

Dell EMC Ready Bundle for Red Hat OpenStack Platform

Hadoop Solution Deployment Guide

Version 10.0.1



Dell EMC Converged Platforms and Solutions

Contents

List of Figures.....	iv
List of Tables.....	v
Trademarks.....	6
Glossary.....	7
Notes, Cautions, and Warnings.....	12
 Chapter 1: Executive Summary.....	 13
About the Dell EMC Ready Bundle for Red Hat OpenStack Platform.....	14
About This Document.....	14
Intended Audience.....	14
Prerequisites.....	15
Dependencies.....	15
Related Documentation.....	15
Deployment Workflow.....	15
 Chapter 2: High-level Architecture.....	 17
Hardware Configuration.....	18
Network Architecture.....	18
 Chapter 3: Cloud Administrator Configurations.....	 19
Configure RAID on Compute Nodes.....	20
Deploy the Undercloud.....	20
Map the Cinder Volume Service.....	21
Deploy the Overcloud.....	21
Create the Cinder Type for LVM Storage.....	22
Create a Hadoop Tenant and Administrator User.....	22
Create a Resource Configuration File for the Hadoop Administrator.....	23
Set the Quota Values for the Hadoop Tenant.....	23
Create an External Network and Subnet.....	24
Create Host Aggregates for the Compute Nodes.....	24
Create Nova Node Instance Flavors.....	25
 Chapter 4: Hadoop Administrator Configurations.....	 26
Deploy Hadoop Clusters with OpenStack Sahara.....	27
Create a Hadoop Image.....	27
Upload the Hadoop Image to the OpenStack Image Service.....	28
Register the Image with the Sahara Image Registry.....	28
Create a Nova Key Pair for Sahara.....	28
Create a Tenant Network, Subnet, and Router.....	29
Create Node Group Templates.....	29

Create a Hadoop Cluster Template.....	34
Create the Hadoop Cluster.....	34
Configure Cloudera Hadoop.....	35
 Appendix A: Configuration Files Notes.....	37
dell_environment.yaml.....	38
first_boot.yaml.....	39
post_install.yaml.....	40
 Appendix B: Configuring a Node or VM for Hadoop Admin Use.....	42
Configure a Node or VM for Hadoop Admin Use.....	43
 Appendix C: Getting Help.....	44
Technical References.....	45
Performance and Benchmarking References.....	45
Cloudera 5.7.0 Documentation.....	45
Apache Hadoop Documentation.....	45
Red Hat Documentation.....	45
Contacting Dell EMC.....	45
References.....	46
To Learn More.....	46

List of Figures

Figure 1: Simple Network Schematic..... 18

List of Tables

Table 1: Deployment Workflow.....	15
Table 2: Hardware Configuration.....	18
Table 3: Nova Node Flavor Resources.....	25
Table 4: NameNode Group Template Configurations.....	29
Table 5: Cloudera Manager Group Template Configuration.....	30
Table 6: Data Node 0 Group Template Configuration.....	31
Table 7: Data Node 1 Group Template Configuration.....	32
Table 8: Data Node 2 Group Template Configuration.....	33
Table 9: Cluster Template Configurations.....	34
Table 10: Hadoop Cluster Parameters.....	35
Table 11: Hadoop Cluster Configuration Values.....	35

Trademarks

Copyright © 2014-2017 Dell Inc. or its subsidiaries. All rights reserved.

Microsoft® and Windows® are registered trademarks of Microsoft Corporation in the United States and/or other countries.

Red Hat®, Red Hat Enterprise Linux®, and Ceph are trademarks or registered trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Oracle® and Java® are registered trademarks of Oracle Corporation and/or its affiliates.

DISCLAIMER: The OpenStack® Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries, and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community.

Glossary

API

Application Programming Interface is a specification that defines how software components can interact.

BMC/iDRAC Enterprise

Baseboard management controller. An on-board microcontroller that monitors the system for critical events by communicating with various sensors on the system board, and sends alerts and log events when certain parameters exceed their preset thresholds.

Bundle

A customer-orderable solution that consists of:

- All server, network, and storage hardware needed to install and operate the solution as outlined
- All necessary solution software licenses needed to install and operate the solution as outlined

CDH

Cloudera Distribution for Apache Hadoop

Cloud Computing

See <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

Cluster

A set of servers dedicated to OpenStack that can be attached to multiple distribution switches.

Compute Node

The hardware configuration that best supports the hypervisor server or Nova compute roles.

DevOps

Development Operations (DevOps) is an operational model for managing data centers using improved automated deployments, shortened lead times between fixes, and faster mean time to recovery. See <https://en.wikipedia.org/wiki/DevOps>.

DIMM

Dual In-line Memory Module

DNS

The domain name system (DNS) defines how Internet domain names are located, and translated into Internet Protocol (IP) addresses.

GUI

Graphical User Interface - A visual interface for human interaction with the software, taking inputs and generating easy to understand visual outputs.

Hypervisor

Software that runs virtual machines (VMs).

IaaS

Infrastructure as a Service.

Infrastructure Node

Systems that handle the control plane and deployment functions.

ISV

Independent Software Vendor.

JBOD

Just a Bunch of Disks

LAG

Link Aggregation Group.

LOM

LAN on motherboard.

LVM

Logical Volume Management

ML2

The Modular Layer 2 plug-in is a framework that allows OpenStack to utilize different layer 2 networking technologies.

