

HPC Cloud On-Demand Data Center

User Guide 7.7.1

June 2024



HPC Cloud On-Demand Data Center™

Contents

Welcome	5
Legal Notices	5
Revision History	6
Related Documentation	6
Chapter 1: On-Demand Data Center Platform Installation	8
1.1 ODDC Server Prerequisites	9
1.1.1 On-Demand Data Center Overview	9
1.1.2 ODDC Server Setup	9
1.1.3 Ingress Traffic	9
1.1.4 Cloud Provider Credentials	10
1.2 Installation Steps	16
1.3 Verifying the Installation	17
1.4 Adding the Server License Key	17
1.5 Modifying the Code	17
1.6 Supported MIG Options	18
Chapter 2: Using the On-Demand Data Center Web UI	19
2.1 Creating an ODDC Account	20
2.2 User Profile	20
2.2.1 Settings - Admins	20
2.2.2 Settings - Users	21
2.3 Dashboard	22
2.3.1 Job Graph	23
2.3.2 ODDC Server	23
2.3.3 Services	24
2.3.4 Terminal	24
2.4 Credentials Manager	25
2.5 Stack Manager	26
2.5.1 Viewing Stacks	26
2.5.2 Creating a Stack	27
2.5.3 Building a Stack	28
2.6 Building Stacks into Images	28
2.6.1 Customizing your Stack with a Provisioner Script	28
2.6.2 Building Customized Stacks	29
2.7 Cluster Manager	30
2.7.1 Viewing Clusters	30

2.7.2 Creating a Cluster	32
2.7.3 Deploying a Cluster	33
2.7.4 Deleting a Cluster	33
2.8 Job Manager	34
2.8.1 Adding a New Job	34
2.9 File Manager	37
2.10 Accounting Manager	38
2.11 Instance Prices	39
2.12 Cloud Budget Analysis	40
Chapter 3: On-Demand Data Center CLI Procedures	41
3.1 Users	41
3.2 Architecture	42
3.3 CLI Usage	43
3.3.1 Jobs	43
3.3.2 Building Images	43
3.4 Cluster Actions	43
3.4.1 Connecting to a Cluster	43
3.4.2 Destroying a Cluster	43
3.4.3 Viewing All Clusters	43
3.5 Job Actions	44
3.6 Common CLI Commands	44
3.6.1 Cluster Management	44
3.6.2 Cluster Monitoring	46
3.6.3 Cluster Jobs	46
3.6.4 Admin Commands	47
3.7 GPU Enabled Clusters	47
3.7.1 GPU Enabled Cluster Restrictions	47
3.7.2 GPU Stack Software	47
3.7.3 Building a GPU Stack	48
3.7.4 Updating a Cluster to Reference the GPU Prefix	48
3.7.5 Creating a GPU Job	48
3.7.6 Submitting a GPU Job	48
Chapter 4: On-Demand Data Center CLI Commands	49
4.1 Actions	49
4.2 Users	51
4.3 Credentials	53
4.4 Stacks	57
4.5 Clusters	76
4.6 Jobs	110
4.7 Providers	128

4.8 Misc	129
Appendix A: Creating Accounts for Cloud Service Providers	133
Appendix B: Obtaining Cloud Service Provider Account Information and Credentials ..	135
Appendix C: Job Preparation	138
Appendix D: Configuring the ODDC to be a Secure Server	140
Appendix E: Credential JSON Examples	141
Appendix F: Writing ODDC Scripts	152
Appendix G: Performing a Clean Installation of the ODDC	154
Appendix H: HPC Cloud On-Demand Data Center Workload Considerations	155
Appendix I: Data Staging to the Cloud	162
Appendix J: Using Client-Server Directory Service Protocol Setups	163
Appendix K: Modifying Provider Information To Support Customer Requirements	164
Appendix L: Creating a Private Oracle Cluster	168
Glossary	169
Index	172

Welcome

Welcome to the *HPC Cloud On-Demand Data Center User Guide* for version 7.7.1

Adaptive Computing's HPC Cloud On-Demand Data Center platform gives companies the ability to spin up temporary or persistent data center infrastructure resources quickly, inexpensively, and on-demand. This intelligent Cloud management platform gives immediate access to all computational resources, whether on-premises or in the Cloud on any leading Cloud provider.

Teams can automatically deploy and build clusters in the Cloud, automatically run applications on those clusters, and then terminate the Cloud resources on a daily, weekly, or even hourly basis.

The HPC Cloud On-Demand Data Center (ODDC) includes all the necessary tools to provision compute power and run workloads in the Cloud or on-premises. Access to all major Cloud providers is preconfigured and built into the interface (GUI or CLI).

This guide will show you how to use the ODDC user interface and command line interface.

Legal Notices

Copyright © 2018, 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, Moab®, Moab HPC Suite, Moab Viewpoint, Moab Grid, 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
February 2018	Moab/NODUS Cloud Bursting 1.1
November 2018	NODUS 3.0
July 2019	NODUS Cloud OS 3.2
December 2019	NODUS Cloud OS 4.0
June 2020	NODUS Cloud OS 5.0
September 2020	NODUS Cloud OS 5.1
January 2021	NODUS Cloud OS 5.2
August 2021	On-Demand Data Center 6.0
December 2021	HPC Cloud On-Demand Data Center 6.1
June 2024	HPC Cloud On-Demand Data Center 7.7.1

Related Documentation

- *Cloud Business Value Process Spreadsheet*
<http://support.adaptivecomputing.com/Cloud-Business-Value-Process-Spreadsheet.xlsx>
- *HPC Cloud On-Demand Data Center Quick Start Guide*
[HPC Cloud On-Demand Data Center Quick Start Guide \(Adaptive Computing Support Portal / Product Documentation / HPC Cloud On-Demand Data Center\)](#)

Chapter 1: On-Demand Data Center Platform Installation


This chapter provides information about the installer, which contains the entire ODDC platform, along with dependencies to install it.

In this chapter:

1.1	ODDC Server Prerequisites	9
1.1.1	On-Demand Data Center Overview	9
1.1.2	ODDC Server Setup	9
1.1.3	Ingress Traffic	9
1.1.4	Cloud Provider Credentials	10
1.2	Installation Steps	16
1.3	Verifying the Installation	17
1.4	Adding the Server License Key	17
1.5	Modifying the Code	17
1.6	Supported MIG Options	18

1.1 ODDC Server Prerequisites

1.1.1 On-Demand Data Center Overview



1.1.2 ODDC Server Setup

You must provision to a clean server with one of the following supported operating systems:

- CentOS 8 Stream
- RHEL 8
- Rocky 8

The recommended minimum specification is as follows:

- 8 Cores
- 16 GB RAM
- 100 GB disk space

1.1.3 Ingress Traffic

To enable access to the ODDC Server, the following ports are allowed for ingress:

- 22 (SSH) - user direct access / Moab job routing
- 80 (HTTP) - connection to the ODDC UI
- 443 (HTTPS) - optional, if you secure your HTTP connection
- 2200 (SSH over HTTPS) - connection to a terminal from the Web UI
- 12345 (ODDC API) - the ODDC platform API server

Note: The ODDC Server installation will enable these ports on the server.

1.1.4 Cloud Provider Credentials

This section contains information about obtaining credentials from the Cloud service providers. These credentials can be registered within ODDC in order to provision resources on a given Cloud service provider.

These are the supported Cloud service providers:

-  Amazon Web Services (AWS)
-  Google Cloud (GCP)
-  HPC CloudMart
-  Microsoft Azure
-  Open Telekom Cloud (OTC)
-  Oracle Cloud
-  Oracle Cloud HPC

Note: For unsupported Cloud service providers, please contact us at oddcsupport@adaptivecomputing.com to inquire about adding the provider into the ODDC.

Important! The following sections are intended to be used for reference only -- refer to the providers' support pages for detailed information. If you cannot get provider credentials / correct admin level credentials, please contact us at oddcsupport@adaptivecomputing.com.

Obtaining Credentials

Continue with the steps for your chosen provider below to obtain credentials:

- [AWS](#)
- [Google Cloud](#)
- [HPC CloudMart](#)
- [Microsoft Azure](#)
- [Open Telekom Cloud](#)
- [Oracle Cloud](#)
- [Oracle Cloud HPC](#)

AWS

AWS has two forms of credentials; both require admin privileges to set up:

- Programmatic access (access key, secret key) is a floating set of coupled values.
- EC2 roles only exist within AWS and are attached to AWS instances.

These accounts must have valid permissions associated in order to use the ODDC. The ODDC can deploy clusters using either programmatic access or EC2 roles. However, stack builds can only be done with programmatic access.

The ODDC requires programmatic access to access the provider's resources. For example, if the ODDC is outside of the AWS environment, access keys are required. If the ODDC is running inside AWS, the best practice is to use IAM roles instead. An IAM role is a defined set of permissions. It is not associated with a specific user or group and any trusted entity can assume the role to perform a specific business task. The ODDC supports both access methods.

When creating a cluster via **nodus cluster:aws:create**, the flag `--role` can be used to set the cluster to use EC2 role instead of programmatic access. If the cluster entry already exists, then **nodus cluster:aws:update** can be used with flag `--role` or `--no-role` to switch on EC2 role use or turn it off, respectfully.

1. Go to <https://aws.amazon.com> and open the **IAM** dashboard.
2. In the navigation pane, select **Users**.
3. Select the name of the user whose access keys you want to create, and then select the **Security Credentials** tab.
4. Add user permissions.
5. In the **Access keys** section, select **Access keys (access key ID and secret key) / Create New Access Key**, then download the file with the access key and secret key.
Note: If you forget the secret key, you will have to delete it and create a new one.

6. To view the new access key pair, select **Show**. You will not have access to the secret access key again after this dialog box closes.
7. To download the key pair, select **Download .csv file**. Store the keys in a secure location. You will not have access to the secret access key again after this dialog box closes.
8. After you download the .csv file, click **Close**. When you create an access key, the key pair is active by default, and you can use the pair right away.

See <https://docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html> for additional information.

Google Cloud

1. Go to <https://cloud.google.com>.
2. Select **IAM & Admin / Service Accounts**.
3. From the **Service account** list, select **New service account**.
4. In the **Service account name** field, enter a name, and then click **Create**.
5. From the **Role** list, select **Project / Owner**.
6. Click **Add Key / Create new key**.
7. Click **Create**. A JSON file that contains your key downloads to your computer.

See <https://cloud.google.com/docs/authentication/production> for additional information.

HPC CloudMart

Contact Adaptive Computing at oddcsupport@adaptivecomputing.com to obtain credentials.

Microsoft Azure

1. Go to <https://azure.microsoft.com/>.
2. Run the script **account-scripts_azure.sh**.

See <https://docs.microsoft.com/en-us/azure/automation/shared-resources/credentials> for additional information.

Open Telekom Cloud

1. Go to <https://open-telekom-cloud.com/en>.
2. Create an account and for the **Access Type**, select **Programmatic Access**.
3. On the Open Telekom Cloud (OTC) account to be used, create a VPC.

4. Within that VPC, create a subnet.
5. Create a Connection, create a Virtual Gateway, and then create a Virtual Interface.
6. Under `Access Control`, on the default `Security Group`, select **Inbound Rules** and then **Allow Common Ports**, and also select **Outbound Rules** and then **Allow Common Ports**.
7. Create a subnet:
 - A. On the management console, click **Service List**, and choose **Network > Virtual Private Cloud** to launch the VPC console.
 - B. In the navigation pane, choose **Subnets**. On the displayed page, click **Create Subnets**.
 - C. After the subnet is created, click its name to obtain the network ID, which will be required to create a stack.

See <https://open-telekom-cloud.com/en/support/first-steps> for additional information.

Oracle Cloud

1. From a terminal, install OpenSSL.
 - A. Run the command `openssl genrsa -out private-key.pem <xxxx>`. Where `xxxx` is the number of bits in the generated key, for example 2048.
 - B. Run the command `openssl rsa -in private-key.pem -pubout -out public-key.pem`.
2. Go to <https://www.oracle.com/cloud/>.
 - A. On the Oracle account to be used, click **Profile / User Settings**.
 - B. On the API Key page, under **Resources**, select **API Keys**, and then click **Add API Key**.
 - C. Select the option **Choose Public Key File**, select the `public-key.pem` generated previously, and then click **Add**.
 - D. Copy the contents of `private-key.pem` into a text editor and format as in the appendix [Credential JSON Examples](#), and save as a JSON file.
 - E. Under the **Capabilities** section, click **View Configuration file**, and then copy the contents into the JSON file created previously and format as in the appendix example. **Note:** Do not include the `=` signs or the text to the left of the `=` signs.
3. From the Oracle Cloud menu, select **Identity & Security**, and then under **Identity**, click **Compartments**.

- A. Create a compartment and then within that compartment, create a VPC (virtual private cloud) network.
4. From the Oracle Cloud menu, select **Networking** and then click **Virtual Cloud Networks**.

Note: All procedures in this step should be done in the `root` compartment.

- A. Create a VCN in the root compartment (under **List Scope / Compartment**).
- B. Within that VCN, create at least one subnet (also referred to as the Subnet Network).
- C. Click the subnet and then in the **Compartment Information**, save the **OCID:** info because it will be needed to create stacks.
- D. Click the VCN, under **Resources**, select **Security Lists**, and then click **Create Security List**. Provide the required information and also add these rules:

Ingress Rules

Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows	Description
No	0.0.0.0/0	TCP	All	7946		TCP traffic for ports: 7946	
No	0.0.0.0/0	TCP	All	12001		TCP traffic for ports: 12001	
No	0.0.0.0/0	TCP	All	15001-15003		TCP traffic for ports: 15001-15003	
No	0.0.0.0/0	TCP	All	1-65535		TCP traffic for ports: 1-65535	
No	0.0.0.0/0	UDP	All	7946		UDP traffic for ports: 7946	
No	0.0.0.0/0	TCP	All	7373		TCP traffic for ports: 7373	
No	0.0.0.0/0	TCP	All	2200		TCP traffic for ports: 2200	

Egress Rules

Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows	Description
No	0.0.0.0/0	TCP	All	All		TCP traffic for ports: All	

- E. Add rules to the default security list, by selecting **Security Lists**, clicking **Default Security List for VCN**, and then adding these rules:

Ingress Rules

Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows	Description
No	0.0.0.0/0	TCP	All	22		TCP traffic for ports: 22 SSH Remote Login Protocol	
No	0.0.0.0/0	ICMP			3, 4	ICMP traffic for: 3, 4 Destination Unreachable: Fragmentation Needed and Don't Fragment was Set	
No	10.0.0.0/16	ICMP			3	ICMP traffic for: 3 Destination Unreachable	

Egress Rules

Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows	Description
No	0.0.0.0/0	All Protocols				All traffic for all ports	

- F. Create another security list by selecting **Security Lists**, clicking **Create Security List**. We recommend naming the security list something like `nodus-server`. Provide the required information, and also add these rules:

Ingress Rules

Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows	Description
No	0.0.0.0/0	TCP	All	7946		TCP traffic for ports: 7946	
No	0.0.0.0/0	TCP	All	12001		TCP traffic for ports: 12001	
No	0.0.0.0/0	TCP	All	15001-15003		TCP traffic for ports: 15001-15003	
No	0.0.0.0/0	TCP	All	1-65535		TCP traffic for ports: 1-65535	

Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows	Description
No	0.0.0.0/0	UDP	All	7946		UDP traffic for ports: 7946	
No	0.0.0.0/0	TCP	All	7373		TCP traffic for ports: 7373	

- G.** Under **Resources**, click **Internet Gateways**, click **Create Internet Gateway**. We recommend naming the internet gateway something like `nodus-cluster`. Provide the required information, and then click **Create Internet Gateway**.
- H.** Under **Resources**, click **NAT Gateways**, click **Create NAT Gateway**. We recommend naming the NAT gateway something like `nat_gateway`. Provide the required information, and then click **Create NAT Gateway**.
- I.** Under **Resources**, click **Route Tables**, click **Create Route Table**. Create two route tables. We recommend naming them something like `nodus-server` and `nodus-cluster`.
- a.** Click the `nodus-server` route table, click **Add Route Rules**, and add these rules:

Destination	Target Type	Target	Description
0.0.0.0/0	Internet Gateway	nodus-cluster	

- b.** Click the `nodus-cluster` route table, click **Add Route Rules**, and add these rules:

Destination	Target Type	Target	Description
0.0.0.0/0	NAT Gateway	nat_gateway	

- J.** Click **Virtual Cloud Networks** click the VCN, click the subnet, click **Add Security List**, and then add the `nodus-server` and default security lists created above.

See <https://docs.oracle.com/en-us/iaas/Content/Identity/Tasks/managingcompartments.htm> for additional information.

Oracle Cloud HPC

Oracle Cloud HPC uses the regular Oracle Cloud Credentials

1.2 Installation Steps

Important! If upgrading from a previous version of ODDC, you must destroy all clusters.

1. Download the install file from adaptivecomputing.com (Support Portal Home / Product Download / HPC Cloud On-Demand Data Center) and then copy the file to the installation user.

1.3 Verifying the Installation

1. Check the UI configuration files by editing `/NODUS/config.json`. This file tells the UI the address and port of the ODDC API, and which providers will be used. You can edit the providers to only show which ones you will supply licenses to. The file looks similar to this:

```
{
  "server": "http://<public_ip>:12345",
  "providers": ["aws", "gcp", "azure", "oracle", "otc"],
  "register": true,
  "idleTimer": false
}
```

2. Verify the UI in a browser by entering **`http://<ODDC_server_address>/`**. You should be prompted to enter a license or to log in. If not, rerun step 1.

1.4 Adding the Server License Key

If your license key has expired or is invalid, follow the procedure below to add a license key from a terminal.

1. SSH in to the `<user_name>@<oddc_server_address>`.
2. Add the license key that was provided to `/NODUS/.license`.
3. Verify that the license is working by running the command `oddc license`.

1.5 Modifying the Code

Modification of the code without the approval of Adaptive Computing is discouraged as this may result in an unstable environment.

1.6 Supported MIG Options

This section describes the supported Multi-Instance GPU (MIG) options.

Providers

Oracle

Instance Shapes

BM.GPU4.8

MIG Instance Configurations

all-disabled: No MIG instances are provisioned

all-1g.5gb: All GPUs on instance are configured to have 7 MIG instances

all-2g.10gb: All GPUs on instance are configured to have 3 MIG instances

all-3g.20gb: All GPUs on instance are configured to have 2 MIG instances

Resource Managers

Only Torque 7 is currently supported.

Stack Build

The stack must be built using Torque 7 w/ MIG support.

Chapter 2: Using the On-Demand Data Center Web UI

This chapter provides information about using the ODDC GUI to run your workloads in the Cloud.

In this chapter:

2.1	Creating an ODDC Account	20
2.2	User Profile	20
2.2.1	Settings - Admins	20
2.2.2	Settings - Users	21
2.3	Dashboard	22
2.3.1	Job Graph	23
2.3.2	ODDC Server	23
2.3.3	Services	24
2.3.4	Terminal	24
2.4	Credentials Manager	25
2.5	Stack Manager	26
2.5.1	Viewing Stacks	26
2.5.2	Creating a Stack	27
2.5.3	Building a Stack	28
2.6	Building Stacks into Images	28
2.6.1	Customizing your Stack with a Provisioner Script	28
2.6.2	Building Customized Stacks	29
2.7	Cluster Manager	30
2.7.1	Viewing Clusters	30
2.7.2	Creating a Cluster	32
2.7.3	Deploying a Cluster	33
2.7.4	Deleting a Cluster	33
2.8	Job Manager	34
2.8.1	Adding a New Job	34
2.9	File Manager	37
2.10	Accounting Manager	38
2.11	Instance Prices	39
2.12	Cloud Budget Analysis	40

2.1 Creating an ODDC Account

1. To create an account, go to `http://<ODDC_server_address>` and click **Create an Account**.
2. Follow the instructions and complete the registration process.

2.2 User Profile

From your user profile at the top right of the screen, the following are available:

- **Settings** - Add new users and/or edit profiles, depending on whether you are an admin or a user. Your user profile shows whether you are an admin or a user.
- **User SSH Key** - Download a pem file to access the head node of a cluster via SSH for that given user.
- **Logout**
- **ODDC Version Number**

2.2.1 Settings - Admins


Admins can maintain user information/roles and create new users. The first user created has an administrator role and all others have a user role. Admins can be created by an existing admin user.

Non-admin users have limited ODDC features permissions:

- Cluster Manager - View clusters.
- Stack Manager - No permissions.
- Credentials Manager - No permissions.
- Job Manager - Add and submit jobs.
- File Manager - View, download, and delete jobs.


2.2.1.A Editing User Profiles

1. To change a user's name/email or change an admin to a user and vice versa, click **Settings**, and then on the **Users** tab, click **Edit**. You can also change a user's password, download the SSH Key, and delete a user.



2.2.1.B Adding a New User

1. To add a new user, on the **New User** tab, complete the fields and then click **Register**.



2.2.2 Settings - Users

Users can edit their name/email and change their password.

user

ID: 60c20533209ac7401b48115eaf

First: jill

Last: Bailey

Email: jbailey@adaptimecomputing.com

Role: User

Verified: Yes

Created: Jun 10, 2023 8:27 AM

Modified: Jun 10, 2023 8:27 AM

UPDATE USER:

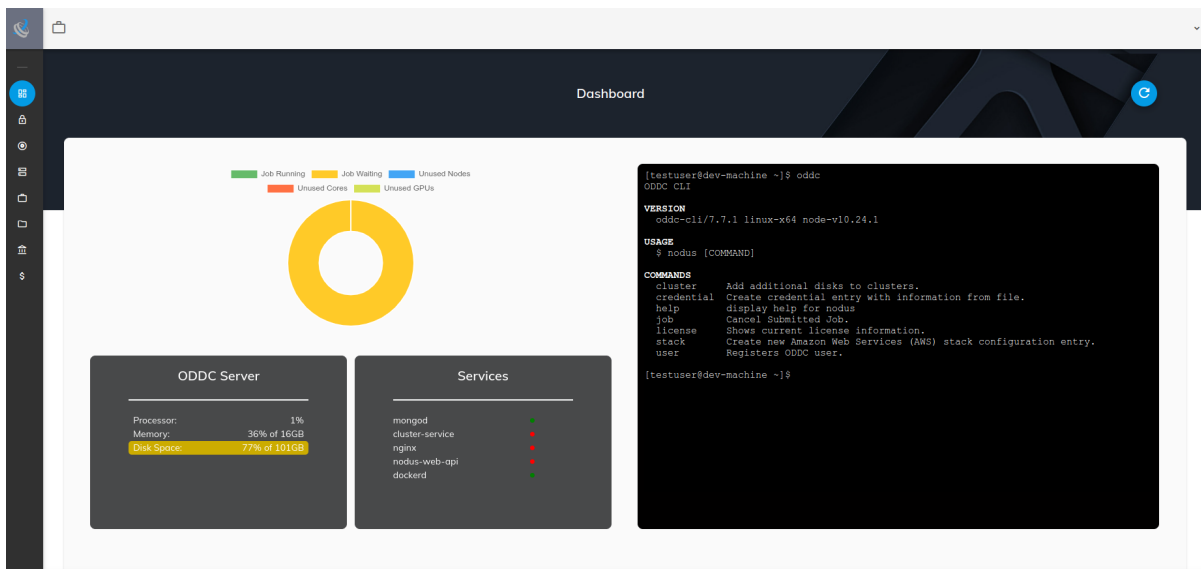
[EDIT](#)

PASSWORD:

[CHANGE](#)

2.3 Dashboard


1. On the **Applications** menu, click **Dashboard**:



This screen shows the Job Graph, ODDC Server, Services, and Terminal.

2.3.1 Job Graph


After you create a cluster and run jobs, the infograph shows this information: Job Running, Job Waiting, Unused Nodes, Unused Cores, and Unused GPUs.



Hover the pointer over the infograph to see the corresponding indicator number.


2.3.2 ODDC Server

This shows the ODDC server resources:



2.3.3 Services

This shows the status of the services that ODDC needs to run:



Green indicates that the services are good. Red indicates that the services are not running.

2.3.4 Terminal

This is an integrated terminal to use the CLI and to also execute Linux server commands:


```

ODDC CLI

VERSION
  oddc-cli/x.x.x linux-x64 node-v14.21.3

USAGE
  $ oddc [COMMAND]

COMMANDS
  action      List action entries from database.
  cluster     Add additional disks to clusters.
  credential  Create credential entry with information from file.
  daemon      Manage Master Bursting Daemon.
  db          Backup ODDC database to file.
  help        Display help for oddc.
  job         Create new Amazon Web Services (AWS) on-demand job entry.
  license     Shows current license information.
  provider    List provider instances.
  stack       Create new Alibaba Cloud stack configuration entry.
  upgrade     Prepare or finish upgrade process.
  user        Registers Linux user as ODDC user.

```

2.4 Credentials Manager

Credentials are required by the Cloud service providers in order to access their environment and resources. Adding and saving them now means that you can easily select them with just one click when creating a cluster.

1. To add credentials, log in to your ODDC user account at `http://<oddc_server_address>`.
2. On the **Applications** menu, click **Credentials Manager**:

In Use	Name	Provider	Date/Time Created	Actions
✓	awscredential		June 29, 2023 8:59 AM	...
✓	googlecredential		June 29, 2023 8:59 AM	...
✓	azurecredential		June 29, 2023 8:59 AM	...
✓	oraclecredential		June 29, 2023 8:59 AM	...
✓	atccredential		June 29, 2023 8:59 AM	...

The **Credentials Manager** screen shows this information: In Use, Name, Provider, Date/Time Created, and Actions (...). Edit Collaborators, Set, & Delete.

Note: From the **Applications** menu, you can also create clusters, build stacks, submit jobs, view a job's output file, view an accounting of all the jobs run, and estimate the cost of running a job. See the sections [Cluster Manager](#), [Stack Manager](#), [Job Manager](#), [File Manager](#), [Accounting Manager](#), and [Instance Prices](#) for information.

3. Click the add icon (+) and then select the desired Cloud service provider.
4. On the **New Credential** form, enter the desired name for the credentials.
5. Enter your account credentials.
6. Click **Save** to finish adding the credentials for this provider.

Note: Clicking the name of a credential shows information about it.


2.5 Stack Manager

This section contains the following information:

- [Viewing Stacks](#)
- [Creating a Stack](#)
- [Building a Stack](#)

2.5.1 Viewing Stacks

1. On the **Applications** menu, click **Stack Manager**:



The screenshot shows the 'Stack Manager' interface. At the top, there's a search bar and a '+ Add' button. Below, there are two tabs: 'My Stacks' (active) and 'Marketplace'. The 'My Stacks' tab displays a table with the following columns: Built, Name, Provider, Owner, Instance Size, Prefix, Type, Credential, Builder, Last Built, and Actions. A single row is visible with the following data: Built: None (with a red cloud icon), Name: aws, Provider: Amazon, Owner: qouser1, Instance Size: t2.nano, Prefix: core, Type: centos-7, Credential: awscredential, Builder: market, Last Built: N/A, and Actions: a three-dot menu icon. At the bottom right of the table, it says 'Rows per page: 10' and '1-1 of 1'.

