

NODUS Cloud OS

Quick Start Guide 3.2.3

October 2019



Contents

Welcome	4
Legal Notices	4
Revision History	4
Chapter 1: NODUS Cloud OS Prerequisites	5
Chapter 2: Onboarding To NODUS Cloud OS	6
2.1 Create An Account	6
2.2 Login To NODUS	6
2.3 Create A New Cluster	7
2.3.1 For Amazon Web Services (AWS)	9
2.3.2 For Other Supported Cloud Providers	10
2.3.3 For Unsupported Cloud Providers	12
2.4 Stacks And Images	14
2.4.1 Building Stacks	15
2.5 Submit Workload	16
2.6 Cloud Bursting	21
2.7 Cloud Budget Analysis	22
Chapter 3: Making Accounts For Cloud Providers	23
3.1 Microsoft Azure	23
3.2 Amazon Web Services (AWS)	24
3.3 Google Cloud Platform	24
Chapter 4: NODUS Platform Overview And Cluster Resources	25
Chapter 5: NODUS Command Line	26
5.1 Setup	26
5.2 Cluster Actions	27
5.2.1 Deploying A Cluster	27
5.2.2 Connecting To A Cluster	27
5.2.3 Destroying A Cluster	27
5.2.4 Viewing All Clusters	27
5.3 Node Actions	28
5.4 Job Actions	28
5.5 Bursting	28
5.6 User Keys	28

5.6.1 List All User Keys	28
5.6.2 Get/Set Values	29
Chapter 6: NODUS Cloud Platform Installation	30
6.1 Prerequisites	30
6.2 Installation Steps	30
Chapter 7: NODUS CLI Commands	32
7.1 Cluster Management	32
7.1.1 Cluster Deploy	32
7.1.2 Cluster Deploy-Node	33
7.1.3 Cluster Connect	33
7.1.4 Cluster Resize	33
7.1.5 Cluster Destroy	33
7.1.6 Cluster Destroy-Node	33
7.2 Cluster Monitoring	33
7.2.1 Cluster Info	33
7.2.2 Cluster Jobs	34
7.2.3 Cluster JSON	34
7.2.4 Cluster List	34
7.2.5 Cluster Members	34
7.2.6 Cluster Monitor	34
7.3 Cluster Jobs	34
7.3.1 Cluster Burst	34
7.3.2 Cluster Run-Job	34
7.4 Cluster Utilities	35
7.4.1 Cluster Run-Task	35
7.5 Resource Management	35
7.5.1 User Set	35
7.5.2 User Get	35
7.5.3 User List	35
7.5.4 User Delete	35
7.5.5 Cluster Add	35
7.5.6 Cluster Delete	36
Chapter 8: NODUS Cloud OS Considerations	37
Appendix A: Glossary	43
Index	44

Welcome

Welcome to the NODUS Cloud OS Quick Start Guide 3.2.3. This guide will show you how to use the NODUS Cloud OS user interface.

NODUS Cloud OS for intelligent cloud management gives immediate access to all computational resources, whether on-premise or in the cloud, on any leading cloud provider. This highly flexible and customizable solution enables HPC or enterprise systems to 'burst' the additional workload to an external cloud on demand. NODUS Cloud OS includes all the necessary tools to facilitate moving HPC and enterprise workloads and applications to the cloud and/or extending on-premise resources. NODUS Cloud OS is scheduler-independent and works with any job scheduler.

Legal Notices

Copyright © 2019 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.

Adaptive Computing, Moab®, Moab HPC Suite, Moab Viewpoint, Moab Grid, NODUS Cloud OS™, 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.

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

Revision History

Date	Description of Change
October 2019	<ul style="list-style-type: none">Initial publication.

Chapter 1: NODUS Cloud OS Prerequisites

Complete the steps below to ensure you are ready to start onboarding to NODUS Cloud OS.

1. From the computer where the job submission will be made, test your internet access from a browser by going to <http://www.nodusplatform.com> and ensure that you can get to the NODUS Cloud OS landing page. Also make sure that your network is configured to allow HTTP requests.
2. Make sure you have the proper cloud credentials and authorization to provision instances in the cloud for your preferred cloud providers.
3. Make sure you are using one of these supported browsers:
 - Google Chrome (latest version)
 - Mozilla Firefox (latest version)
 - Microsoft Edge (two latest major versions)
 - Apple Safari (two latest major versions)

Chapter 2: Onboarding to NODUS Cloud OS

This chapter provides information about onboarding to NODUS Cloud OS and running your workloads in the cloud.

In this chapter:

2.1	Create an Account	6
2.2	Login to NODUS	6
2.3	Create a New Cluster	7
2.3.1	For Amazon Web Services (AWS)	9
2.3.2	For other Supported Cloud Providers	10
2.3.3	For Unsupported Cloud Providers	12
2.4	Stacks and Images	14
2.4.1	Building Stacks	15
2.5	Submit Workload	16
2.6	Cloud Bursting	21
2.7	Cloud Budget Analysis	22

2.1 Create an Account

1. Go to <http://www.nodusplatform.com> and click **Create an Account**.

Note: For most installations, nodusplatform.com refers to the server that the NODUS platform is installed on, for example, the Web server address of the machine that you installed NODUS on

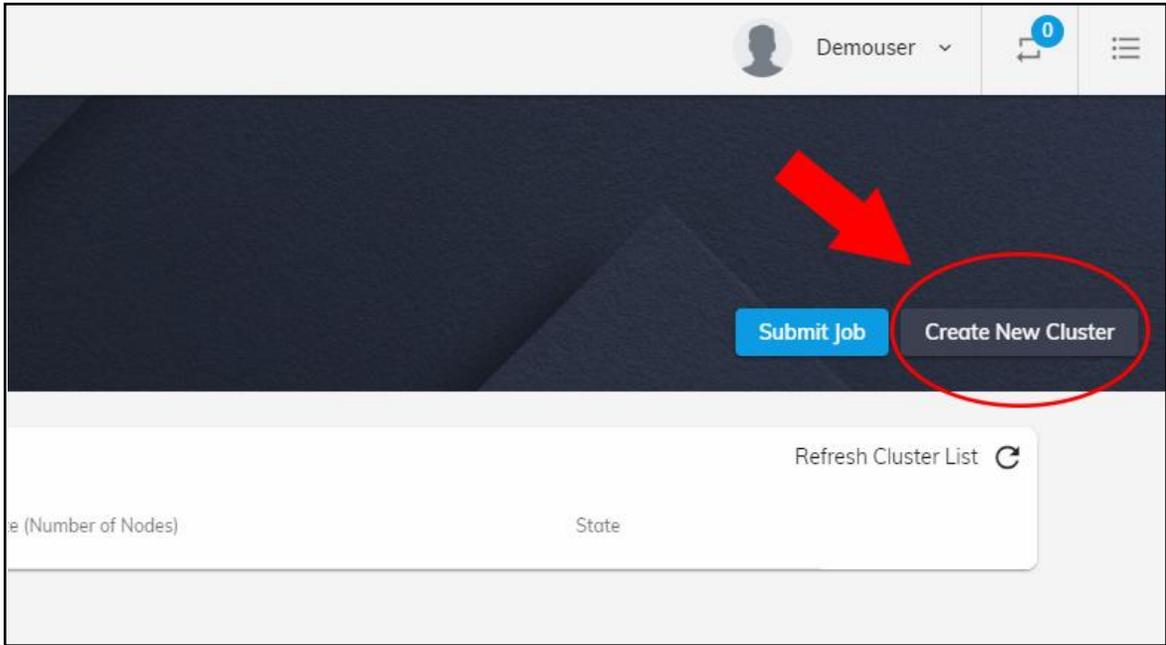
2. Follow the instructions and complete the registration process. At the end of this step, you will have your NODUS Cloud OS user ID and password. Write them down and/or store them in a secure location. Upon successful registration, you will receive a confirmation email from nodus@adaptivecomputing.com and be redirected to the login screen at <http://www.nodusplatform.com>.

2.2 Login to NODUS

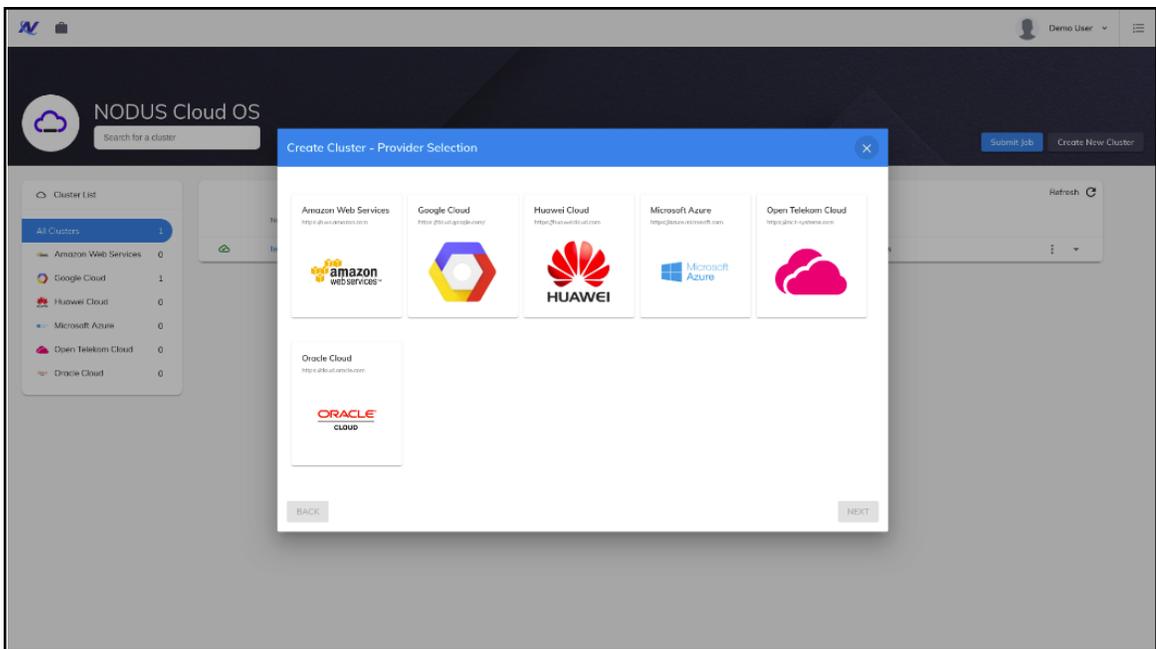
3. Login to your NODUS Cloud OS user account using your username and password. From the main NODUS Cluster management page you can easily deploy and destroy clusters, as well as submit jobs to either an existing cluster or an on-demand cluster. Additionally, you can see a list of all your current clusters once created.

