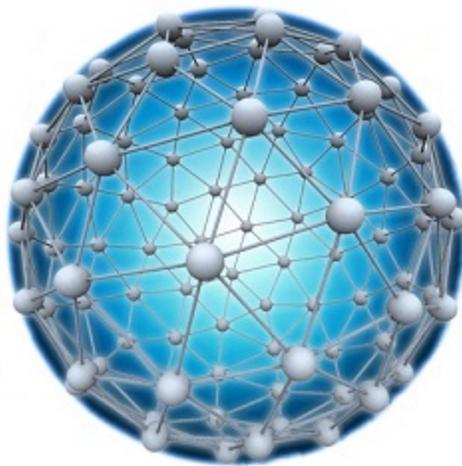


Moab ODDC Connect

User Guide 4.0.0

April 2020



 **Adaptive**
COMPUTING
On-Demand Data Center™
Powered by NODUS Cloud OS

Contents

Welcome	3
Legal Notices	3
Revision History	4
Related Documentation	4
Chapter 1: Moab ODDC Connect Cloud Bursting	5
1.1 Moab ODDC Connect Bursting Overview	5
1.2 Moab ODDC Connect Bursting Architecture	6
1.3 Securing Data Transmissions	7
Chapter 2: Installing Moab, Viewpoint, and ODDC	8
2.1 Installing Moab	8
2.2 Installing Viewpoint	8
2.3 Installing ODDC	8
Chapter 3: Moab ODDC Connect Configuration	9
3.1 Deploying the Bursting Cluster	9
3.2 Enabling MOAB to use ODDC as a Remote Resource Manager	10
3.3 Configuration	10
3.4 Configuring Moab to Connect to ODDC	11
3.5 ODDC Use Cases	11
3.6 Starting Moab ODDC Connect	13
Glossary	14
Index	17

Welcome

Adaptive Computing's On-Demand Data Center (ODDC) platform, powered by NODUS Cloud OS, gives companies the ability to spin up temporary or persistent data center infrastructure resources quickly, inexpensively, and on-demand.

Moab ODDC Connect is a highly flexible and extendable solution that enables HPC systems to burst scheduled workloads to an external cloud on-demand. It can provision a full stack across multiple cloud environments, and can be customized for multiple use cases and scenarios. Access to unlimited HPC compute resources is available in the cloud from multiple cloud service providers.

This *Moab ODDC Connect User Guide* will show you how to install and use the Moab ODDC Connect environment. This guide is intended for Moab ODDC Connect system administrators.

For pricing information, contact us at sales@adaptivecomputing.com.

For support, contact us at support@adaptivecomputing.com.

Legal Notices

Copyright © 2020, 2021. Adaptive Computing Enterprises, Inc. All rights reserved.

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, Moab®, Moab HPC Suite, Moab Viewpoint, Moab Wide Area Grid, NODUS Cloud OS™, and other Adaptive Computing products are either registered trademarks or trademarks of Adaptive Computing Enterprises, Inc. The Adaptive

Revision History

Computing logo is a trademark of Adaptive Computing Enterprises, Inc. All other company and product names may be trademarks of their respective companies.

The information contained herein is subject to change without notice and is not warranted to be error free. If you find any errors, please report them to us in writing.

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

Revision History

Date	Release
April 2020	Moab NODUS Cloud OS 1.0.0
September 2020	Moab NODUS Cloud OS 2.0.0
November 2020	Moab NODUS Cloud OS 3.0.0
April 2021	Moab ODDC Connect 4.0.0

Related Documentation

- *Moab HPC Suite Installation and Configuration Guide for SUSE 12-Based Systems* (see the link below)
- *Moab HPC Suite Installation and Configuration Guide for Red Hat 6-Based Systems* (see the link below)
- *Moab HPC Suite Installation and Configuration Guide for Red Hat 7-Based Systems* (see the link below)
- *Moab Viewpoint Reference Guide* (see the link below)
- *Moab Workload Manager Administrator Guide* (see the link below)
- *Torque Resource Manager Administrator Guide* (see the link below)
- *On-Demand Data Center User Guide* (see the link below)

[Adaptive Computing Support Portal / Product Documentation](#)

Chapter 1: Moab ODDC Connect Cloud Bursting

This chapter provides information about Moab ODDC Connect cloud bursting.