Built	Name	Provider	Owner	Instance Size	Prefix	Type	Credential	Builder	Last Built	Actions
None	aws	Amazon	qouser1	t2.nano	core	centos-7	awscredential	market	N/A	...


On the **My Stacks** tab, you can create new stacks and see a list of all your current stacks and information such as: Built (🟢 Success / 🛑 Failed / 🛑 None), Name, Provider, Owner, Instance Size, Prefix, Type, Credential, Builder, Last Built, and Actions (Edit Collaborators, Edit, Build, Build GPUs, Build Server, Build Node, Logs, & Delete).

Notes:

- To edit collaborators' stacks, the License Type must be Site.
- Build GPUs creates a stack with a node with a GPU.
- Build Server creates a stack with just a server.
- Build Node creates a stack with just a node.

2.5.2 Creating a Stack

1. To create a new stack, on the **Stack Manager** screen, click the add icon (+). The **New Stack** form shows the Cloud service providers that you have credentials for:



2. Select the Cloud service provider that you want to create a stack in.
3. Click the **Advanced** button to show additional fields based on the provider.
4. Click **Save**. The `Built` state shows `None` until it is built.

2.5.3 Building a Stack

1. To build a stack, under **Actions**, click **Build**. Once the build process finishes, if it is successful, the state changes to `Success`; otherwise, the state changes to `Failed`.

2.6 Building Stacks into Images

This section contains the following information:

- [Customizing your Stack with a Provisioner Script](#)
- [Building Customized Stacks](#)

2.6.1 Customizing your Stack with a Provisioner Script

This step adds custom user provisioning to the ODDC stack building process. You can create your own custom scripts and run them as a user-defined provisioner during the server and node stack build. Supplemental dataFiles can be added to expand and simplify the provisioning scripts.

Example Usage

```
$ nodus stack:<provider>:create myStack centos-7 myCred -p
myPrefix --serverProvisioner ~/myServerScript --nodeProvisioner
~/myNodeScript --dataFile ~/myDataFile1 --dataFile ~/myDataFile2
```

- | | |
|------------------------|---|
| -D --dataFile | These files are copied and written to the provisioning root directory and can be called by the main provisioning scripts. Multiple files are supported. |
| -N --nodeProvisioner | User-defined provisioning script to be run on the head node of the cluster. This script is run from the present working directory (pwd) directory, so all dataFiles can be called from the \$(pwd). |
| -S --serverProvisioner | User-defined provisioning script to be run on the compute node of the cluster. The script is run from the present working directory (pwd) directory, so all dataFiles can be called from the \$(pwd). |

During the stack build, all provisioner files are located in the \$(pwd) directory on the server or node instances.

Note: If any user provisioning script fails, the stack build process will abort and the image artifacts will be deleted. It is a best practice to ensure the scripts run successfully before testing them on a stack build.

2.6.2 Building Customized Stacks

ODDC stack images need to be built for your provider and in the regions you want to deploy an ODDC cluster in.

If you just want to deploy another cluster in the same region using the same credentials, then stack rebuilds should not be necessary.

By default, ODDC images are built on top of the marketplace images.

Successful Build

A successful stack build looks like this from the logs:

```
==> Builds finished. The artifacts of successful builds are:
--> market-server: AMIs were created:
us-east-1: ami-00d17452c46306a94

--> market-server:
--> market-node: AMIs were created:
us-east-1: ami-082631c41bb6b0761

--> market-node:
```

```
=== BUILD COMPLETE :: Wed Jun 17 2023 09:19:34 GMT-0400 ===
```

2.7 Cluster Manager

This section contains the following information:

- [Viewing Clusters](#)
- [Creating a Cluster](#)
- [Deploying a Cluster](#)
- [Deleting a Cluster](#)

2.7.1 Viewing Clusters

1. On the **Applications** menu, click **Cluster Manager**:

State	Name	Provider	Owner	Bursting	Nodes	Credential	Uptime	Actions
Available	google		qauser1	Off	0	googlecredential	4 minutes	
Down	aws		qauser1	Off	0	awscredential	N/A	
Down	azure		qauser1	Off	0	azurecredential	N/A	
Down	otc		qauser1	Off	0	otccredential	N/A	
Down	oracle		qauser1	Off	0	oraclecredential	N/A	

On the **Cluster Manager** screen, you can configure new clusters and see a list of all your current clusters and information such as their State (Available / Down / Failed), Name, Provider, Bursting, Nodes, Credential, Uptime, and Actions (Edit / Edit Collaborators, Deploy, Destroy, Configuration, Logs, Plan, Save Config, Delete, and Cluster SSH Key). The Actions vary based on the State and License Type (see the notes below for additional information).


Notes:

- Clicking **User SSH Key** from the drop-down on your profile will download a `.pem` file that can be used to access the head node of a cluster as your user via SSH. This provides you full control over the cluster.
 - Clicking the refresh icon (🔄) at the top right of the cluster list, refreshes the list.
 - Clicking the add icon (+) at the top right opens the **New Cluster** form.
 - Clicking the name of an available cluster shows cluster information.
 - Clicking the information icon (i) at the end of the row shows additional details about the cluster, including: Cluster ID, SSH Username, Manager, Created, Modified, Expected Deploy Time, Actual Deploy Time, Uptime, Head Node Public IP, Prefix, Region, Head Node Size, and Compute Node Sizes. The buttons Deploy/**Destroy** and Logs also show.
 - Clicking the horizontal ellipsis (⋮) at the end of the cluster row shows actions that you can take regarding the cluster:
 - **Edit** - Modifies the configuration of the cluster. The cluster must be in a `Down` state to edit.
 - **Configuration** - Opens the **Configuration** screen to view/modify Cluster Information, Cluster Compute Nodes, Bursting Service, and Disks Information.
 - **Deploy** - Deploys the cluster into an available state.
 - **Destroy** - Destroys the head node and the linked compute nodes so that they are no longer in a state to do work. This also means that you will not be billed further for these Cloud resources. After a cluster is destroyed, the state shows as `Down`. Make sure you back up any files before proceeding because you will lose all files in File Manager once the cluster is destroyed.
 - **Logs** - Shows information about the provisioning. A screen shows information in real-time about the cluster's state. Additionally, these logs contain information for deploy and destroy.
 - **Delete** - Permanently deletes the cluster. The cluster must be down before it can be deleted. Removing a failed cluster before destroying it from ODDC removes all cluster data from the ODDC database, not the provider account. Doing this may leave behind abandoned resources that require manual removal (in the provider account). To avoid this, destroy the cluster first and then perform this action.
 - **Cluster SSH Key** - Downloads the cluster SSH key.
 - **Edit Collaborators** - Shares the cluster (the License Type must be `Site`).
-

2.7.2 Creating a Cluster

If you have already added your Cloud credentials for the chosen provider, follow the steps below to create a new cluster, which can then be deployed at any time. If you have not yet added your credentials, you must first follow the steps in the section [Credentials Manager](#).

1. To create a new cluster, on the **Cluster Manager** screen, click the add icon (+). The **New Cluster** screen shows the Cloud service providers that you have credentials for:



2. Select the Cloud service provider that you want to deploy a cluster in. You can also click the **Import** button to import a cluster configuration file that was previously saved, to easily create a new cluster.

 The screenshot shows the 'New Cluster' configuration form. At the top, there's a blue header with the text 'New Cluster', a price indicator '\$0.00 per Hour', and three buttons: 'ADVANCED', 'SAVE', and 'CLOSE'. The form contains several input fields and dropdown menus:

- Name ***: A text input field.
- Credential ***: A dropdown menu.
- Region ***: A dropdown menu.
- OS Type ***: A dropdown menu with 'centos-7' selected.
- Prefix ***: A text input field with 'core'.
- Head Node Size ***: A dropdown menu with 'VM.Standard2.1 - vCPU: 2, Mem (GB): ...' selected.
- Manager ***: A dropdown menu with 'Torque' selected.
- Availability Zone ***: A text input field.
- Bursting Configuration**: Radio buttons for 'Off' (selected), 'Min', 'Max', and 'All'.
- Compute Nodes**: A section with a blue plus icon and an information icon.
- Node Prefix Label ***: A text input field.
- Group Node Prefix ***: A text input field.
- Size ***: A dropdown menu with 'VM.Standard2.1 - vCPU: 2, Mem (GB): 15' selected.
- Count ***: A text input field with '0'.
- Description**: A large text area at the bottom.

3. Enter the desired information for the new cluster. Clicking the **Advanced** button shows additional fields.

Notes:

- The EFA field will not be visible if the AWS region, availability zone, and instance size do not allow EFA.
 - Multiple Compute Node Instances of different sizes can be added.
 - When entering the Count, if you exceed the specified number, a message similar to this displays: The quota of CPUs/Instances established by the provider is exceeded. This must first be configured in the `providers.json` file (`/NODUS/providers.json`) inside of `regions/zones/quotas/instance` according to the provider's quota.
 - The available Regions, Availability Zones, Server Instance Sizes, Compute Node Instance Sizes, and quantity of instances may vary based on your provider account status.
 - When creating a cluster with a greater number of nodes, we recommend using large instance types to ensure that the cluster deploys successfully.
-

4. Click **Save**. The state shows `Down` until it is deployed.

2.7.3 Deploying a Cluster

1. To deploy a cluster, click the ellipsis and then click **Deploy**. When successfully deployed, the state changes to `Available`.

While deploying, a % completion bar shows an estimate of the progress, and a timer shows in yellow. After deployment, the deploy time shows in green.

While destroying, a timer shows in yellow. After destroying, the destroy time shows in red.

The deploy and destroy timer works as a countdown in descending mode. For example (color yellow) 5:00, 4:59, 4:58... But also if a cluster has an estimated time of 5:00 minutes, it will run in descending mode until reaching 0 and after that if the deploy or destroy takes more than the estimated time, for this example 5:00, the timer will start counting in ascending mode with a negative sign in front (e.g., -1, -2, -3...)

2.7.4 Deleting a Cluster

To delete a cluster, the state must first be `Down`.


1. Click the ellipsis for the desired cluster and then click **Destroy**.
2. After the state changes to `Down`, click **Delete**.

2.8 Job Manager

You can add jobs and submit them to an existing cluster by using Job Manager.

2.8.1 Adding a New Job

1. To add a new job, on the **Applications** menu, click **Job Manager**:




The screenshot shows the Job Manager interface. At the top, there is a 'Job Manager' header with a search bar and a blue plus icon. Below this is a table with the following columns: Name, Nodes, Cores, Data Files, Walltime, Submit, and Actions. The table contains one row with the job name 'stdaut-stderr', 1 Node, 1 Core, Data Files as '--', Walltime as '01:00:00', a submit icon, and an actions menu (three dots). At the bottom right of the table, there is a pagination control showing 'Rows per page: 10' and '1-1 of 1'.

Name	Nodes	Cores	Data Files	Walltime	Submit	Actions
stdaut-stderr	1	1	--	01:00:00		...

This screen shows this information: Name, Nodes, Cores, Data Files, Walltime, Submit, and Actions (Edit, Copy, & Delete).

2. Click the add icon (+) at the top right to open the **New Job** screen.
 - A. On the **Job Info** panel, enter the desired information for the new job: Name, Walltime, Number of Nodes, Cores per Node, GPUs per Node, On Demand, and Description. This can be edited after the job is saved by clicking the name of the job.
 - B. On the **Script** panel, either edit the script for this job or click **Upload Script**. This can be edited after the job is saved by clicking the name of the job.
 - C. Click **Save**.
3. Click the submit (▶) icon, then select the cluster to run the job on. If there are no created clusters or ones that meet the hardware requirements of the job, then certain clusters may not appear.

4. Click **Submit**. The job is queued and runs on the selected cluster. The cluster details screen shows the submitted job information:



- **Cluster Info** - Jobs Submitted (Total, Today, This Month, This Year), Queue (Running, Queued), Nodes (Available, Busy, Down, Offline), Activities Log (All, Job), and Nodes/Workload graph.

The graph shows the status of Available Nodes, Busy Nodes, & Down Nodes, and the status of the Running Jobs, Queued Jobs, & Blocked Jobs. Click to deselect nodes and jobs to narrow the display results. Hover over the graph to see the number of nodes or jobs.

- **Queue** - ID, Name, State (Queued, Running, Completed), Submitted, Cores, Nodes, and Actions (⊗).
- **Note:** After a job completes, it only shows here temporarily, but can be viewed in [File Manager](#).
- **Nodes** - ☐ (Deploy, Destroy), Name, State (see descriptions below), Cores, Load Avg, Actions (Details, Deploy, Destroy), and MIG Mode.
 - All Nodes - For the selected cluster, all nodes are listed, no matter if they are available, busy, deploying, destroying, offline or down.
 - Busy - For the selected cluster, all nodes that are in use by workload/jobs are listed. If all available cores on the nodes are being used, you will see the nodes that are being fully utilized. If the workload/jobs are only using part of the available cores on a node, the node will not be shown in the Busy state, but will be listed under the Available nodes, as there are parts of the available cores on the node still available and can be used for more workload/jobs being sent to the cluster.
 - Available - All nodes that are available (not being used for workload/jobs, are not offline, are not being deployed or not being destroyed are listed (for the selected

cluster), and nodes that are partially used (workload/jobs that are only using part of the available cores on a node) are listed in the Available state.

- **Deploying** - For the selected cluster, all nodes that are currently being deployed within a Cloud Provider are listed.
- **Offline** - For the selected cluster, all nodes that were used (busy), and have become available after being used, and the Idle Purge Time of the cluster configuration has expired will be placed in an Offline state before being destroyed. After the next burst cycle has run, an offline node will be moved to be destroyed. The main reason for an Offline state is, if a node goes into an Offline state, and before being destroyed, more workloads are sent to the cluster, the node can very easily and quickly be moved back into an Available state, rather than needing to be deployed again after it was destroyed.
- **Destroying** - For the selected cluster, all nodes that are currently being destroyed within a Cloud Provider are listed.
- **Down** - For the selected cluster, all nodes that have gone through the full termination process and are terminated (completely down) on a Cloud Provider, are listed.

These are the Actions/Details that display: State, Power State, Processors, Type, Service Port, Manager Port, Total Sockets, Total Nodes, Total Cores, Total Threads, Dedicated Sockets, Dedicated Nodes, Dedicated Cores, Dedicated Threads, Operating System, User, Sessions, N Sessions, Idle Time, Total Memory, Available Memory, Physical Memory, CPUs, Load Average, Net Load, State, CPU Clock, MAC Address, Torque Version, REC Time, and Jobs.

Notes:

- Nodes are available by default, but can be taken down by selecting **Destroy**. A node that is down has been deprovisioned in the Cloud service provider and does not count as a billable resource. A node that is offline is still provisioned in the Cloud service provider and isn't being used to run jobs at the moment, but counts as a billable resource.
 - MIG Mode only appears if you have a MIG-capable instance in your cluster. When a cluster is deployed, you can change the MIG mode cluster-wide if a worker node has a MIG capable instance ID. Navigate to the deployed cluster view, and click the **Nodes** tab. From the MIG selector drop-down at the bottom left, select the desired MIG mode and then a loading indicator will appear. When you perform a MIG mode change during a deployment, the newly selected MIG mode will be stored as your desired mode, even if you destroy and redeploy your cluster.
-

- **Configuration - Cluster Information** (Cluster ID, Head Node Size, Image Name, Cluster IP, SSH Username, Created, Download Cluster SSH Key, and Install/Uninstall OpenVPN), **Cluster Compute Nodes** (Size, OCPUs, Memory, Count, and Resize Cluster), and **Bursting Service** (Min/Max, View Logs, Burst Once), and Disks Information (Manage Disks).

The Manage Disks feature enables ODDC Administrators to add extra drives to the head node. When the cluster is deployed, another drive is attached to the head node that can be shared. When the cluster is destroyed, the drive is also destroyed.

Notes: Resize requires bursting service to be disabled. See the sections [Bursting Configurations](#) and [Bursting Service](#) for information about bursting.

Note: Jobs can also quickly be run from any screen at any time by clicking the **Quick Launch / Briefcase** icon (📁) near the top left of the screen.

5. (optional) To add extra drives to the head node, click **Configuration / Manage Disks**, and then click the add icon (+). Select the desired information and then click **Close/Do**.

2.9 File Manager

You can view, download, and delete a job's standard output file or error file using File Manager. **Note:** If a cluster is destroyed, these contents will be erased.

1. On the **Applications** menu, click **File Manager**, and then from the **Select Cluster** drop-down list at the top right, select a cluster:

Manager ID	Nodus ID	Name	Status	Owner	Size(B)	Modified	Actions
40	40	test-stdout-stderr	Running		70	June 16, 2021 9:22 AM	...

There are no Files...

The **File Manager** screen shows this information: Manager ID, Nodus ID, Name, Status, Owner, Size(B), Modified, and Actions.

- Click the ellipsis at the end of a row and select the desired option: View STD-Out, Download STD-Out, View STD-Err, Download STD-Err, or Delete.

2.10 Accounting Manager

Accounting shows an accounting of all the jobs run by each user over a period of time specified, and also the clusters. You must be an Admin to access the Accounting feature.

- On the **Applications** menu, click **Accounting**:

User	Jobs	Time	Cost	Details	Download
ctaylor	192	6.483h	\$2.24		
qgauser	172	2.883h	\$0		

The Jobs panel shows this information: Users, Job, Time, Cost, Details, and Download.

- To see a user's jobs that ran, click **Details** on the desired user row:

Job	Start Time	Provider	Cluster	Node	Cores	Run Time
mpi-1x8	30 Jun 2023 9:29	otc	otc-moab	nodus-s2-8xlarge-1-4797-0	8	1.067h
random	30 Jun 2023 9:55	otc	otc-moab	nodus-s2-8xlarge-1-4797-0	1	0.017h
Moab.4309	30 Jun 2023 11:57	otc	otc-moab	nodus-s2-8xlarge-1-4797-2	16	0.8h
Moab.4310	30 Jun 2023 11:57	otc	otc-moab	nodus-s2-8xlarge-1-4797-3	24	1.2h
Moab.4311	30 Jun 2023 11:57	otc	otc-moab	nodus-s2-8xlarge-1-4797-1	32	1.067h
Moab.4308	30 Jun 2023 11:57	otc	otc-moab	nodus-s2-8xlarge-1-4797-2	8	1.067h
Moab.4312	30 Jun 2023 11:57	otc	otc-moab	nodus-s2-8xlarge-1-4797-2 nodus-s2-8xlarge-1-4797-3	8 8	2.4h
Moab.4320	30 Jun 2023 12:01	otc	otc-moab	nodus-s2-8xlarge-1-4797-0	8	1.067h
Moab.4321	30 Jun 2023 12:01	otc	otc-moab	nodus-s2-8xlarge-1-4797-0	16	1.067h
Moab.4313	30 Jun 2023 11:59	otc	otc-moab	nodus-s2-8xlarge-1-4797-3 nodus-s2-8xlarge-1-4797-1	16 16	1.067h

This screen shows this information: Job, Start Time, Provider, Cluster, Node, Cores, and Run Time.

- To see cluster information, click the **Clusters** tab:

Jobs	Clusters						
Name	User	Deploy	Destroy	Uptime	Costs	Details	Download
otc-moab	qauser1	12 Aug 2023 8:44	27 Jan 2024 10:24 ...	10.40h	\$3.40		
aws-moab	nodus	19 Aug 2023 11:50	27 Jan 2024 10:24 ...	187.20h	\$72.93		
aws	nodus	19 Aug 2023 17:30	27 Jan 2024 10:24 ...	62.40h	\$22.93		
oracle	nodus	19 Aug 2023 17:30	27 Jan 2024 10:24 ...	10.40h	\$2.97		

This screen shows this information: Name, User, Deploy, Destroy, Uptime, Costs, Details, and Download.

2.11 Instance Prices

This feature helps you estimate the cost of running a job in different Cloud providers.

1. On the **Applications** menu, click **Instance Prices**:

\$

Instance Prices

Cloud Providers

Oracle Cloud HPC

Oracle Cloud

Amazon Web Services

Google Cloud

Microsoft Azure

Open Telekom Cloud

Instance	US Virginia (US EAST)	US Silicon Valley (US WEST)	EU Germany (Frankfurt)	CN China (Hong Kong)	AP Singapore	AP Australia (Sydney)	AP Japan (Tokyo)	AP India (Mumbai)
ecs.c6.xlarge - CPU: 4, Mem (GB): 8	0.1520	0.1520	0.1720	0.1920	0.1740	0.1980	0.1900	0.1520
ecs.c6.2xlarge - CPU: 8, Mem (GB): 16	0.3040	0.3040	0.3440	0.3840	0.3480	0.3960	0.3800	0.3040
ecs.c6.4xlarge - CPU: 16, Mem (GB): 32	0.6080	0.6080	0.6880	0.7680	0.6960	0.7920	0.7600	0.6080
ecs.c6.8xlarge - CPU: 32, Mem (GB): 64	1.2160	1.2160	1.3760	1.5360	1.3920	1.5840	1.5200	1.2160
ecs.c6.13xlarge - CPU: 52, Mem (GB): 96	1.9760	1.9760	2.2360	2.4960	2.2620	2.5740	2.4700	1.9760
ecs.c6.26xlarge - CPU: 104, Mem (GB): 192	3.9520	3.9520	4.4720	4.9920	4.5240	5.1480	4.9400	3.9520

This panel shows this information: Cloud Providers, Instance, Region, and job cost per hour estimation.

Notes:

- To change the price per region estimation to match the cost from your provider, open the `providers.json` file (`/NODUS/providers.json`) and change the `cost-region` amount. For example, for `"instance_id": "t2.nano"`, change `"us-east-2": 0.0030` to `"us-east-2": 0.0045`.
- If you get a message The quota of CPUs established by the provider is exceeded, you need to manually specify the quotas established by the provider in the `providers.json` file, inside the `zones` (`zones: {cpu:99, instances:99, network:10}`).

2.12 Cloud Budget Analysis

Now that you have successfully run one job in the Cloud by completing the steps above, you can run all your jobs in the Cloud and begin developing your Cloud budget. See the appendix [HPC Cloud On-Demand Data Center Workload Considerations](#) for additional information.

Chapter 3: On-Demand Data Center CLI Procedures

In this chapter:


- 3.1 Users 41
- 3.2 Architecture 42
- 3.3 CLI Usage 43
 - 3.3.1 Jobs 43
 - 3.3.2 Building Images 43
- 3.4 Cluster Actions 43
 - 3.4.1 Connecting to a Cluster 43
 - 3.4.2 Destroying a Cluster 43
 - 3.4.3 Viewing All Clusters 43
- 3.5 Job Actions 44
- 3.6 Common CLI Commands 44
 - 3.6.1 Cluster Management 44
 - 3.6.2 Cluster Monitoring 46
 - 3.6.3 Cluster Jobs 46
 - 3.6.4 Admin Commands 47
- 3.7 GPU Enabled Clusters 47
 - 3.7.1 GPU Enabled Cluster Restrictions 47
 - 3.7.2 GPU Stack Software 47
 - 3.7.3 Building a GPU Stack 48
 - 3.7.4 Updating a Cluster to Reference the GPU Prefix 48
 - 3.7.5 Creating a GPU Job 48
 - 3.7.6 Submitting a GPU Job 48

3.1 Users


An ODDC user is just a registered Linux system user that enables using the ODDC in various capacities, permissions, and license permitting.

See the CLI Commands section [Users](#) for information about user commands.

3.2 Architecture



Cluster Resources



3.3 CLI Usage

3.3.1 Jobs

Show standard error messages: `$ nodus job:output ID CLUSTER [OPTIONS]`

3.3.2 Building Images

After a cluster is associated with a set of credentials and a stack file, use the command `$ nodus stack:build STACK [OPTIONS]` to create the image.

3.4 Cluster Actions

3.4.1 Connecting to a Cluster

1. To connect to a deployed cluster via SSH, use the command `nodus cluster:ssh CLUSTER`.

3.4.2 Destroying a Cluster

1. Assuming the cluster was deployed properly through the above CLI commands, a cluster can be destroyed using the command `$ nodus cluster:destroy CLUSTER [OPTIONS]`.
2. The cluster should be destroyed before you delete it, or you will have to destroy it manually. To remove all information about the cluster, use the command `nodus cluster:remove CLUSTER`.

3.4.3 Viewing All Clusters

1. To view all saved clusters, deployed or not, use the command `nodus cluster:list`.

3.5 Job Actions

1. To submit a job to a cluster, use the command `nodus job:submit JOB CLUSTER`.

3.6 Common CLI Commands

This section contains common ODDC CLI commands:

- [Cluster Management](#)
- [Cluster Monitoring](#)
- [Cluster Jobs](#)
- [Admin Commands](#)

See the chapter [On-Demand Data Center CLI Commands](#) for detailed information.

3.6.1 Cluster Management

This section contains the following information:

- [Cluster Deploy](#)
- [Cluster Destroy](#)
- [Cluster SSH](#)

3.6.1.A Cluster Deploy

This deploys nodes on a cluster: `$ nodus cluster:deploy CLUSTER --target HOSTNAME`

See [cluster:deploy CLUSTER](#) for additional information.

3.6.1.B Cluster Destroy

This destroys nodes on a cluster: `$ nodus cluster:destroy CLUSTER --target HOSTNAME`

See [cluster:destroy CLUSTER](#) for additional information.

3.6.1.C Cluster SSH

This connects to a cluster: `$ nodus cluster:ssh CLUSTER [OPTIONS]`

See [cluster:ssh CLUSTER](#) for additional information.

3.6.2 Cluster Monitoring

This section contains the following information:

- [Cluster Show](#)
- [Cluster Jobs Queue](#)
- [Cluster List](#)
- [Cluster Members](#)
- [Cluster Nodes](#)

3.6.2.A Cluster Show

This gets general information about the cluster (i.e., cluster username, public IP, node list...): `$ nodus cluster:show CLUSTER`

3.6.2.B Cluster Jobs Queue

This displays JSON information about the jobs queue: `$ nodus cluster:queue CLUSTER`

3.6.2.C Cluster List

This lists all user clusters: `$ nodus cluster:list`

3.6.2.D Cluster Members

This displays serf members from the head node: `$ nodus cluster:members CLUSTER [OPTIONS]`

3.6.2.E Cluster Nodes

This displays a detailed description of compute nodes on the cluster: `$ nodus cluster:nodes CLUSTER [OPTIONS]`

3.6.3 Cluster Jobs

This section contains the following information:

- [Job Submit](#)

3.6.3.A Job Submit

This remotely runs a local job script on a cluster: `$ nodus job:submit JOB CLUSTER [OPTIONS]`

3.6.4 Admin Commands

This section contains the following information:

- [Delete User](#)

3.6.4.A Delete User

This deletes a user from the system: `sudo nodus admin delete-users <username>`

3.7 GPU Enabled Clusters

This section describes the process for running GPU jobs:

- [GPU Enabled Cluster Restrictions](#)
- [GPU Stack Software](#)
- [Building a GPU Stack](#)
- [Updating a Cluster to Reference the GPU Prefix](#)
- [Creating a GPU Job](#)
- [Submitting a GPU Job](#)

3.7.1 GPU Enabled Cluster Restrictions

These are the restrictions for creating and running GPU enabled clusters:

- Must be Nvidia-based GPUs
- Must use Torque
- Must be non-bursting clusters

3.7.2 GPU Stack Software

The GPU stack comes with the following software configured for Nvidia GPU support:

- Open MPI 4
- Torque 7
- Docker (Only the use of Docker's container management system is supported; tools like Podman are not supported.)