2.3 Create a New Cluster

4. On the NODUS Cluster Management view, click **Create New Cluster** to open the Create Cluster dialog to configure a new cluster:



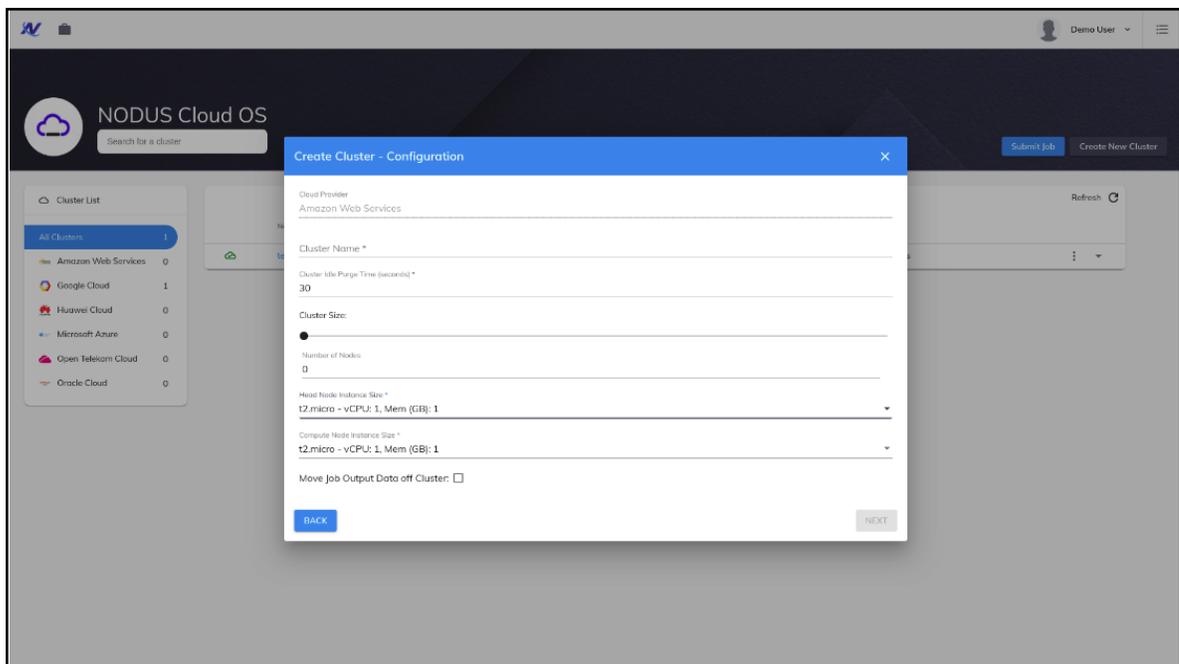
- A. On the Provider Selection dialog, select the cloud provider that you want to deploy a cluster in:



These are the supported cloud providers:

- Amazon Web Services (AWS)
- Google Cloud
- Huawei Cloud
- Microsoft Azure
- Open Telekom Cloud
- Oracle Cloud

B. Click **Next** to continue, and then on the Configuration dialog, choose an appropriate name for your cluster:

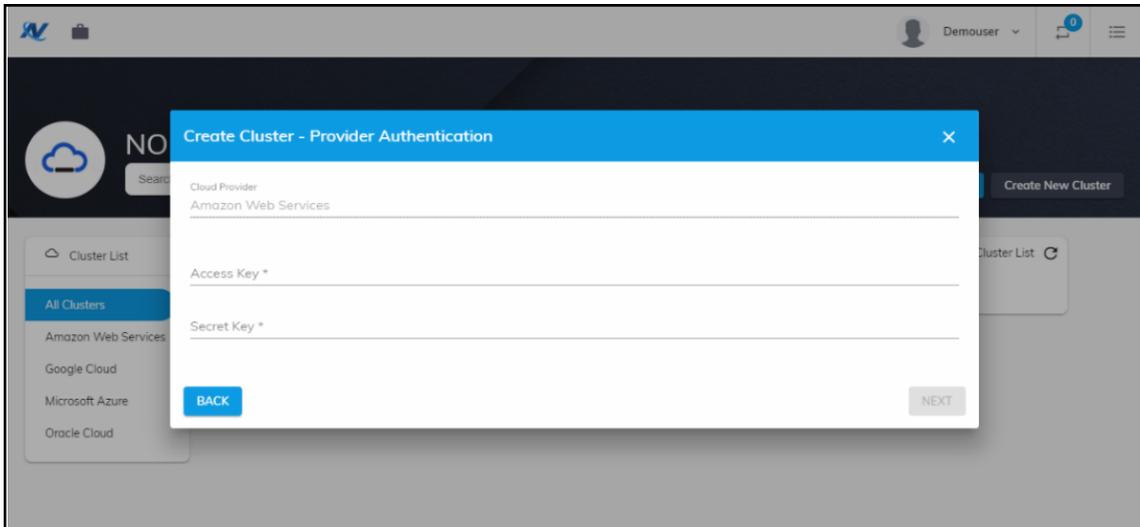


C. Define the cluster size and node size, and also select instance sizes for the head and compute nodes. To have all your job output data moved automatically off the cluster by NODUS, select the check box.

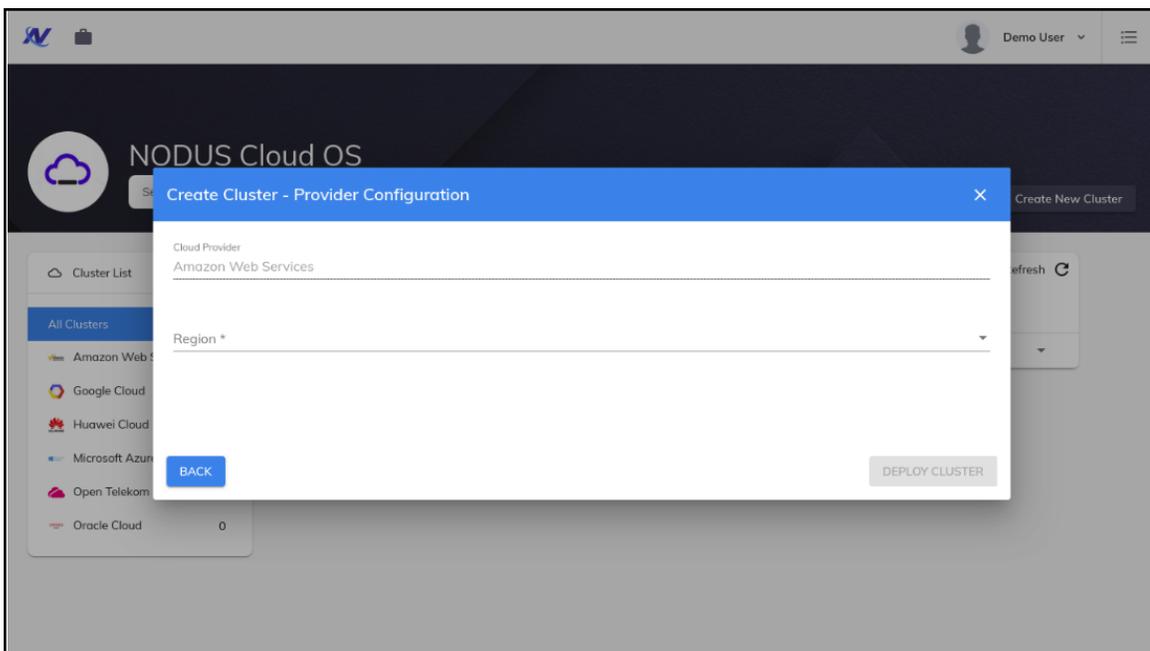
D. Click **Next** to continue, and then add your cloud credentials for the chosen provider.

2.3.1 For Amazon Web Services (AWS)

- a. Enter your AWS account access key and secret key:



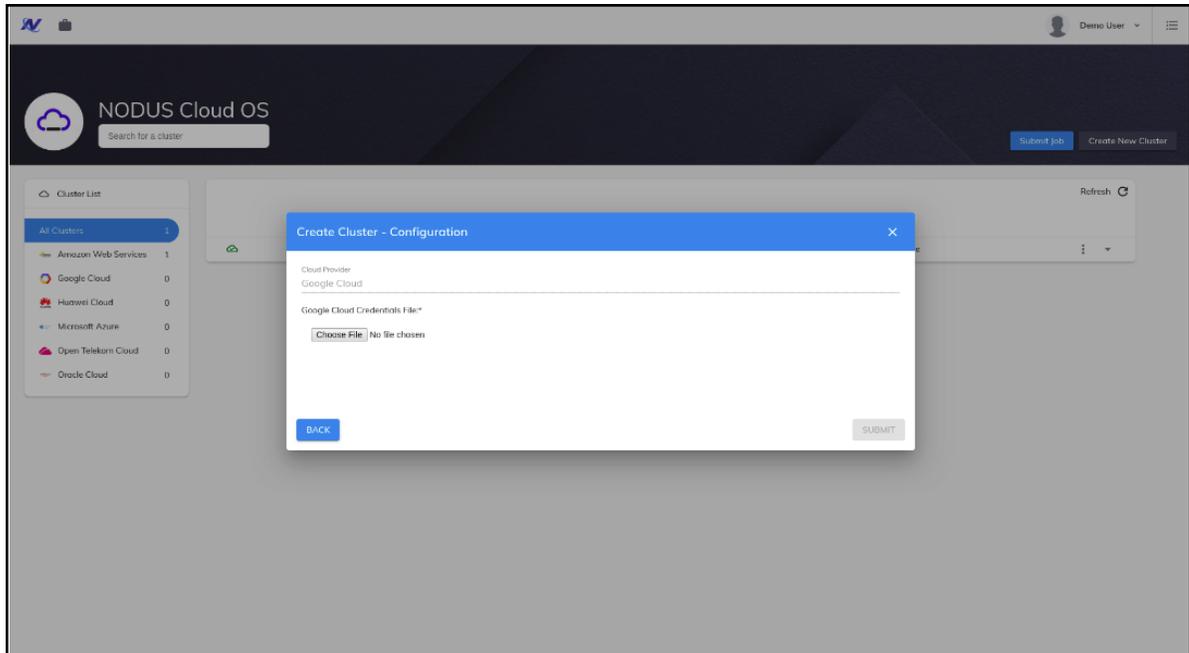
- b. Click **Next** to continue, and then select the region you want to deploy the cluster in:



- c. Then click **Deploy Cluster** to finish cluster creation.