In this chapter:

1.1 Moab ODDC Connect Bursting Overview	5
1.2 Moab ODDC Connect Bursting Architecture	6
1.3 Securing Data Transmissions	7

1.1 Moab ODDC Connect Bursting Overview

During the course of operation, the number of job submissions will increase and decrease. Under some circumstances, the job backlog may increase to the point where additional resources are required to complete the job backlog in a reasonable time frame. In this scenario, jobs are held until resources become available.

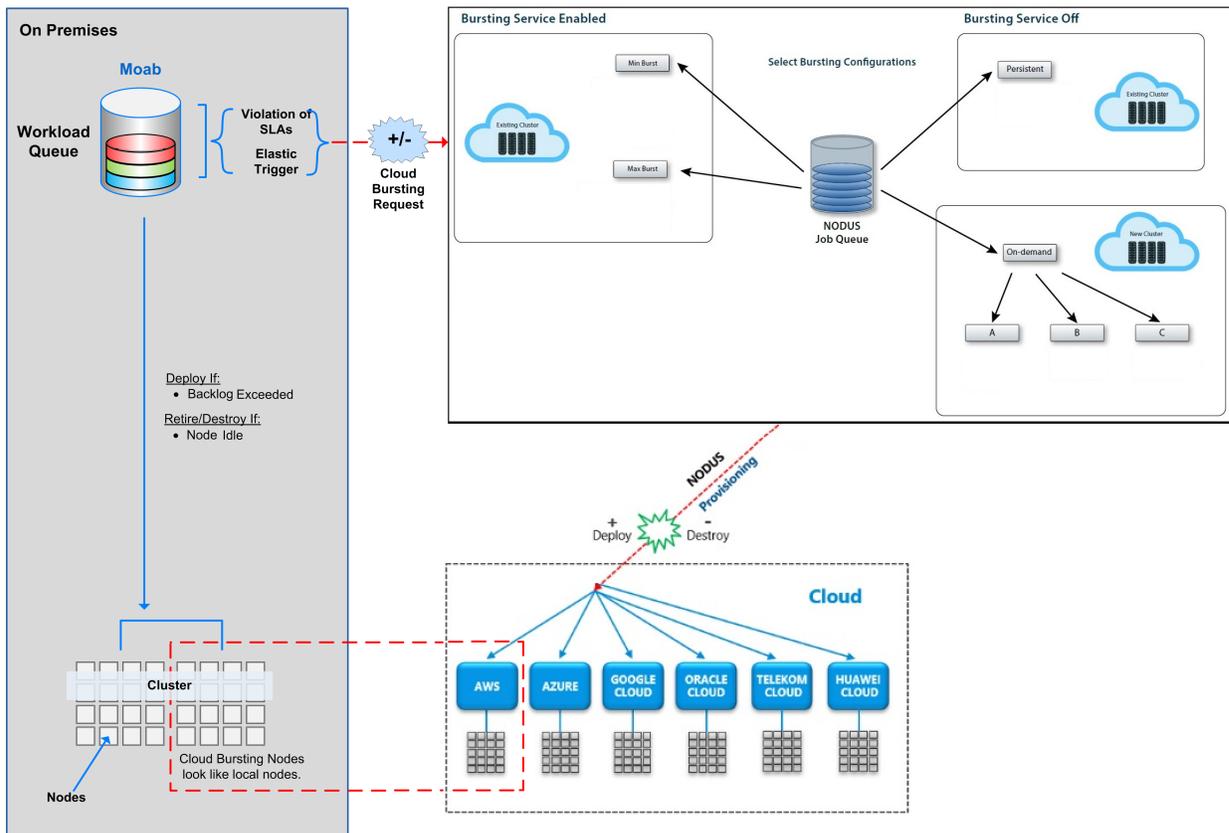
The Moab ODDC Connect solution enables MOAB to use ODDC as a Remote Resource Manager (Native Resource Manager). The Native Resource Manager communicates with ODDC directly. Moab controls job submission as before, using on-premises nodes first and then passes on jobs submitted to available ODDC nodes when a backlog occurs.