3.7.3 Building a GPU Stack

```
$ nodus stack:build myGpuStack -s gpu
```

Special Consideration

Use a volume greater than 10 GB for the GPU stack.

Note: See [stack:build STACK](#) for additional information.

3.7.4 Updating a Cluster to Reference the GPU Prefix

```
$ nodus cluster:<provider>:update myGpuCluster --node-prefix  
myGpuPrefix
```

3.7.5 Creating a GPU Job

```
$ nodus job:create /path/to/my/job/script myGpuJob -c numCores  
-n numNodes -g numGpus
```

3.7.6 Submitting a GPU Job

```
$ nodus job:submit myGpuJob myGpuCluster
```

Note: To reserve GPU resources, the cluster cannot be in bursting mode. If you want to use bursting with GPU instances, you can, but GPU resource isolation is not supported.

Special Consideration

Queue polling of GPU jobs is not currently supported. If you submit a job requiring GPU resources, but the resources are in use (isolated by Torque), the submission will fail and a 500 response code will be returned.

To release the job from the queue, run this command: `$ nodus cluster:qrun myGpuCluster myGpuJobNumber`

Chapter 4: On-Demand Data Center CLI Commands

This chapter contains a list of CLI commands.

In this chapter:

4.1

Actions

.....

49

4.2

Users

.....

51

4.3

Credentials

.....

53

4.4

Stacks

.....

57

4.5

Clusters

.....

76

4.6

Jobs

.....

110

4.7

Providers

.....

128

4.8

Misc

.....

129

4.1 Actions

This section contains the following commands:

- `action:list`
- `action:remove ACTION`
- `action:show ACTION`

action:list

Description

List action entries from database.

Usage

\$ nodus action:list [OPTIONS]

Arguments

NAME	Name Identifier
TYPE	(centos-7) Image Type
CLUSTER	Cluster Name or ID

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

--id=id	Filter by Affected ID
--type=action cluster credential job provider stack user	Filter by Action Type
Examples	
\$ nodus action:list	
\$ nodus action:list --format json	

action:remove ACTION

Description

Remove action entries from database.

Usage

```
$ nodus action:remove ACTION [[ACTION] ...] [OPTIONS]
```

Argument

ACTION	Action ID
--------	-----------

Option

-h	--help	Show Help
----	--------	-----------

Example

```
action:remove 5e8f6a8ac86619500ab573da
```

action:show ACTION

Description

Display details of a particular action entry from the database.

Usage

```
$ nodus action:show ACTION [OPTIONS]
```

Argument

ACTION	Action ID
--------	-----------

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus action:show 5ea06014df12871f264716eb
```

```
$ nodus action:show 5ea06014df12871f264716eb --format json
```

4.2 Users

This section contains the following commands:

- `user:add USERNAME`
- `user:list`
- `user:remove USERNAME`
- `user:show USER`
- `user:update USERNAME`

user:add USERNAME

Description

Registers linux user as an ODDC user.

Usage

```
$ nodus user:add USERNAME [OPTIONS]
```

Argument

USERNAME	Linux Username
----------	----------------

Options

-e	--email=email	Email Address
-f	--first=first	First Name (Given Name)
-h	--help	Show Help
-l	--last=last	Last Name (Family Name)
	--[no-]admin	Make User Admin
	--[no-]ui	Allow User to Login via Web App / UI

Examples

```
$ nodus user:add jdoe
```

```
$ nodus user:add jdoe --admin
```

Note: If a specific SSH key is to be used for a pre-existing Linux user, the key must be named `oddc_id_rsa` and reside in the user's `.ssh` home directory.

user:list

Description

List user entries from database.

Usage

```
$ nodus user:list [OPTIONS]
```

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus user:list
```

```
$ nodus user:list --format json
```

user:remove USERNAME**Description**

Remove ODDC user(s).

Usage

```
$ nodus user:remove USERNAME [OPTIONS]
```

Argument

USERNAME	Username of the system user
----------	-----------------------------

Options

-h	--help	Show Help
	--clean	Remove User Home Directory from Filesystem
	--force	Force Remove User Home Directory from Filesystem

Examples

```
$ nodus user:remove jdoe
```

```
$ nodus user:remove jdoe --clean
```

```
$ nodus user:remove jdoe --force
```

user:show USER**Description**

Display details of a particular user entry from the database.

Usage

```
$ nodus user:show USER [OPTIONS]
```

Argument

USERNAME	User Name or ID
----------	-----------------

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus user:show jdoe
```

```
$ nodus user:show jdoe --format json
```

user:update USERNAME**Description**

Update ODDC user details.

Usage

```
$ nodus user:update USERNAME [OPTIONS]
```

Argument

USERNAME	Linux Username
----------	----------------

Options

-e	--email=email	Email Address
-f	--first=first	First Name (Given Name)
-h	--help	Show Help
-l	--last=last	Last Name (Family Name)
-u	--[no-]ui	Allow User to Login via Web App / UI
	--[no-]admin	Make User Admin
	--pass	Change Password
	--private=private	Path to Private SSH Key
	--public=public	Path to Public SSH Key

Example

```
$ nodus user:update jdoe --admin
```

4.3 Credentials

This section contains the following commands:

- [credential:add FILE NAME](#)
- [credential:list](#)
- [credential:list CREDENTIAL](#)
- [credential:remove CREDENTIAL](#)
- [credential:share CREDENTIAL USERNAME](#)

- [credential:show CREDENTIAL](#)
- [credential:unshare CREDENTIAL](#)

credential:add FILE NAME

Description

Create credential entry with information from file.

Usage

```
$ nodus credential:add FILE NAME --provider PROVIDER
```

Arguments

FILE	Credential File (Absolute / Relative) Path
NAME	Credential Name

Options

-d	--description=description	Description of Credential
-h	--help	Show Help
-p	--provider=aws cloudmart gcp azure oracle otc	(required) Provider Code

Examples

```
$ nodus credential:add /home/user/Downloads/credential.json myCredential --provider aws
```

```
$ nodus credential:add ~/Downloads/credential.json myCredential --provider aws
```

credential:list

Description

List credential entries from database.

Usage

```
$ nodus credential:list [OPTIONS]
```

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus credential:list [OPTIONS]
```

```
$ nodus credential:list --format json
```

credential:list CREDENTIAL

Description

Update credential entry.

Usage

```
$ nodus credential:list CREDENTIAL [OPTIONS]
```

Argument

CREDENTIAL	Credential Name or ID
------------	-----------------------

Options

-d	--description=description	Description
-h	--help	Show Help
-n	--name=name	Name Identifier

Examples

```
$ nodus credential:list myCredential --name newCredential
```

credential:remove CREDENTIAL

Description

Remove credential entries from database.

Usage

```
$ nodus credential:remove CREDENTIAL [OPTIONS]
```

Argument

CREDENTIAL	Credential Name or ID
------------	-----------------------

Options

-f	--force	Force Removal from Database
-h	--help	Show Help

Example

```
credential:remove myCredential
```

credential:share CREDENTIAL USERNAME

Description

Share your credential to another user.

Usage

```
$ nodus credential:share CREDENTIAL USERNAME
```

Arguments

CREDENTIAL	Credential Name or ID
USERNAME	Linux Username

Options

-h	--help	Show Help
----	--------	-----------

Examples

```
$ nodus credential:share myCredential otherUser
```

credential:show CREDENTIAL**Description**

Display details of a particular credential entry from the database.

Usage

```
$ nodus credential:show CREDENTIAL [OPTIONS]
```

Argument

CREDENTIAL	Credential Name or ID
------------	-----------------------

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus credential:show myCredential
```

```
$ nodus credential:show myCredential --format json
```

credential:unshare CREDENTIAL**Description**

Unshare your credential to another user.

Usage

```
$ nodus credential:unshare CREDENTIAL [USERNAME] [OPTIONS]
```

Arguments

CREDENTIAL	Credential Name or ID
USERNAME	Linux Username

Options

-h	--help	Show Help
	--all	Remove all shared users

Examples


```
$ nodus credential:unshare myCredential otherUser  
$ nodus credential:unshare myCredential --all
```

4.4 Stacks

This section contains the following commands:

- `stack:aws:create NAME TYPE`
- `stack:aws:update STACK`
- `stack:azure:create NAME TYPE`
- `stack:azure:update STACK`
- `stack:build STACK`
- `stack:cloudmart:create NAME TYPE`
- `stack:cloudmart:update STACK`
- `stack:gcp:create NAME TYPE`
- `stack:gcp:update STACK`
- `stack:list`
- `stack:log CLUSTER TYPE`
- `stack:ocihpc:create NAME TYPE`
- `stack:ocihpc:update STACK`
- `stack:oracle:create NAME TYPE`
- `stack:oracle:update STACK`
- `stack:otc:create NAME TYPE`
- `stack:otc:update STACK`
- `stack:remove STACK`
- `stack:set STACK CREDENTIAL`
- `stack:share STACK USERNAME`
- `stack:show STACK`
- `stack:unshare STACK`

stack:aws:create NAME TYPE

Description

Create new Amazon Web Services (AWS) stack configuration entry.

Usage

```
$ nodus stack:aws:create NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Associated Name
TYPE	(centos-7) Image Type
CREDENTIAL	Set Credential

Options

-D	--dataFile=dataFile	Path (absolute or relative) of user input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: t2.micro] Instance Type
-p	--prefix=prefix	(required) [default: nodus] Output Image Prefix
-r	--region=region	(required) Region ID
-s	--source=source	Source Image ID
-u	--username=username	SSH Username
	--private	Provision via Private IP
	--proxy=proxy	HTTP(S) Proxy URL
	--security=security	Security Group ID
	--subnet=subnet	Subnet ID
	--volume=volume	Storage Volume (GB)

Examples

```
$ nodus stack:aws:create myStack centos-7 myCred --region=us-east-2
```

```
$ nodus stack:aws:create myStack centos-7 --prefix=myNodus --region=us-east-2 --instance=t2.large
```

```
$ nodus stack:aws:create myStack centos-7 --region us-east-2 --proxy $http_proxy
```

```
$ nodus stack:aws:create myStack centos-7 --region=us-east-2 --custom --source=ami-04152c3a27c49a944 --username=ec2-user --instance=t2.xlarge
$ nodus stack:aws:create myStack centos-7 --region=us-east-2 --instance=t2.xlarge --serverProvisioner ~/myServerScript --nodeProvisioner ~/myNodeScript --dataFile ~/myDataFile2
```

stack:aws:update STACK

Description

Update Google Cloud Platform (GCP) cluster configuration entry.

Usage

```
$ nodus stack:aws:update STACK [OPTIONS]
```

Arguments

STACK	Stack Name or ID
-------	------------------

Options

-D	--dataFile=dataFile	Input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	User provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	User provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Instance Type
-n	--name=name	Name Identifier
-p	--prefix=prefix	Output Image Prefix
-r	--region=region	Region ID
-s	--source=source	Source Image ID
-t	--type=centos-7	OS Type
-u	--username=username	SSH Username
	--credential=credential	Credential
	--no-proxy	Remove HTTP(S) Proxy
	--no-security	Remove Security Group
	--no-subnet	Remove Subnet
	--no-volume	Remove Extended Root Storage Volume
	--[no-]private	Provision via Private IP
	--proxy=proxy	HTTP(S) Proxy URL

Options

<code>--security=security</code>	Security Group ID
<code>--subnet=subnet</code>	Subnet ID
<code>--volume=volume</code>	Extended Root Storage Volume (GB)

Examples

```
$ nodus stack:aws:update myStack --name newStack
$ nodus stack:aws:update myStack --region=us-east-2
$ nodus stack:aws:update myStack --prefix=myNodus --region=us-east-2 --
instance=t2.large
$ nodus stack:aws:update myStack --region us-east-2 --proxy $http_proxy
$ nodus stack:aws:update myStack --region=us-east-2 --custom --source=ami-
04152c3a27c49a944 --username=ec2-user --instance=t2.xlarge
```

stack:azure:create NAME TYPE**Description**

Create new Microsoft Azure stack configuration entry.

Usage

```
$ nodus stack:azure:create NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Associated Name
TYPE	(centos-7) Image Type
CREDENTIAL	Set Credential

Options

-D	<code>--dataFile=dataFile</code>	Path (absolute or relative) of user input files for user provisioning process (multiple files supported)
-N	<code>--nodeProvisioner=nodeProvisioner</code>	Path (absolute or relative) of user provisioning script to run during stack build on Compute node only
-S	<code>--serverProvisioner=serverProvisioner</code>	Path (absolute or relative) of user provisioning script to run during stack build on Head node only
-c	<code>--custom</code>	Use Managed Image
-d	<code>--description=description</code>	Description
-h	<code>--help</code>	Show Help
-i	<code>--instance=instance</code>	(required) [default: Standard_DS2_v2] Instance Type ID

-p	--prefix=prefix	(required) [default: nodus] Output Image Prefix
-s	--source=source	Source Image ID
-u	--username=username	SSH Username
	--offer=offer	Source Image Offer
	--os-type=os-type	Source Image OS Type
	--publisher=publisher	Source Image Publisher
	--resource=resource	(required) Managed Resource Group Name
	--sku=sku	Source Image SKU
	--source-resource=source-resource	Source Resource Group Name

Examples

```
$ nodus stack:azure:create myStack centos-7 myCred --resource=myResourceGroup --instance=Standard_DS2_v2
```

```
$ nodus stack:azure:create myStack centos-7 myCred --resource=myResourceGroup -S ~/myServerScript -N ~/myNodeScript -D ~/myDataFile2 -D ~/myDataFile2
```

```
$ nodus stack:azure:create myStack centos-7 --prefix=myNodus --resource=myResourceGroup --instance=Standard_DS2_v2
```

stack:azure:update STACK

Description

Update Microsoft Azure stack configuration entry.

Usage

```
$ nodus stack:azure:update STACK [OPTIONS]
```

Arguments

STACK	Stack Name or ID
-------	------------------

Options

-D	--dataFile=dataFile	Input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	User provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	User provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: Standard_DS2_v2] Instance Type ID
-n	--name=name	Name Identifier

Options

-p	--prefix=prefix	(required) [default: nodus] Output Image Prefix
-s	--source=source	Source Image ID
-t	--type=centos-7	OS Type
-u	--username=username	SSH Username
	--credential=credential	Credential
	--os-type=os-type	Source Image OS Type
	--publisher=publisher	Source Image Publisher
	--resource=resource	(required) Managed Resource Group Name
	--sku=sku	Source Image SKU
	--source-resource=source-resource	Source Resource Group Name

Examples

```
$ nodus stack:azure:update myStack --name newStack
```

```
$ nodus stack:azure:update myStack --resource=myResourceGroup --instance=Standard_DS2_v2
```

```
$ nodus stack:azure:update myStack --prefix=myNodus --resource=myResourceGroup --instance=Standard_DS2_v2
```

stack:build STACK**Description**

Generates images compatible with ODDC clusters. Generate two images: 'nodus-server'; 'nodus-node'.

Usage

```
$ nodus stack:build STACK [OPTIONS]
```

Argument

STACK	Stack Name or ID
-------	------------------

Options

-S	--slow	Sequential Builds
-h	--help	Show Help
-s	--stack=server node	Sub-Stack

Examples

```
$ nodus stack:build myStack
```

```
$ nodus stack:build myStack --stack server
```

stack:cloudmart:create NAME TYPE

Description

Create new HPC CloudMart stack configuration entry.

Usage

```
$ nodus stack:cloudmart:create NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Associated Name
TYPE	(centos-7) Image Type
CREDENTIAL	Set Credential

Options

-D	--dataFile=dataFile	Path (absolute or relative) of user input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: VM.Standard2.1] Instance Type ID
-n	--subnet=subnet	(required) Subnet OCID
-p	--prefix=prefix	(required) [default: nodus] Output Image Prefix
-r	--region=region	(required) Region ID
-s	--source=source	Source Image ID
-u	--username=username	SSH Username
-z	--availability-domain=availability-domain	(required) Availability Domain for OCI tenancy
	--min	Minimum required stack build
	--private	Provision via Private IP

Example

```
$ nodus stack:cloudmart:create myStack centos-7 myCred --prefix=myNodus --region=us-ashburn-1 --subnet=ocid1.subnet.oc1.iad.aaaaaaa7v2znysxlclowk7rhgkgldvjaxmkkjpiktnahnusgknhw5pcvi6q
```

```
$ nodus stack:cloudmart:create myStack centos-7 myCred --prefix=myNodus --region=us-ashburn-1 --subnet=mySubnet -S ~/myServerScript -N ~/myNodeScript -D ~/myDataFile1 -D ~/myDataFile2
```

stack:cloudmart:update STACK

Description

Update HPC CloudMart stack configuration entry.

Usage

```
$ nodus stack:cloudmart:update STACK [OPTIONS]
```

Arguments

STACK	Stack Name or ID
-------	------------------

Options

-D	--dataFile=dataFile	Input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	User provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	User provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Instance Type ID
-n	--name=name	Name Identifier
-n	--subnet=subnet	(required) Subnet OCID
-p	--prefix=prefix	Output Image Prefix
-r	--region=region	Region ID
-s	--source=source	Source Image ID
-t	--type=centos-7 centos-7-bm centos-7-latest oel-7 oel-7-bm	OS Type
-u	--username=username	SSH Username
-z	--availability-domain=availability-domain	Availability Domain for OCI tenancy
	--credential=credential	Credential
	--[no-]min	Minimum required stack build
	--[no-]private	Provision via Private IP

Examples

```
$ nodus stack:cloudmart:update myStack --name newStack
```


Examples

```
$ nodus stack:cloudmart:update myStack --prefix=myNodus --region=us-ashburn-1 --
subnet=ocid1.subnet.oc1.iad.aaaaaaa7v2znysxlclowk7rhgkgldvjaxmkkjpiktnahnusgknhw
5pcvi6q
```

stack:gcp:create NAME TYPE**Description**

Create new Google Cloud Platform (GCP) stack configuration entry.

Usage

```
$ nodus stack:gcp:create NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Associated Name
TYPE	(centos-7) Image Type
CREDENTIAL	Set Credential

Options

-D	--dataFile=dataFile	Path (absolute or relative) of user input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Head node only
-c	--custom	Managed Image
-d	--disk=disk	Disk Size (GB)
-h	--help	Show Help
-i	--instance=instance	(required) [default: n1-standard-1] Instance Type ID
-p	--prefix=prefix	(required) [default: nodus] Output Image Prefix
-s	--source=source	Source AMI ID
-u	--username=username	SSH Username
-z	--zone=zone	(required) Availability Zone ID
	--[no-]account	Auth via Service Account
	--disk=disk	Extended Root Disk Size (GB)
	--proxy=proxy	HTTP(S) Proxy URL
	--subnet=subnet	Subnet Self Link ID

Examples

```
$ nodus stack:gcp:create myStack centos-7 myCred --zone=us-east1-b --instance=n1-standard-1
```

```
$ nodus stack:gcp:create myStack centos-7 --prefix=myNodus --zone=us-east1-b --instance=n1-standard-1
```

```
$ nodus stack:gcp:create myStack centos-7 myCred -z us-east4-b -p myPrefix --serverProvisioner ~/myServerScript --nodeProvisioner ~/myNodeScript --dataFile ~/myDataFile1 --dataFile ~/myDataFile2
```

stack:gcp:update STACK**Description**

Update Google Cloud Platform (GCP) stack configuration entry.

Usage

```
$ nodus stack:gcp:update STACK [OPTIONS]
```

Arguments

STACK	Stack Name or ID
-------	------------------

Options

-D	--dataFile=dataFile	Input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	User provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	User provisioning script to run during stack build on Head node only
-c	--custom	Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Instance Type ID
-n	--name=name	Name Identifier
-p	--prefix=prefix	Output Image Prefix
-s	--source=source	Source AMI ID
-t	--type=centos-7	OS Type
-u	--username=username	SSH Username
-w	--walltime=walltime	Job walltime
-z	--zone=zone	Availability Zone ID
	--credential=credential	Credential
	--[no-]account	Auth via Service Account
	--disk=disk	Extended Root Disk Size (GB)

Options

<code>--no-disk</code>	Remove Disk
<code>--no-proxy</code>	Remove HTTP(S) Proxy
<code>--no-subnet</code>	Remove Subnet
<code>--proxy=proxy</code>	HTTP(S) Proxy URL
<code>--subnet=subnet</code>	Subnet Self Link ID

Examples

```
$ nodus stack:gcp:update myStack --name newStack
$ nodus stack:gcp:update myStack --zone=us-east1-b --instance=n1-standard-1
$ nodus stack:gcp:update myStack --prefix=myNodus --zone=us-east1-b --instance=n1-standard-1
```

stack:list**Description**

List stack entries from database.

Usage

```
$ nodus stack:list [OPTIONS]
```

Options

<code>-f</code>	<code>--format=text json</code>	[default: text] Output Format
<code>-h</code>	<code>--help</code>	Show Help

Examples

```
$ nodus stack:list
$ nodus stack:list --format json
```

stack:log CLUSTER TYPE**Description**

Display stack build logs.

Usage

```
$ nodus stack:log CLUSTER TYPE [OPTIONS]
```

Arguments

STACK	Stack Name or ID
TYPE	(master build) [default: master] Log Type

Options

<code>-d</code>	<code>--debug</code>	Debug Logs
-----------------	----------------------	------------

-f	--follow	Follow File Output Stream
-h	--help	Show Help
-n	--lines=lines	Last N Lines
	--follow	Follow File Output Stream

Example

```
$ nodus stack:log myCluster
```

stack:ocihpc:create NAME TYPE

Description

Create new Oracle Cloud stack configuration entry.

Usage

```
$ nodus stack:ocihpc:create NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Associated Name
TYPE	(centos-7-bm oel-7 oel-7-gpu) Image Type
CREDENTIAL	Set Credential

Options

-D	--dataFile=dataFile	Path (absolute or relative) of user input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: BM.Optimized3.36] Instance Type ID
-n	--subnet=subnet	(required) Subnet OCID
-p	--prefix=prefix	(required) [default: nodus] Output Image Prefix
-r	--region=region	(required) Region ID
-s	--source=source	Source Image ID
-u	--username=username	SSH Username
-z	--availability-domain=availability-	(required) Availability Domain for OCI tenancy

domain	
--min	Minimum required stack build
--private	Provision via Private IP

Example

```
$ nodus stack:ocihpc:create myStack centos-7 myCred --prefix=myNodus --region=us-ashburn-1 --subnet=ocid1.subnet.oc1.iad.aaaaaaa7v2znysxlclowk7rhgkgldvjaxmkkjpiktnahnusgknhw5pcvi6q
```

```
$ nodus stack:ocihpc:create myStack centos-7 myCred --prefix=myNodus --region=us-ashburn-1 --subnet=mySubnet -S ~/myServerScript -N ~/myNodeScript -D ~/myDataFile1 -D ~/myDataFile2
```

stack:ocihpc:update STACK**Description**

Update Oracle Cloud stack configuration entry.

Usage

```
$ nodus stack:ocihpc:update STACK [OPTIONS]
```

Arguments

STACK	Stack Name or ID
-------	------------------

Options

-N	--name=name	Name Identifier
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Instance Type ID
-n	--subnet=subnet	Subnet OCID
-p	--prefix=prefix	Output Image Prefix
-r	--region=region	Region ID
-s	--source=source	Source Image ID
-t	--type=centos-7-bm oel-7 oel-7-gpu	OS Type
-u	--username=username	SSH Username
-z	--availability-domain=availability-domain	Availability Domain for OCI tenancy
	--credential=credential	Credential
	--[no-]min	Minimum required stack build
	--[no-]private	Provision via Private IP

Example

```
$ nodus stack:ocihpc:update myStack --name newStack
$ nodus stack:ocihpc:update myStack --prefix=myNodus --region=us-ashburn-1 --
subnet=ocid1.subnet.oc1.iad.aaaaaaa7v2znysxlclowk7rhgkgldvjaxmkkjpiktnahnusgknhw
5pcvi6q
```

stack:oracle:create NAME TYPE

Description

Create new Oracle Cloud stack configuration entry.

Usage

```
$ nodus stack:oracle:create NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Associated Name
TYPE	(centos-7) Image Type
CREDENTIAL	Set Credential

Options

-D	--dataFile=dataFile	Path (absolute or relative) of user input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: VM.Standard2.1] Instance Type ID
-n	--subnet=subnet	(required) Subnet OCID
-p	--prefix=prefix	(required) [default: nodus] Output Image Prefix
-r	--region=region	(required) Region ID
-s	--source=source	Source Image ID
-u	--username=username	SSH Username
-z	--availability-domain=availability-domain	(required) Availability Domain for OCI tenancy
	--min	Minimum required stack build
	--private	Provision via Private IP

Example

```
$ nodus stack:oracle:create myStack centos-7 myCred --prefix=myNodus --region=us-ashburn-1 --subnet=ocid1.subnet.oc1.iad.aaaaaaaa7v2znysxlclowk7rhgkgldvjaxmkkjpiktnahnusgknhw5pcvi6q
$ nodus stack:oracle:create myStack centos-7 myCred --prefix=myNodus --region=us-ashburn-1 --subnet=mySubnet -S ~/myServerScript -N ~/myNodeScript -D ~/myDataFile1 -D ~/myDataFile2
```

stack:oracle:update STACK**Description**

Update Oracle Cloud stack configuration entry.

Usage

```
$ nodus stack:oracle:update STACK [OPTIONS]
```

Arguments

STACK	Stack Name or ID
-------	------------------

Options

-D	--dataFile=dataFile	Input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	User provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	User provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Instance Type ID
-n	--name=name	Name Identifier
-n	--subnet=subnet	(required) Subnet OCID
-p	--prefix=prefix	Output Image Prefix
-r	--region=region	Region ID
-s	--source=source	Source Image ID
-t	--type=centos-7 centos-7-bm centos-7-latest oel-7 oel-7-bm	OS Type
-u	--username=username	SSH Username
-z	--availability-domain=availability-domain	Availability Domain for OCI tenancy

Options

--credential=credential	Credential
--[no-]min	Minimum required stack build
--[no-]private	Provision via Private IP

Examples

```
$ nodus stack:oracle:update myStack --name newStack
```

```
$ nodus stack:oracle:update myStack --prefix=myNodus --region=us-ashburn-1 --
subnet=ocid1.subnet.oc1.iad.aaaaaaaa7v2znysxlclowk7rhgkgldvjaxmkkjpiktnahnusgknhw
5pcvi6q
```

stack:otc:create NAME TYPE**Description**

Create new Open Telekom Cloud stack configuration entry.

Usage

```
$ nodus stack:otc:create NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Associated Name
TYPE	(centos-7) Image Type
CREDENTIAL	Set Credential

Options

-D	--dataFile=dataFile	Path (absolute or relative) of user input files for user provisioning process (multiple files supported)
-N	--nodeProvisioner=nodeProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Compute node only
-S	--serverProvisioner=serverProvisioner	Path (absolute or relative) of user provisioning script to run during stack build on Head node only
-c	--custom	Use Managed Image
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: s2.medium.2] Instance Type ID
-p	--prefix=prefix	(required) [default: nodus] Output Image Prefix
-s	--source=source	Source Image ID
-u	--username=username	SSH Username

<code>--private</code>	Provision via Private IP
<code>--proxy=proxy</code>	HTTP(S) Proxy URL
<code>--security=security</code>	Security Group ID
<code>--subnet=subnet</code>	(required) Subnet Network ID
<code>--volume=volume</code>	Extended Root Storage Volume (GB)

Examples

```
$ nodus stack:otc:create myStack centos-7 myCred --subnet 3bc69be8-3ebc-434c-b828-9aeb5a8ec78a --instance s2.medium.2
```

```
$ nodus stack:otc:create myStack centos-7 --prefix myNodus --subnet 3bc69be8-3ebc-434c-b828-9aeb5a8ec78a --instance s2.medium.2
```

```
$ nodus stack:otc:create myStack centos-7 --subnet 3bc69be8-3ebc-434c-b828-9aeb5a8ec78a --proxy $http_proxy
```

```
$ nodus stack:otc:create myStack centos-7 --subnet mySubnet --serverProvisioner ~/myServerScript --nodeProvisioner ~/myNodeScript --dataFile ~/myDataFile2
```

stack:otc:update STACK**Description**

Update Open Telekom Cloud stack configuration entry.