2.3.2 For other Supported Cloud Providers

- a. Upload your credentials file for your selected cloud provider:



- b. Click **Submit** to finish cluster creation.

2.3.2.A Credential File Templates

Below are credential file templates for the cloud providers.

AWS

```
{  
  "access_key": "",  
  "secret_key": "",  
  "region" : ""  
}
```

Microsoft Azure

```
{  
  "client_id": "",  
  "client_secret": "",  
  "object_id": "",  
}
```

```
"subscription_id": "",  
"tenant_id": "",  
"region": ""  
}
```

Google

```
{  
"type" : "",  
"project_id" : "",  
"private_key_id" : "",  
"private_key" : "",  
"client_email" : "",  
"client_id" : "",  
"auth_uri" : "https://accounts.google.com/o/oauth2/auth",  
"token_uri" : "https://oauth2.googleapis.com/token",  
"auth_provider_x509_cert_url": "https://www.googleapis.com/oauth2/v1/certs",  
"client_x509_cert_url" : "https://www.googleapis.com/robot/v1/metadata/x509/937257091049-compute%40developer.gserviceaccount.com"  
}
```

Huawei

```
{  
"username" : "",  
"password" : "",  
"domain_name": "",  
"tenant_name": "",  
"region" : "",  
"auth_url" : "https://iam.myhwclouds.com:443/v3"  
}
```

Open Telekom

```
{  
"username" : "",  
"password" : "",  
"domain_name": "",  
"tenant_name": "",
```

```
"auth_url" : "https://iam.eu-de.otc.t-systems.com:443/v3"  
}
```

Oracle

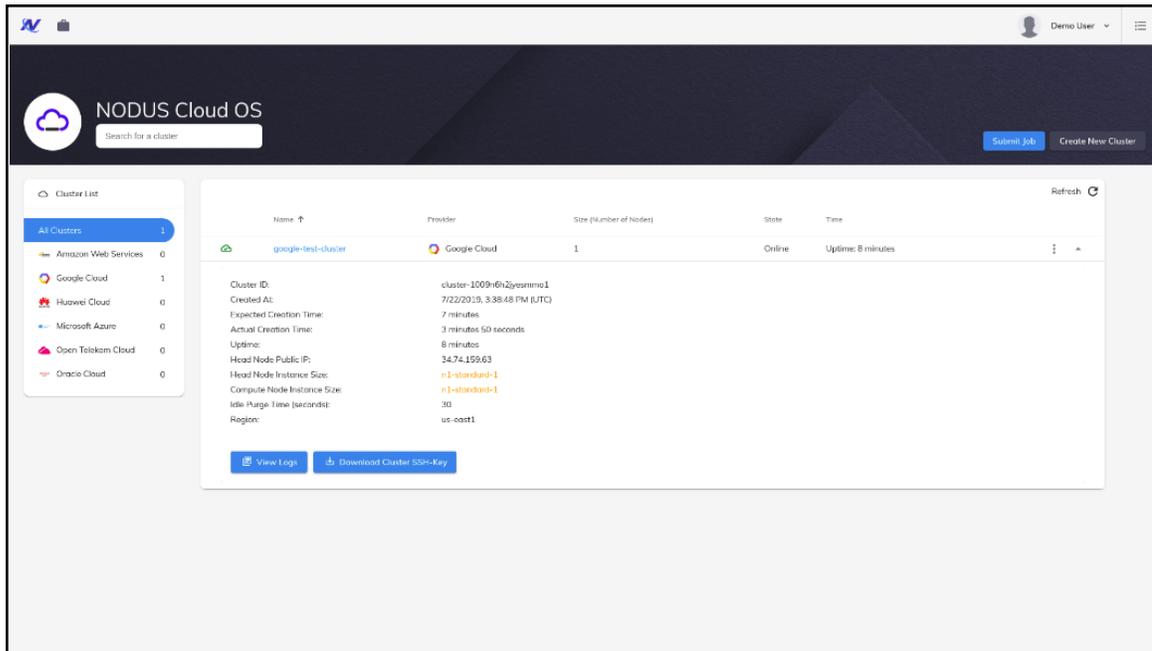
```
{  
  "compartment_ocid" : "",  
  "tenancy_ocid" : "",  
  "user_ocid" : "",  
  "region" : "",  
  "availability_domain_number": "",  
  "oci_api_key" : "",  
  "oci_api_key_fingerprint" : "",  
  "ssh_public_key" : "",  
  "ssh_private_key" : ""  
}
```

2.3.3 For Unsupported Cloud Providers

- a. Please contact nodus@adaptivecomputing.com to inquire about adding the provider into NODUS Cloud OS.

- E. Click Refresh Cluster List (↻) at the top right of the cluster list to see the added cluster. From here you can manage the cluster and see information regarding the cluster.

F. To explore cluster details, click the down-arrow at the end of the row:



Additional details regarding the expanding cluster show, including: ID; date of creation; expected creation time; actual creation time; head node IP; instance sizes of head and compute nodes; and region.

View Logs shows additional information regarding the provisioning. A dialog presents information in real-time about the cluster's state. Additionally, these logs contain information for deploy, re-deploy, resize, and destroy.

Download SSH-Key downloads the SSH key to gain SSH access to the cluster. A '.pem' file is downloaded and can be used to access the head node of the cluster via SSH. This provides you full control over the cluster.

Clicking the vertical ellipsis at the end of the cluster row shows actions that you can take regarding the cluster:

- **Re-Deploy** re-provisions the entire cluster into an online state as if it were newly deployed.
- **Destroy** removes the entire cluster. Meaning the head node and the linked compute nodes are destroyed, removed from the list, and no longer in a state to do work. This also means that you will not be billed further for these cloud resources.

Deploying Clusters

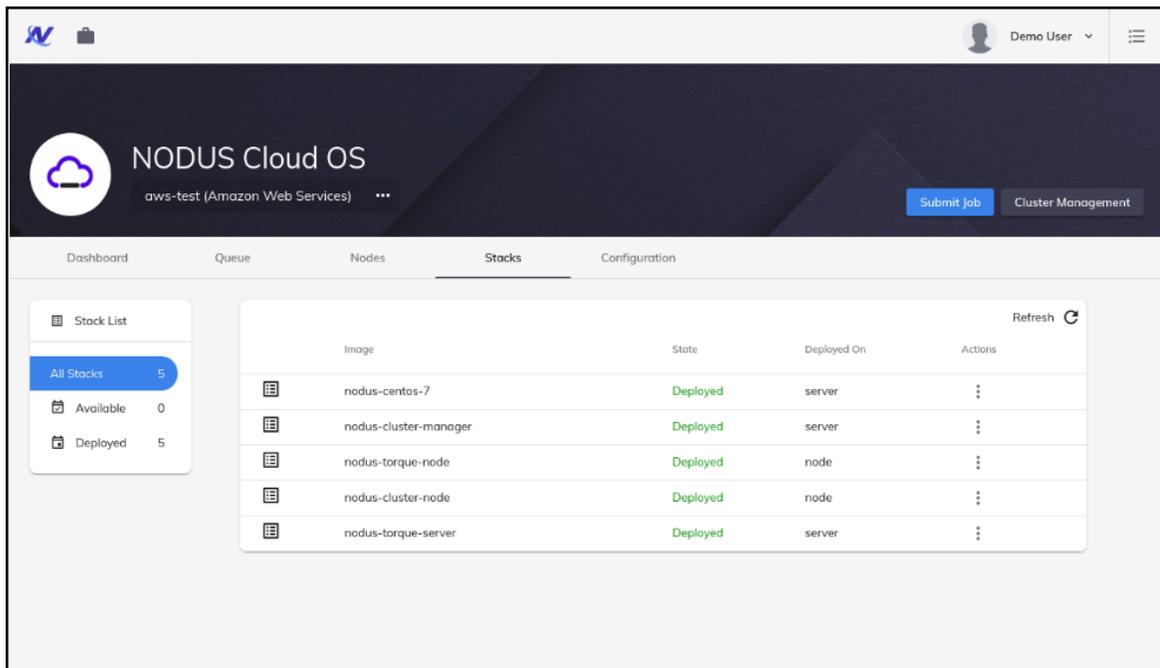
This is the command to deploy a new cluster (or redeploy an existing cluster): **\$ nodus cluster deploy <cluster-name>**

where <cluster-name> is the name of a cluster from the cluster.json file that has been added to NODUS using the command: **\$ nodus cluster add <cluster.json> [credentials.json]**

You can see the list of all clusters added to NODUS using the cluster list command: **\$ nodus cluster list**

2.4 Stacks and Images

5. On the image management screen, click the cluster from the list, then click the **Stacks** tab to view the predefined stacks:



Deployed On - These comprise the software packages loaded on to the NODUS cluster: server or node.