All functions are being managed by ODDC after receiving instruction from Moab that there is a job awaiting a resource to be able to run. If there is already an available ODDC node, it uses this node first to run a job against. If there are more jobs waiting to run (backlog), the ODDC Remote Resource Manager communicates directly with ODDC Cloud OS to assist in making resources available for a waiting job on Moab, and manages all bursting requirements on ODDC. When jobs complete using an ODDC node, based on the configuration set on ODDC (bursting options), ODDC manages the removal of unused nodes. With the available ODDC functionality (being heterogeneous), ODDC can be configured to access multiple cloud service providers.

1.2 Moab ODDC Connect Bursting Architecture

Moab ODDC Connect is a solution based on Moab's Remote Resource Manager framework, using the ODDC to launch instances in your cloud service provider account. Elastic Computing with Moab ODDC Connect is triggered based on the size of the job backlog.

This diagram depicts the Moab ODDC Connect solution:



When a job backlog occurs, Moab uses ODDC as a Remote Resource Manager to utilize the ODDC bursting capabilities to spin up the required nodes for jobs to run being delayed due to a backlog. Using the ODDC bursting functionality, ODDC automatically spins up, takes offline, or shuts down nodes depending on the total requirements for the job queue. If there are not enough online nodes to run all jobs, bursting brings on as many nodes as needed. If there are more nodes than needed, the excess nodes are taken offline. If the job queue is empty, all nodes are shut down after a specified period of time.

Min Burst spins up the minimum number of compute nodes required to complete all jobs in the queue, which is ideal for budgeting and controlling cloud costs.

Max Burst spins up enough compute nodes to complete all the jobs in the queue immediately; this gets results as fast as possible. Max Burst is limited by the size of the cluster and will not create new nodes.

Persistent bursting spins up all or a portion of the licensed instances in a cluster that remain persistent for a period of time and brings nodes online or shuts them down as needed.

On-demand bursting spins up the number of nodes required to run one job now; this is an isolated cluster, not for sharing with other jobs. The on-demand types are: Destroy Compute Nodes (the head node stays active and the compute nodes are destroyed), Offline Compute Nodes (the head node stays active and the compute nodes go offline), and Destroy Full Cluster (the full cluster is destroyed including the head node).

1.3 Securing Data Transmissions

Because jobs can be run off-site, it is paramount that the transfer of data is done securely.

There are three ways to create a secure connection from your network to your cloud service provider account:

- Create a point-to-point VPN connection from your network to a Virtual Private Cloud (VPC) in your cloud service provider account. Your cloud service provider does charge for this capability, but this enables one single secure line of communication between your network and your cloud service provider account. Transfer speeds are based on your internet bandwidth and the distance from your network to the Region and Availability Zone you are bursting to. See VPN connection information in your cloud service provider documentation for how this is to be configured.
- Use your cloud service provider's direct connection options to configure a dedicated network connection from your network directly to your cloud service provider account. Because it is a private connection that does not go over the public internet, there is no need for a VPN, and speeds between 1 Gbps to 10 Gbps are possible. This solution comes at a considerably higher price than a cloud service provider VPN, but with the advantage of a more predictable network speed.
- Install a VPN client into your cloud service provider image to connect to a VPN server on your network when the cloud service provider's instance is launched. Although this solution requires the most up-front configuration, it enables the cloud service provider's instances to connect to your network securely with no cost from your cloud service provider. This solution uses one VPN connection per cloud service provider instance, so there is a potential for excessive network overhead. Transfer speeds are also based on your internet bandwidth and the distance from your network to the Region and Availability Zone to which you are bursting.

