




HPC Cloud On-Demand Data Center

Quick Start Guide 7.7.1

June 2024



HPC Cloud On-Demand Data Center™

Legal Notices	3
Chapter 1: On-Demand Data Center Platform Installation	5
1.1 ODDC Server Prerequisites	5
1.1.1 On-Demand Data Center Overview	5
1.1.2 ODDC Server Setup	6
1.1.3 Ingress Traffic	6
1.1.4 Cloud Provider Credentials	6
1.2 Installation Steps	13
1.3 Verifying the Installation	13
1.4 Adding the Server License Key	14
Chapter 2: On-Demand Data Center CLI Procedures	15
2.1 Stacks	15
2.2 Users	34
2.3 Architecture	35
Chapter 3: Using the On-Demand Data Center Web UI	36
3.1 Creating an ODDC Account	36
3.2 User Profile	37
3.2.1 Settings - Admins	37
3.2.2 Settings - Users	38
3.3 Credentials Manager	39
3.4 Stack Manager	40
3.4.1 Viewing Stacks	40
3.4.2 Creating a Stack	41
3.4.3 Building a Stack	41
3.5 Cluster Manager	42
3.5.1 Viewing Clusters	42
3.5.2 Creating a Cluster	43
3.5.3 Deploying a Cluster	45
3.5.4 Deleting a Cluster	45
3.6 Job Manager	46
3.6.1 Adding a New Job	46
3.7 File Manager	49
3.8 Accounting Manager	50
3.9 Instance Prices	51

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

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	5
1.1.1 On-Demand Data Center Overview	5
1.1.2 ODDC Server Setup	6
1.1.3 Ingress Traffic	6
1.1.4 Cloud Provider Credentials	6
1.2 Installation Steps	13
1.3 Verifying the Installation	13
1.4 Adding the Server License Key	14

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	

Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows	Description
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	
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`.

Chapter 2: On-Demand Data Center CLI Procedures

In this chapter:

2.1 Stacks	15
2.2 Users	34
2.3 Architecture	35

2.1 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

<code>--subnet=subnet</code>	Subnet ID
<code>--volume=volume</code>	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	<code>--dataFile=dataFile</code>	Input files for user provisioning process (multiple files supported)
-N	<code>--nodeProvisioner=nodeProvisioner</code>	User provisioning script to run during stack build on Compute node only
-S	<code>--serverProvisioner=serverProvisioner</code>	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
-n	<code>--name=name</code>	Name Identifier
-p	<code>--prefix=prefix</code>	Output Image Prefix
-r	<code>--region=region</code>	Region ID
-s	<code>--source=source</code>	Source Image ID
-t	<code>--type=centos-7</code>	OS Type
-u	<code>--username=username</code>	SSH Username

Options

--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
--security=security	Security Group ID
--subnet=subnet	Subnet ID
--volume=volume	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	--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	--	Path (absolute or relative) of user provisioning

	serverProvisioner=serverProvisioner	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
-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	--	User provisioning script to run during stack

Options		
	serverProvisioner=serverProvisioner	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
-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.aaaaaaaa7v2znysxlclowk7rhgkgldvjaxmkkjpiktnahnusgknhw5pcvi6q
$ 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

Options

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

Examples

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

```
$ 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

Options

-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)
	--no-disk	Remove Disk
	--no-proxy	Remove HTTP(S) Proxy
	--no-subnet	Remove Subnet
	--proxy=proxy	HTTP(S) Proxy URL
	--subnet=subnet	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

-f	--format=text json	[default: text] Output Format
-h	--help	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		
-d	--debug	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-domain	(required) Availability Domain for OCI tenancy
	--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.aaaaaaa7v2znysxlclowk7rhgkglvjaxmkkjpiktnahnusgknhw
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.aaaaaaa7v2znysxlclowk7rhgkgldvjaxmkkjpiktnahnusgknhw5pcvi6q
```

```
$ 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

Options		
-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:oracle:update myStack --name newStack		
\$ nodus stack:oracle:update myStack --prefix=myNodus --region=us-ashburn-1 --subnet=ocid1.subnet.oc1.iad.aaaaaaa7v2znysxlclowk7rhgkglvjaxmkkjpiktnahnusgknhw5pcvi6q		

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
	--private	Provision via Private IP
	--proxy=proxy	HTTP(S) Proxy URL
	--security=security	Security Group ID
	--subnet=subnet	(required) Subnet Network ID
	--volume=volume	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	--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	Instance Type ID
-n	--name=name	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
```