2.4.1 Building Stacks

The only time you need to build stacks is when switching your Cloud Credentials (AWS access/secret key), Cloud Provider, and/or moving regions within AWS (us-east -> us-west). You can see the list of stack images that have been built for a particular region in the AWS Console under the Images section of the sidebar.

NODUS should detect when new stack images need to be built and should build stacks for new cluster deployments automatically.

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

This is the command to build/rebuild a stack for a cluster: **\$ nodus stack build <cluster-name> <stack-name>**

Building Images

The command for building images is: **nodus image build <cluster-name> <image-name>**

By default, NODUS images are built on top of the marketplace image CentOS 7. Alternatively, it can be configured to build NODUS images on top of an existing image.

Note: See <https://wiki.centos.org/About/Product> for minimum requirements for CentOS 7.

Build Order

Images must be built in this order:

0. [user-image]
1. nodus-centos-7
2. nodus-server
3. nodus-node

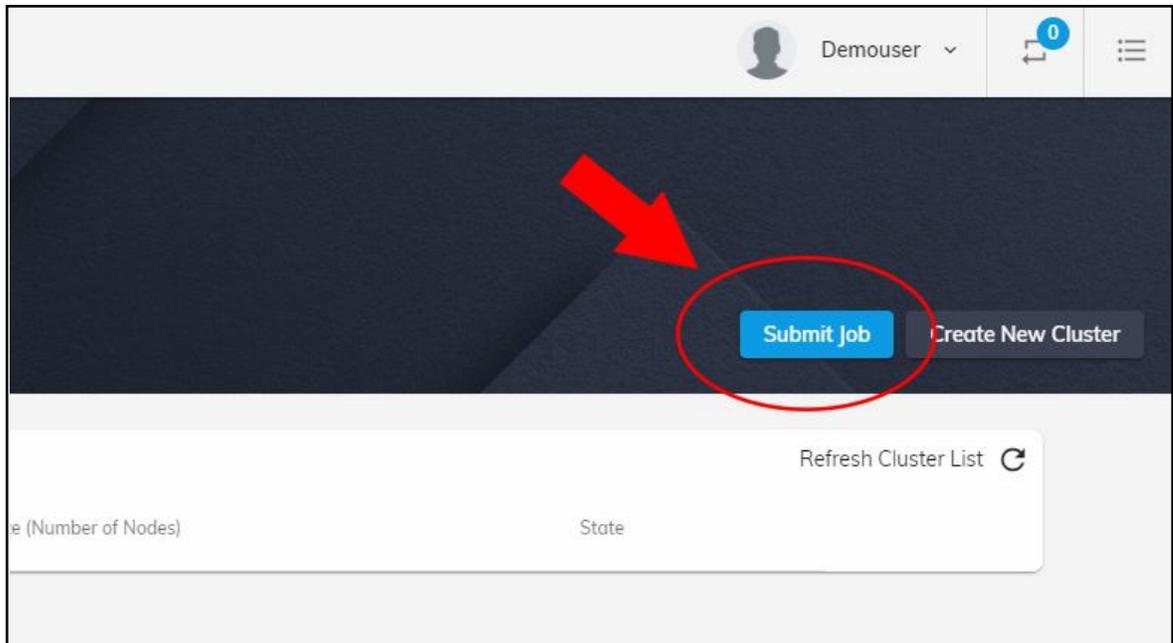
Note: 0. [user-image] is optional.

A successful build should end with the following:

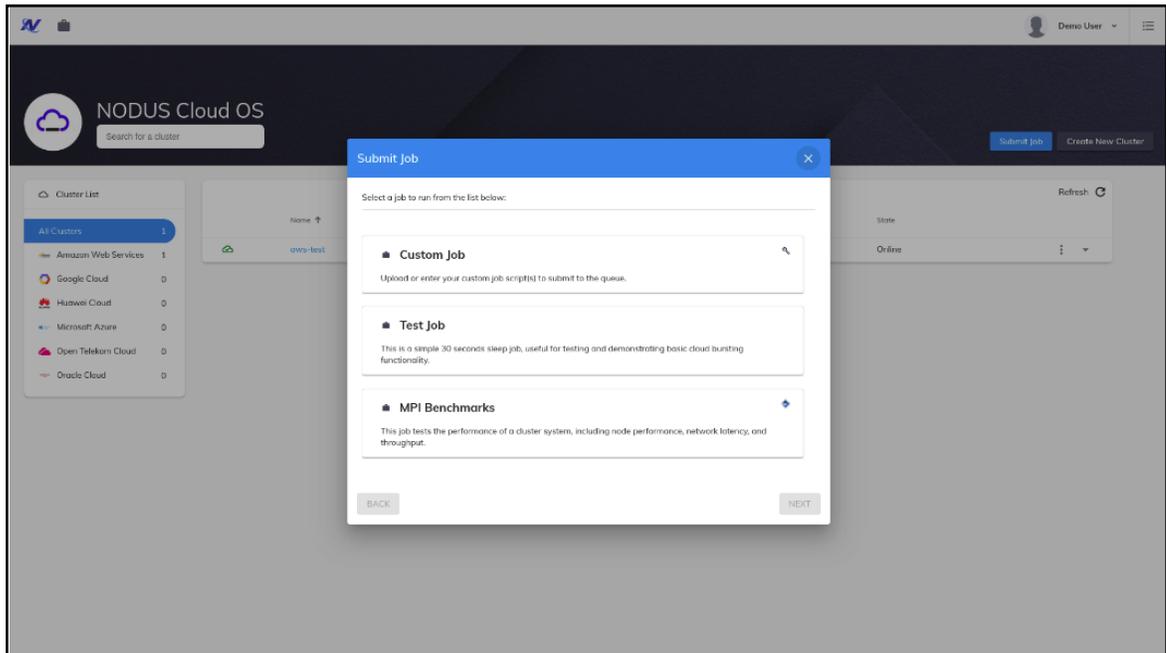
```
Build 'NODUS Image Builder - nodus-centos-7-1566325733 - ' finished.
==> Builds finished. The artifacts of successful builds are:
--> NODUS Image Builder - nodus-centos-7-1566325733 - : AMIs were created:
us-east-1: ami-07a14470d92be82a2
/home/name/.nodus/resources/aws/cluster/aws-cluster/nodus-centos-7.image
```

2.5 Submit Workload

6. Click **Submit Job** to open the job submission dialog:



A. On the Submit Job dialog, select the type of job you want to submit:



Custom Job: A job that is customizable and configurable.

Test Job: A job that is best used to test bursting functionality.

MPI Benchmarks: A job that tests performance of the cluster.

- a. Enter a job name, select how you want to enter the job script (either upload or enter your custom job script), upload any additional data files if the job requires it, and scroll down to specify the job's resource requirements:

Submit Job - Configuration

Upload or enter a custom script and configure the node requirements for the job to run:

BASIC JOB SETTINGS

Job Name *

Job Script *:

Upload Script Enter Script

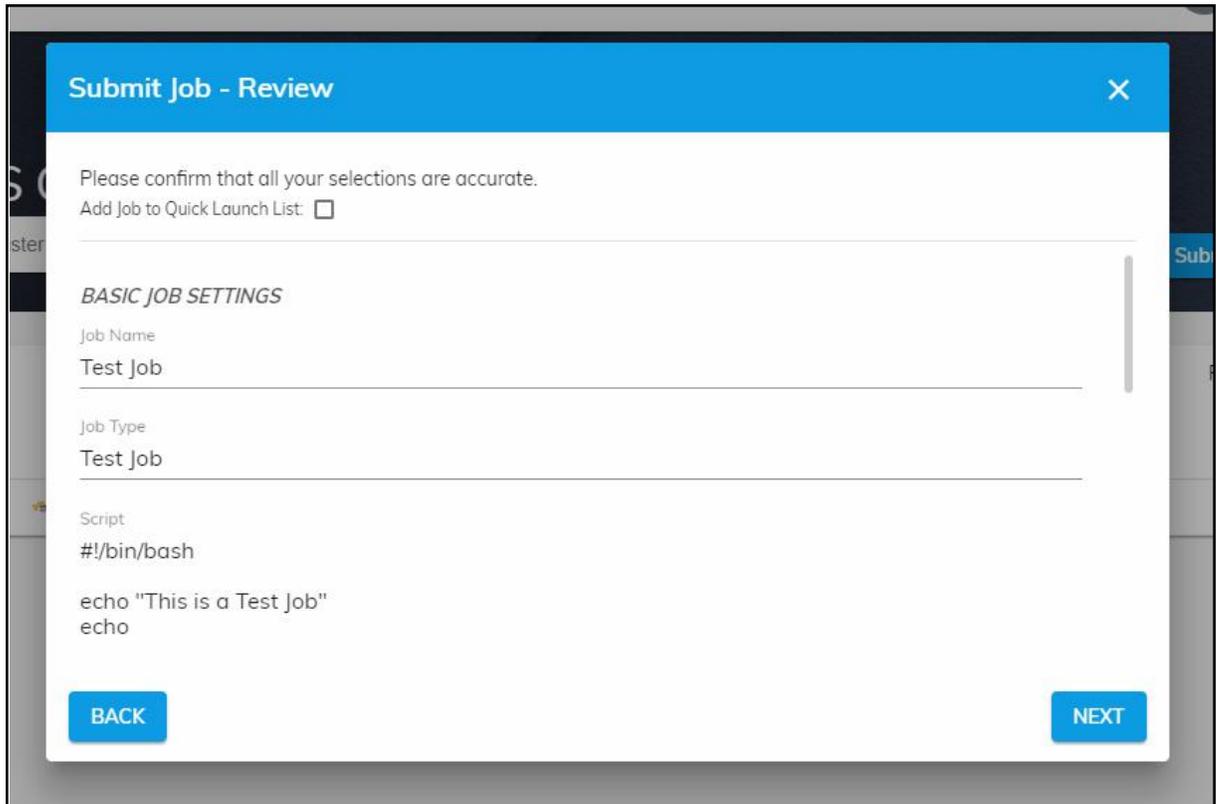
Choose File No file chosen

Data Files (optional):