Usage

```
$ nodus stack:otc:update STACK [OPTIONS]
```

Arguments

STACK	Stack Name or ID
CREDENTIAL	Set Credential

Options

-D	<code>--dataFile=dataFile</code>	Path (absolute or relative) of user input files for user provisioning process (multiple files supported)
-N	<code>--nodeProvisioner=nodeProvisioner</code>	Path (absolute or relative) of user provisioning script to run during stack build on Compute node only
-S	<code>--serverProvisioner=serverProvisioner</code>	Path (absolute or relative) of user provisioning script to run during stack build on Head node only
-c	<code>--custom</code>	Use Managed Image
-d	<code>--description=description</code>	Description
-h	<code>--help</code>	Show Help
-i	<code>--instance=instance</code>	Instance Type ID
-n	<code>--name=name</code>	Name Identifier

-p	--prefix=prefix	Output Image Prefix
-s	--source=source	Source Image ID
-t	--type=centos-7	OS Type
-u	--username=username	SSH Username
	--credential=credential	Credential
	--no-proxy	Remove HTTP(S) Proxy
	--no-security	Remove Security Group
	--no-volume	Remove Extended Root Storage Volume
	--[no-]private	Provision via Private IP
	--proxy=proxy	HTTP(S) Proxy URL
	--security=security	Security Group ID
	--subnet=subnet	Subnet Network ID
	--volume=volume	Extended Root Storage Volume (GB)

Examples

```
$ nodus stack:otc:create myStack myCred --name newStack
```

```
$ nodus stack:otc:create myStack --subnet 3bc69be8-3ebc-434c-b828-9aeb5a8ec78a --instance s2.medium.2
```

```
$ nodus stack:otc:create myStack --prefix myNodus --subnet 3bc69be8-3ebc-434c-b828-9aeb5a8ec78a --instance s2.medium.2
```

```
$ nodus stack:otc:create myStack --subnet 3bc69be8-3ebc-434c-b828-9aeb5a8ec78a --proxy $http_proxy
```

stack:remove STACK

Description

Remove stack entries from database.

Usage

```
$ nodus stack:remove STACK [OPTIONS]
```

Argument

STACK	Stack Name or ID
-------	------------------

Options

-h	--help	Show Help
	--[no-]clean	Remove Files Created for Cluster

Examples

```
$ nodus stack:remove myStack
```

```
$ nodus stack:remove myStack --no-clean
```

```
$ nodus stack:remove myStack --force
```

stack:set STACK CREDENTIAL**Description**

Associate stack entry with credential entry.

Usage

```
$ nodus stack:set STACK CREDENTIAL [OPTIONS]
```

Arguments

STACK	Stack Name or ID
CREDENTIAL	Credential Name or ID

Option

-h	--help	Show Help
----	--------	-----------

Example

```
$ nodus stack:set myStack myCredential
```

stack:share STACK USERNAME**Description**

Share your stack to another user.

Usage

```
$ nodus stack:share STACK USERNAME
```

Argument

STACK	Stack Name or ID
USERNAME	Linux Username

Options

-h	--help	Show Help
----	--------	-----------

Examples

```
$ nodus stack:share myStack otherUser
```

stack:show STACK**Description**

Display details of a particular stack entry from the database.

Usage

```
$ nodus stack:show STACK [OPTIONS]
```

Argument

STACK	Stack Name or ID
-------	------------------

Options

-f	--format=text json xml	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus stack:show myStack
```

```
$ nodus stack:show myStack --format json
```

stack:unshare STACK**Description**

Unshare your stack to another user.

Usage

```
$ nodus stack:unshare STACK [USERNAME] [OPTIONS]
```

Argument

STACK	Stack Name or ID
USERNAME	Linux Username

Options

-h	--help	Show Help
	--all	Remove all shared users

Examples

```
$ nodus stack:unshare myStack otherUser
```

```
$ nodus stack:unshare myStack --all
```

4.5 Clusters

This section contains the following commands:

- [cluster:add-disk](#)
- [cluster:aws:create NAME TYPE MANAGER](#)
- [cluster:aws:update CLUSTER](#)
- [cluster:azure:create NAME TYPE MANAGER](#)
- [cluster:azure:update CLUSTER](#)
- [cluster:burst CLUSTER](#)
- [cluster:cloudmart:create NAME TYPE MANAGER](#)

- `cluster:cloudmart:update CLUSTER`
- `cluster:cores CLUSTER`
- `cluster:deploy CLUSTER`
- `cluster:destroy CLUSTER`
- `cluster:gcp:create NAME TYPE MANAGER`
- `cluster:gcp:update CLUSTER`
- `cluster:get-mig-mode`
- `cluster:list`
- `cluster:local:create`
- `cluster:local:update CLUSTER`
- `cluster:log`
- `cluster:members CLUSTER`
- `cluster:nodes CLUSTER`
- `cluster:node-create CLUSTER LABEL PREFIX INSTANCE COUNT`
- `cluster:node-remove CLUSTER`
- `cluster:node-update CLUSTER LABEL`
- `cluster:nodes-list CLUSTER`
- `cluster:ocihpc:create NAME TYPE MANAGER`
- `cluster:ocihpc:update CLUSTER`
- `cluster:oracle:create NAME TYPE MANAGER`
- `cluster:oracle:update CLUSTER`
- `cluster:otc:create NAME TYPE MANAGER`
- `cluster:otc:update CLUSTER`
- `cluster:ovpn CLUSTER`
- `cluster:qrun`
- `cluster:queue CLUSTER`
- `cluster:releasehold CLUSTER`
- `cluster:remove CLUSTER`

- `cluster:remove-output CLUSTER ID`
- `cluster:service CLUSTER`
- `cluster:set CLUSTER CREDENTIAL`
- `cluster:set-mig-mode`
- `cluster:share CLUSTER USERNAME`
- `cluster:show CLUSTER`
- `cluster:ssh CLUSTER`
- `cluster:sshfs CLUSTER`
- `cluster:sync CLUSTER TYPE`
- `cluster:unshare CLUSTER`

cluster:add-disk

Description

Add additional disks to clusters.

Usage

```
$ nodus cluster:add-disk CLUSTER [[NAME TYPE VOLUME] ...] [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
DISK-NAME	Name of disk
DISK-TYPE	Disk type ID code
VOLUME	Size of disk (GB)

Options

-f	--force	Ignore Checks
-h	--help	Show Help

Examples

```
$ nodus cluster:add-disk myCluster ssd1 pd-ssd 25 balanced1 pd-balanced 20
```

```
$ nodus cluster:add-disk myCluster my-ssd1 gp3 10 my-ssd2 gp2 15
```

```
$ nodus cluster:add-disk myOracleCluster my-ssd1 paravirtualized 30 (in Oracle-cluster, TYPE is 'paravirtualized', MANDATORY!!!)
```

cluster:aws:create NAME TYPE MANAGER

Description

Create new Amazon Web Services (AWS) cluster configuration entry.

Usage

```
$ nodus cluster:aws:create NAME TYPE MANAGER [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Name Identifier
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Set Credential

Options

-P	--package=package	Packages
-b	--[no-]burstOnDeploy	Initiate Burst Service post Deploy
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) Server Instance Type ID
-m	--mode=min max	[default: min] Bursting Mode
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (example: nodus)
-r	--region=region	(required) Region ID
-z	--zone=zone	Availability Zone ID
	--eip=eip	Elastic IP ID
	--idle-purge=idle-purge	[default: 600] Node Idle Purge Time(s)
	--job-completed-dropoff	Completed Jobs Queue Time
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--private	Provision via Private IP
	--role	Auth via EC2 Role
	--security=security	Security Group ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)
	--subnet=subnet	Subnet ID
	--vpc=vpc	VPC Network Self Link ID

Examples

```
$ nodus cluster:aws:create myCluster centos-7 torque myCred --prefix myNodus --region us-east-2 --instance=t2.large --job-completed-dropoff=60
```

```
$ nodus cluster:aws:create myCluster centos-7 slurm --region us-east-2 --package git --package python3 --package perl
```

```
$ nodus cluster:aws:create myCluster centos-7 torque --region us-east-2 --instance=t2.large --vpc=vpc-03ceecbf61bb37c99 --subnet=subnet-05aa121e98148d0ce --security=sg-09cda11597d6fd926
```

cluster:aws:update CLUSTER

Description

Update Amazon Web Services (AWS) cluster configuration entry.

Usage

\$ nodus cluster:aws:update CLUSTER [OPTIONS]

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-P	--package=package	Add Package
-R	--no-package=no-package	Remove Package
-b	--[no-]burstOnDeploy	Initiate Burst Service post Deploy
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Server Instance Type ID
-m	--mode=min max	Bursting Mode
-n	--name=name	Name Identifier
-p	--prefix=prefix	Source Image Root Prefix (Ex: nodus)
-r	--region=region	Region ID
-t	--type=centos-7	OS Type
-z	--zone=zone	Availability Zone ID
	--credential=credential	Credential
	--eip=eip	Elastic IP ID
	--force	Skip Checks
	--idle-purge=idle-purge	Node Idle Purge Time(s)
	--job-completed-dropoff	Completed Jobs Queue Time
	--manager=torque slurm	HPC Workload Manager
	--no-eip	Remove Elastic IP
	--no-node-volume	Remove Expanded Node Root Volume
	--no-security	Remove Security Group
	--no-server-volume	Remove Expanded Server Root Volume
	--no-subnet	Remove Subnet
	--no-vpc	Remove VPC
	--no-zone	Remove Zone
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--[no-]private	Provision via Private IP
	--[no-]role	Auth via EC2 Role

Options

<code>--security=security</code>	Security Group ID
<code>--server-volume=server-volume</code>	Expanded Server Root Volume (GB)
<code>--subnet=subnet</code>	Subnet ID
<code>--vpc=vpc</code>	VPC ID

Examples

```
$ nodus cluster:aws:update myCluster --name newCluster
$ nodus cluster:aws:update myCluster --job-completed-dropoff=60
$ nodus cluster:aws:update myCluster --instance=t2.large
$ nodus cluster:aws:update myCluster --vpc=vpc-03ceecbf61bb37c99 --subnet=subnet-05aa121e98148d0ce --security=sg-09cda11597d6fd926
```

cluster:azure:create NAME TYPE MANAGER**Description**

Create new Microsoft Azure cluster configuration entry.

Usage

```
$ nodus cluster:azure:create NAME TYPE MANAGER [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Name Identifier
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Set Credential

Options

<code>-P --package=package</code>	Packages
<code>-R --resource=resource</code>	(required) Image Resource Group Name
<code>-b --[no-]burstOnDeploy</code>	Initiate Burst Service post Deploy
<code>-d --description=description</code>	Description
<code>-h --help</code>	Show Help
<code>-i --instance=instance</code>	(required) [default: Standard_DS2_v2] Server Instance Type ID
<code>-m --mode=min max</code>	[default: min] Bursting Mode
<code>-p --prefix=prefix</code>	(required) [default: nodus] Source Image Root Prefix (example: nodus)
<code>-r --region=region</code>	(required) Region ID
<code>--idle-purge=idle-purge</code>	[default: 600] Node Idle Purge Time(s)
<code>--job-completed-dropoff</code>	Completed Jobs Queue Time

<code>--node-volume=node-volume</code>	Expanded Node Root Volume (GB)
<code>--server-volume=server-volume</code>	Expanded Server Root Volume (GB)

Example

```
$ nodus cluster:azure:create myCluster centos-7 torque myCred --region=eastus --instance=Standard_DS2_v2 --resource=myResourceGroup --job-completed-dropoff=60
```

```
$ nodus cluster:azure:create myCluster centos-7 slurm --prefix=myNodus --region=westus --instance=Standard_DS2_v2 --resource=myResourceGroup
```

cluster:azure:update CLUSTER**Description**

Update Microsoft Azure cluster configuration entry.

Usage

```
nodus cluster:azure:update CLUSTER [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

<code>-P</code>	<code>--package=package</code>	Add Package
<code>-R</code>	<code>--no-package=no-package</code>	Remove Package
<code>-R</code>	<code>--resource=resource</code>	Image Resource Group Name
<code>-b</code>	<code>--[no-]burstOnDeploy</code>	Initiate Burst Service post Deploy
<code>-d</code>	<code>--description=description</code>	Description
<code>-h</code>	<code>--help</code>	Show Help
<code>-i</code>	<code>--instance=instance</code>	Server Instance Type ID
<code>-m</code>	<code>--mode=min max</code>	Bursting Mode
<code>-n</code>	<code>--name=name</code>	Name Identifier
<code>-p</code>	<code>--prefix=prefix</code>	Prefix Source Image Root Prefix (Ex: nodus)
<code>-r</code>	<code>--region=region</code>	Region ID
<code>-t</code>	<code>--type=centos-7</code>	OS Type
	<code>--credential=credential</code>	Credential
	<code>--force</code>	Skip Checks
	<code>--idle-purge=idle-purge</code>	Node Idle Purge Time(s)
	<code>--job-completed-dropoff</code>	Completed Jobs Queue Time
	<code>--manager=torque slurm</code>	HPC Workload Manager
	<code>--no-node-volume</code>	Remove Expanded Node Root Volume

Options

<code>--no-server-volume</code>	Remove Expanded Server Root Volume
<code>--node-volume=node-volume</code>	Expanded Node Root Volume (GB)
<code>--server-volume=server-volume</code>	Expanded Server Root Volume (GB)

Examples

```
$ nodus cluster:azure:update myCluster --name newCluster
```

```
$ nodus cluster:azure:update myCluster --job-completed-dropoff=60
```

```
$ nodus cluster:azure:update myCluster --prefix=myNodus --region=eastus --resource=myResourceGroup
```

cluster:burst CLUSTER**Description**

Manage cluster bursting state or manual burst.

Usage

```
$ nodus cluster:burst CLUSTER [OPTIONS]
```

Argument

CLUSTER	Cluster Name or ID
---------	--------------------

Options

<code>-f</code>	<code>--force</code>	Skip State Checks
<code>-h</code>	<code>--help</code>	Show Help
<code>-m</code>	<code>--mode=min max</code>	Bursting Mode
	<code>--raw</code>	Raw Details
	<code>--start</code>	Start bursting service Daemon
	<code>--status</code>	Check Service Status
	<code>--stop</code>	Stop bursting service Daemon

Examples

```
$ nodus cluster:burst myCluster
```

```
$ nodus cluster:burst myCluster --start
```

```
$ nodus cluster:burst myCluster --stop
```

```
$ nodus cluster:burst myCluster --status
```

```
$ nodus cluster:burst myCluster --mode max --start
```

cluster:cloudmart:create NAME TYPE MANAGER**Description**

Create new HPC CloudMart cluster configuration entry.

Usage

\$ nodus cluster:cloudmart:create NAME TYPE MANAGER [CREDENTIAL] [OPTIONS]

Arguments

NAME	Name Identifier
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Set Credential

Options

-P	--package=package	Packages
-b	--[no-]burstOnDeploy	Initiate Burst Service post Deploy
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: VM.Standard2.1] Server Instance Type ID
-m	--mode=min max	[default: min] Bursting Mode
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (example: nodus)
-r	--region=region	(required) Region ID
-s	--share=nfs beegfs	[default: nfs] Shared mnt type (Ex: beegfs or nfs)
-z	--availability-domain=availability-domain	(required) Availability Domain for OCI tenancy
	--LPG=LPG	LPG for Provision via Private IP --idle-purge=idle-purge [default: 600] Node Idle Purge Time (s)
	--idle-purge=idle-purge	[default: 600] Node Idle Purge Time(s)
	--job-completed-dropoff	Completed Jobs Queue Time
	--local_cidr=local_cidr	cidr of ODDC vcn
	--node-prefix=node-prefix	Source Image Root Prefix (Ex: nodus) - Node Only
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--private	Provision via Private IP
	--script=script	Provisioning Script
	--security=security	Security Group ID
	--server-cores=server-cores	Server OCPUs - flex shapes only ¹
	--server-mem=server-mem	Server memory (in GBs) - flex shapes only ¹
	--server-volume=server-	Expanded Server Root Volume (GB)

volume	
--vcn_cidr=vcn_cidr	vcn_cidr for cluster to be created

Example

```
$ nodus cluster:cloudmart:create myCluster centos-7 torque myCred --prefix=myNodus --region=us-ashburn-1 --instance=VM.Standard2.1 --job-completed-dropoff=60
```

```
$ nodus cluster:cloudmart:create myCluster centos-7 slurm --region=us-ashburn-1 --instance=VM.Standard2.1 --share=beegfs
```

¹Flex shapes (flexible shape) enable you to use flexible instance shapes on the OCI platform for changing the number of cores and memory for compute nodes and the head node. This can be done via flags at cluster creation by using the `cluster:oracle:create` command for the head node, and by using the `cluster:oracle:update` command for compute nodes. Additionally, in the CLI, when there are no specified values for the server cores or memory, the default values (specified by Oracle) are used. If these flags are used for a non-flexible shape, the values are ignored. In the UI, however, it will be set to 1 instead.

cluster:cloudmart:update CLUSTER

Description

Update HPC CloudMart cluster configuration entry.

Usage

```
nodus cluster:cloudmart:update CLUSTER [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-P	--package=package	Add Package
-R	--no-package=no-package	Remove Package
-b	--[no-]burstOnDeploy	Initiate Burst Service post Deploy
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Server Instance Type ID
-m	--mode=min max	Bursting Mode
-n	--name=name	Name Identifier
-p	--prefix=prefix	Prefix Source Image Root Prefix (Ex: nodus)
-r	--region=region	Region ID
-s	--share=nfs beegfs	Shared mnt type (Ex: beegfs or nfs)
-t	--type=centos-7	OS Type
-z	--availability-domain=availability-	Availability Domain for OCI tenancy

Options	
domain	
--credential=credential	Credential
--force	Skip Checks
--idle-purge=idle-purge	Node Idle Purge Time(s)
--job-completed-dropoff	Completed Jobs Queue Time
--manager=torque slurm	HPC Workload Manager
--no-node-volume	Remove Expanded Node Root Volume
--no-script	Remove Provisioning Script
--no-security	Remove Security Group
--no-server-volume	Remove Expanded Server Root Volume
--node-prefix=node-prefix	Source Image Root Prefix (Ex: nodus) - Node Only
--node-volume=node-volume	Expanded Node Root Volume (GB)
--script=script	Provisioning Script
--server-cores=server-cores	Server OCPUs - flex shapes only ¹
--server-mem=server-mem	Server memory (in GBs) - flex shapes only ¹
--server-volume=server-volume	Expanded Server Root Volume (GB)
--private	Create a cluster with no public IP
--LPG	OCID of local peering gateway associated with ODDC VCN
--vcn_cidr=vcn_cidr	vcn_cidr for cluster to be created
--local_cidr	cidr of ODDC server VCN

Examples

```
$ nodus cluster:cloudmart:update myCluster --name newCluster
```

```
$ nodus cluster:cloudmart:update myCluster --job-completed-dropoff=60
```

```
$ nodus cluster:cloudmart:update myCluster --prefix=myNodus
```

¹ Flex shapes (flexible shape) enable you to use flexible instance shapes on the OCI platform for changing the number of cores and memory for compute nodes and the head node. This can be done via flags at cluster creation by using the `cluster:oracle:create` command for the head node, and by using the `cluster:oracle:update` command for compute nodes. Additionally, in the CLI, when there are no specified values for the server cores or memory, the default values (specified by Oracle) are used. If these flags are used for a non-flexible shape, the values are ignored. In the UI, however, it will be set to 1 instead.

cluster:cores CLUSTER

Description

Get cluster core information.

Usage

\$ nodus cluster:cores CLUSTER [OPTIONS]

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

\$ nodus cluster:cores myCluster

cluster:deploy CLUSTER

Description

Modify cloud resources based on cluster configuration entry.

Usage

\$ nodus cluster:deploy CLUSTER [OPTIONS]

Argument

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-h	--help	Show Help
-t	--target=target	Node Hostname
	--[no-]plan	Show Plan

Examples

\$ nodus cluster:deploy myCluster

\$ nodus cluster:deploy myCluster --no-plan

cluster:destroy CLUSTER

Description

Destroy cloud resources based on cluster configuration entry.

Usage

\$ nodus cluster:destroy CLUSTER [OPTIONS]

Argument

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-f	--force	Skip Checks
-h	--help	Show Help
-t	--target=target	Node Hostname
	--plan	Show Plan

Examples

```
$ nodus cluster:destroy myCluster
```

```
$ nodus cluster:destroy myCluster --no-plan
```

cluster:gcp:create NAME TYPE MANAGER**Description**

Create new Google Cloud Platform (GCP) cluster configuration entry.

Usage

```
$ nodus cluster:gcp:create NAME TYPE MANAGER [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Name Identifier
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Set Credential

Options

-P	--package=package	Packages
-b	--[no-]burstOnDeploy	Initiate Burst Service post Deploy
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: n1-standard-1] Server Instance Type ID
-m	--mode=min max	[default: min] Bursting Mode
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (example: nodus)
-r	--region=region	(required) Region ID
-z	--zone=zone	(required) Availability Zone ID
	--account	Auth via Service Account
	--disk-type=ssd balanced standard	[default: standard] Storage class for boot disk Options correspond to the GCP persistent disk types [ssd persistent disks (pd-ssd), balanced persistent disks (pd-balanced), and standard persistent disks (pd-standard)]
	--firewall=firewall	Firewall Self Link ID

<code>--idle-purge=idle-purge</code>	[default: 600] Node Idle Purge Time(s)
<code>--job-completed-dropoff</code>	Completed Jobs Queue Time
<code>--node-volume=node-volume</code>	Expanded Node Root Volume (GB)
<code>--server-volume=server-volume</code>	Expanded Server Root Volume (GB)
<code>--subnet=subnet</code>	Subnet Self Link ID
<code>--vpc=vpc</code>	VPC Network Self Link ID

Example

```
$ nodus cluster:gcp:create myCluster centos-7 torque myCred --region=us-east1 --zone=e=us-east1-b --instance=n1-standard-1 --job-completed-dropoff=60
```

```
$ nodus cluster:gcp:create myCluster centos-7 slurm --prefix=myNodus --region=us-east1 --zone=us-east1-b --instance=n1-standard-1
```

cluster:gcp:update CLUSTER**Description**

Update Google Cloud Platform (GCP) cluster configuration entry.

Usage

```
nodus cluster:gcp:update CLUSTER [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

<code>-P --package=package</code>	Add Package
<code>-R --no-package=no-package</code>	Remove Package
<code>-b --[no-]burstOnDeploy</code>	Initiate Burst Service post Deploy
<code>-d --description=description</code>	Description
<code>-h --help</code>	Show Help
<code>-i --instance=instance</code>	Server Instance Type ID
<code>-m --mode=min max</code>	Bursting Mode
<code>-n --name=name</code>	Name Identifier
<code>-p --prefix=prefix</code>	Prefix Source Image Root Prefix (Ex: nodus)
<code>-r --region=region</code>	Region ID
<code>-t --type=centos-7</code>	OS Type
<code>-z --zone=zone</code>	Availability Zone ID
<code>--credential=credential</code>	Credential
<code>--[no-]account</code>	Auth via Service Account
<code>--disk-</code>	Storage class for boot disk

Options	
type=ssd balanced standard	Options correspond to the GCP persistent disk types [ssd persistent disks (pd-ssd), balanced persistent disks (pd-balanced), and standard persistent disks (pd-standard)]
--firewall=firewall	Firewall Self Link ID
--force	Skip Checks
--idle-purge=idle-purge	Node Idle Purge Time(s)
--job-completed-dropoff	Completed Jobs Queue Time
--manager=torque slurm	HPC Workload Manager
--no-firewall	Remove Firewall
--no-node-volume	Remove Expanded Node Root Volume
--no-server-volume	Remove Expanded Server Root Volume
--no-subnet=no-subnet	Remove Subnet
--no-vpc=no-vpc	Remove VPC
--node-volume=node-volume	Expanded Node Root Volume (GB)
--server-volume=server-volume	Expanded Server Root Volume (GB)
--subnet=subnet	Subnet Self Link ID
--vpc=vpc	VPC Network Self Link ID
Examples	
\$ nodus cluster:gcp:update myCluster --name newCluster	
\$ nodus cluster:gcp:update myCluster --job-completed-dropoff=60	
\$ nodus cluster:gcp:update myCluster --prefix=myNodus --region=us-east1 --zone=us-east1-b	

cluster:get-mig-mode

Description

Returns cluster MIG configuration.

Usage

```
$ nodus cluster:get-mig-mode CLUSTER
```

Arguments

CLUSTERID	Cluster Name or ID
-----------	--------------------

Options

-h	--help	Show Help
----	--------	-----------

Example

```
$ nodus cluster:get-mig-mode myCluster
```

cluster:list**Description**

List cluster entries from database.

Usage

\$ nodus cluster:list [OPTIONS]

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

\$ nodus cluster:list

\$ nodus cluster:list --format json

cluster:local:create**Description**

Create new On Premises cluster configuration entry.

Usage

\$ nodus cluster:local:create NAME [OPTIONS]

Arguments

NAME	Name Identifier
------	-----------------

Options

-d	--description=description	Description
-h	--help	Show Help

Examples

\$ nodus cluster:local:create myCluster

cluster:local:update CLUSTER**Description**

Update On Premises cluster configuration entry.