Chapter 2: Installing Moab, Viewpoint, and ODDC

To interface Moab with ODDC, the installation of Moab, Viewpoint, and ODDC is required.

2.1 Installing Moab

For instructions, see the following documents for your operating system:

- *Moab HPC Suite Installation and Configuration Guide 9.1.3 for Red Hat 6-Based Systems* (see the link below)
- *Moab HPC Suite Installation and Configuration Guide 9.1.3 for Red Hat 7-Based Systems* (see the link below)
- *Moab HPC Suite Installation and Configuration Guide 9.1.3 for SUSE 12-Based Systems* (see the link below)

[Adaptive Computing Support Portal / Product Documentation](#)

2.2 Installing Viewpoint

For instructions, see the *Moab Viewpoint Reference Guide*.

[Adaptive Computing Support Portal / Product Documentation](#)

2.3 Installing ODDC

For instructions, see the *On-Demand Data Center User Guide*.

[Adaptive Computing Support Portal / Product Documentation](#)

Chapter 3: Moab ODDC Connect Configuration

Now that you have Moab, Viewpoint, and ODDC installed and accessible, complete the steps in this chapter to configure your Moab ODDC Connect environment. The configuration steps below apply to any cloud service provider.

Note: All commands must be run as `root`.

In this chapter:

3.1	Deploying the Bursting Cluster	9
3.2	Enabling MOAB to use ODDC as a Remote Resource Manager	10
3.3	Configuration	10
3.4	Configuring Moab to Connect to ODDC	11
3.5	ODDC Use Cases	11
3.6	Starting Moab ODDC Connect	13

3.1 Deploying the Bursting Cluster

1. Deploy your bursting cluster from the ODDC GUI. Refer to the *On-Demand Data Center User Guide* for instructions.
2. After successful deployment of your cluster, complete the sections below. The steps are done on ODDC or Moab as indicated.

3.2 Enabling MOAB to use ODDC as a Remote Resource Manager

1. Untar `moab-nodus-connect.tar.gz` (supplied by Adaptive Computing) at your home directory and install Moab ODDC Connect.

- A. We suggest using the following commands:

```
cp moab-nodus-connect.tar.gz / ~
```

```
tar -xvf moab-nodus-connect.tar.gz
```

```
cd /release
```

```
sudo ./install.sh
```

- B. If this is an upgrade, you need to stop the ODDC monitor daemon first (**`systemctl stop nodus-monitor`**), remove/delete the file `.tracking.json` (**`rm/opt/moab/contrib/nodus/.tracking.json`**), and start the ODDC monitor daemon again (**`systemctl start nodus-monitor`**), as the ODDC monitor daemon creates a new `.tracking.json` file on start up.
- C. Also important to note, if this is a new install, you need to restart the ODDC monitor daemon (**`systemctl restart nodus-monitor`**) after you have created and updated the `.env` file (see the section [Configuring Moab to Connect to ODDC](#)) as it reads the content of the file to correctly connect to ODDC and read the cluster you want your Moab jobs to run on.