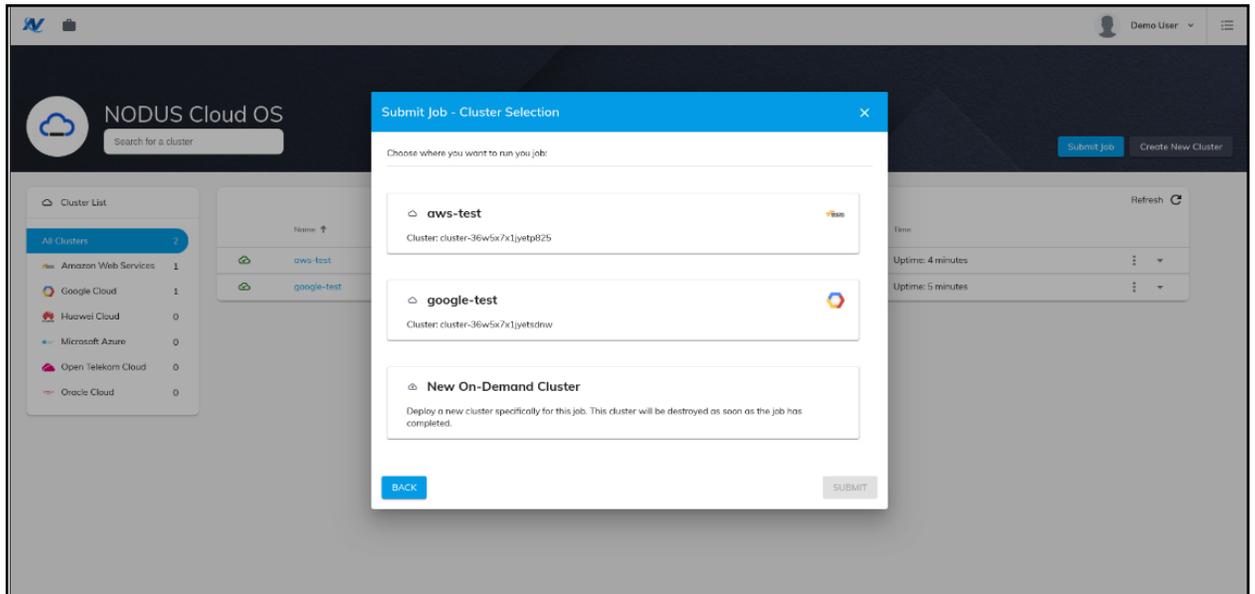
Choose Files No file chosen

BACK NEXT

- B. Click **Next** to continue, and then on the Review dialog, review the entered information and configurations of the job:



- C. Click **Next** to continue, and then on the Cluster Selection dialog, select the cluster to run the job on:



Named Cluster: If there are no created clusters or ones that meet the hardware requirements of the job, then certain named clusters may not appear at all.

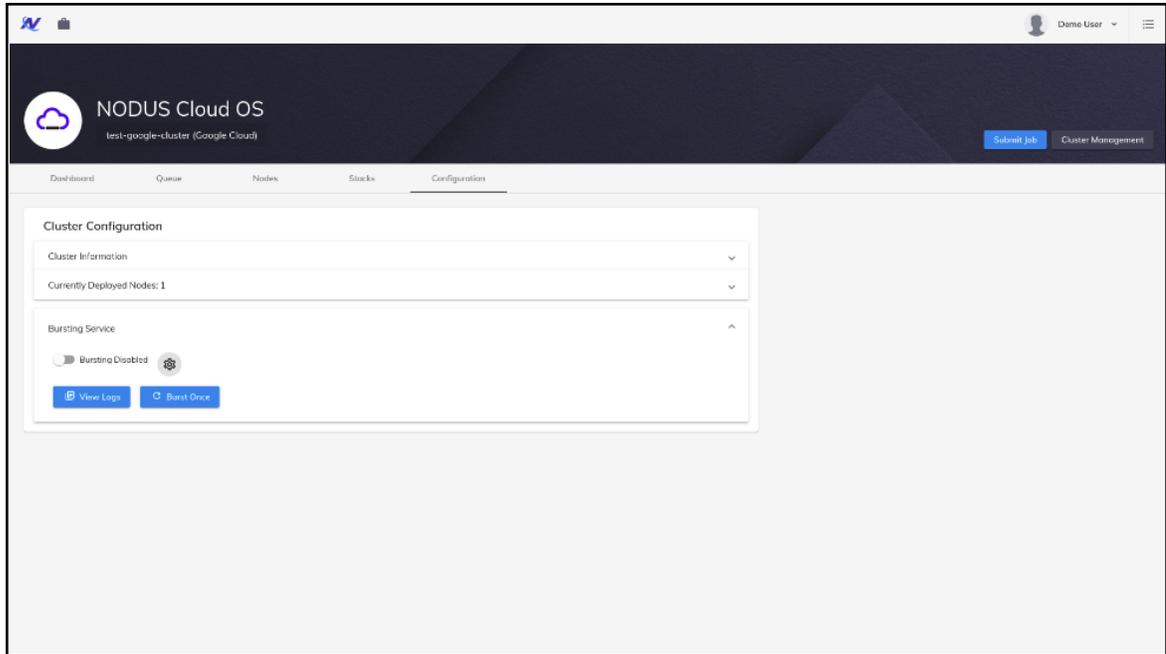
New On-Demand Cluster: A new temporary cluster that is tailored to the parameters of your job will be deployed to run your job and then destroyed when the job completes.

- D. Click **Submit** to finish job submission. The job is queued and runs on the selected cluster from the list of created clusters.

After you have made all the selections above, your workload will run given the chosen cloud provider. When the workload completes, the output will be attached to the job ID on the User Interface. The cluster nodes will be destroyed when there is no more work.

2.6 Cloud Bursting

7. Select a cluster and click the **Configuration** tab to manage cluster cloud bursting configurations:

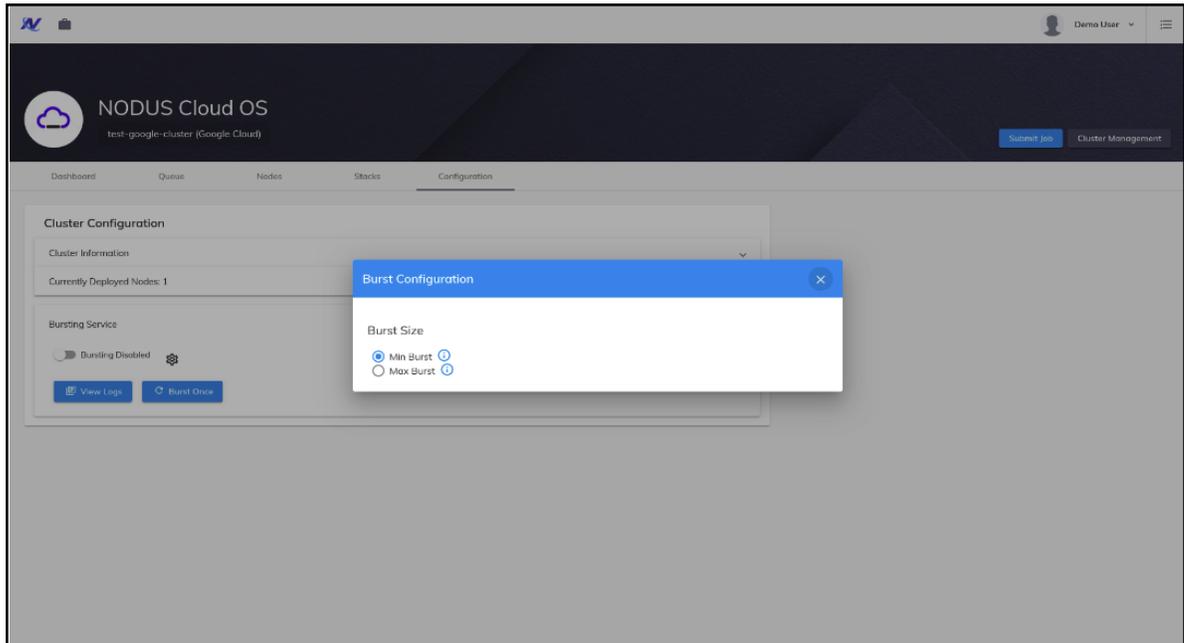


Disable or enable bursting functionality for a single cluster:

- **Bursting Enabled** - Jobs are run and clusters provisioned to handle the jobs
- **Bursting Disabled** - The cluster remains in its current state.

The Settings/Gear icon enables selecting the Burst Size:

- **Min Burst** completes the workload with the minimum number of nodes required. This is great for keeping to budgets and assessing costs.
- **Max Burst** completes the workload immediately with the maximum number of required nodes. This is great for getting results as fast as possible.



2.7 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 Cloud Business Value Process Spreadsheet.