Usage

\$ nodus cluster:local:update CLUSTER [OPTIONS]

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-d	--description=description	Description
-h	--help	Show Help
-n	--name=name	Name Identifier

Examples

```
$ nodus cluster:local:update myCluster --name newCluster
```

cluster:log**Description**

Display cluster logs.

Usage

```
$ nodus cluster:log CLUSTER TYPE [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
TYPE	(init plan deploy destroy master burst) [default: master] Log Type

Options

-d	--debug	Debug Logs
-f	--follow	Follow File Output Stream
-h	--help	Show Help
-n	--lines=lines	Last N Lines

Examples

```
$ nodus cluster:log myCluster
```

```
$ nodus cluster:log myCluster burst
```

cluster:members CLUSTER**Description**

Get cluster information as reported by 'serf members'.

Usage

```
$ nodus cluster:members CLUSTER [OPTIONS]
```

Argument

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Example

```
$ nodus cluster:members myCluster
```

cluster:nodes CLUSTER**Description**

Get cluster node information.

Usage

```
$ nodus cluster:nodes CLUSTER [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus cluster:nodes myCluster
```

cluster:node-create CLUSTER LABEL PREFIX INSTANCE COUNT**Description**

Adds node to cluster.

Usage

```
$ nodus cluster:node-create CLUSTER LABEL PREFIX INSTANCE COUNT [...[OPTIONS]]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-c	--count=count	(required) Number of nodes
-f	--force	Ignore Checks
-h	--help	Show Help
-i	--instance=instance	(required) Instance ID
-l	--label=label	(required) Flag for the nodes
-m	--mig-mode=mig-mode	Sets MIG mode for given instance shape. (Use set-mig-mode once node has been deployed.)
-p	--purge	Purge All Existing Node Groups Except Given Args
-s	--server=server	New Instance ID Code For Server
	--[no-]efa	Enable EFA

<code>--node-cores=node-cores</code>	Node OCPUs - flex shapes only
<code>--node-mem=node-mem</code>	Node memory (in GBs) - flex shapes only
<code>--prefix=prefix</code> (required)	Group node prefix

Example

```
$ nodus cluster:node-create myCluster --label=test --prefix=core --instance=t2.micro --count=4
```

```
$ nodus cluster:node-create myCluster --label=test --prefix=core --instance=t2.micro --count=4 --purge
```

```
$ nodus cluster:node-create myCluster --label=test --prefix=core --instance=BM.GPU4.8 --count=1 --mig-mode=all
```

cluster:node-remove CLUSTER**Description**

Delete a group of nodes from a cluster, or delete all cluster nodes using '--purge' flag.

Usage

```
$ nodus cluster:node-remove CLUSTER [LABEL]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-h	--help	Show Help
-l	--label=label	Label of the node to be deleted
-p	--purge	Delete all the nodes of a cluster

Example

```
$ nodus cluster:nodus node-remove myCluster -l test 1
```

```
$ nodus cluster:node-remove myCluster --label test
```

```
$ nodus cluster:node-remove myCluster --purge
```

The command can remove either a collection of nodes from a cluster by using the --label flag followed by the label of the group, or it can remove all of the nodes in the cluster by utilizing the --purge flag.

¹ This example command will remove the node(s) with the 'test' label from the myCluster cluster.

cluster:node-update CLUSTER LABEL

Description

Update cluster nodes.

Usage

```
$ nodus cluster:node-update CLUSTER LABEL [[COUNT][INSTANCE ID] ...] [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
LABEL	Label for the group of nodes

Options

-c	--count=count	(required) Number of nodes
-f	--force	Ignore Checks
-h	--help	Show Help
-i	--instance=instance	Instance ID
-l	--label=label	Flag for the nodes
-m	--mig-mode=mig-mode	Sets MIG mode for given instance shape. (Use set-mig-mode once node has been deployed.)
-p	--purge	Purge All Existing Node Groups Except Given Args
-s	--server=server	New Instance ID Code For Server
	--efa=efa	Enable EFA
	--node-cores=node-cores	Node OCPUs - flex shapes only
	--node-mem=node-mem	Node memory (in GBs) - flex shapes only
	--prefix=prefix (required)	Group node prefix

Example

```
$ nodus cluster:node-update myCluster test --prefix=core --instance=t2.micro --count=4 1
```

```
$ nodus cluster:node-update myCluster test --count=4 --purge
```

¹ This command will modify the node with the name 'test' that belongs to cluster myCluster and it will change the instance to t2.micro and the count to 4.

cluster:nodes-list CLUSTER

Description

List cluster nodes.

Usage

```
$ nodus cluster:nodes-list CLUSTER [OPTIONS]
```

Arguments

NAME	Name of the cluster to list nodes from
------	--

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Example

```
$ nodus cluster:nodes-list
```

```
$ nodus cluster:nodes-list --format json
```

cluster:ocihpc:create NAME TYPE MANAGER

Description

Create new Oracle Cloud cluster configuration entry.

Usage

```
$ nodus cluster:ocihpc:create NAME TYPE MANAGER [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Name Identifier
TYPE	(centos-7 centos-7-bm centos-7-last) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Set Credential

Options

-P	--package=package	Packages
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: VM.Standard2.1] Server Instance Type ID
-m	--mode=min max	[default: min] Bursting Mode
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (example: nodus)
-r	--region=region	(required) Region ID
-s	--share=nfs beegfs	[default: nfs] Shared mnt type (Ex: beegfs or nfs)
-z	--availability-domain=availability-domain	(required) Availability Domain for OCI tenancy
	--idle-purge=idle-purge	[default: 600] Node Idle Purge Time(s)
	--job-completed-dropoff=job-completed-dropoff	Completed Jobs Drop Off Duration (s)

--node-prefix=node-prefix	Source Image Root Prefix (Ex: nodus) - Node Only
--node-volume=node-volume	Expanded Node Root Volume (GB)
--script=script	Provisioning Script
--security=security	Security Group ID
--server-volume=server-volume	Expanded Server Root Volume (GB)

Example

```
$ nodus cluster:ocihpc:create myCluster centos-7 torque myCred --prefix=myNodus --region=us-ashburn-1 availability-domain=YNZz:US-ASHBURN-AD-1 --instance=BM.Optimized3.36 --job-completed-dropoff=60
```

```
$ nodus cluster:ocihpc:create myCluster centos-7 slurm --region=us-ashburn-1 -z=YNZz:US-ASHBURN-AD-3 --instance=BM.Optimized3.36 --share=beegfs
```

cluster:ocihpc:update CLUSTER**Description**

Update Oracle Cloud cluster configuration entry.

Usage

```
$ nodus cluster:ocihpc:update NAME TYPE MANAGER [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Name Identifier
TYPE	(centos-7 centos-7-bm centos-7-last) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Set Credential

Options

-P	--package=package	Packages
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: VM.Standard2.1] Server Instance Type ID
-m	--mode=min max	[default: min] Bursting Mode
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (example: nodus)
-r	--region=region	(required) Region ID
-s	--share=nfs beegfs	[default: nfs] Shared mnt type (Ex: beegfs or nfs)
-z	--availability-domain=availability-domain	(required) Availability Domain for OCI tenancy
	--credential=credential	Credential
	--idle-purge=idle-purge	[default: 600] Node Idle Purge Time(s)

--job-completed-dropoff=job-completed-dropoff	Completed Jobs Drop Off Duration (s)
--node-prefix=node-prefix	Source Image Root Prefix (Ex: nodus) - Node Only
--node-volume=node-volume	Expanded Node Root Volume (GB)
--script=script	Provisioning Script
--security=security	Security Group ID
--server-volume=server-volume	Expanded Server Root Volume (GB)

Example

```
$ nodus cluster:ocihpc:update myCluster centos-7 torque myCred --prefix=myNodus --
region=us-ashburn-1 availability-domain=YNZz:US-ASHBURN-AD-1 --
instance=BM.Optimized3.36 --job-completed-dropoff=60
```

```
$ nodus cluster:ocihpc:update myCluster centos-7 slurm --region=us-ashburn-1 -
z=YNZz:US-ASHBURN-AD-3 --instance=BM.Optimized3.36 --share=beegfs
```

cluster:oracle:create NAME TYPE MANAGER**Description**

Create new Oracle Cloud cluster configuration entry.

Usage

```
$ nodus cluster:oracle:create NAME TYPE MANAGER [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Name Identifier
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Set Credential

Options

-P	--package=package	Packages
-b	--[no-]burstOnDeploy	Initiate Burst Service post Deploy
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	(required) [default: VM.Standard2.1] Server Instance Type ID
-m	--mode=min max	[default: min] Bursting Mode
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (example: nodus)
-r	--region=region	(required) Region ID
-s	--share=nfs beegfs	[default: nfs] Shared mnt type (Ex: beegfs or nfs)
-z	--availability-	(required) Availability Domain for OCI tenancy

domain=availability-domain	
--LPG=LPG	LPG for Provision via Private IP --idle-purge=idle-purge [default: 600] Node Idle Purge Time (s)
--idle-purge=idle-purge	[default: 600] Node Idle Purge Time(s)
--job-completed-dropoff	Completed Jobs Queue Time
--local_cidr=local_cidr	cidr of ODDC vcn
--node-prefix=node-prefix	Source Image Root Prefix (Ex: nodus) - Node Only
--node-volume=node-volume	Expanded Node Root Volume (GB)
--private	Provision via Private IP
--script=script	Provisioning Script
--security=security	Security Group ID
--server-cores=server-cores	Server OCPUs - flex shapes only ¹
--server-mem=server-mem	Server memory (in GBs) - flex shapes only ¹
--server-volume=server-volume	Expanded Server Root Volume (GB)
--vcn_cidr=vcn_cidr	vcn_cidr for cluster to be created

Example

```
$ nodus cluster:oracle:create myCluster centos-7 torque myCred --prefix=myNodus --region=us-ashburn-1 --instance=VM.Standard2.1 --job-completed-dropoff=60
```

```
$ nodus cluster:oracle:create myCluster centos-7 slurm --region=us-ashburn-1 --instance=VM.Standard2.1 --share=beegfs
```

¹ Flex shapes (flexible shape) enable you to use flexible instance shapes on the OCI platform for changing the number of cores and memory for compute nodes and the head node. This can be done via flags at cluster creation by using the `cluster:oracle:create` command for the head node, and by using the `cluster:oracle:update` command for compute nodes. Additionally, in the CLI, when there are no specified values for the server cores or memory, the default values (specified by Oracle) are used. If these flags are used for a non-flexible shape, the values are ignored. In the UI, however, it will be set to 1 instead.

cluster:oracle:update CLUSTER

Description

Update Oracle Cloud cluster configuration entry.

Usage		
nodus cluster:oracle:update CLUSTER [OPTIONS]		
Arguments		
CLUSTER	Cluster Name or ID	
Options		
-P	--package=package	Add Package
-R	--no-package=no-package	Remove Package
-b	--[no-]burstOnDeploy	Initiate Burst Service post Deploy
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Server Instance Type ID
-m	--mode=min max	Bursting Mode
-n	--name=name	Name Identifier
-p	--prefix=prefix	Prefix Source Image Root Prefix (Ex: nodus)
-r	--region=region	Region ID
-s	--share=nfs beegfs	Shared mnt type (Ex: beegfs or nfs)
-t	--type=centos-7	OS Type
-z	--availability-domain=availability-domain	Availability Domain for OCI tenancy
	--credential=credential	Credential
	--force	Skip Checks
	--idle-purge=idle-purge	Node Idle Purge Time(s)
	--job-completed-dropoff	Completed Jobs Queue Time
	--manager=torque slurm	HPC Workload Manager
	--no-node-volume	Remove Expanded Node Root Volume
	--no-script	Remove Provisioning Script
	--no-security	Remove Security Group
	--no-server-volume	Remove Expanded Server Root Volume
	--node-prefix=node-prefix	Source Image Root Prefix (Ex: nodus) - Node Only
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--script=script	Provisioning Script
	--server-cores=server-cores	Server OCPUs - flex shapes only ¹
	--server-mem=server-mem	Server memory (in GBs) - flex shapes only ¹
	--server-volume=server-volume	Expanded Server Root Volume (GB)
	--private	Create a cluster with no public IP
	--LPG	OCID of local peering gateway associated with ODDC VCN

Options

<code>--vcn_cidr=vcn_cidr</code>	vcn_cidr for cluster to be created
<code>--local_cidr</code>	cidr of ODDC server VCN

Examples

```
$ nodus cluster:oracle:update myCluster --name newCluster
```

```
$ nodus cluster:oracle:update myCluster --job-completed-dropoff=60
```

```
$ nodus cluster:oracle:update myCluster --prefix=myNodus
```

¹ Flex shapes (flexible shape) enable you to use flexible instance shapes on the OCI platform for changing the number of cores and memory for compute nodes and the head node. This can be done via flags at cluster creation by using the `cluster:oracle:create` command for the head node, and by using the `cluster:oracle:update` command for compute nodes. Additionally, in the CLI, when there are no specified values for the server cores or memory, the default values (specified by Oracle) are used. If these flags are used for a non-flexible shape, the values are ignored. In the UI, however, it will be set to 1 instead.

cluster:otc:create NAME TYPE MANAGER**Description**

Create new Open Telekom Cloud cluster configuration entry.

Usage

```
$ nodus cluster:otc:create NAME TYPE MANAGER [CREDENTIAL] [OPTIONS]
```

Arguments

NAME	Name Identifier
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Set Credential

Options

-P	<code>--package=package</code>	Packages
-b	<code>--[no-]burstOnDeploy</code>	Initiate Burst Service post Deploy
-d	<code>--description=description</code>	Description
-h	<code>--help</code>	Show Help
-i	<code>--instance=instance</code>	(required) [default: s2.medium.2] Server Instance Type ID
-m	<code>--mode=min max</code>	[default: min] Bursting Mode
-p	<code>--prefix=prefix</code>	(required) [default: nodus] Source Image Root Prefix (example: nodus)
	<code>--fip=fip</code>	Floating IP Address
	<code>--idle-purge=idle-purge</code>	[default: 600] Node Idle Purge Time(s)

--job-completed-dropoff	Completed Jobs Queue Time
--network=network	Subnet Network ID
--node-volume=node-volume	Expanded Node Root Volume (GB)
--private	Provision via Private IP
--security=security	Security Group ID
--server-volume=server-volume	Expanded Server Root Volume (GB)
--subnet=subnet	Subnet ID
--zone=zone	Availability Zone ID

Example

```
$ nodus cluster:otc:create myCluster centos-7 torque myCred --prefix=myNodus --instance=e=s2.medium.2 --job-completed-dropoff=60
```

```
$ nodus cluster:otc:create myCluster centos-7 slurm --instance=s2.medium.2
```

cluster:otc:update CLUSTER**Description**

Update Open Telekom Cloud cluster configuration entry.

Usage

```
$ nodus cluster:otc:update CLUSTER [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-P	--package=package	Add Package
-R	--no-package=no-package	Remove Package
-b	--[no-]burstOnDeploy	Initiate Burst Service post Deploy
-d	--description=description	Description
-h	--help	Show Help
-i	--instance=instance	Server Instance Type ID
-m	--mode=min max	Bursting Mode
-n	--name=name	Name Identifier
-p	--prefix=prefix	Prefix Source Image Root Prefix (Ex: nodus)
-t	--type=centos-7	OS Type
	--credential=credential	Credential
	--fip=fip	Floating IP Address
	--force	Skip Checks

Options		
	--idle-purge=idle-purge	Node Idle Purge Time(s)
	--job-completed-dropoff	Completed Jobs Queue Time
	--manager=torque slurm	HPC Workload Manager
	--network=network	Subnet Network ID
	--no-fip	Remove Floating IP Address
	--no-network	Remove Subnet Network
	--no-node-volume	Remove Expanded Node Root Volume
	--no-security	Remove Security Group
	--no-server-volume	Remove Expanded Server Root Volume
	--no-subnet=no-subnet	Remove Subnet
	--no-zone	Remove Availability Zone
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--[no-]private	Provision via Private IP
	--security=security	Security Group ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)
	--subnet=subnet	Subnet Self Link ID
	--zone=zone	Availability Zone ID

Examples

```
$ nodus cluster:otc:update myCluster --name newCluster
$ nodus cluster:otc:update myCluster --job-completed-dropoff=60
$ nodus cluster:otc:update myCluster --prefix=myNodus --instance=s2.medium.2
```

cluster:ovpn CLUSTER

Description

Installs OpenVPN on head node and can download the ovpn file that is generated.

Usage

```
$ nodus cluster:ovpn CLUSTER [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-h	--help	Show Help
----	--------	-----------

Examples

```
$ nodus cluster:ovpn myCluster
```

cluster:qrun

Description

Manually override queued job (Torque only).

Usage

```
$ nodus cluster:qrun CLUSTER JOBID [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
JOBID	Id of job to run

Options

-h	--help	Show Help
----	--------	-----------

Examples

```
$ nodus cluster:qrun myCluster myJobId
```

```
$ nodus cluster:qrun myCluster 1
```

cluster:queue CLUSTER

Description

Get workload queue information.

Usage

```
$ nodus cluster:queue CLUSTER [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus cluster:queue myCluster
```

cluster:releasehold CLUSTER

Description

Release all job holds.

Usage

```
$ nodus cluster:releasehold CLUSTER [OPTIONS]
```

Argument

CLUSTER	Cluster Name or ID
---------	--------------------

Option		
-h	--help	Show Help
Example		
\$ nodus cluster:releasehold myCluster		

cluster:remove CLUSTER

Description		
Remove cluster entries from database.		
Usage		
\$ nodus cluster:remove CLUSTER [OPTIONS]		
Argument		
CLUSTER		Cluster Name or ID
Options		
-h	--help	Show Help
	--[no-]clean	Remove Files Created for Cluster
	--force	Force Removal from Database
Examples		
cluster:remove myCluster		
cluster:remove myCluster --no-clean		
cluster:remove myCluster --force		

cluster:remove-output CLUSTER ID

Description		
Delete job directory from cluster.		
Usage		
\$ nodus cluster:remove-output CLUSTER ID [ID]... [OPTIONS]		
Arguments		
CLUSTER		Cluster Name or ID
JOB		NODUS Job ID
Options		
-h	--help	Show Help
Examples		
\$ nodus cluster:remove-output myCluster 0		
\$ nodus cluster:remove-output myCluster 0 1 2		

cluster:service CLUSTER

Description

Restarts service on the cluster head node.

Usage

```
$ nodus cluster:service CLUSTER SERVICE
```

Arguments

CLUSTER Cluster Name or ID

SERVICE (serf|pbs_server|pbs_mom|trqauthd|maui|mariadb|munge|slurmdbd|slurmctld|slurmd|slurmrestd)
Service Name

Options

-h	--help	Show Help
----	--------	-----------

Examples

```
$ nodus cluster:service myCluster pbs_server
```

cluster:set CLUSTER CREDENTIAL

Description

Associate cluster entry with credential entry.

Usage

```
$ nodus cluster:set CLUSTER CREDENTIAL [OPTIONS]
```

Arguments

CLUSTER Cluster Name or ID

CREDENTIAL Credential Name or ID

Option

-h	--help	Show Help
----	--------	-----------

Example

```
$ nodus cluster:set myCluster myCredential
```

cluster:set-mig-mode

Description

Alter cluster MIG configuration by modifying MIG instance counts.

Usage

```
$ nodus cluster:set-mig-mode CLUSTER INSTANCE_ID MIG-MODE
```

Arguments

CLUSTERID	Cluster Name or ID
INSTANCE_ID	Instance ID for MIG mode
MIG_MODE	MIG Partition Schema
Options	
-h	--help Show Help
Examples	
\$ nodus cluster:set-mig-mode myCluster BM.GPU4.8 all-disabled	
\$ nodus cluster:set-mig-mode myCluster BM.GPU4.8 all-3g.20g	

Note: To edit the GPU MIG configuration on an already deployed cluster, use `nodus cluster:set-mig-mode`.

cluster:share CLUSTER USERNAME

Description	
Share your cluster to another user.	
Usage	
\$ nodus cluster:share CLUSTER USERNAME	
Arguments	
CLUSTER	Cluster Name or ID
USERNAME	Linux Username
Options	
-h	--help Show Help
Examples	
\$ nodus cluster:share myCluster otherUser	

cluster:show CLUSTER

Description	
Display details of a particular cluster entry from the database.	
Usage	
\$ nodus cluster:show CLUSTER [OPTIONS]	
Argument	
CLUSTER	Cluster Name or ID
Options	
-b	--backup Show TF State Backup
-f	--format=text json [default: text] Output Format

-h	--help	Show Help
-k	--keys	Show Keys
-s	--state	Show TF State

Example

```
$ nodus cluster:show myCluster
```

cluster:ssh CLUSTER

Description

SSH into the head node of cluster.

Usage

```
$ nodus cluster:ssh CLUSTER [OPTIONS]
```

Argument

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-d	--details	Show SSH Details Instead
-f	--force	Ignore Checks
-h	--help	Show Help
-p	--private	Force Private IP Use
	--admin	Use Default Sudoer Instead

Example

```
$ nodus cluster:ssh myCluster
```

```
$ nodus cluster:ssh myCluster --details
```

```
$ nodus cluster:ssh myCluster --admin
```

```
$ nodus cluster:ssh myCluster --admin --details
```

cluster:sshfs CLUSTER

Description

Mount/Unmount a remote file system over SSH.

Usage

```
$ nodus cluster:sshfs CLUSTER [OPTIONS]
```

Argument

CLUSTER	Cluster Name or ID
---------	--------------------

Options

-d	--destination	Where to Mount the Directory
----	---------------	------------------------------

-h	--help	Show Help
-m	--mount	Mount the Directory
-u	--unmount	Unmount the Directory

Example

```
$ nodus cluster:sshfs myCluster --mount /mnt/beegfs --destination ~/mnt
```

```
$ nodus cluster:sshfs myCluster --unmount ~/mnt
```

cluster:sync CLUSTER TYPE

Description

Synchronizes data with cluster. Used for version upgrades, maintenance, and development purposes.

Usage

```
$ nodus cluster:sync CLUSTER TYPE [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
TYPE	(api scripts users) Type

Options

-h	--help	Show Help
	--force	Skip State Check

Examples

```
$ nodus cluster:sync myCluster api
```

```
$ nodus cluster:sync myCluster scripts
```

```
$ nodus cluster:sync myCluster users
```

cluster:unshare CLUSTER

Description

Unshare your cluster to another user.

Usage

```
$ nodus cluster:unshare CLUSTER [USERNAME] [OPTIONS]
```

Arguments

CLUSTER	Cluster Name or ID
USERNAME	Linux Username

Options

-h	--help	Show Help
	--all	Remove all shared users

Examples

```
$ nodus cluster:unshare myCluster otherUser
```

```
$ nodus cluster:unshare myCluster --all
```

4.6 Jobs

This section contains the following commands:

- `job:aws:create SCRIPT NAME TYPE`
- `job:aws:update NAME`
- `job:azure:create SCRIPT NAME TYPE`
- `job:azure:update NAME`
- `job:cancel ID CLUSTER`
- `job:cloudmart:create SCRIPT NAME TYPE`
- `job:cloudmart:update NAME`
- `job:create SCRIPT NAME`
- `job:gcp:create SCRIPT NAME TYPE`
- `job:gcp:update NAME`
- `job:list`
- `job:oracle:create SCRIPT NAME TYPE`
- `job:oracle:update NAME`
- `job:otc:create SCRIPT NAME TYPE`
- `job:otc:update NAME`
- `job:output ID CLUSTER`
- `job:remove JOB`
- `job:show JOB`
- `job:submit JOB CLUSTER`
- `job:submit-od JOB`
- `job:update NAME`

job:aws:create SCRIPT NAME TYPE

Description

Create new AWS on-demand job entry.

Usage

```
$ nodus job:aws:create SCRIPT NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

SCRIPT	(Absolute / Relative) File path to job script
NAME	Job Name
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Credential Name or ID

Options

-P	--package=package	Add Package
-c	--cores=cores	(required) [default: 1] Number of required cores per node
-d	--description=description	Description of the job
-f	--dataFile=dataFile	(Absolute / Relative) File path to data file
-h	--help	Show Help
-n	--nodes=nodes	(required) [default: 1] Number of required nodes
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (Ex: nodus)
-r	--region=region	(required) Region ID
-w	--walltime=walltime	[default: 24:00:00] Job maximum wall-time
-z	--zone=zone	Availability Zone ID
	--node-instance=node-instance	(required) Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	(required) On-Demand Behavior
	--server-instance=server-instance	(required) Server Instance Type ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)

Examples

```
$ nodus job:aws:create ~/Documents/myScript.sh myJob centos-7 --nodes=2 --cores=16 --on-demand destroy
```

```
$ nodus job:aws:create ~/Documents/myScript.sh myJob centos-7 --dataFile=~/Documents/data0.dat --dataFile=~/Documents/data1.dat --on-demand down
```

job:aws:update NAME

Description

Update AWS on-demand job entry.

Usage

\$ nodus job:aws:update NAME [OPTIONS]

Arguments

JOB	Job Name
-----	----------

Options

-P	--package=package	Add Package
-R	--no-package=no-package	Remove Package
-a	--add-df=add-df	Absolute Path to any data files to be added
-c	--cores=cores	Number of required cores per node
-d	--description=description	Description of the job
-h	--help	Show Help
-n	--nodes=nodes	Number of required nodes
-p	--prefix=prefix	Source Image Root Prefix (Ex: nodus)
-r	--remove-df=remove-df	Absolute Path to any data files to be removed
-s	--script=script	Absolute File path to job script
-t	--type=centos-7	OS Type
-w	--walltime=walltime	Job walltime
	--credential=credential	Add Credential
	--manager=torque slurm	HPC Workload Manager
	--name=name	Name Identifier
	--no-credential=no-credential	Remove Credential
	--no-node-volume	Remove Expanded Node Root Volume
	--no-server-volume	Remove Expanded Server Root Volume
	--no-zone	Remove Zone
	--node-instance=node-instance	Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	On-Demand Behavior
	--region=region	Region ID
	--server-instance=server-instance	Server Instance Type ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)
	--zone=zone	Availability Zone ID

Examples

```
$ nodus job:aws:update myJob --description="My Job" --
script=/home/jdoe/Documents/myScript.sh
```

```
$ nodus job:aws:update myJob --description="My Job" --
script=/home/jdoe/Documents/myScript.sh --nodes=2 --cores=16 --walltime="00:30:00" -
-remove-df=/home/jdoe/Documents/file0.dat --add-df=/home/jdoe/Documents/file1.dat
```

job:azure:create SCRIPT NAME TYPE**Description**

Create new AWS on-demand job entry.