3.3 Configuration

1. Add the following lines to your `$MOABHOMEDIR/etc/moab.cfg` file:

```
RMCFG [nodus] TYPE=NATIVE
RMCFG [nodus] CLUSTERQUERYURL=exec://$HOME/contrib/nodus/clusterquery
RMCFG [nodus] WORKLOADQUERYURL=exec://$HOME/contrib/nodus/workloadquery
RMCFG [nodus] JOBSUBMITURL=exec://$HOME/contrib/nodus/jobsubmit
RMCFG [nodus] JOBSTARTURL=exec://$HOME/contrib/nodus/jobstart
RMCFG [nodus] JOBCANCELURL=exec://$HOME/contrib/nodus/jobcancel
RMCFG [nodus] JOBMODIFYURL=exec://$HOME/contrib/nodus/jobmodify
RMCFG [nodus] JOBQUEUEURL=exec://$HOME/contrib/nodus/jobqueue
RMCFG [nodus] JOBSUSPENDURL=exec://$HOME/contrib/nodus/jobsuspend
RMCFG [nodus] JOBRESUMEURL=exec://$HOME/contrib/nodus/jobresume
RMCFG [nodus] NODEPOWERURL=exec://$HOME/contrib/nodus/nodepower
RMCFG [nodus] NODEMODIFYURL=exec://$HOME/contrib/nodus/nodemodif
RMCFG [nodus] SYSTEMMODIFYURL=exec://$HOME/contrib/nodus/systemmodify
RMCFG [nodus] RESOURCECREATEURL=exec://$HOME/contrib/nodus/resourcecreate
RMCFG [nodus] JOBMIGRATEURL=exec://$HOME/contrib/nodus/jobmigrate
```

JOBNODEMATCHPOLICY EXACTNODE

3.4 Configuring Moab to Connect to ODDC

Perform the following operations as the root user.

1. Run the command `cd $MOABHOMEDIR/contrib/nodus/`.
2. Create an environment file by running the command `cp template.env .env`.
3. Edit the `.env` file:
 - A. The ODDC Host (for example: `HOST=http://xxxx`)
 - B. The ODDC User on Host (for example: `USERNAME=xxxx`)
 - C. The registered ODDC Cluster Name (for example: `CLUSTER_NAME=xxxx`)
 - D. The ODDC Web API Port (for example: `API_PORT=12345`)
 - E. The ODDC-monitor Daemon Port (for example: `MONITOR_PORT=12346`)
 - F. The Monitor Poll Frequency [in ms] (for example: `(MONITOR_POLL_MS=30000)`)
 - G. Restart the ODDC monitor daemon (`systemctl restart nodus-monitor`) to read the new `.env` details.

3.5 ODDC Use Cases

Select one of the following two options to manage backlog jobs:

- [Backlog Management Option 1](#)
- [Backlog Management Option 2](#)

Backlog Management Option 1

Use all available on-premise Moab nodes, then use available ODDC nodes and when on-premise Moab nodes becomes available after jobs complete (while available ODDC nodes are still in use), start using on-premise Moab nodes again. To achieve the above use case, follow the two steps below on ODDC.

1. Spin up a persistent cluster - Spin up an ODDC cluster with the desired hardware, and Moab will schedule jobs as if they are on-premise. Viewpoint will see the jobs and return `stdout` and `stderr` (job output). Steps to spin up an ODDC cluster in ODDC are explained in the *On-Demand Data Center User Guide*.
2. Use your spun up cluster in the step above (backlog cluster) - Activate bursting service on the spun up ODDC cluster to enable the cluster to be used to take only backlog jobs with full bursting functionality. Steps to activate bursting service on an ODDC cluster in ODDC are explained in the *On-Demand Data Center User Guide*.

There are two suggested ODDC bursting service configuration options:

- Min Burst - Spins up the minimum number of compute nodes required to complete all jobs in the queue, which is ideal for budgeting and controlling cloud costs. Using this configuration, if you have two nodes allocated to a cluster, as an example, when a backlog occurs, it only uses one node.
- Max Burst - Spins up enough compute nodes to complete all the jobs in the queue immediately; this gets results as fast as possible. Max Burst is limited by the size of the cluster and will not create new nodes. When using this configuration, backlog jobs run against all allocated nodes.

Backlog Management Option 2

Use all available on-premise Moab nodes, then use available ODDC nodes and when on-premise Moab nodes becomes available after jobs complete, continue using only available ODDC nodes and not available on-premise Moab nodes. To achieve the above use case, follow the configuration step below on Moab.