Chapter 3: Making Accounts for Cloud Providers

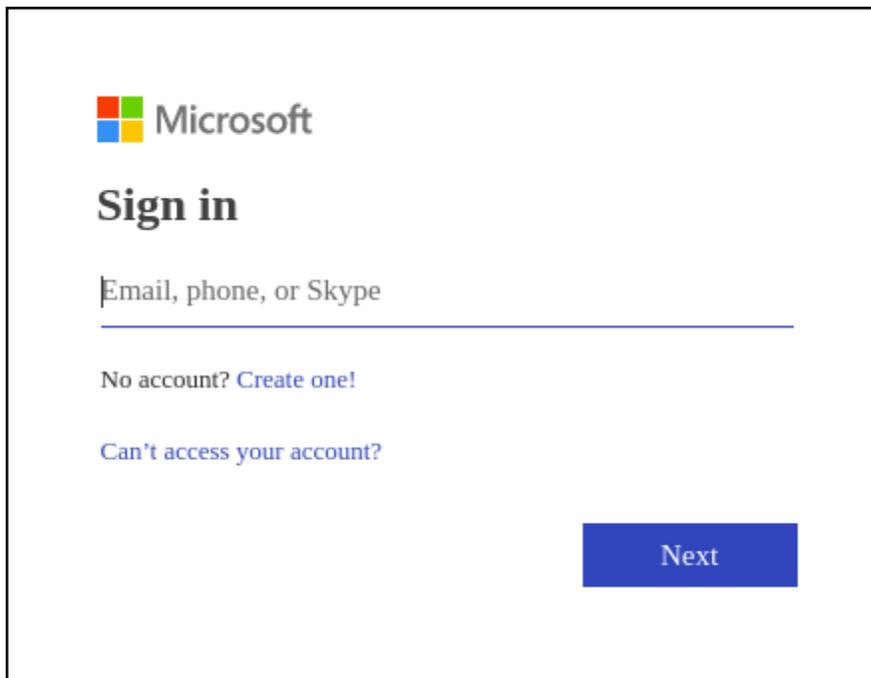
If you don't have credentials through your company, you can get them by following the procedures below.

In this chapter:

3.1	Microsoft Azure	23
3.2	Amazon Web Services (AWS)	24
3.3	Google Cloud Platform	24

3.1 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 input the rest of your user information (i.e., username, passwords, credit/debit, etc.).

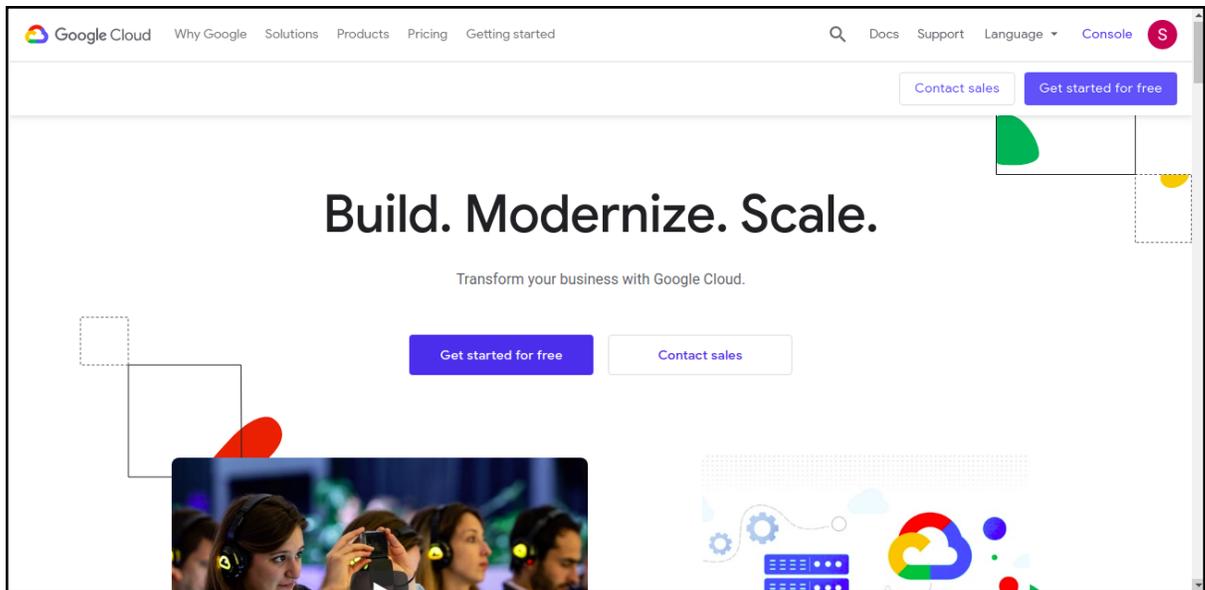
3.2 Amazon Web Services (AWS)

1. Go to <https://aws.amazon.com/> and click **Create an AWS account** (if you already have an account, it may say **Sign in to**).
2. Follow the steps and enter your account information and then click **Continue**.
3. Choose **Personal** or **Professional**.
4. Enter your company information and accept the customer agreement.
5. Choose **Create Account and Continue**.
6. Provide a payment method and verify your phone number.
7. Click **My Security Credentials > Access keys (access key ID and secret key) > Create New Access Key** and create a new one, 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.

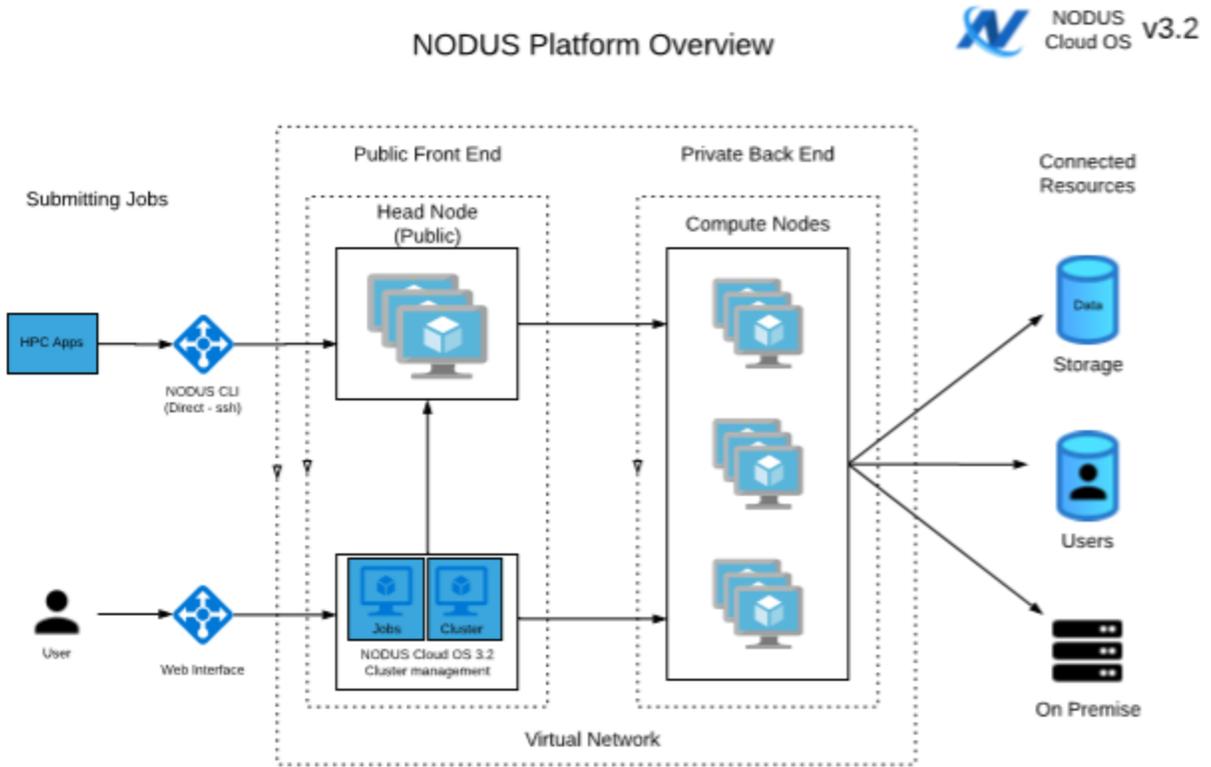
3.3 Google Cloud Platform

1. Go to <https://cloud.google.com/> and click **Get started for free**:

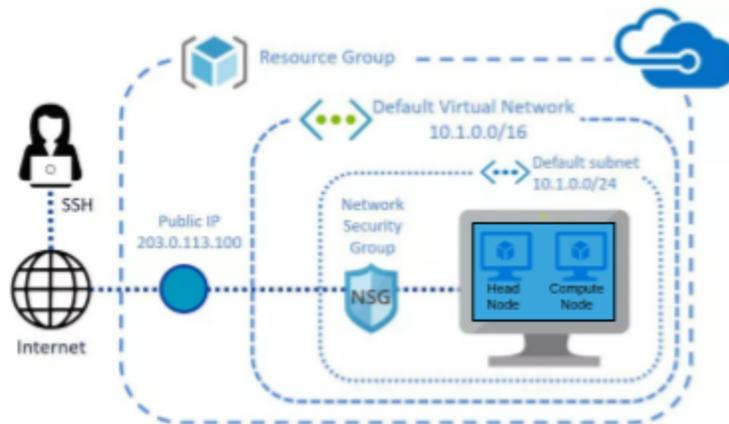


2. Follow the steps on the screen to complete account registration.

Chapter 4: NODUS Platform Overview and Cluster Resources



NODUS Cluster Resources



Chapter 5: NODUS Command Line

After NODUS is installed, set up the command line interface (CLI) by following the steps below.

In this chapter:

5.1	Setup	26
5.2	Cluster Actions	27
5.2.1	Deploying a Cluster	27
5.2.2	Connecting to a Cluster	27
5.2.3	Destroying a Cluster	27
5.2.4	Viewing All Clusters	27
5.3	Node Actions	28
5.4	Job Actions	28
5.5	Bursting	28
5.6	User Keys	28
5.6.1	List all User Keys	28
5.6.2	Get/Set Values	29

5.1 Setup

1. Ensure the NODUS CLI is available by running **nodus --help** from a terminal. **Note:** If this is unsuccessful, follow these steps to link the CLI:

```
cd /NODUS/nodus-cli
```

```
npm i
```

```
sudo npm link
```

```
nodus -help
```

2. Set the location of nodus-cluster-providers for the CLI to use (this folder is located in /NODUS by default):

```
nodus user set packages/nodus-cluster-providers /NODUS/nodus-cluster-providers
```

5.2 Cluster Actions

5.2.1 Deploying a Cluster

1. Obtain credentials from the cloud provider you intend to deploy this cluster on.
2. Create a cluster file for the cluster you are deploying.
3. Add the cluster via **nodus cluster add <cluster.json> [credentials.json]**. Cluster add optionally let's you add your credentials when you add a cluster. If you don't, you can still set them using **nodus user set cluster/<cluster-name>/credentials**.
4. Add credentials for the cluster using **nodus user set cluster/<cluster-name>/credentials credentials.json**, providing the full path to the file.
5. Deploy the cluster with **nodus cluster deploy <cluster-name>**.