Usage

```
$ nodus job:azure:create SCRIPT NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

SCRIPT	(Absolute / Relative) File path to job script
NAME	Job Name
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Credential Name or ID

Options

-P	--package=package	Package
-c	--cores=cores	(required) [default: 1] Number of required cores per node
-d	--description=description	Description of the job
-f	--dataFile=dataFile	(Absolute / Relative) File path to data file
-g	--resource=resource	(required) Image Resource Group Name
-h	--help	Show Help
-n	--nodes=nodes	(required) [default: 1] Number of required nodes
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (Ex: nodus)
-r	--region=region	(required) Region ID
-w	--walltime=walltime	[default: 24:00:00] Job maximum wall-time
	--node-instance=node-instance	(required) Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	(required) On-Demand Behavior
	--server-instance=server-instance	(required) Server Instance Type ID

Options

<code>--server-volume=server-volume</code>	Expanded Server Root Volume (GB)
--	----------------------------------

Examples

```
$ nodus job:azure:create ~/Documents/myScript.sh myJob centos-7 --nodes=2 --cores=16
--on-demand destroy
```

```
$ nodus job:azure:create ~/Documents/myScript.sh myJob centos-7 --
dataFile=~ /Documents/data0.dat --dataFile=~ /Documents/data1.dat --on-demand down
```

job:azure:update NAME**Description**

Update Azure on-demand job entry.

Usage

```
$ nodus job:azure:update NAME [OPTIONS]
```

Arguments

JOB	Job Name
------------	----------

Options

<code>-P</code>	<code>--package=package</code>	Add Package
<code>-R</code>	<code>--no-package=no-package</code>	Remove Package
<code>-a</code>	<code>--add-df=add-df</code>	Absolute Path to any data files to be added
<code>-c</code>	<code>--cores=cores</code>	Number of required cores per node
<code>-d</code>	<code>--description=description</code>	Description of the job
<code>-g</code>	<code>--resource=resource</code>	(required) Image Resource Group Name
<code>-h</code>	<code>--help</code>	Show Help
<code>-n</code>	<code>--nodes=nodes</code>	Number of required nodes
<code>-p</code>	<code>--prefix=prefix</code>	Source Image Root Prefix (Ex: nodus)
<code>-r</code>	<code>--remove-df=remove-df</code>	Absolute Path to any data files to be removed
<code>-s</code>	<code>--script=script</code>	Absolute File path to job script
<code>-t</code>	<code>--type=centos-7</code>	OS Type
<code>-w</code>	<code>--walltime=walltime</code>	Job walltime
<code>-z</code>	<code>--zone=zone</code>	Availability Zone ID
	<code>--[no-]account</code>	Auth via Service Account
	<code>--firewall=firewall</code>	Firewall Self Link ID
	<code>--force</code>	Skip Checks
	<code>--idle-purge=idle-purge</code>	Node Idle Purge Time(s)
	<code>--manager=torque slurm</code>	HPC Workload Manager
	<code>--no-firewall</code>	Remove Firewall
	<code>--no-node-volume</code>	Remove Expanded Node Root Volume

Options

<code>--no-server-volume</code>	Remove Expanded Server Root Volume
<code>--no-subnet=no-subnet</code>	Remove Subnet
<code>--no-vpc=no-vpc</code>	Remove VPC
<code>--node-volume=node-volume</code>	Expanded Node Root Volume (GB)
<code>--server-volume=server-volume</code>	Expanded Server Root Volume (GB)
<code>--subnet=subnet</code>	Subnet Self Link ID
<code>--vpc=vpc</code>	VPC Network Self Link ID

Examples

```
$ nodus cluster:gcp:update myCluster --name newCluster
```

```
$ nodus cluster:gcp:update myCluster --prefix=myNodus --region=us-east1 --zone=us-east1-b
```

job:cancel ID CLUSTER**Description**

Cancel Submitted Job.

Usage

```
$ nodus job:cancel ID CLUSTER [OPTIONS]
```

Arguments

ID	Torque Job ID
CLUSTER	Cluster Name or ID

Option

<code>-h</code>	<code>--help</code>	Show Help
-----------------	---------------------	-----------

Example

```
$ nodus job:cancel 0.nodus-server myCluster
```

job:cloudmart:create SCRIPT NAME TYPE**Description**

Create new HPC CloudMart on-demand job entry.

Usage

```
$ nodus job:cloudmart:create SCRIPT NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

SCRIPT	(Absolute / Relative) File path to job script
NAME	Job Name
TYPE	(centos-7) Image Type

Arguments		
MANAGER		(torque slurm) HPC Workload Manager
CREDENTIAL		Credential Name or ID
Options		
-P	--package=package	Packages
-c	--cores=cores	(required) [default: 1] Number of required cores per node
-d	--description=description	Description of the job
-f	--dataFile=dataFile	(Absolute / Relative) File path to data file
-h	--help	Show Help
-n	--nodes=nodes	(required) [default: 1] Number of required nodes
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (Ex: nodus)
-r	--region=region	(required) Region ID
-w	--walltime=walltime	[default: 24:00:00] Job maximum wall-time
	--node-instance=node-instance	(required) Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	(required) On-Demand Behavior
	--server-instance=server-instance	(required) Server Instance Type ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)
Examples		
\$ nodus job:cloudmart:create ~/Documents/myScript.sh myJob centos-7 --nodes=2 --cores=16 --on-demand destroy		
\$ nodus job:cloudmart:create ~/Documents/myScript.sh myJob centos-7 --dataFile=~/Documents/data0.dat --dataFile=~/Documents/data1.dat --on-demand down		

job:cloudmart:update NAME

Description

Update HPC CloudMart on-demand job entry.

Usage

\$ nodus job:cloudmart:update NAME [OPTIONS]

Arguments

JOB Job Name

Options

-P --package=package Add Package

Options		
-R	--no-package=no-package	Remove Package
-a	--add-df=add-df	Absolute Path to any data files to be added
-c	--cores=cores	Number of required cores per node
-d	--description=description	Description of the job
-h	--help	Show Help
-n	--nodes=nodes	Number of required nodes
-p	--prefix=prefix	Source Image Root Prefix (Ex: nodus)
-r	--remove-df=remove-df	Absolute Path to any data files to be removed
-s	--script=script	Absolute File path to job script
-t	--type=centos-7	OS Type
-w	--walltime=walltime	Job walltime
	--credential=credential	Add Credential
	--manager=torque slurm	HPC Workload Manager
	--name=name	Name Identifier
	--no-credential=no-credential	Remove Credential
	--no-node-volume	Remove Expanded Node Root Volume
	--no-server-volume	Remove Expanded Server Root Volume
	--no-zone	Remove Zone
	--node-instance=node-instance	Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	On-Demand Behavior
	--region=region	Region ID
	--server-instance=server-instance	Server Instance Type ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)
	--zone=zone	Availability Zone ID

Examples

```
$ nodus job:cloudmart:update myJob --description="My Job" --script=/home/jdoe/Documents/myScript.sh
```

```
$ nodus job:cloudmart:update myJob --description="My Job" --script=/home/jdoe/Documents/myScript.sh --nodes=2 --cores=16 --walltime="00:30:00" --remove-df=/home/jdoe/Documents/file0.dat --add-df=/home/jdoe/Documents/file1.dat
```

job:create SCRIPT NAME

Description

Create new job entry. When a job is created, the location of the script will be absolute. If the script and data files are modified/deleted, the changes will be automatically included in the next running of the job.

Usage

```
$ nodus job:create SCRIPT NAME [OPTIONS]
```

Arguments

SCRIPT	(Absolute/Relative) File Path to Job Script
NAME	Job Name

Options

-c	--cores=cores	(required) [default: 1] Number of required cores per node
-d	--description=description	Description of the job
-f	--dataFile=dataFile	(Absolute/Relative) File path to data file
-h	--help	Show Help
-n	--nodes=nodes	(required) [default: 1] Number of required nodes
-w	--walltime=walltime	[default: 24:00:00] Job maximum wall-time
	--override=override	Override 'nodes', 'cores', and 'walltime' with raw flag string

Examples

```
$ nodus job:create ~/Documents/myScript.sh myJob
```

```
$ nodus job:create ~/Documents/myScript.sh myJob --nodes=2 --cores=16
```

```
$ nodus job:create ~/Documents/myScript.sh myJob --dataFile=~/Documents/data0.dat -  
-dataFile=~/Documents/data1.dat
```

```
$ nodus job:create ~/Documents/myScript.sh myJob --override "-N myJobName --  
procs=16" --dataFile=~/Documents/data0.dat
```

job:gcp:create SCRIPT NAME TYPE

Description

Update Google Cloud Platform (GCP) cluster configuration entry.

Usage

```
$ nodus job:gcp:create SCRIPT NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

SCRIPT	(Absolute / Relative) File path to job script
NAME	Job Name
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Credential Name or ID

Options

-P	--package=package	Add Package
----	-------------------	-------------

Options		
-c	--cores=cores	(required) [default: 1] Number of required cores per node
-d	--description=description	Description of the job
-f	--dataFile=dataFile	(Absolute / Relative) File path to data file
-h	--help	Show Help
-n	--nodes=nodes	(required) [default: 1] Number of required nodes
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (Ex: nodus)
-r	--region=region	(required) Region ID
-w	--walltime=walltime	[default: 24:00:00] Job maximum wall-time
-z	--zone=zone	Availability Zone ID
	--node-instance=node-instance	(required) Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	(required) On-Demand Behavior
	--server-instance=server-instance	(required) Server Instance Type ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)

Examples

```
$ nodus job:gcp:create ~/Documents/myScript.sh myJob centos-7 --nodes=2 --cores=16 --on-demand destroy
```

```
$ nodus job:gcp:create ~/Documents/myScript.sh myJob centos-7 --dataFile=~/Documents/data0.dat --dataFile=~/Documents/data1.dat --on-demand down
```

job:gcp:update NAME

Description

Update Google Cloud Platform (GCP) on-demand job entry.

Usage

```
$ nodus job:gcp:update NAME [OPTIONS]
```

Arguments

JOB	Job Name
-----	----------

Options

-P	--package=package	Add Package
-R	--no-package=no-package	Remove Package
-a	--add-df=add-df	Absolute Path to any data files to be added
-c	--cores=cores	Number of required cores per node

Options		
-d	--description=description	Description
-h	--help	Show Help
-n	--nodes=nodes	Number of required nodes
-p	--prefix=prefix	Source Image Root Prefix (Ex: nodus)
-r	--remove-df=remove-df	Absolute Path to any data files to be removed
-s	--script=script	Absolute File path to job script
-t	--type=centos-7	OS Type
-w	--walltime=walltime	Job walltime
	--credential=credential	Add Credential
	--manager=torque slurm	HPC Workload Manager
	--name=name	Name Identifier
	--no-credential=no-credential	Remove Credential
	--no-node-volume	Remove Expanded Node Root Volume
	--no-server-volume	Remove Expanded Server Root Volume
	--no-zone	Remove Zone
	--node-instance=node-instance	Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	On-Demand Behavior
	--region=region	Region ID
	--server-instance=server-instance	Server Instance Type ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)
	--zone=zone	Availability Zone ID

Examples

```
$ nodus job:gcp:update myJob --description="My Job" --script=/home/jdoe/Documents/myScript.sh
```

```
$ nodus job:gcp:update myJob --description="My Job" --script=/home/jdoe/Documents/myScript.sh --nodes=2 --cores=16 --walltime="00:30:00" --remove-df=/home/jdoe/Documents/file0.dat --add-df=/home/jdoe/Documents/file1.dat
```

job:list

Description

List job entries from database.

Usage

```
$ nodus job:list [CLUSTER] [OPTIONS]
```


Argument		
CLUSTER		Cluster Name or ID
Options		
-f	--format=text json	[default: text] Output Format
-h	--help	Show Help
Examples		
\$ nodus job:list		
\$ nodus job:list --format json		
\$ nodus job:list myCluster		

job:oracle:create SCRIPT NAME TYPE

Description		
Create new Oracle on-demand job entry.		
Usage		
\$ nodus job:oracle:create SCRIPT NAME TYPE [CREDENTIAL] [OPTIONS]		
Arguments		
SCRIPT		(Absolute / Relative) File path to job script
NAME		Job Name
TYPE		(centos-7) Image Type
MANAGER		(torque slurm) HPC Workload Manager
CREDENTIAL		Credential Name or ID
Options		
-P	--package=package	Packages
-c	--cores=cores	(required) [default: 1] Number of required cores per node
-d	--description=description	Description of the job
-f	--dataFile=dataFile	(Absolute / Relative) File path to data file
-h	--help	Show Help
-n	--nodes=nodes	(required) [default: 1] Number of required nodes
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (Ex: nodus)
-r	--region=region	(required) Region ID
-w	--walltime=walltime	[default: 24:00:00] Job maximum wall-time
	--node-instance=node-instance	(required) Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)

Options

<code>--on-demand=persist offline down destroy</code>	(required) On-Demand Behavior
<code>--server-instance=server-instance</code>	(required) Server Instance Type ID
<code>--server-volume=server-volume</code>	Expanded Server Root Volume (GB)

Examples

```
$ nodus job:oracle:create ~/Documents/myScript.sh myJob centos-7 --nodes=2 --cores=16 --on-demand destroy
```

```
$ nodus job:oracle:create ~/Documents/myScript.sh myJob centos-7 --dataFile=~/Documents/data0.dat --dataFile=~/Documents/data1.dat --on-demand down
```

job:oracle:update NAME**Description**

Update Oracle on-demand job entry.

Usage

```
$ nodus job:oracle:update NAME [OPTIONS]
```

Arguments

JOB	Job Name
-----	----------

Options

<code>-P --package=package</code>	Add Package
<code>-R --no-package=no-package</code>	Remove Package
<code>-a --add-df=add-df</code>	Absolute Path to any data files to be added
<code>-c --cores=cores</code>	Number of required cores per node
<code>-d --description=description</code>	Description of the job
<code>-h --help</code>	Show Help
<code>-n --nodes=nodes</code>	Number of required nodes
<code>-p --prefix=prefix</code>	Source Image Root Prefix (Ex: nodus)
<code>-r --remove-df=remove-df</code>	Absolute Path to any data files to be removed
<code>-s --script=script</code>	Absolute File path to job script
<code>-t --type=centos-7</code>	OS Type
<code>-w --walltime=walltime</code>	Job walltime
<code>--credential=credential</code>	Add Credential
<code>--manager=torque slurm</code>	HPC Workload Manager
<code>--name=name</code>	Name Identifier
<code>--no-credential=no-credential</code>	Remove Credential

Options

<code>--no-node-volume</code>	Remove Expanded Node Root Volume
<code>--no-server-volume</code>	Remove Expanded Server Root Volume
<code>--no-zone</code>	Remove Zone
<code>--node-instance=node-instance</code>	Node Instance Type ID
<code>--node-volume=node-volume</code>	Expanded Node Root Volume (GB)
<code>--on-demand=persist offline down destroy</code>	On-Demand Behavior
<code>--region=region</code>	Region ID
<code>--server-instance=server-instance</code>	Server Instance Type ID
<code>--server-volume=server-volume</code>	Expanded Server Root Volume (GB)
<code>--zone=zone</code>	Availability Zone ID

Examples

```
$ nodus job:oracle:update myJob --description="My Job" --script=/home/jdoe/Documents/myScript.sh
```

```
$ nodus job:oracle:update myJob --description="My Job" --script=/home/jdoe/Documents/myScript.sh --nodes=2 --cores=16 --walltime="00:30:00" -remove-df=/home/jdoe/Documents/file0.dat --add-df=/home/jdoe/Documents/file1.dat
```

job:otc:create SCRIPT NAME TYPE**Description**

Create new Open Telekom Cloud (OTC) on-demand job entry.

Usage

```
$ nodus job:otc:create SCRIPT NAME TYPE [CREDENTIAL] [OPTIONS]
```

Arguments

SCRIPT	(Absolute / Relative) File path to job script
NAME	Job Name
TYPE	(centos-7) Image Type
MANAGER	(torque slurm) HPC Workload Manager
CREDENTIAL	Credential Name or ID

Options

-P	<code>--package=package</code>	Packages
-c	<code>--cores=cores</code>	(required) [default: 1] Number of required cores per node
-d	<code>--description=description</code>	Description of the job
-f	<code>--dataFile=dataFile</code>	(Absolute / Relative) File path to data file

Options		
-h	--help	Show Help
-n	--nodes=nodes	(required) [default: 1] Number of required nodes
-p	--prefix=prefix	(required) [default: nodus] Source Image Root Prefix (Ex: nodus)
-w	--walltime=walltime	[default: 24:00:00] Job maximum wall-time
-z	--zone=zone	Availability Zone ID
	--node-instance=node-instance	(required) Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	(required) On-Demand Behavior
	--server-instance=server-instance	(required) Server Instance Type ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)
Examples		
\$ nodus job:otc:create ~/Documents/myScript.sh myJob centos-7 --nodes=2 --cores=16 --on-demand destroy		
\$ nodus job:otc:create ~/Documents/myScript.sh myJob centos-7 --dataFile=~/Documents/data0.dat --dataFile=~/Documents/data1.dat --on-demand down		

job:otc:update NAME

Description		
Update Open Telekom Cloud (OTC) on-demand job entry.		
Usage		
\$ nodus job:otc:update NAME [OPTIONS]		
Arguments		
JOB	Job Name	
Options		
-P	--package=package	Add Package
-R	--no-package=no-package	Remove Package
-a	--add-df=add-df	Absolute Path to any data files to be added
-c	--cores=cores	Number of required cores per node
-d	--description=description	Description of the job
-h	--help	Show Help
-n	--nodes=nodes	Number of required nodes
-p	--prefix=prefix	Source Image Root Prefix (Ex: nodus)
-r	--remove-df=remove-df	Absolute Path to any data files to be

Options		
		removed
-s	--script=script	Absolute File path to job script
-t	--type=centos-7	OS Type
-w	--walltime=walltime	Job walltime
	--credential=credential	Add Credential
	--manager=torque slurm	HPC Workload Manager
	--name=name	Name Identifier
	--no-credential=no-credential	Remove Credential
	--no-node-volume	Remove Expanded Node Root Volume
	--no-server-volume	Remove Expanded Server Root Volume
	--no-zone	Remove Zone
	--node-instance=node-instance	Node Instance Type ID
	--node-volume=node-volume	Expanded Node Root Volume (GB)
	--on-demand=persist offline down destroy	On-Demand Behavior
	--server-instance=server-instance	Server Instance Type ID
	--server-volume=server-volume	Expanded Server Root Volume (GB)
	--zone=zone	Availability Zone ID

Examples

```
$ nodus job:otc:update myJob --description="My Job" --script=/home/jdoe/Documents/myScript.sh
```

```
$ nodus job:otc:update myJob --description="My Job" --script=/home/jdoe/Documents/myScript.sh --nodes=2 --cores=16 --walltime="00:30:00" --remove-df=/home/jdoe/Documents/file0.dat --add-df=/home/jdoe/Documents/file1.dat
```

job:output ID CLUSTER

Description

Display stdout or stderr of from.

Usage

```
$ nodus job:output ID CLUSTER [OPTIONS]
```

Arguments

JOB	NODUS Job ID
CLUSTER	Cluster ID or Name

Options

```
-f --format=text|json [default: text] Output Format
```

-h	--help	Show Help
-t	--type=stdout stderr	[default: stdout] Output Type

Examples

```
$ nodus job:output 0 myCluster
```

```
$ nodus job:output 0 myCluster --format json
```

job:remove JOB**Description**

Remove job entries from database.

Usage

```
$ nodus job:remove JOB [OPTIONS]
```

Argument

JOB	Job Name or ID
-----	----------------

Option

-h	--help	Show Help
----	--------	-----------

Example

```
job:remove myJob
```

job:show JOB**Description**

Display job entry.

Usage

```
$ nodus job:show JOB [OPTIONS]
```

Argument

JOB	Job ID or Name
-----	----------------

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus job:show myJob
```

```
$ nodus job:show myJob --format json
```

job:submit JOB CLUSTER

Description

Submit a job to the cluster.

Usage

```
$ nodus job:submit JOB CLUSTER [OPTIONS]
```

Arguments

JOB	Job Name or ID
CLUSTER	Cluster Name or ID

Option

-h	--help	Show Help
----	--------	-----------

Example

```
$ nodus job:submit myJob myCluster
```

job:submit-od JOB

Description

Submit on-demand job.

Usage

```
$ nodus job:submit-od JOB [CREDENTIAL] [OPTIONS]
```

Arguments

JOB	Job Name or ID
CREDENTIAL	Credential Name or ID

Option

-h	--help	Show Help
----	--------	-----------

Example

```
$ nodus job:submit-od myJob myCredential
```

job:update NAME

Description

Update job configuration entry.

Usage

```
$ nodus job:update NAME [OPTIONS]
```

Argument

JOB	Job Name
-----	----------

Options

-a	--add-df=add-df	Absolute Path to any data files to be added
-c	--cores=cores	Number of required cores per node
-d	--description=description	Description of the job
-h	--help	Show Help
-n	--nodes=nodes	Number of required nodes
-r	--remove-df=remove-df	Absolute Path to any data files to be removed
-s	--script=script	Absolute File path to job script
-w	--walltime=walltime	Job walltime

Examples

```
$ nodus job:update myJob --description="My Job" --
script=/home/jdoe/Documents/myScript.sh
```

```
$ nodus job:update myJob --description="My Job" --
script=/home/jdoe/Documents/myScript.sh --nodes=2 --cores=16 --walltime="00:30:00" -
-remove-df=/home/jdoe/Documents/file0.dat --add-df=/home/jdoe/Documents/file1.dat
```

4.7 Providers

This section contains the following commands:

- [provider:instance](#)
- [provider:region](#)

provider:instance

Description

List provider instances.

WARNING Non-exhaustive list.

WARNING Availability is dependent on region and demand.

WARNING Eligibility is strictly dependent on your account.

Usage

```
$ nodus provider:instance [OPTIONS]
```

Argument

PROVIDER (aws|gcp|azure|oracle|otc) Cloud Provider

Options

-f	--format=text json	[default: text] Output Format
-h	--help	Show Help

Examples

```
$ nodus provider:instance
```

```
$ nodus provider:instance --format json
```

provider:region**Description**

List provider regions and zones.

WARNING Non-exhaustive list.

WARNING Availability and eligibility is strictly dependent on your account.

Usage

```
$ nodus provider:region [OPTIONS]
```

Argument

PROVIDER (aws|gcp|azure|oracle|otc) Cloud Provider

Options

-f --format=text|json [default: text] Output Format

-h --help Show Help

Examples

```
$ nodus provider:region
```

```
$ nodus provider:region --format json
```

4.8 Misc

This section contains the following commands:

- [daemon](#)
- [db:backup](#)
- [db:drop](#)
- [db:restore](#)
- [help](#)
- [license](#)
- [upgrade](#)

daemon

Description

Manage Master Bursting Daemon. Run 'daemon --init' to initialize the bursting daemon.

Usage

```
$ nodus daemon [OPTIONS]
```

Options

-h	--help	Show Help
	--init	Initializes and Starts Bursting Service
	--restart	Restart Bursting Service
	--start	Start Bursting Service
	--status	Check Bursting Service Status
	--stop	Stop Bursting Service

Examples

```
$ nodus daemon --init
```

```
$ nodus daemon --status
```

db:backup

Description

Backup ODDC database To File.

Usage

```
$ nodus db:backup [OPTIONS]
```

Options

-h	--help	Show Help
-u	--user=user	Username

Examples

```
$ nodus db:backup
```

```
$ nodus db:backup --user myUser
```

db:drop

Description

Drop a ODDC database.

Usage

```
$ nodus db:drop [OPTIONS]
```

Options

-h	--help	Show Help
-u	--user=user	Username

Examples

```
$ nodus db:drop
```

```
$ nodus db:drop --user myUser
```

db:restore**Description**

Restore ODDC database From Backup.

Usage

```
$ nodus db:restore [OPTIONS]
```

Options

-h	--help	Show Help
-u	--user=user	Username

Examples

```
$ nodus db:restore
```

help**Description**

Display help for ODDC.

Usage

```
$ nodus help [COMMAND]
```

Argument

COMMAND	Command to show help for
---------	--------------------------

Option

--all	See all commands in CLI
-------	-------------------------

license**Description**

Show current license information.

Usage

```
$ nodus license [OPTIONS]
```

Options

-h	--help	Show Help
	--main	Check main license

Example

```
$ nodus license
```

upgrade**Description**

Prepare or Finish Upgrade Process.

Usage

```
$ nodus upgrade [OPTIONS]
```

Options

-h	--help	Show Help
----	--------	-----------

Example

```
$ nodus upgrade
```

Appendix A: Creating Accounts for Cloud Service Providers

If you don't have credentials through your company, you can get them by following the procedures below:

- [AWS](#)
- [Google Cloud](#)
- [HPC CloudMart](#)
- [Microsoft Azure](#)
- [Open Telekom Cloud \(OTC\)](#)
- [Oracle Cloud](#)

AWS

1. Go to <https://aws.amazon.com/> and click **Create an AWS account**.
2. Follow the steps and enter your account information and then click **Continue**.
3. Select **Professional**.
4. Enter your company information and accept the customer agreement.
5. Select **Create Account and Continue**.
6. Provide a payment method and verify your phone number.

Google Cloud

1. Go to <https://cloud.google.com/> and click **Get started for free**.
2. Follow the steps on the screen to complete account registration.

HPC CloudMart

Contact Adaptive Computing at oddcsupport@adaptivecomputing.com for CloudMart account information.

Microsoft Azure

1. Go to <https://azure.microsoft.com/en-us/free/> and click **Start Free**.
2. On the **Sign in** page, click **Create one!**.
3. Click **Next** and then provide the rest of your user information (i.e., username, passwords, credit/debit, etc.).
4. Run the script **Azure.sh** to configure the Azure account.

Open Telekom Cloud (OTC)

1. Go to <https://open-telekom-cloud.com/en> and click **Login / Console**.
2. Click **Register** and follow the steps on the screen to complete account registration.

Oracle Cloud

1. Go to <https://www.oracle.com/cloud/> and sign up for a cloud account.
2. Follow the steps to complete account registration.

See the appendix [Obtaining Cloud Service Provider Account Information and Credentials](#) below for additional information.

Appendix B: Obtaining Cloud Service Provider Account Information and Credentials

Follow the checklists below to prepare accounts and get credentials for the chosen Cloud service provider:

- [AWS](#)
- [Google Cloud](#)
- [Microsoft Azure](#)
- [Open Telekom Cloud \(OTC\)](#)
- [Oracle Cloud](#)

AWS

- Create Key: IAM / Security Status / Manage Security Credentials / Access Keys / Create New Access Key
- Copy and fill Credentials File Template

Google Cloud

- Select Project / IAM / Service Accounts / Create Key (JSON)

Microsoft Azure

Azure Template Credential File

This is a template credential file for Azure:

```
{
  "client_id": " ",
  "client_secret": " ",
  "subscription_id": " ",
  "tenant_id": " "
  "region": " "
}
```

Obtaining Azure Credential Information

Follow these steps to obtain credential information.

client_id/tenant_id

1. On the Azure homepage (<https://portal.azure.com/>), under the **Azure services** section, click **App registrations**.
2. Under **Display name**, click the requested application (e.g., oddc) and copy/paste the **Application (client) ID** into the template between the quotes for the **client_id** variable, and copy/paste the **Directory (tenant) ID** into the template between the quotes for the **tenant_id** variable.
Note: If you have not previously registered an application, for this step, first click **Register an application** and enter the required information.

client_secret

Since Azure hides the client key, we can make another one that coexists with the one that is hidden by Azure.

