ODDC Connect User Guide 6.0.0

September 2021





Contents

Welcome	
Legal Notices	
Revision History	
Related Documentation	
Chapter 1: ODDC Connect Cloud Bursting	
1.1 ODDC Connect Bursting Overview	
1.2 ODDC Connect Bursting Architecture	
1.3 Securing Data Transmissions	
Chapter 2: Installing Moab and ODDC	
2.1 Installing Moab	
2.2 Installing ODDC	
Chapter 3: ODDC Connect Configuration	
3.1 Deploying the Bursting Cluster	
3.2 Installing ODDC Connect	
Glossary	14
Index	

Welcome

Adaptive Computing's Moab HPC Workload and Resource Orchestration Platform, already a world leader in dynamically optimizing large-scale HPC computing environments, has been enhanced to automatically extend on-premise HPC workloads into public clouds. Moab is now fully integrated with HPC Cloud ODDC, which works with Moab to deliver cloud services, enabling true scalability and elasticity.

ODDC Connect can provision a full stack across multiple cloud environments, and can be customized to satisfy multiple use cases and scenarios. Access to unlimited HPC compute resources is available in the cloud from multiple cloud providers. For less than the cost of one employee, peak workload demands can be met rapidly and accurately.

This *ODDC Connect User Guide* will show you how to install and use the ODDC Connect environment. This guide is intended for 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[™], On-Demand Data Center[™], and other Adaptive Computing products are either registered trademarks or trademarks of Adaptive Computing Enterprises, Inc. The

Adaptive 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	
June 2021	ODDC Connect 5.0.0	
September 2021	ODDC Connect 6.0.0	

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 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

Related Documentation

Chapter 1: ODDC Connect Cloud Bursting

This chapter provides information about ODDC Connect cloud bursting.

In tl	his chapter:	
1.1	ODDC Connect Bursting Overview	6
1.2	ODDC Connect Bursting Architecture	. 7
1.3	Securing Data Transmissions	. 8

1.1 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 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 ODDC Connect Bursting Architecture

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 ODDC Connect is triggered based on the size of the job backlog.



This diagram depicts the 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).



See the diagram below for details.

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.

Chapter 1: ODDC Connect Cloud Bursting

• 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 and ODDC

To interface Moab with ODDC, the installation of Moab 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 ODDC

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

Adaptive Computing Support Portal / Product Documentation

Chapter 3: ODDC Connect Configuration

Now that you have Moab and ODDC installed and accessible, complete the steps in this chapter to configure your 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	11
3.2	Installing ODDC Connect	12

3.1 Deploying the Bursting Cluster

- **1.** Deploy your cluster (bursting or persistent) from the ODDC GUI. Refer to the *On-Demand Data Center User Guide* for instructions. The deployed cluster info will be needed in the next section for the env file.
- **2.** After successful deployment of your cluster, complete the steps below. The steps are done on ODDC or Moab as indicated.

3.2 Installing ODDC Connect

If this is an upgrade, start at step 1. If this is a new install, start at step 2.

- 1. Run the command yum upgrade .*.rpm.
- Download the RPM file (link supplied by Adaptive Computing), transfer the files to the Moab server, and then install the rpm with the command rpm -i moab-oddc-connect-<version>.rpm.

Most of the user files are found in /opt/moab/contrib/oddc.

3. Copy env.template to .env in the /opt/moab/contrib/oddc directory, and make modifications. The file looks like this:

URL of the NODUS Server URL=https://<oddc_server_ip> # Log level for python logger # Options: DEBUG, INFO, WARNING, ERROR LOG_LEVEL=INFO # Linux User Who is a Registered NODUS User on HOST USERNAME=<oddc User> # Registered NODUS Cluster Name CLUSTER_NAME=<cluster_name_deployed_on_oddc> # NODUS Web API Port API_PORT=12345 # oddc-monitor Daemon Port MONITOR_PORT=12346 # Monitor Poll Frequency (s) MONITOR_POLL_SEC=30

Where:

<oddc_server_ip> is the URL to the ODDC Server. <oddc User> is the user name where the clusters reside. <cluster_name_deployed_on_oddc> is a comma separated list of cluster names to run jobs on.

- 4. Copy the oddc.cfg file to /opt/moab/etc as it needs to be included in Moab. Make modifications to the file as needed.
- 5. Edit the file /opt/moab/etc/moab.cfg by adding #include oddc.cfg to the end of the file.
- **6.** Restart the Connect monitor daemons with the command **systemctl restart oddc-monitor oddc-job-monitor**.
- **7.** Restart Moab with the command **mschedctl** -**R**.

Connect logs are written to /opt/moab/contrib/oddc/log/oddc-connect.log. LOG_LEVEL in .env controls how much is written to the files.

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/oddc/./clusterquery** provides GMETRIC information about each compute node defined on the cluster that is being monitored on ODDC from Moab.

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

5 x 5

.env 12

- A -

Adaptive Computing Sales 3 Support 3 Availability Zone 8

- B -

Bursting Architecture 7 Bursting Cluster 11 Bursting Overview 6

- C -

canceljob (command) 14 checknode (command) 14 cluster 7, 11-12 CLUSTER_NAME= 12 Configuration 4, 10-11

- D -

Data Transmissions 8
Deploying the Bursting Cluster 11
Destroy Compute Nodes (on-demand bursting type) 7
Destroy Full Cluster (on-demand bursting type) 7

- E -

elastic 14 Elastic Computing 7, 14 elastic trigger 14

- F -

FQDN 14

- G -

GMETRIC 13

- H -

hardware 14 Head Node 14 hostname 16

- I -

Installing Moab 10 Installing ODDC 10 instance type 14

- L -

Legal Notices 3 License 15 Linux 12 LOG_LEVEL 13 logs 13

- M -

Max Burst 7 mdiag -n (command) 13 mdiag -t -v (command) 13 mdiag (command) 13, 15 Min Burst 7 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 server 12 Moab Workload Manager 4, 16 Moab Workload Manager Administrator Guide 4 moab.cfg 13 mschedctl -R (command) 13 mschedctl -R (connect) 13

- N -

Native Resource Manager 6 NODUS Cluster 12

- 0 -

ODDC 11
oddc.cfg 13
Offline Compute Nodes (on-demand bursting type) 7
On-Demand Bursting 7, 15
On-Demand Data Center User Guide 10

- P -

pbsnodes (command) 15Persistent bursting 7

- Q -

qdel (command) 15 qmgr (command) 15

- R -

Red Hat 7 4, 10 Region 8 Related Documentation 4 Remote Resource Manager 6 Resource Manager 4, 6, 15 Revision History 4 root 11 rpm -i moab-oddc-connect-<version>.rpm (command) 12

- S -

scheduler 14
Securing Data Transmissions 8
shell 15
showq (command) 15
Suite Installation (Automated
Installer) 4, 10
SUSE 4, 10
systemctl restart oddc-monitor oddcjob-monitor (command) 13

- T -

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

- U -

USERNAME= 12

- V -

Viewpoint 16 Virtual Private Cloud 8 VPC 8 VPN 8

- Y -

yum upgrade .*.rpm (command) 12