_server>:8081/login/	viewpoint-admin
/opt/viewpoint/etc/viewpoint.cfg	password	https://<viewpoint_server>	changeme!

Declared Parameter		Used By	Default Value
File	Parameter Name		
		>:8081/login/	

3.2 Port Reference

The following tables contains the port numbers for the various products in the Moab HPC Suite.

RLM Server

Location	Ports	Functions	When Needed
RLM Server Host	5053	RLM Server Port	Always
RLM Server Host	5054	RLM Web Interface Port	Always
RLM Server Host	5135	ISV adaptiveco Port (for the Adaptive license-enabled products)	For Moab Workload Manager <i>and</i> if Nitro is part of your configuration

Torque Resource Manager

Location	Ports	Functions	When Needed
Torque Server Host	15001	Torque Client and MOM communication to Torque Server	Always
Torque MOM Host (Compute Nodes)	15002	Torque Server communication to Torque MOMs	Always
Torque MOM Host (Compute Nodes)	15003	Torque MOM communication to other Torque MOMs	Always

Moab Workload Manager

Location	Ports	Functions	When Needed
Moab Server Host	42559	Moab Server Port	If you intend to run client commands on a host different from the Moab Server Host <i>or</i> if you will be using Moab in a grid

Moab Accounting Manager

Location	Ports	Functions	When Needed
MAM Server Host	7112	MAM Server Port	If you will be installing the MAM Server on a different host from where you installed the Moab Server <i>or</i> you will be installing the MAM Clients on other hosts
MAM Web Server Host	443	HTTPS Port	If using the MAM GUI or MAM Web Services
MAM Database Host	5432	MAM PostgreSQL Server Port	If you will be installing the MAM Database on a different host from the MAM Server

Moab Web Services

Location	Ports	Functions	When Needed
MWS Server Host	8080	Tomcat Server Port	Always
MWS Database Host	27017	MWS MongoDB Server Port	If you will be installing the MWS Database on a different host from the MWS Server

Moab Insight

Location	Ports	Functions	When Needed
Insight Server Host	5568	Insight Server Port	Always
Moab MongoDB Database Host	27017	Moab MongoDB Server Port	Always

Location	Ports	Functions	When Needed
Moab Server Host	5574	Moab Data Port	Always
Moab Server Host	5575	Moab Reliability Port	Always

Moab Viewpoint

Location	Ports	Functions	When Needed
Viewpoint Server Host	8081	Viewpoint Web Server Port	Always
Moab Server Host	8443	Viewpoint File Manager Port	Always
Viewpoint Database Host	5432	Viewpoint PostgreSQL Database Port	If you will be installing the Viewpoint Database on a different host from the Viewpoint Server

Nitro

i The listed ports are for configurations that have only one coordinator. If multiple coordinators are run on a single compute host, then sets of ports (range of 4) must be opened for the number of expected simultaneous coordinators.

Location	Ports	Functions	When Needed
Compute Hosts (Nitro Coordinator)	47000	Coordinator/Worker communication	Always
Compute Hosts (Nitro Coordinator)	47001	Coordinator PUB/SUB channel - publishes status information	Always
Compute Hosts (Nitro Coordinator)	47002	Reserved for future functionality	
Compute Hosts (Nitro Coordinator)	47003	API communication channel	Always

Nitro Web Services

Location	Ports	Functions	When Needed
Nitro Web Services Host	9443	Tornado Web Port	Always
Nitro Web Services Host	47100	ZMQ Port	Always
Nitro Web Services Database Host	27017	Nitro Web Services MongoDB Server Port	If you will be installing the Nitro Web Services Database on a different host from Nitro Web Services

Reporting

Suggested Host	Service	Ports	Function	When Needed
Reporting Master	HDFS name node	8020	HDFS communication	Always
Reporting Master	HDFS name node	50070	HDFS web interface	Always
Reporting Master	Spark Master	6066, 7077	Spark communication	Always
Reporting Master	Spark Master	8082	Spark Master web interface	Always
Reporting Master	Apache Kafka	9092	Kafka communication	Always
Reporting Master	Apache Zookeeper	2181	Zookeeper communication with Kafka and Drill	Always
Insight Server	Apache Drill	8047	Drill HTTP interface	Always
Reporting Worker	HDFS data node	50075, 50010, 50020	HDFS communication	Always
Reporting Worker	Spark	4040	Spark communication	Always

Suggested Host	Service	Ports	Function	When Needed
	Worker			
Reporting Worker	Spark worker	8083	Spark worker web interface	Always
MWS Host	Tomcat	8080	Reporting Web Services HTTP interface	Always
MWS Host	MongoDB	27017	MongoDB communication	Always

3.3 Moab Workload Manager Issues

This topic details some common problems and general solutions for Moab Workload Manager. See also 'Troubleshooting and System Maintenance' in the *Moab Workload Manager Administrator Guide*.

In this topic:

[3.3.1 Moab error: cannot determine local hostname - page 81](#)

[3.3.2 Moab error: Moab will now exit due to license file not found - page 82](#)

3.3.1 Moab error: cannot determine local hostname

```
# systemctl start moab.service
Starting moab: ERROR:   cannot determine local hostname - node is misconfigured
                        [FAILED]
```

```
...
SCHEDCFG [Moab]                SERVER=<moab-hostname>:42559
...
```

Also check `/etc/hosts` to be sure the host name resolves, at least with localhost:

```
...
127.0.0.1  <moab-hostname> localhost localhost.localdomain localhost4
localhost4.localdomain4
...
```

3.3.2 Moab error: Moab will now exit due to license file not found

```
# systemctl start moab.service
Starting moab: Moab will now exit due to license file not found
Please contact Adaptive Computing (sales@adaptivecomputing.com) to get a license for
your system
[FAILED]
```

If you encounter this error when starting Moab HPC Suite, make sure your Moab HPC Suite license file is named `moab.lic` and is located in the `/opt/moab/etc/` directory.

Also make sure the license is not expired. The expiration date is listed in the license file, for example:

```
# cat /opt/moab/etc/moab.lic
...
# Expires after Tue Dec 31 10:43:46 2024
...
```