3. Under the **Azure services** section, click **App registrations**.
4. Under **Display name**, click the requested application (e.g., oddc).
5. Under the **Manage** section on the left side of the page, click **Certificates & secrets**.
6. In the **Client secrets** section, click **New client secret**.
7. Add a **Description** and specify an **Expiration** time frame.
8. Copy/paste the **Value** into the template between the quotes for the **client_secret** variable.

subscription_id

9. Under the **Navigate** section, click **Subscriptions**.
10. Copy/paste the **Subscription ID** into the template between the quotes for the **subscription_id** variable.

region

11. Enter the desired region into the template between the quotes for the **region** variable.
For example, the region can be **eastus** or **westus**.

Open Telekom Cloud (OTC)

Create Images (see <https://open-telekom-cloud.com/en/products-services/virtual-private-cloud>):

- Create VPC network
- Create subnet on VPC
- Create stack file (where: network = {network_id} of the subnet)

Oracle Cloud

Create Images (see <https://docs.cloud.oracle.com/iaas/Content/home.htm>):

- Create VPC network
- Create subnet on VPC

Appendix C: Job Preparation

The Job script can be run using standard shell scripts such as batch, Perl, Geant4, or Python. The Job script files are stored in the `Input` folder and the results are stored in the `Output` (head node) folder.

Note: Any files that need to be accessed in multiple compute nodes must be copied or moved to the shared file system (under the user's home directory, or under `/mnt/nfs` or `/mnt/beegfs`) before attempting to reference those files with applications such as MPI.

Example Job Script

```
#!/bin/bash

set -x

cd input

PROCESSORS=1

TASKNAME=geant4

cp Dockerfile.template Dockerfile

sed -i -e "s/___FILE___/$FILE/g" Dockerfile

sed -i -e "s/___PROCESSORS___/$PROCESSORS/g" Dockerfile

sudo docker build -t $TASKNAME .

sudo docker run -v $PWD/../output:/app/output $TASKNAME
```

Example Input Dockerfile.template Contents

```
FROM ifurther/geant4

RUN mkdir BUILD

RUN mkdir output

RUN sed -i -e "s/g4root.hh/g4csv.hh/g"
/src/geant4.10.05.p01/examples/extended/radioactivedecay/rdecay01/
include/HistoManager.hh
```

```
RUN cd BUILD; cmake -DGeant4_DIR=/app/geant4.10.5.1-  
install/lib/Geant4-10.5.1/  
/src/geant4.10.05.p01/examples/extended/radioactivedecay/rdecay01/  
  
RUN cd BUILD; make -j __PROCESSORS__ rdecay01  
  
RUN cd BUILD; make install  
  
RUN . /app/geant4.10.5.1-install/bin/geant4.sh ;  
/usr/local/bin/rdecay01 ./BUILD/__FILE__.mac  
  
CMD mv *.csv output
```

Note: This program creates CSV files and moves them to the Output folder according to the script.

Appendix D: Configuring the ODDC to be a Secure Server

Follow the steps below to configure the ODDC to be a secure server.

1. Edit the file `/NODUS/config.json`:

- a.** Replace `http://` with `https://`.
- b.** Save the file.

2. Edit the file `nodus-web-api.json`:

- A.** Add the following lines to identify the server's private key location and the certificate file (this is an example):

```
{
  "https_key_file" : "/example/directory/privkey.pem",
  "https_cert_file": "/example/directory/fullchain.pem"
}
```

- B.** Save the file.

3. Enter `systemctl restart nodus-web-api` to restart the ODDC web API.

Appendix E: Credential JSON Examples

Credential JSON files should look like the below:

- [AWS](#)
- [Google Cloud](#)
- [HPC CloudMart](#)
- [Microsoft Azure](#)
- [Open Telekom Cloud \(OTC\)](#)
- [Oracle Cloud](#)

AWS

```
{  
  "access_key": "",  
  "secret_key": ""  
}
```

nodus_cluster_permissions.json

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Sid": "VisualEditor0",  
      "Effect": "Allow",  
      "Action": [  
        "ec2:AuthorizeSecurityGroupIngress",  
        "ec2:DescribeInstances",  
        "ec2:ModifyVolumeAttribute",  
        "ec2:ReplaceRouteTableAssociation",  
        "aws-marketplace:*",  
      ]  
    }  
  ]  
}
```

```

"ec2:DeleteVpcEndpoints",
"ec2:CreateKeyPair",
"ec2:AcceptTransitGatewayVpcAttachment",
"ec2:AttachInternetGateway",
"ec2:ReportInstanceState",
"route53:ListHostedZonesByName",
"ec2:UpdateSecurityGroupRuleDescriptionsIngress",
"ec2:DeleteRouteTable",
"ec2:DescribeVolumeStatus",
"ec2:StartInstances",
"ec2:CreateNetworkInterfacePermission",
"ec2:RevokeSecurityGroupEgress",
"ec2:CreateRoute",
"ec2:CreateInternetGateway",
"ec2:DescribeScheduledInstanceAvailability",
"ec2:DescribeVolumes",
"ec2>DeleteInternetGateway",
"ec2:UnassignPrivateIpAddresses",
"ec2:RejectTransitGatewayVpcAttachment",
"ec2:DescribeKeyPairs",
"ec2:DescribeReservedInstancesListings",
"ec2:DisassociateTransitGatewayRouteTable",
"ec2:ImportKeyPair",
"ec2:CreateTags",
"ec2:DescribeVpcClassicLinkDnsSupport",
"ec2:ModifyNetworkInterfaceAttribute",
"route53:ListVPCAssociationAuthorizations",
"ec2:RunInstances",
"ec2:StopInstances",
"ec2:AssignPrivateIpAddresses",
"ec2:DisassociateRouteTable",
"route53:CreateVPCAssociationAuthorization",

```

```

"ec2:DescribeVolumeAttribute",
"ec2:CreateVolume",
"ec2:ReplaceNetworkAclAssociation",
"ec2:CreateVpcEndpointServiceConfiguration",
"ec2:RevokeSecurityGroupIngress",
"ec2:CreateNetworkInterface",
"ec2:DescribeVpcEndpointServicePermissions",
"ec2:CreateTransitGatewayVpcAttachment",
"ec2:DescribeScheduledInstances",
"ec2:DescribeImageAttribute",
"ec2>DeleteNatGateway",
"ec2:DescribeReservedInstancesModifications",
"ec2:EnableTransitGatewayRouteTablePropagation",
"ec2:DescribeSubnets",
"ec2:CreateSubnet",
"ec2:ModifyVpcEndpoint",
"ec2:AttachVolume",
"ec2:DisassociateAddress",
"ec2:ModifyVpcEndpointServicePermissions",
"ec2:MoveAddressToVpc",
"ec2:CreateNatGateway",
"ec2:CreateVpc",
"ec2:DescribeVpcEndpointServices",
"ec2:DescribeSpotInstanceRequests",
"ec2:DescribeVpcAttribute",
"ec2:ModifySubnetAttribute",
"ec2:DescribeTransitGatewayRouteTables",
"ec2:DescribeAvailabilityZones",
"ec2:DescribeNetworkInterfaceAttribute",
"ec2:DescribeVpcEndpointConnections",
"ec2:DescribeInstanceStatus",
"ec2:ReleaseAddress",

```

```

"ec2:RebootInstances",
"ec2:ModifyInstanceMetadataOptions",
"ec2:AssignIpv6Addresses",
"ec2:AcceptVpcEndpointConnections",
"route53:ListHostedZones",
"ec2:DescribeClassicLinkInstances",
"ec2:DisassociateSubnetCidrBlock",
"ec2:DescribeVpcEndpointConnectionNotifications",
"ec2:DescribeSecurityGroups",
"ec2>DeleteVpcEndpointConnectionNotifications",
"ec2:RestoreAddressToClassic",
"ec2:DescribeVpcs",
"ec2:DisableVpcClassicLink",
"ec2:DisableVpcClassicLinkDnsSupport",
"ec2:ModifyVpcTenancy",
"route53:AssociateVPCWithHostedZone",
"ec2:DescribeStaleSecurityGroups",
"ec2>DeleteSubnet",
"ec2:ModifyVpcEndpointServiceConfiguration",
"ec2:UnmonitorInstances",
"ec2:DetachClassicLinkVpc",
"ec2:MonitorInstances",
"ec2:CreateTransitGatewayRouteTable",
"route53:GetHostedZone",
"ec2:DescribeVolumesModifications",
"ec2:AssociateVpcCidrBlock",
"ec2:ReplaceRoute",
"ec2:AssociateRouteTable",
"ec2:DisassociateVpcCidrBlock",
"ec2:DescribeInternetGateways",
"ec2>DeleteVolume",
"ec2>DeleteTransitGatewayVpcAttachment",

```



```

"ec2:ReplaceNetworkAclEntry",
"ec2:AssociateTransitGatewayRouteTable",
"ec2:UnassignIpv6Addresses",
"ec2:DescribeNetworkInterfacePermissions",
"ec2:DescribeReservedInstances",
"route53:UpdateHostedZoneComment",
"ec2:RejectVpcEndpointConnections",
"ec2:DescribeNetworkAcls",
"ec2:DescribeRouteTables",
"ec2:EnableVpcClassicLink",
"ec2:DescribeEgressOnlyInternetGateways",
"ec2:DetachVolume",
"ec2:ModifyVolume",
"ec2:DisableTransitGatewayRouteTablePropagation",
"ec2:UpdateSecurityGroupRuleDescriptionsEgress",
"ec2:CreateVpcEndpointConnectionNotification",
"route53:DisassociateVPCFromHostedZone",
"ec2:ResetNetworkInterfaceAttribute",
"ec2:DescribeReservedInstancesOfferings",
"ec2:CreateRouteTable",
"ec2>DeleteNetworkInterface",
"ec2:DescribeFleetInstances",
"ec2:DetachInternetGateway",
"ec2:DescribeVpcEndpointServiceConfigurations",
"ec2:ModifyVpcEndpointConnectionNotification",
"ec2:DescribeInstanceCreditSpecifications",
"ec2:DescribeVpcClassicLink",
"ec2>DeleteTransitGatewayRouteTable",
"route53>DeleteVPCAssociationAuthorization",
"ec2:GetTransitGatewayRouteTablePropagations",
"ec2:AssociateSubnetCidrBlock",
"route53:ChangeTagsForResource",

```

```

"ec2:DeleteVpc",
"ec2:CreateEgressOnlyInternetGateway",
"ec2:DescribeVpcEndpoints",
"ec2:AssociateAddress",
"ec2:DeleteKeyPair",
"ec2:DescribeAddresses",
"ec2:DeleteTags",
"ec2:DescribeInstanceAttribute",
"ec2:DeleteVpcEndpointServiceConfigurations",
"ec2:DeleteNetworkInterfacePermission",
"ec2:DescribeNetworkInterfaces",
"ec2:CreateSecurityGroup",
"ec2:CreateNetworkAcl",
"ec2:ModifyVpcAttribute",
"ec2:ModifyInstanceAttribute",
"ec2:GetTransitGatewayRouteTableAssociations",
"ec2:AuthorizeSecurityGroupEgress",
"ec2:AttachClassicLinkVpc",
"ec2:ModifyTransitGatewayVpcAttachment",
"ec2>DeleteEgressOnlyInternetGateway",
"ec2:TerminateInstances",
"ec2:DetachNetworkInterface",
"ec2:DescribeIamInstanceProfileAssociations",
"ec2:DescribeTags",
"ec2>DeleteRoute",
"ec2:DescribeNatGateways",
"ec2:AllocateAddress",
"ec2:DescribeImages",
"ec2:DescribeSpotFleetInstances",
"ec2:CreateVpcEndpoint",
"ec2>DeleteSecurityGroup",
"ec2:AttachNetworkInterface",

```

```

        "ec2:EnableVpcClassicLinkDnsSupport",
        "ec2:DescribeTransitGatewayVpcAttachments",
        "ec2:CreateNetworkAclEntry"
    ],
    "Resource": "*"
}
]
}

```

nodus_stack_permissions.json

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "ec2:ModifyVolumeAttribute",
        "ec2:DescribeInstances",
        "ec2:UnmonitorInstances",
        "aws-marketplace:*",
        "ec2:MonitorInstances",
        "ec2:CreateKeyPair",
        "ec2:DescribeVolumesModifications",
        "ec2:CreateImage",
        "ec2:DescribeSnapshots",
        "ec2>DeleteVolume",
        "ec2:DescribeVolumeStatus",
        "ec2:ModifySnapshotAttribute",
        "ec2:StartInstances",
        "ec2:DescribeVolumes",

```

```
"ec2:DescribeKeyPairs",
"ec2:DetachVolume",
"ec2:ResetImageAttribute",
"ec2:DisableFastSnapshotRestores",
"ec2:ImportKeyPair",
"ec2:CreateTags",
"ec2:RegisterImage",
"ec2:RunInstances",
"ec2:DescribeFastSnapshotRestores",
"ec2:StopInstances",
"ec2:CreateVolume",
"ec2:DescribeVolumeAttribute",
"ec2:CreateSnapshots",
"ec2:DescribeImageAttribute",
"ec2>DeleteKeyPair",
"ec2:AttachVolume",
"ec2:DeregisterImage",
"ec2>DeleteSnapshot",
"ec2>DeleteTags",
"ec2:DescribeInstanceAttribute",
"ec2:DescribeRegions",
"ec2:ModifyImageAttribute",
"ec2:EnableFastSnapshotRestores",
"ec2:CreateSnapshot",
"ec2:DescribeInstanceStatus",
"ec2:RebootInstances",
"ec2:TerminateInstances",
"ec2:DescribeTags",
"ec2:ResetSnapshotAttribute",
"ec2:DescribeSecurityGroups",
"ec2:DescribeImages",
"ec2:DescribeStaleSecurityGroups"
```

```

    ],
    "Resource": "*"
  }
]
}

```

Google Cloud

```

{
  "type": "",
  "project_id": "",
  "private_key_id": "",
  "private_key": ""
}

```

HPC CloudMart

```

{
  "user_ocid" : "<user_content_of_View_Configuration_file>",
  "tenancy_ocid" : "<tenancy_content_of_View_Configuration_file>",
  "compartment_ocid" : "<content_of_subnet_OCID>",
  "availability_domain": "<content_of_Availability_Domain>",
  "key_file" : "-----BEGIN RSA PRIVATE KEY-----\n<content_of_private-key.pem>\n-----END RSA PRIVATE KEY-----",
  "fingerprint" : "<fingerprint_content_of_View_Configuration_file>"
}

```

Note: The key_file must not have any spaces and each line must start with \n, and the key_file must end with \n.

Microsoft Azure

```
{
  "client_id": "",
  "client_secret": "",
  "subscription_id": "",
  "tenant_id": ""
}
```

Open Telekom Cloud (OTC)

```
{
  "username" : "",
  "password" : "",
  "domain_name": "",
  "tenant_name": ""
}
```

Oracle Cloud

```
{
  "user_ocid" : "<user_content_of_View_Configuration_file>",
  "tenancy_ocid" : "<tenancy_content_of_View_Configuration_file>",
}
```

```
"compartment_ocid" : "<content_of_subnet_OCID>",  
  
"availability_domain": "<content_of_Availability_Domain>",  
  
"key_file" : "-----BEGIN RSA PRIVATE KEY-----\n<content_of_  
private-key.pem>\n-----END RSA PRIVATE KEY-----",  
  
"fingerprint" : "<fingerprint_content_of_View_Configuration_file>"  
  
}
```

Note: The key_file must not have any spaces and each line must start with \n, and the key_file must end with \n.

Appendix F: Writing ODDC Scripts

The environment variables below can be used to write ODDC scripts.

Variable	Description
ENVIRONMENT	BATCH
HISTCONTROL	ignoredups
HISTSIZE	1000
HOME	/home/linux
HOSTNAME	nodus-c3-15xlarge-2-node-0
LANG	en_US.UTF-8
LESSOPEN	/usr/bin/lesspipe.sh %s
LOADED_MODULES	
LOGNAME	linux
MAIL	/var/spool/mail/linux
MODULEPATH	/usr/share/Modules/modulefiles:/etc/modulefiles
MODULESHOME	/usr/share/Modules
NO_PROXY	127.0.0.1,localhost,169.254.169.254
PATH	/usr/local/bin:/usr/local/sbin:/bin:/usr/bin:/usr/local/sbin:/usr/sbin:/home/linux/.local/bin:/home/linux/bin
PBS_ENVIRONMENT	PBS_BATCH
PBS_GPUFILE	/var/spool/torque/aux//63.nodus-servergpu
PBS_JOBCOOKIE	E36EE984E318A76438F6A28EB31CC5DB
PBS_JOBID	63.nodus-server
PBS_JOBNAME	env
PBS_MICFILE	/var/spool/torque/aux//63.nodus-servermic
PBS_MOMPORT	15003
PBS_NODEFILE	/var/spool/torque/aux//63.nodus-server
PBS_NODENUM	0
PBS_NP	1
PBS_NUM_NODES	1
PBS_NUM_PPN	1
PBS_O_HOME	/home/linux
PBS_O_HOST	nodus-server
PBS_O_INITDIR	/NODUS/jobs/63
PBS_O_LANG	en_US.UTF-8

Variable	Description
PBS_O_LOGNAME	linux
PBS_O_MAIL	/var/spool/mail/linux
PBS_O_PATH	/usr/local/bin:/usr/local/sbin:/usr/local/maui/bin:/usr/local/maui/sbin:/usr/local/bin:/bin:/usr/bin:/usr/local/sbin:/usr/sbin
PBS_O_QUEUE	batch
PBS_O_SERVER	nodus-server
PBS_O_SHELL	/bin/sh
PBS_O_WORKDIR	/NODUS/jobs/63
PBS_QUEUE	batch
PBS_TASKNUM	1
PBS_VERSION	TORQUE-6.1.3
PBS_VNODENUM	0
PBS_WALLTIME	3600
PWD	/NODUS/jobs/63
SHELL	/bin/bash
SHLVL	2
USER	linux

Appendix G: Performing a Clean Installation of the ODDC

Follow the steps below to clean the ODDC off the server and perform a fresh install.

1. On the server, run `sudo /NODUS/nodus cluster:list` to see all the clusters that are still available for all the users.
2. Destroy all available clusters, because this process will clean out all clusters. Otherwise, you will have instances running on the Cloud service provider that you cannot destroy from the ODDC.
3. Save the `/NODUS/.license`, `/NODUS/config.json`, and the `/NODUS/nodus-web-api.json` files (if they exist), in your home directory for possible later use.

Appendix H: HPC Cloud On-Demand Data Center Workload Considerations

Listed below are some of the considerations and questions that need to be answered when running workloads in the Cloud. All of these may not be appropriate for each of your workloads, or you may have additional considerations.

1. Not all workloads can be effectively run in the Cloud, so each workload should be tested for Cloud appropriateness by running them in the Cloud and collecting statistics.
 - The ODDC has multi-cloud capabilities and will run workloads in the Cloud from the UI on one or more of the major Cloud service providers.
 - Run your workload on each Cloud service provider and record the statistics for each one in the Cloud Business Value Process Benchmark Report (see Figure 1 Benchmarks).
 - HTC (High-Throughput Computing) workloads and small scale HPC (High-Performance Computing) workloads that are not memory-, communication-, large-data-, or bandwidth-intensive are ideal for migrating to Cloud environments.

2. Cloud OpEx costs are difficult to model without actually running the workload in the Cloud.
 - Build a cost and performance model based on the statistics from (Figure 1) below. This will help you develop your Cloud budget. See the attached [Cloud Business Value Process Spreadsheet](#). It is also available on adaptivecomputing.com (Support Portal Home / Product Download / Adaptive HPC Cloud On-Demand Data Center).

 Adaptive COMPUTING Cloud Business Value Process-Benchmarks Your Company									
WORKLOAD NAME	CLOUD PROVIDER	INSTANCE TYPE	NUMBER OF NODES	NUMBER OF CPUs PER NODE	PERFORMANCE	COST PER CPU HOUR	JOB RUN FREQUENCY PER MONTH	MONTHLY COST	TOTAL COST
	AWS								
	AWS								
	AWS								
	Google Cloud								
	Google Cloud								
	Google Cloud								
	Azure								
	Azure								
	Azure								
	Oracle Cloud								
	Oracle Cloud								
	Oracle Cloud								
WORKLOAD NAME	CLOUD PROVIDER	INSTANCE TYPE	NUMBER OF NODES	NUMBER OF CPUs PER NODE	PERFORMANCE	COST PER CPU HOUR	JOB RUN FREQUENCY PER MONTH	MONTHLY COST	TOTAL COST
	AWS								
	AWS								
	AWS								
	Google Cloud								
	Google Cloud								
	Google Cloud								
	Azure								
	Azure								
	Azure								
	Oracle Cloud								
	Oracle Cloud								
	Oracle Cloud								

Figure 1: Benchmarks

- When using the ODDC, expenditures move from capital expenses (CapEx) to operational expenses (OpEx) because you are gaining capacity without buying more hardware. This allows for the immediate availability of resources and the ability to scale up clusters instantaneously.
3. Choose the appropriate delivery model for each of your workloads.
- The ODDC accommodates on-premises, even your laptop, SaaS, and cloud-hosted delivery models, which are very easy to implement because of the low barriers to entry using the ODDC. The delivery model for each workload could be different.

4. Do you want to run your workload independent from or as part of your on-premises cluster?
- The ODDC offers superior usability when devising Cloud and on-premises integration strategies.
 - The ODDC works with any HPC or enterprise job scheduler or without a workload scheduler.
 - In the diagram below (Figure 2) the workload has burst to the Cloud automatically based on backlog. The job is seen by the scheduler as part of the on-premises cluster. The ODDC has the capability to seamlessly integrate on-premises and Cloud resources. See the section [Bursting Service](#) for additional information.




Figure 2: Cloud Becoming On-Premise

- In the diagram below (Figure 3) the workload has burst to the Cloud on demand and is running in the Cloud independently of the on-premises cluster.




Figure 3: Bursting Off-Premise

5. Which workloads can be run in the Cloud on VMs and which need bare metal?
 - Can you get the performance you need from VMs?
 - Virtualization layers can often cause slower performance, although VMs can be easily moved to and from similar Clouds and are less expensive.
 - You may have to run your workload on bare metal to get the performance results required.

6. Which workloads require additional licenses to run in the Cloud and which do not?
 - Determine if your workload requires additional licensing to run in the Cloud.
 - License sharing can be difficult in international organizations and will increase costs. License sharing is easier to manage with regional workloads, making these better suited to run in the Cloud and are less expensive.
 - Use vendors or open-source software with cloud-friendly terms.
7. Which of your workloads require moving large amounts of data to the Cloud and which do not?
 - Look at each job, the amount of data it requires, and then make a determination.
 - Consider compute, memory, and storage when assessing data requirements.
 - Test your data transfer line.
 - Move your data first and then spin up the nodes to run your workload in the Cloud.
 - The ODDC has a data movement feature, which can automatically move the data and proliferate it.
8. Which public Cloud service providers match which of your workloads? One size does not fit all.
 - Using the ODDC, you can test your workloads on each one of the Cloud service providers. Within the ODDC interface, select the Cloud service provider that your job will run on. The nodes will shut down when the job completes, and you can then select another Cloud service provider to test your workload on. For example, one workload should run on AWS and another should run on Google Cloud.
 - The ODDC supports all major public Clouds, and can be configured to support any public Cloud of any size.
 - After calculating the cost for a workload, multiply that cost by the number of times the workload runs each month. Repeat this for each Cloud service provider (Figure 1).

- A blank worksheet is provided if you want to test other Cloud service providers (Figure 4). See the attached *Cloud Business Value Process Spreadsheet*.

[illegible]

Figure 4: Worksheet

- A budget sheet is included for you to summarize your testing results and prepare your Cloud Business Value Process Budget Report for management (Figure 5). See the attached [Cloud Business Value Process Spreadsheet](#).

 Cloud Business Value Process - Budget Your Company									
WORKLOAD NAME	CLOUD PROVIDER	INSTANCE TYPE	NUMBER OF NODES	NUMBER OF CPUS PER NODE	PERFORMANCE	COST PER CPU HOUR	JOB RUN FREQUENCY PER MONTH	MONTHLY COST	TOTAL COST

Figure 5: Budget

Test your workloads in the Cloud one at a time (the most cloud-friendly ones first) and begin collecting statistics.

Enter the gathered statistical data into the attached spreadsheet. In a very short period, you will know which of your workloads can be run in the Cloud, and on which Cloud service provider using which instance type. We recommend that you run each workload using at least three different instance types for each Cloud service provider to determine the best cost/performance for each workload (Figure 1).

When testing is concluded, your Cloud budget will be determined.

Appendix I: Data Staging to the Cloud

Data Staging with ODDC can be accomplished in multiple ways. See three examples below.

1. When a cluster is created or shared with a user within the ODDC, an SSH key pair is created and available to the user to transfer files via the SSL protocol to the head node of the cluster. ¹
2. The ODDC allows for the creation of an OpenVPN server on the head node. An OVPN file can be retrieved and used to create a VPN connection between on-premises resources and the cluster head node. ¹
3. The compute nodes are created on a private network behind the head node. Depending on the configuration within the Cloud provider, these nodes have the ability to pull data from external sources, such as provider-created storage services, or any other IP address.

¹ A directory `/NODUS/share` exists on the head node and is shared across all compute nodes via NFS or BeeGFS (default: NFS). Placing data files in this directory will make them visible to all compute nodes when the job is executed.

Please contact us at oddcsupport@adaptivecomputing.com if you need a different Data Staging method.

Appendix J: Using Client–Server Directory Service Protocol Setups

When adding ODDC to client–server directory service protocol networks (such as NIS/YP, LDAP, etc.), you need to make operating system modifications. ODDC uses the Linux command `useradd` to create any new accounts.

Make the following changes to `/etc/default/useradd`:

1. Add **UID_MAX (number)** and **UID_MIN (number)** to the file so that any ODDC created user does not conflict with existing UIDs.
2. Change the **HOME (string)** information so that the `HOME` directory is local and does not conflict with the existing home directories. The **`useradd`** command defaults to `/HOME`.

Appendix K: Modifying Provider Information To Support Customer Requirements

A Provider file, which is included in the ODDC installation file, enables you to populate your specific needs to list providers, regions, and instance types.