1. Configure the `moab.cfg` file by adding the following lines:

```
NODEALLOCATIONPOLICY PRIORITY
NODECFG[DEFAULT] PRIORITYF='APROCS - LOAD - 1000*GMETRIC[nodus]'
```

Note: Each reported ODDC cluster node is tagged with `GMETRIC[nodus]=1`.

The above is a sample of one way to configure Moab to respect ODDC as a backlog only cluster. In this configuration, we are using priority to ensure the least priority to ODDC by nature of not having a priority value less than the on-premises hardware.

3.6 Starting Moab ODDC Connect

1. Restart Moab using the command **systemctl restart moab** (must be run as root).
2. Restart the ODDC monitor daemon using the command **systemctl restart nodus-monitor** (must be run as root).

The following commands will assist in confirming the communication between Moab and ODDC is correct:

- **mdiag -n** provides a compute node summary, where both the on-prem (Moab) and ODDC nodes must display.
 - **mdiag -t -v** provides partition status, where it must show the system partition settings (including pbs and ODDC).
 - **/opt/moab/contrib/nodus/./clusterquery** provides GMETRIC information about each compute node defined on the cluster that is being monitored on ODDC from Moab.
2. Check that the nodes appear on Viewpoint, which should show both the on-prem (Moab) as well as the ODDC nodes.

Glossary

Bursting: The event of clusters and nodes being deployed to run jobs, then being destroyed.

canceljob: Command that selectively cancels the specified jobs (active, idle, or nonqueued) from the queue.

checknode: Command that shows detailed state information and statistics for nodes that run jobs.

Cluster: A collection of compute instances consisting of a head node and compute nodes.

Cluster Size: The number of compute nodes.

Compute Nodes: The servers, typically designed for fast computations and large amounts of I/O, that provide the storage, networking, memory, and processing resources.

Compute Node Size: An instance type or hardware configuration (for example, n1-standard-2 - vCPU: 2, Mem (GB): 7.50).

Core: An individual hardware-based execution unit within a processor that can independently execute a software execution thread and maintain its execution state separate from the execution state of all other cores within the processor.

Credentials: Authentication information required to access the respective cloud service provider from code.

Custom Job: A job that is customizable and configurable.

Elastic Computing: A feature in Moab that enables the Moab scheduler to take advantage of systems that can temporarily provide additional nodes (for example, to create new virtual machines or borrow physical nodes from another system) to fulfill the workload demand in a timely manner. The Elastic Computing framework can be configured to access multiple cloud providers either on-demand or based on a job backlog.

Elastic Trigger: Enables the Moab scheduler to take advantage of systems that can temporarily provide additional nodes to fulfill the backlog in a reasonable time frame. The elastic trigger calls a script to request additional nodes (dynamic nodes) from an external service.

FQDN (Fully qualified domain name): The most complete domain name that identifies a host or server.

Head Node: The server that manages the delegation of jobs.

Image: A snapshot of an OS.

Job: A workload submitted to a scheduler for the purpose of scheduling resources on which the workload executes when started up by the scheduler. Typically, a user creates a script that executes the workload (one or more applications) and submits the script to the scheduler where it becomes a job.

Job Script: A program to be run on a cluster (generally a shell script).

mdiag: Command that displays information about various aspects of the cluster and the results of internal diagnostic tests.

mpi-benchmarks: A job that tests the performance of a cluster system, including node performance, network latency, and throughput.

mschedctl: Command that controls various aspects of scheduling behavior, including managing scheduling activity, shutting down the scheduler, and creating resource trace files. It can also evaluate, modify, and create parameters, triggers, and messages.

On-Demand Bursting: Unlike backlog bursting, on-demand bursting is not triggered by the amount of work in the job backlog. You can simply request that the job run in the cloud service provider by specifying the preconfigured job template in the Moab config file that contains the trigger defined to call out to ODDC.