Summary

nodus cluster add cluster-file.json

nodus user set cluster/<cluster-name>/credentials /path/to/credentials.json

5.2.2 Connecting to a Cluster

1. To connect to a deployed cluster, run **nodus cluster connect <cluster-name>**.

5.2.3 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-name>**.
2. The cluster should be destroyed before you delete it or you will have to delete it manually. To remove all information about the cluster, run **nodus user delete cluster/<cluster-name>**.

5.2.4 Viewing All Clusters

1. To view all saved clusters, deployed or not, run **nodus cluster list**.

5.3 Node Actions

1. To deploy/destroy a cluster's individual nodes manually, run:

nodus cluster deploy-node <cluster-name> <node-name>

nodus cluster destroy-node <cluster-name> <node-name>

Note: Node names are always nodus-cluster-node-x, where x is a number from 0 to (cluster size - 1).

5.4 Job Actions

1. To submit a job to a cluster, run **nodus cluster run-job** <cluster-name> <job-script>.

5.5 Bursting

The NODUS burst function looks at what jobs are queued on your cluster and automatically spins up or tears down nodes depending on the total requirements for the queue. If there are not enough online nodes to run all jobs, bursting will bring on as many nodes as needed. If there are more nodes than needed, the excess nodes will be destroyed. If the job queue is empty, all nodes will be destroyed.

To run a single burst cycle, the command is: **nodus cluster burst** <cluster-name>

5.6 User Keys

The NODUS CLI comes with user key, value storage. These resources are used by the CLI commands.

nodus cluster add <cluster-name> creates a user key, **cluster/<cluster-name>**, containing the cluster file.

5.6.1 List all User Keys

1. To view all of your user keys, run **nodus user list**.

5.6.2 Get/Set Values

1. To retrieve the value of a key, run **nodus user get <key>**.
2. To set the value of a key or create a new key, run **nodus user set <key> <value>**.

Chapter 6: NODUS Cloud Platform Installation

This chapter provides information about the installer, which contains the entire NODUS Cloud Platform, along with dependencies to install it, and setup scripts to simplify/automate the process.

In this chapter:

6.1 Prerequisites	30
6.2 Installation Steps	30

6.1 Prerequisites

The NODUS installation requires either CentOS 7 or Oracle 7.7 as the operating system.

Note: Refer to the operating system's product specifications for minimum requirements.

Ensure the device running the API (e.g., your in-house CentOS 7 NODUS Platform Server) has the following ports open:

- 22 (SSH) - User direct access / Moab job routing
- 80 (HTTP) - Connection to the NODUS UI
- 443 (HTTPS) - Optional, if you secure your HTTP connection
- 12021 (NODUS-web-API) - The NODUS Platform API server that your UI will communicate with

6.2 Installation Steps

1. Create a nodus user on the NODUS Platform. **Note:** This account must have sudo access to the root account.
2. Login as the nodus user.
3. Transfer the following file to the nodus user account: **Nodus_Cloud_Platform-3.2.x.tgz**
Where x is the current release.
4. Extract the file using this command: **tar xf Nodus_Cloud_Platform-3.2.x.tgz**

5. Run one of the following installation scripts:

For CentOS

```
cd Nodus-3.2.x/  
./centos7-install.sh
```

For Oracle

```
cd Nodus-3.2.x/  
./oracle7-install.sh
```

6. When prompted, enter the nodus user password. Depending on your system, the installation may take up to 45 minutes to complete. If it does not complete successfully, refer to the **install.log** for troubleshooting.

Chapter 7: NODUS CLI Commands

After NODUS is installed, set up the command line interface by following the steps below.

In this chapter:

7.1	Cluster Management	32
7.1.1	Cluster Deploy	32
7.1.2	Cluster Deploy-Node	33
7.1.3	Cluster Connect	33
7.1.4	Cluster Resize	33
7.1.5	Cluster Destroy	33
7.1.6	Cluster Destroy-Node	33
7.2	Cluster Monitoring	33
7.2.1	Cluster Info	33
7.2.2	Cluster Jobs	34
7.2.3	Cluster JSON	34
7.2.4	Cluster List	34
7.2.5	Cluster Members	34
7.2.6	Cluster Monitor	34
7.3	Cluster Jobs	34
7.3.1	Cluster Burst	34
7.3.2	Cluster Run-Job	34
7.4	Cluster Utilities	35
7.4.1	Cluster Run-Task	35
7.5	Resource Management	35
7.5.1	User Set	35
7.5.2	User Get	35
7.5.3	User List	35
7.5.4	User Delete	35
7.5.5	Cluster Add	35
7.5.6	Cluster Delete	36

7.1 Cluster Management

7.1.1 Cluster Deploy

Deploys a named cluster using credentials from bash environment variable `$CREDENTIALS_FILE` or from nodus user get cluster/cluster-name/credentials.

\$ nodus cluster deploy <cluster name>

7.1.2 Cluster Deploy-Node

Deploys a single node on a named cluster.

\$ nodus cluster deploy-node <cluster name> <node name>

7.1.3 Cluster Connect

Connects to a named cluster.

\$ nodus cluster connect <cluster name>

7.1.4 Cluster Resize

Resizes a cluster.

\$ nodus cluster resize <cluster name> <new node count>

7.1.5 Cluster Destroy

Destroys a cluster.

\$ nodus cluster destroy <cluster name>

7.1.6 Cluster Destroy-Node

Destroys one node on a cluster.

\$ nodus cluster destroy-node <cluster name> <node name>

7.2 Cluster Monitoring

7.2.1 Cluster Info

Get general information about the cluster (i.e., cluster username, public IP, node list...).

\$ nodus cluster info <cluster name>

7.2.2 Cluster Jobs

Displays JSON information about the jobs queue.

\$ nodus cluster jobs <cluster name>

7.2.3 Cluster JSON

Displays the cluster definition JSON file for a named cluster.

\$ nodus cluster json <cluster name>

7.2.4 Cluster List

Lists all user clusters.

\$ nodus cluster list

7.2.5 Cluster Members

Displays serf members from the head node.

\$ nodus cluster members <cluster name>

7.2.6 Cluster Monitor

Displays serf logging information.

\$ nodus cluster monitor <cluster name>

7.3 Cluster Jobs

7.3.1 Cluster Burst

Runs a single burst cycle that spins up and tears down nodes depending on workload.

\$ nodus cluster burst <cluster name>

7.3.2 Cluster Run-Job

Remotely runs a local job script on a cluster.

\$ nodus cluster run-job <cluster name> <job script>

7.4 Cluster Utilities

7.4.1 Cluster Run-Task

Runs NODUS Utility functions for deployed clusters.

```
$ nodus cluster run-task <cluster name> <task name>
```

7.5 Resource Management

7.5.1 User Set

Sets a user's key, value pair for resources.

```
$ nodus user set <name of resource> <path to resource>
```

7.5.2 User Get

Gets the value of a user's resource.

```
nodus user get <name of resource>
```

7.5.3 User List

Lists all resource keys for a user.

```
nodus user list
```

7.5.4 User Delete

Deletes a user's resource.

```
nodus user delete <name of resource>
```

7.5.5 Cluster Add

Adds a cluster using the provided cluster definition file.

```
nodus cluster add <cluster file>
```

7.5.6 Cluster Delete

Deletes the cluster from user storage (does not destroy the cluster).

\$ nodus cluster delete <cluster name>

Chapter 8: NODUS Cloud OS 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.
 - NODUS Cloud OS has multi-cloud capabilities and will run workloads in the cloud from the UI on one or more of the major cloud providers.
 - Run your workload on each cloud provider and record the statistics for each one in the Cloud Business Value Process Benchmark Report. See Figure 2.1 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 2.1) below. This will help you develop your cloud budget.

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 2.1: Benchmarks

- When using NODUS Cloud OS, expenditures move from CapEx to 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.
 - NODUS Cloud OS accommodates on-premise, even your laptop, SaaS, and cloud-hosted delivery models, which are very easy to implement because of the low barriers to entry using NODUS Cloud OS. 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-premise cluster?
 - NODUS Cloud OS offers superior usability when devising cloud and on-premise integration strategies.
 - In the diagram below (Figure 2.2) the workload has burst to the cloud automatically based on backlog. The job is seen by the scheduler as part of the on-premise cluster. NODUS Cloud OS has the capability to seamlessly integrate on-premise and cloud resources.