Example

```
{
  "providers": [
    {
      "id": "local",
      "regions": [],
      "instanceSizes": []
    },
    {
      "id": "aws",
      "regions": [
        {
          "name": "US East 1",
          "region_id": "us-east-1",
          "zones": [
            {
              "name": "US East 1a",
              "zone_id": "us-east-1a"
            },
            {
              "name": "US East 1b",
              "zone_id": "us-east-1b"
            },
            {
              "name": "US East 1c",
              "zone_id": "us-east-1c"
            },
            {
              "name": "US East 1d",
              "zone_id": "us-east-1d"
            },
            {
              "name": "US East 1e",
              "zone_id": "us-east-1e"
            },
            {
              "name": "US East 1f",
              "zone_id": "us-east-1f"
            }
          ]
        },
        {
          "name": "US West 1",
          "region_id": "us-west-1",
          "zones": [
            {
              "name": "US West 1a",
```

```

        "zone_id": "usw1-az3"
      },
      {
        "name": "US West 1b",
        "zone_id": "usw1-az1"
      }
    ]
  },
  {
    "name": "US Gov East 1",
    "region_id": "us-gov-east-1",
    "zones": [
      {
        "name": "US Gov East 1a",
        "zone_id": "usge1-az1"
      },
      {
        "name": "US Gov East 1b",
        "zone_id": "usge1-az2"
      },
      {
        "name": "US Gov East 1c",
        "zone_id": "usge1-az3"
      }
    ]
  },
  {
    "name": "CA Central 1",
    "region_id": "ca-central-1",
    "zones": [
      {
        "name": "CA Central 1a",
        "zone_id": "cac1-az1"
      },
      {
        "name": "CA Central 1b",
        "zone_id": "cac1-az2"
      },
      {
        "name": "CA Central 1d",
        "zone_id": "cac1-az4"
      }
    ]
  },
  {
    "name": "EU Central 1",
    "region_id": "eu-central-1",
    "zones": [
      {
        "name": "EU Central 1a",
        "zone_id": "eucl-az2"
      },
      {
        "name": "EU Central 1b",
        "zone_id": "eucl-az3"
      },
      {
        "name": "EU Central 1c",

```

```

        "zone_id": "eucl-az1"
      }
    ],
    {
      "name": "EU West 1",
      "region_id": "eu-west-1",
      "zones": [
        {
          "name": "EU West 1a",
          "zone_id": "euw1-az3"
        },
        {
          "name": "EU West 1b",
          "zone_id": "euw1-az1"
        },
        {
          "name": "EU West 1c",
          "zone_id": "euw1-az2"
        }
      ]
    }
  ],
  "instanceSizes": [
    {
      "name": "t2.nano - vCPU: 1, Mem (GB): 0.50",
      "instance_id": "t2.nano",
      "cores": "1"
    },
    {
      "name": "t2.micro - vCPU: 1, Mem (GB): 1",
      "instance_id": "t2.micro",
      "cores": "1"
    },
    {
      "name": "t2.small - vCPU: 1, Mem (GB): 2",
      "instance_id": "t2.small",
      "cores": "1"
    },
    {
      "name": "t2.medium - vCPU: 2, Mem (GB): 4",
      "instance_id": "t2.medium",
      "cores": "2"
    },
    {
      "name": "t2.large - vCPU: 2, Mem (GB): 8",
      "instance_id": "t2.large",
      "cores": "2"
    },
    {
      "name": "t2.xlarge - vCPU: 4, Mem (GB): 16",
      "instance_id": "t2.xlarge",
      "cores": "4"
    },
    {
      "name": "t2.2xlarge - vCPU: 8, Mem (GB): 32",
      "instance_id": "t2.2xlarge",
      "cores": "8"
    }
  ]
}

```

```
    }  
  ]  
}
```

Appendix L: Creating a Private Oracle Cluster

1. Install ODDC onto an instance in OCI.
2. Identify the CIDR of the ODDC VCN and also the CIDR of the cluster to be created.
3. Create an LPG (Local Peering Gateway) where the VCN ODDC is installed using the CIDR of the cluster to be created.
4. Create and attach a route to the LPG.
5. Create and deploy your cluster with the OCID of the LPG, CIDR of the cluster to be created for VCN_CIDR, and CIDR of the VCN ODDC.

If the cluster needs to be redeployed for any reason, undo and then redo steps 3 and 4.

Glossary

BeeGFS: A parallel file system developed and optimized for HPC and includes a distributed metadata architecture for scalability and flexibility.

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

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, which 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.

Custom Job: A job that is customizable and configurable.

Flexible Shape: A flexible shape (flex shape) is a shape that lets you customize the number of OCPUs and the amount of memory when launching or resizing your VM. When you create a VM instance using a flexible shape, you select the number of OCPUs and the amount of memory that you need for the workloads that run on the instance. The network bandwidth and number of VNICs scale proportionately with the number of OCPUs. This flexibility lets you build VMs that match your workload, enabling you to optimize performance and minimize cost.

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

High Performance Computing: The use of highly parallel and/or specialized supercomputers for executing parallel workloads such as large simulations, solving problems that require very complex and extensive calculations, computations that require very long running calculations, etc. Such workloads are characterized by their use of many compute nodes, often in the thousands, to work on a single problem and have execution times ranging from minutes to months. HPC systems often execute from one to a few dozen or hundreds of simultaneous workloads and have a job (workload) queue with a few hundred to several thousands of pending jobs. In HPC systems, the performance of individual workloads within a time interval is the primary objective and therefore HPC schedulers attempt to optimize their use of an HPC system's resources regardless of the scheduling overhead incurred to do so (within reason).

High Throughput Computing: Describes workloads that often execute on just a single core and may have execution times ranging from sub-seconds to minutes and perhaps

hours. HTC systems often execute hundreds to tens of thousands of simultaneous workloads. In HTC systems, the quantity of workloads processed per time interval is the primary objective and therefore HTC schedulers attempt to minimize scheduling overhead in order to maximize workload throughput.

Hyper-Threading: Term used by Intel for its Simultaneous Multi-Threading (SMT) capability in its Atom, Core, Itanium, Pentium 4, Xeon, and Xeon Phi processor families. See also Simultaneous Multi-Threading.

Idle Purge Time: Specifies how long to keep a compute node available after a job (workload) is processed, before destroying the compute node when using bursting.

Image: A snapshot of an OS.

Job: HPC term for workload submitted by a user to a scheduler for the purpose of scheduling resources on which the workload executes when started up by the scheduler. It identifies the workload, in whatever form, submitted to a scheduler that schedules the workload for execution on a system (HPC cluster or commercial datacenter). 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. The user also gives information identifying the types of resources, typically one or more nodes and optionally other hardware (such as GPU or MIC accelerators) or software and/or software licenses, required by the workload to execute, either in the job script itself or at the time of job submission (for example, via command line options or web portal form). The scheduler schedules the job for the requested resources and when they are available allocates them to the job and then starts the job by executing the script on one of the allocated nodes. The script executes the workloads/applications that then use the resources allocated to the job by the scheduler.

Job Scheduler: HPC term for a scheduler that manages submitted workloads (called jobs) for an HPC cluster.

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

Multi-threading: The use of multiple software threads, which may or may not be pinned to hardware threads (core affinity), to implement processing in parallel. There are multiple implementations of multi-threading, such as Linux pthreads, etc.

OCPU (Oracle CPU): Provides the CPU capacity equivalent of one physical core of an Intel Xeon processor with hyperthreading enabled. Each OCPU corresponds to two hardware execution threads, known as vCPUs.

Process: An individual executing program managed by an operating system. It has its own resources and memory address space, independent of all other executing processes managed by the operating system. A process may itself be multi-threaded, which means the operating system can execute simultaneously different software execution threads of the process.

Processor: A physical hardware chip (sometimes called a socket); regardless of whether it supports a single core or multiple cores (multi-core processor). Socket is a strong, unambiguous synonym for processor while CPU is an ambiguous synonym individuals and/or processor vendor documentation may use. In addition, people and literature sometimes use the term processor to refer to a hardware core or hardware thread.

Provisioning: The event of configuring a node or cluster with its stack and getting it into a ready-to-work state.

Scheduler: The specialized software between the user and the HPC cluster/datacenter system that manages submitted workloads or jobs. This includes queuing jobs, prioritizing queued jobs for execution, scheduling and allocating requested resources for each job, and starting jobs when their requested resources become available and the jobs have the highest priority.

Shape: A template that determines the number of OCPUs, amount of memory, and other resources that are allocated to an instance.

Socket: A placeholder for the CPUs that are attached with a series of pins to a computer's motherboard.

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

Thread: In SMT or hyper-threading, refers to a hardware-based thread execution capability. For example, the consumer-oriented Intel Core i7 processor has four cores, each of which has hardware that can simultaneously track two software execution thread states and execute the other thread when one thread blocks waiting on a memory access; thus increasing the utilization of each core's computational capability and yielding 8 hardware-based threads for the entire processor. The server-oriented Intel Xeon processor documentation refers to threads as logical processors. Cray documentation uses the term threads relative to the SMT capability of the Intel Xeon processors in its XC systems. Regardless of vendor terminology, one thread in the context of SMT refers to the hardware capability for tracking and executing one software execution thread. The term threads used in the context of a processor core refers to the quantity of software execution threads the core can simultaneously track; for example, 2 threads per Intel Xeon E5-2650 v3 core and 20 threads per Intel Xeon E5-2650 v3 processor. BIOS settings enable or disable the SMT capability of SMT-capable processors and therefore determine at boot time whether a processor has only one thread per core or multiple threads per core.

Torque (Terascale Open-source Resource and Queue manager): An industry-standard resource manager solution with a built-in scheduler.

vCPU (virtual centralized processing unit): Represents a portion or share of the underlying physical CPU that is assigned to a particular VM.

Walltime: The job's time limit (HH:MM:SS).

Workload: Jobs to be run and/or jobs in the queue.

Index

- % -

% completion bar 33

- 5 -

500 response code 48

- A -

Access key 11

access_key 141

Account 20, 65, 88, 114, 135

account-scripts_azure.sh 12

Account Information and
Credentials 135

Accounting 38

Accounting Manager 38

Accounts 12, 133, 135

Actions (CLI commands) 49

Actions (cluster) 30

Actions (credential) 26

Actions (file) 37

Actions (stack) 27

Activities Log 35

Adding a New Job 34

Adding a New User 21

Adding the Server License Key 17

Admin 12, 38, 51

Admin Commands 47

Administrator role 20

Advanced (button) 28

All Nodes (node state) 35

Amazon Web Services (AWS) 10, 58,
78

API 10, 17, 140

API Key 13

Applications 22, 25-26, 30, 34, 37-39

Architecture 42

Availability Domain 63, 84

Availability Zone 65, 79, 111

availability_domain 149

Available (cluster state) 30

Available (node state) 35

Available Memory 36

Available Nodes 35

AWS 10, 33, 58, 78, 111, 159

AWS Account 133

Azure 10, 60, 81, 114

Azure Account 134

Azure services 136

Azure Template Credential File 135

- B -

Balanced persistent disks 88

BeeGFS 162, 169

Blocked Jobs 35

Briefcase 37

Budget 40, 161

Build GPUs 27

Build Node 27

Build Server 27

Builder 27

Building a GPU Stack 48

Building a Stack 28

Building Customized Stacks 29

Building Images 43

Building Stacks into Images 28

Burst cycle 36

Burst Once 37

Bursting 6, 30, 37, 79, 130, 157, 169

Bursting (cluster) 30

Bursting Configurations 37

Bursting Daemon 130

Bursting Service 31, 37, 130, 157

Busy (node state) 35

Busy Nodes 35

- C -

CapEx 156

CentOS 9

CIDR 168

CLI 5, 24, 41, 43-44, 49, 85, 131

CLI Commands 41, 44, 49

CLI Procedures 41

CLI Usage 43

client_id 135, 150

client_secret 135, 150

Cloud Budget Analysis 40

Cloud Business Value Process
Benchmark Report 155

Cloud Business Value Process
Spreadsheet 6, 156

Cloud Service Providers 10, 25, 27,
32, 133, 155

CloudMart 10, 54

Cluster 20, 30, 35, 37-38, 43, 49, 168-
169

Cluster Actions 43

Cluster Compute Nodes 37

Cluster Deploy (command) 44

Cluster Destroy (command) 44

Cluster ID (Cluster Information) 37

Cluster Info 35

Cluster Information 31, 37

Cluster IP (Cluster Information) 37

Cluster Jobs 46

Cluster Jobs Queue (command) 46

Cluster List (command) 46

Cluster Management 44

Cluster Manager 20, 30

Cluster Manager screen 30, 32

Cluster Members (command) 46

Cluster Monitoring 46

Cluster pbsnodes (command) 46

Cluster Show (command) 46

Cluster SSH (command) 45

Cluster SSH Key 31

Cluster SSH Key (cluster action) 30

Clusters (CLI commands) 76

Clusters tab 38

Collaborators 26-27, 30

Commands 44

Connect to a cluster 45

Delete a user from the system 47

Deploy nodes on a cluster 44

Display a description of compute
nodes on the cluster 46

Display JSON information about the
jobs queue 46

Display serf members from the
head node 46

Get general information about the
cluster 46

List all user clusters 46

Run a local job script on a
cluster 47

Show standard error messages 43

Compartment [14, 149](#)
compartment_ocid [149](#)
Completed jobs [35, 79](#)
Compute node [29, 31, 58](#)
Compute Node Instance Sizes [33](#)
Compute Node Sizes [31](#)
Compute Nodes [31, 37, 46, 85, 138, 162, 169](#)
config.json [17, 140, 154](#)
Configuration [13, 30, 37](#)
Configuring the ODDC to be a Secure Server [140](#)
Connecting to a Cluster [43](#)
Copyright [5](#)
Cores [9, 23, 34, 38](#)
Cores (job) [34](#)
Cost [38](#)
Costs [39](#)
Count (Cluster Compute Nodes) [37](#)
CPU Clock [36](#)
CPUs [33, 36, 39](#)
Create an account [20](#)
Create Images [136](#)
Created (Cluster Information) [37](#)
Creating a Cluster [32](#)
Creating a Google Cloud Account [133](#)
Creating a GPU Job [48](#)
Creating a Microsoft Azure Account [134](#)
Creating a Private Oracle Cluster [168](#)
Creating a Stack [27](#)
Creating Accounts for Cloud Service Providers [133](#)
Creating an AWS Account [133](#)
Creating an HPC CloudMart

Account [133](#)
Creating an ODDC Account [20](#)
Creating an Oracle Cloud Account [134](#)
Creating an OTC Account [134](#)
Credential JSON Examples [141](#)
Credentials [10, 20, 25, 27, 29, 32, 43, 53, 133, 135](#)
Credentials (CLI commands) [53](#)
Credentials Manager [20, 25](#)
Credentials Manager screen [26](#)
Credentials Path [53-54, 58, 112](#)
CSV [139](#)
Custom Job [169](#)

- D -

daemon (CLI command) [130](#)
Dashboard [22](#)
Data Files (job) [34](#)
Data Files (panel) [34](#)
Data Staging [162](#)
dataFiles [28](#)
Date/Time Created (credential) [26](#)
Dedicated Cores [36](#)
Dedicated Nodes [36](#)
Dedicated Sockets [36](#)
Dedicated Threads [36](#)
Delete [26-27, 30, 34, 38, 94](#)
Delete (file) [38](#)
Delete (job) [34](#)
Delete User (command) [47](#)
Deleting a Cluster [33](#)
Deploy [35, 39](#)
Deploy (cluster action) [30](#)
Deploy button [31](#)
Deploy Cluster [44](#)

Deploy Time 31
Deploying (node state) 36
Deploying a Cluster 33
Description 14, 34, 49, 51, 54, 58, 78,
111, 128, 130, 136, 152
Destroy 30, 35, 39, 87, 154
Destroy button 31
Destroying (node state) 36
Destroying a Cluster 43
Details 35, 38, 83
disks 78
Disks Information 31, 37
docker 138
Docker 48
Dockerfile 138
domain_name 150
Down (cluster state) 30
Down (node state) 36
Down Nodes 35
Download 12, 17, 20, 37-38, 156
Download standard error 38
Download standard output 38
Download STD-Err (file) 38
Download STD-Out (file) 38

- E -

EC2 11
Edit 26-27, 30, 34, 140
Edit Collaborators 31
Editing User Profiles 20
Egress Rules 14
ellipsis 31, 38
Email 51
environment variables 152
Example Input Dockerfile.template

Contents 138
Example Job Script 138
Expiration 136

- F -

Failed (stack built) 27
File Manager 20, 31, 37
File Manager screen 37
fingerprint 149
firewall 88, 114
First Name 51
flex shapes 84
flexible shape 85
Floating IP Address 101

- G -

Geant4 138
Glossary 169
Google Cloud 10, 59, 88, 118, 159
Google Cloud Account 133
Google Cloud Platform (GCP) 59, 88,
118
GPU 18, 47, 170
GPU Enabled Cluster Restrictions 47
GPU Enabled Clusters 47
GPU Stack Software 47
GPUs per Node 34
Group ID 58, 79

- H -

Head Node Public IP 31
Head Node Size 31, 37
Head Node Size (Cluster
Information) 37
help (CLI command) 131

High-Performance Computing [155](#)
High-Throughput Computing [155](#)
HOME [152, 163](#)
HPC [1, 5-6, 10, 17, 40, 63, 79, 111, 155, 169](#)
HPC Cloud On-Demand Data Center
CLI Commands [49](#)
HPC Cloud On-Demand Data Center
CLI Procedures [41](#)
HPC Cloud On-Demand Data Center
Platform Installation [8](#)
HPC Cloud On-Demand Data Center
Web UI [19](#)
HPC Cloud On-Demand Data Center
Workload Considerations [155](#)
HPC CloudMart [12, 63, 83, 115](#)
HPC CloudMart Account [133](#)
HTC [155, 169](#)
HTTP [10, 58](#)
HTTPS [10](#)

- I -

IAM [11, 135](#)
ICMP [15](#)
ID [11, 31, 35, 37, 43, 49, 52, 58, 136](#)
Idle Purge Time [36, 79, 114, 170](#)
Idle Time [36](#)
idleTimer [17](#)
Image Name [31, 37](#)
Image Resource Group Name [81, 113](#)
Images [28, 43, 136](#)
In Use [26](#)
Information (cluster action) [30](#)
Ingress Rules [14](#)
Ingress Traffic [9](#)
Install OpenVPN (Cluster

Information) [37](#)
Installation Steps [16](#)
Instance Prices [39](#)
Instance Shapes [18, 85](#)
Instance sizes [33](#)
IP [14, 31, 37, 46, 58, 79, 162](#)

- J -

Job [20, 22, 34, 44, 66, 138, 169](#)
Job Actions [44](#)
job cost [39](#)
Job Graph [23](#)
Job Info (panel) [34](#)
Job Manager [20, 34](#)
Job Manager screen [34](#)
Job Preparation [138](#)
Job Running [23](#)
Job script [47, 111, 138, 170](#)
Job Submit (command) [47](#)
Job Waiting [23](#)
Jobs [35, 38, 43, 79, 110, 171](#)
Jobs (CLI commands) [110](#)
Jobs panel [38](#)
Jobs Submitted [35](#)
JSON [12, 46, 135, 141](#)

- K -

Key File [13](#)
key_file [140, 149](#)
Keys [13, 108, 135](#)

- L -

Last Built [27](#)
Last Name [51](#)
LDAP [163](#)

Legal Notices 5

license (CLI command) 131

License Key 17

License Type 27, 30

Linux 24, 41, 51, 56, 75, 107, 163, 170

Load Average 36

Load Avg 35

Local Peering Gateway 168

Logs 27, 30, 37, 67, 92

Logs (cluster action) 30

LPG 84, 168

- M -

MAC Address 36

Manage Disks 37

Manager ID (file) 37

Manager Port 36

Managers 18

market-node 29

Max Burst 37

Microsoft Azure 10, 60, 81

Microsoft Azure Account 134

MIG 18, 35, 90

MIG Mode 35

Min Burst 37

Misc (CLI commands) 129

Moab 6, 10

Modified (date) 31, 37

Modified (file) 37

Modifying Provider Information To Support Customer Requirements 164

Modifying the Code 17

Multi-Instance GPU 18

My Stacks tab 27

- N -

N Sessions 36

Name (credential) 26

Name (file) 37

Name (job) 34

Name (stack) 27

Net Load 36

Network ID 73, 102

New Cluster screen 31

New Credential screen 26

New Job screen 34

NFS 162

NIS 163

Node 20, 27-28, 38, 48, 111, 138, 152, 162

Node Size 31, 37, 169

Nodes 23, 30, 34, 46, 169

Nodes (cluster) 30

Nodes (job) 34

Nodus ID (file) 37

None (stack built) 27

Number of Nodes 34

Nvidia 47

- O -

Obtaining Azure Credential Information 136

Obtaining Cloud Service Provider Account Information and Credentials 135

Obtaining Credentials 10

OCID 14, 63, 86, 149, 168

OCPUs 37, 84, 169

ODDC server [86](#)
ODDC Server [23](#)
ODDC web API [140](#)
Offline (node state) [36](#)
Offline nodes [35](#)
On-Demand Data Center Overview [9](#)
On-premises [5](#), [156](#), [162](#)
On Demand Cluster [34](#)
Open MPI [48](#)
Open Telekom Cloud [10](#), [72](#), [101](#), [123](#)
OpenSSL [13](#)
OpenVPN [37](#), [103](#), [162](#)
Operating System [36](#), [163](#), [170](#)
Operating Systems [9](#)
OpEx [156](#)
Oracle [10](#), [18](#), [68](#), [121](#), [168](#), [170](#)
Oracle Cloud Account [134](#)
OS Type [59](#), [80](#), [112](#)
OTC [10](#), [123](#)
OTC Account [134](#)
OVPN file [162](#)
Owner [12](#), [27](#), [37](#)
Owner (file) [37](#)

- P -

Packages [79](#), [116](#)
Password [53](#)
pd-balanced [78](#)
pd-ssd [78](#)
pd-standard [88](#)
Performing a Clean Installation of the
ODDC [154](#)
Perl [138](#)
Persistent [5](#), [88](#)
Physical Memory [36](#)

Plan (cluster) [30](#)
Ports [9](#)
Power State [36](#)
Prerequisites [9](#)
Private IP [58](#), [79](#)
Private Oracle Cluster [168](#)
Private SSH Key [53](#)
private_key [149](#)
private_key_id [149](#)
Processors [36](#)
programmatic access [11](#)
project_id [149](#)
Provider [10](#), [26-27](#), [30](#), [36](#), [38](#), [54](#),
[135](#), [164](#)
Provider (cluster) [30](#)
Provider (credential) [26](#)
Provider (stack) [27](#)
Provider file [164](#)
Providers (CLI commands) [128](#)
Provisioning [28](#), [31](#), [58](#), [84](#), [171](#)
Proxy [58](#)
Public IP [31](#), [46](#), [86](#)
Public SSH Key [53](#)
Python [138](#)

- Q -

Queue [35](#), [48](#), [79](#), [171](#)
Queued Jobs [35](#)
Quick Launch [37](#)

- R -

REC Time [36](#)
region [33](#), [39](#), [58](#), [79](#), [111](#), [135](#)
Region ID (option) [79](#)
Related Documentation [6](#)

Remove [50, 52, 55, 59, 80, 112](#)

Resize [37](#)

Resize Cluster [37](#)

Resource Managers [18](#)

Revision History [6](#)

Root [59, 79, 111](#)

Route Rules [16](#)

Run Time [38](#)

Running Jobs [35](#)

- S -

Save Config [30](#)

Script [12, 48, 134, 138, 170](#)

Script (panel) [34](#)

Scripts [152](#)

Secret key [11](#)

secret_key [141](#)

Secure Server [140](#)

Security Group ID (option) [79](#)

Select Cluster [37](#)

serf [46, 92](#)

Server [9, 17, 20, 22, 25, 27-28, 33, 79, 111, 140, 154, 162-163](#)

Server Instance Type ID (option) [79](#)

Server Setup [9](#)

Service Port [36](#)

Services [24](#)

Sessions [36](#)

Settings [13, 20](#)

Share [55, 75](#)

shell scripts [138](#)

Show standard error messages
(command) [43](#)

Site (License Type) [27, 31](#)

Size [27, 31, 37, 65, 78, 169](#)

Size (Cluster Compute Nodes) [37](#)

Size (file) [37](#)

slurm [79, 111](#)

Source Image Root Prefix (option) [79](#)

SSD persistent disks [88](#)

SSH [10, 17, 20, 30, 37, 43, 51, 58, 108, 162](#)

SSH Key [20, 31, 37, 53](#)

SSH Username (Cluster
Information) [37](#)

SSL [162](#)

Stack Manager [20, 26](#)

Stack Manager screen [26](#)

Stacks [28, 57](#)

Stacks (CLI commands) [57](#)

Standard persistent disks [88](#)

Start Time [38](#)

State [30, 35, 83](#)

State (cluster) [30](#)

Status [37, 83, 130, 135](#)

Status (file) [37](#)

stderr [125](#)

Submit (job) [34](#)

Submit a job to a cluster
(command) [44](#)

Submitted [35, 115](#)

Submitting a GPU Job [48](#)

Subnet [13, 58, 79, 115, 137, 149](#)

Subnet ID (option) [79](#)

Subnet Network [14, 73, 102](#)

Subscription ID [136](#)

subscription_id [135, 150](#)

sudo [47, 138, 154](#)

Supported MIG Options [18](#)

- T -

TCP [14](#)
Telekom [10, 72, 101, 123](#)
tenancy_ocid [149](#)
tenant_id [135, 150](#)
tenant_name [150](#)
Terminal [24](#)
The quota of CPUs established by the provider is exceeded [33, 39](#)
The quota of Instances established by the provider is exceeded [33](#)
This Year (jobs submitted) [35](#)
Time [26, 31, 36, 38, 79, 114, 170](#)
Today (jobs submitted) [35](#)
Torque [18, 36, 47, 104, 115, 171](#)
Torque Version [36](#)
Total Cores [36](#)
Total Memory [36](#)
Total Nodes [36](#)
Total Sockets [36](#)
Total Threads [36](#)
Traffic [9](#)
Type [12, 27, 30, 36, 49, 58, 79, 111](#)

- U -

UDP [14](#)
UID_MAX [163](#)
Uninstall OpenVPN (Cluster Information) [37](#)
Unused Cores [23](#)
Unused GPUs [23](#)
Unused Nodes [23](#)
Updating a Cluster to Reference the GPU Prefix [48](#)
Upload Script [34](#)

Uptime [30, 39](#)
Uptime (cluster) [30](#)
User [1, 5, 20, 29, 31, 36](#)
User Profile [20](#)
User role [20](#)
User SSH Key [20, 31](#)
user_ocid [149](#)
useradd (command) [163](#)
Username [31, 37, 51, 56, 58, 107, 130](#)
Users [11, 20, 38, 41, 51](#)
Users (CLI commands) [51](#)
Using Client–Server Directory Service Protocol Setups [163](#)
Using the Adaptive HPC Cloud On-Demand Data Center Web UI [19](#)

- V -

variables [152](#)
VCN [14, 86, 168](#)
Verifying the Installation [17](#)
View Logs (bursting) [37](#)
View standard error [38](#)
View standard output [38](#)
View STD-Err (file) [38](#)
View STD-Out (file) [38](#)
Viewing All Clusters [43](#)
Viewing Clusters [30](#)
Viewing Stacks [26](#)
VM [63, 84, 169](#)
VPC (virtual private cloud) [12, 79, 115, 137](#)
VPC Network Self Link ID (option) [79](#)

- W -

Walltime [34, 171](#)

Walltime (job) [34](#)
Web UI [19](#)
Workload [35](#), [40](#), [79](#), [111](#), [155](#)
Workload Considerations [155](#)
Writing ODDC Scripts [152](#)

- Y -

YP [163](#)