On-Demand Cluster: A cluster that carries out a specific job then is removed.

PAM (Pluggable Authentication Module): Used to perform various types of tasks involving authentication, authorization, and some modification.

pbsnodes: Command that marks nodes down, free, or offline. It also lists nodes and their state.

qdel: Command that deletes jobs in the order in which the job identifiers are presented to the command.

qmgr (Queue Manager): Command that provides an administrator interface to query and configure batch system parameters

Reprise License Manager Server (RLM): A flexible and easy-to-use license manager with the power to serve enterprise users.

Resource Manager: Provides the low-level functionality to start, hold, cancel, and monitor jobs. Without these capabilities, a scheduler alone cannot control jobs.

showq: Command that displays information about active, eligible, blocked, and/or recently completed jobs.

showstats: Command that shows various accounting and resource usage statistics for the system.

Stack: An instance of software packages that defines the operating system components.

test-job: A job that is best used to test bursting functionality; it echoes the time and the hostname of the executing machine.

Viewpoint: A web application that interacts with Moab Workload Manager, that enables users to manage jobs and resources without the complexities of maintaining Moab via the command line. Viewpoint uses a customizable portal that enables users to view and configure jobs and to compute node resources, principals, and roles. Viewpoint permissions enable system administrators to specify which pages, tools, and settings that certain users or groups are permitted to use, manage, and view.

Index

- . -

.env 10

- A -

Adaptive Computing

 Sales 3

 Support 3

Availability Zone 7

- B -

Backlog Management Options 11

Bursting Architecture 6

Bursting Cluster 9

Bursting Overview 5

- C -

canceljob (command) 14

checknode (command) 14

cluster 7, 9-10, 12

Configuration 4, 8-10

cp template.env .env (command) 11

- D -

Data Transmissions 7

Deploying the Bursting Cluster 9

Destroy Compute Nodes (on-demand
bursting type) 7

Destroy Full Cluster (on-demand
bursting type) 7

- E -

elastic 14

Elastic Computing 6, 14

elastic trigger 14

Environment file 11

- F -

FQDN 14

- H -

hardware 12, 14

Head Node 14

hostname 16

- I -

Installing Moab 8

instance type 14

- L -

Legal Notices 3

License 15

- M -

Max Burst 7, 12

mdiag -n (command) 13

mdiag -t -v (command) 13

mdiag (command) 13, 15

Min Burst 6, 12

Moab HPC Suite Installation and
Configuration Guide for Red Hat 6-
Based Systems 4

Moab HPC Suite Installation and
Configuration Guide for Red Hat 7-
Based Systems 4

Moab scheduler 14

Moab Viewpoint Reference Guide 8

Moab Workload Manager 4, 16

Moab Workload Manager
Administrator Guide 4
moab.cfg 10, 12
monitor daemon 10-11, 13

- N -

Native Resource Manager 5

- O -

Offline Compute Nodes (on-demand
bursting type) 7
On-Demand Bursting 7, 15

- P -

pbsnodes (command) 15
Persistent bursting 7

- Q -

qdel (command) 15
qmgr (command) 15

- R -

Red Hat 7 4, 8
Region 7
Related Documentation 4
Remote Resource Manager 5, 10
Resource Manager 4-5, 10, 15
Revision History 4
root 9, 11, 13

- S -

scheduler 14
Securing Data Transmissions 7
shell 15
showq (command) 15
stderr 12

stdout 12
Suite Installation (Automated
Installer) 4, 8
SUSE 4, 8
systemctl restart moab (command) 13

- T -

Torque 4
Torque Resource Manager 4
Torque Resource Manager
Administrator Guide 4

- U -

Use Cases 11

- V -

Viewpoint 8, 12-13, 16
Viewpoint Reference Guide 8
Virtual Private Cloud 7
VPC 7
VPN 7