2.2 Users

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

2.3 Architecture



Chapter 3: 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:

3.1	Creating an ODDC Account	36
3.2	User Profile	37
3.2.1	Settings - Admins	37
3.2.2	Settings - Users	38
3.3	Credentials Manager	39
3.4	Stack Manager	40
3.4.1	Viewing Stacks	40
3.4.2	Creating a Stack	41
3.4.3	Building a Stack	41
3.5	Cluster Manager	42
3.5.1	Viewing Clusters	42
3.5.2	Creating a Cluster	43
3.5.3	Deploying a Cluster	45
3.5.4	Deleting a Cluster	45
3.6	Job Manager	46
3.6.1	Adding a New Job	46
3.7	File Manager	49
3.8	Accounting Manager	50
3.9	Instance Prices	51

3.1 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.

3.2 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**

3.2.1 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.

3.2.1.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.



3.2.1.B Adding a New User

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

The screenshot shows the 'New User' form in the On-Demand Data Center Web UI. The form is titled 'New User' and has a 'Register' button at the bottom. The form fields are: Username *, First Name *, Last Name *, Email *, Password *, and Confirm Password *. Each field has a small icon to its right: a person icon for the name fields, an envelope icon for the email field, and a key icon for the password fields.

3.2.2 Settings - Users

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

user

ID 60c2053329bc7401b48115edf

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

3.3 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**:

Credentials Manager				
Search				
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	...
	otccredential		June 29, 2023 8:59 AM	...

Rows per page: 10 1-5 of 5 < >

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.

3.4 Stack Manager

This section contains the following information:

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

3.4.1 Viewing Stacks

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

Built	Name	Provider	Owner	Instance Size	Prefix	Type	Credential	Builder	Last Built	Actions
None	aws	Amazon	qauser1	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.

3.4.2 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.

3.4.3 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`.


3.5 Cluster Manager

This section contains the following information:

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

3.5.1 Viewing Clusters

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



State	Name	Provider	Owner	Bursting	Nodes	Credential	Uptime	Actions
Available	google	Google	qauser1	Off	0	googlecredential	4 minutes	ⓘ ...
Down	aws	Amazon	qauser1	Off	0	awscredential	N/A	ⓘ ...
Down	azure	Microsoft Azure	qauser1	Off	0	azurecredential	N/A	ⓘ ...
Down	etc	etc	qauser1	Off	0	otccredential	N/A	ⓘ ...
Down	oracle	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`).
-

3.5.2 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. The header bar is blue and contains the title 'New Cluster', the price '\$0.00 per Hour', and three buttons: 'ADVANCED', 'SAVE', and 'CLOSE'. The form fields are as follows:

- Name ***: A text input field.
- OS Type ***: A dropdown menu with 'centos-7' selected.
- Prefix ***: A text input field with 'core' entered.
- Credential ***: A dropdown menu.
- Head Node Size ***: A dropdown menu with 'VM.Standard2.1 - vCPU: 2, Mem (GB): ...' selected.
- Manager ***: A dropdown menu with 'Torque' selected.
- Region ***: A dropdown menu.
- Availability Zone ***: A text input field.
- Bursting Configuration**: Radio buttons for 'Off' (selected), 'Min', 'Max', and 'All'.
- Compute Nodes**: A section with a blue '+' icon and an information icon 'i'. It contains:
 - 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' entered.
- 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.

3.5.3 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...)

3.5.4 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**.

3.6 Job Manager

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


3.6.1 Adding a New Job

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

Name	Nodes	Cores	Data Files	Walltime	Submit	Actions
stdout-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**.

3.7 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.

2. 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.

3.8 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.

1. 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.

- A. 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.

2. 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.

3.9 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**:

Cloud Providers		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)
- Oracle Cloud HPC - Oracle Cloud - Amazon Web Services - Google Cloud - Microsoft Azure - Open Telekom Cloud		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}`).