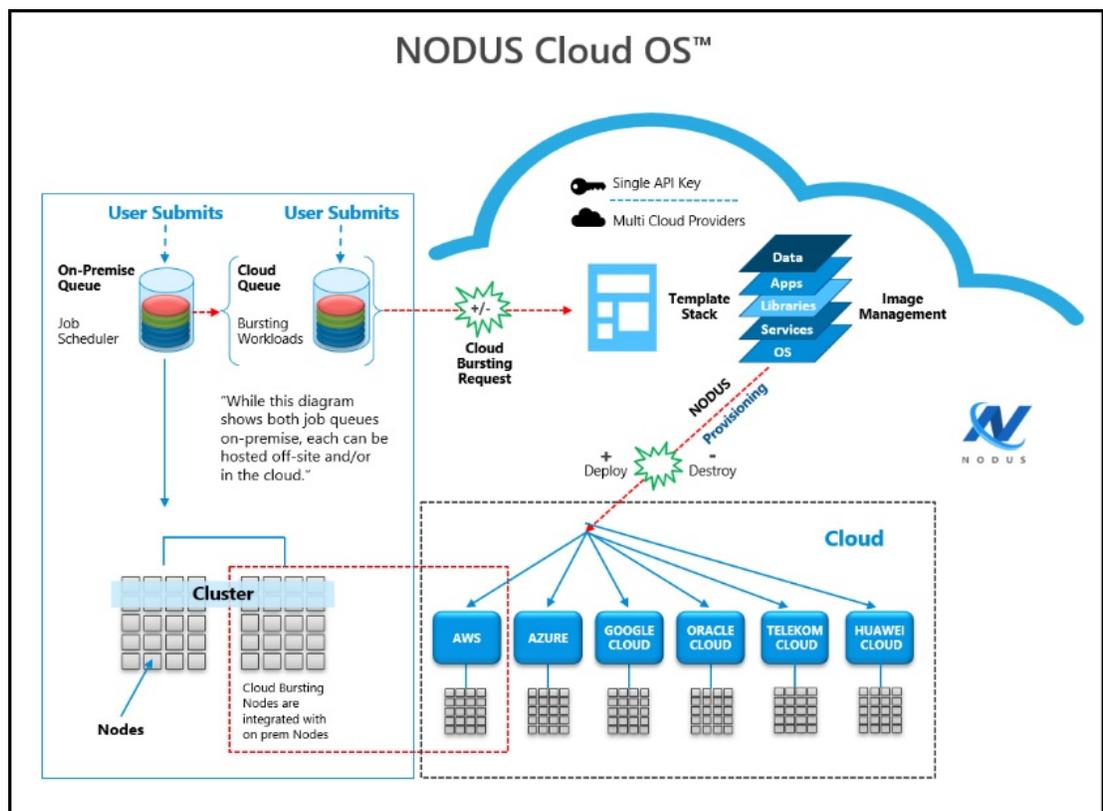


Figure 2.2: NODUS Cloud OS Bursting Becoming On-Premise

- In the diagram below (Figure 2.3) the workload has burst to the cloud on-demand and is running in the cloud independently of the on-premise cluster.

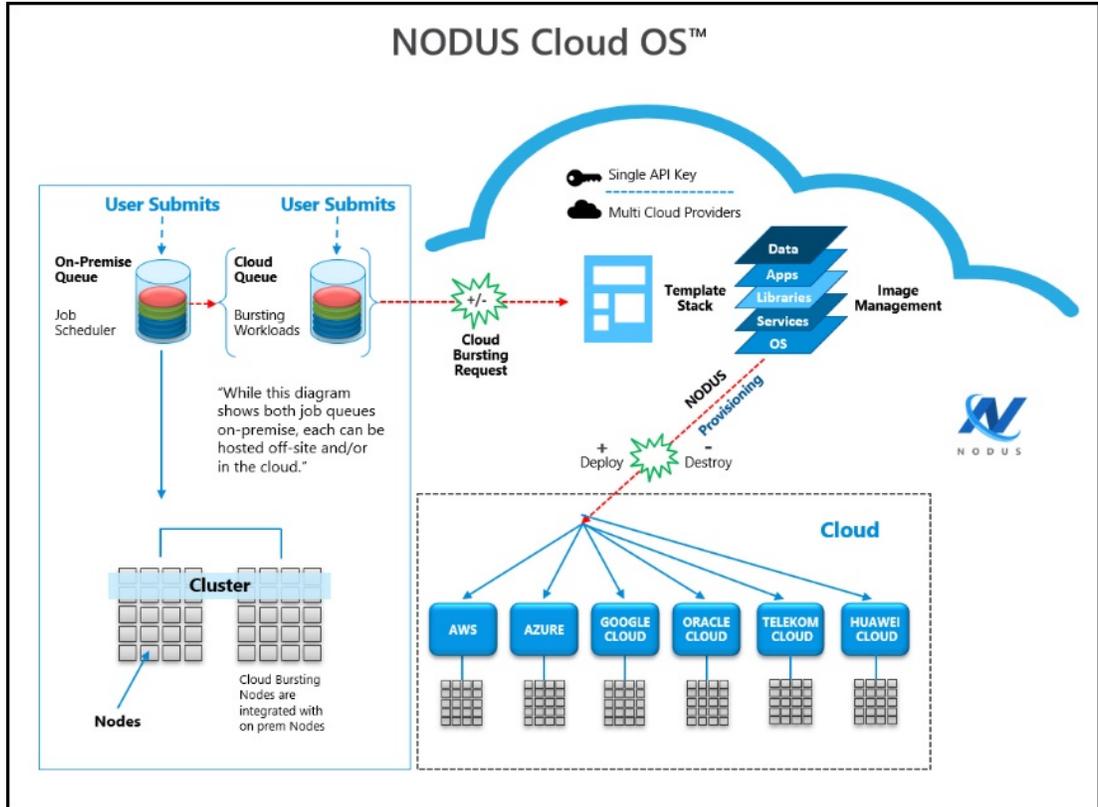


Figure 2.3: NODUS Cloud OS Bursting Off-Premise

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.
 - NODUS Cloud OS has a Data Grid feature, which can automatically move the data and proliferate it.

8. Which public cloud providers match which of your workloads? One size does not fit all.
 - Using NODUS Cloud OS, you can test your workloads on each one of the cloud providers. Within the NODUS Cloud OS interface, select the cloud provider that your job will run on. The nodes will shut down when the job completes, and you can then select another cloud provider to test your workload on. For example, one workload should run on AWS and another should run on Google Cloud.
 - 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 provider. (Figure 2.1)

Appendix A: Glossary

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

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

Cluster Size: The number of compute nodes.

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

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

Custom Job: A job that is customizable and configurable.

Image: A snapshot of an OS.

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

MPI Test Job: A job that tests performance of the cluster.

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

Provisioning: The event of configuring a node or cluster with its stack and getting it into a ready-to-work state.

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

Test Job: A job that is best used to test bursting functionality.

Workload: Jobs to be run; jobs in the queue.

Index

- A -

Access key 9
Actual creation time 13
Amazon Web Services (AWS) 8, 24
Apple Safari 5
AWS 8, 15, 24, 40
Azure 8, 23

- B -

Browsers 5
Build Order 15
Burst Size 22
Bursting 21, 28, 38, 43
Bursting Disabled 21
Bursting Enabled 21

- C -

CapEx 38
CentOS 7 15, 30
Chrome 5
CLI 26-28, 32
CLI Commands 32
Cloud Budget Analysis 22
Cloud Bursting 21
Cloud Business Value Process Benchmark Report 37
Cloud Providers 5, 8, 23, 37
Cluster 6-7, 20, 25, 27, 32-35, 43
Cluster Actions 27
Cluster Add 35
Cluster Burst 34
Cluster Connect 33
Cluster Delete 36

Cluster Deploy 32
Cluster Deploy-Node 33
Cluster Destroy 33
Cluster Destroy-Node 33
Cluster Info 33
Cluster Jobs 34
Cluster JSON 34
Cluster List 12, 34
Cluster Management 7, 32
Cluster Members 34
Cluster Monitoring 33
Cluster Resize 33
Cluster Run-Job 34
Cluster Run-Task 35
Cluster Selection dialog 20
Cluster Utilities 35
Command Line 26
Compute Nodes 8, 43
Connecting to a cluster 27
Copyright 4
Create an account 6
Creating a new cluster 7
Creation time 13
Credential file templates 10
Credentials 5, 8, 15, 23-24, 27, 32, 43
Custom Job 17, 43

- D -

Data Grid 40
Date of creation 13
Deploy a new cluster 14
Deploy Cluster 9

Deployed On column 14
Deploying a cluster 27
Destroy 13, 33
Destroying a cluster 27
Download SSH-Key 13

- E -

Edge 5
Expected creation time 13

- F -

File Templates 10
Firefox 5

- G -

Get Values 29
Google Chrome 5
Google Cloud 8, 24, 40

- H -

Head node IP 13
High-Performance Computing 37
High-Throughput Computing 37
HPC 4, 37
HTC 37
HTTP 5, 30
Huawei Cloud 8

- I -

ID 6, 13, 20, 24
Images 14
Instance sizes 8
IP 13, 33

- J -

Job 16, 28, 43
Job Actions 28
Job script 18, 34, 43

- L -

Legal Notices 4
Logs 13

- M -

Making accounts for cloud providers 23
Max Burst 22
Microsoft Azure 8, 23
Microsoft Edge 5
Min Burst 22
Moab 4, 30
Mozilla Firefox 5
MPI Benchmarks 17

- N -

Named Cluster 20
New On-Demand Cluster 20
Node 8, 14, 28, 33
Node Actions 28
nodus-cluster-providers 26
NODUS-web-API 30
NODUS Cloud OS considerations 37
NODUS Cloud Platform Installation 30
NODUS Cluster Resources Overview 25
NODUS Command Line 26
NODUS Platform Overview 25

- O -

On-Demand Cluster 6, 20, 43
On-premise 4, 38
Onboarding 6
Open Telekom Cloud 8
Operating System 1, 4-6, 12, 30, 37, 43
OpEx 37
Oracle 8, 30-31
Oracle 7 30

- P -

Prerequisites 5, 30

Provisioning 13, 43

- R -

Re-Deploy 13

Redeploy an existing cluster 14

Refresh Cluster List 12

Region 9

Resource management 35

Resource requirements 18

Review dialog 19

Revision History 4

- S -

Safari 5

Script 18, 28, 34, 43

Secret key 9, 15, 24

Server 6, 14, 30

Set Values 29

SSH-Key 13

Stacks 14

Submit Job 16

Submit Workload 16

sudo 26, 30

- T -

Telekom 8

Templates 10

Test Job 17, 43

Troubleshooting 31

- U -

Unsupported cloud providers 12

User Delete 35

User Get 35

User Keys 28

User List 35

User Set 35

- V -

View Logs 13

Viewing all clusters 27

- W -

Workload 4, 16, 22, 34, 43