NFS

The Network File System (NFS) is a distributed filesystem that allows a computer user to access, manipulate, and store files on a remote computer, as though they resided on a local file directory.

NIC

Network Interface Card

Node

One of the servers in the cluster.

NUMA

Non-Uniform Memory Access

Overcloud

The functional cloud that is available to run guest VMs and workloads.

Pod

An installation comprised of three racks, and consisting of servers, storage, and networking.

REST

REST - Representational State Transfer (also ReST). Relies upon stateless, client-server, cacheable communications protocol to access the API.

RHOSP

Red Hat OpenStack Platform

RPC

Remote Procedure Call

SAH

The Solution Admin Host (SAH) is a physical server that supports VMs for the Undercloud machines needed for the cluster to be deployed and operated.

SDS

Software-defined storage (SDS) is an approach to computer data storage in which software is used to manage policy-based provisioning and management of data storage, independent of the underlying hardware.

SDN

Software-defined Network (SDN) is where the software will define, create, use and destroy different networks as needed.

Storage Node

The hardware configuration that best supports SDS functions such as Red Hat Ceph Storage.

ToR

Top-of-rack switch/router.

U

U used in the definition of the size of server, example 1U or 2U. A "U" is a unit of measure equal to 1.75 inches in height.

Undercloud

The Undercloud is the system used to control, deploy, and monitor the Overcloud - it is a single node OpenStack deployment completely under the administrators control. The Undercloud is *not* HA configured.

VLT

A Virtual Link Trunk (VLT) is the combined port channel between an attached device (ToR switch) and the VLT peer switches.




VLTi

A Virtual Link Trunk Interconnect (VLTi) is an interconnect used to synchronize states between the VLT peer switches. Both endpoints must be the same speed, i.e. 40Gb → 40Gb; 1G interfaces are not supported.

VM

Virtual Machine - a simulation of a computer system.

Notes, Cautions, and Warnings

-  A **Note** indicates important information that helps you make better use of your system.
-  A **Caution** indicates potential damage to hardware or loss of data if instructions are not followed.
-  A **Warning** indicates a potential for property damage, personal injury, or death.

This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

Chapter 1

Executive Summary

Topics:

- [*About the Dell EMC Ready Bundle for Red Hat OpenStack Platform*](#)
- [*About This Document*](#)
- [*Intended Audience*](#)
- [*Prerequisites*](#)
- [*Related Documentation*](#)
- [*Deployment Workflow*](#)

Hadoop, a project of the Apache Software Foundation, is a software framework that allows for the distributed processing of large data sets across clusters of computers using simple programming models. It is designed to scale up from single servers to thousands of machines, each offering local computation and storage. Rather than rely on hardware alone to deliver high-availability, the library itself is designed to detect and handle failures at the application layer, delivering a highly-available service on top of a cluster of computers.

Cloudera Distribution for Apache Hadoop (CDH) is a commercially-supported and 100% open source platform distribution of Hadoop.

This document describes the procedures for deploying Cloudera Hadoop on OpenStack cloud infrastructure, with the Dell EMC Ready Bundle for Red Hat OpenStack Platform. Dell EMC validated this procedure using CDH 5.7.0, with RHEL 7.3 as the base operating system and Red Hat OpenStack Platform 10.

About the Dell EMC Ready Bundle for Red Hat OpenStack Platform

The Dell EMC Ready Bundle for Red Hat OpenStack Platform is an integrated hardware and software solution for OpenStack cloud that has been jointly designed and validated by Dell EMC and Red Hat. The Ready Bundle enables organizations to easily and rapidly deploy a highly reliable, optimized, and scalable Infrastructure as a Service (IaaS) cloud solution.

The Ready Bundle architecture is built on Dell EMC PowerEdge servers for the Controller, Compute, and Storage nodes, with Red Hat OpenStack Platform and Red Hat Ceph Storage software, plus a number of validated extensions to the core architecture to enable capabilities such as Platform as a Service (PaaS), Containers as a Service (CaaS), and cloud native development and operations. This release of the Ready Bundle is based on Red Hat OpenStack Platform 10 (RHOSP 10), which is the OpenStack release called Newton.

About This Document

This document describes the procedures for deploying Hadoop in a Red Hat OpenStack Platform cluster. It does not address life-cycle management of either OpenStack or Hadoop.

This document describes the recommended servers and OpenStack configuration procedures for a Hadoop workload in the Dell EMC Ready Bundle for Red Hat OpenStack Platform. With these recommendations, Hadoop performance is expected to be close to a bare-metal Hadoop solution. These include:

- Installation and configuration of Sahara, the OpenStack Data Processing Service
- Configuration needed on OpenStack Compute nodes for:
 - OpenStack Cinder Volume Service
 - CPU Pinning using NUMA Awareness

There are two different sets of configuration procedures involved:

- [Cloud Administrator Configurations](#) on page 19 — Performed by the *Cloud Admin* (responsible for the OpenStack Cloud deployment, and has physical access to the hardware)
- [Hadoop Administrator Configurations](#) on page 26 — Performed by the *Hadoop Administrator* (responsible for the Hadoop Cluster deployment on OpenStack using Sahara)

This document contains code and configuration samples in mono-space fonts. While it is tempting for the user to copy and paste those values from this document into their system, it is inadvisable and not supported. While we make every effort to ensure that the documentation is correct and complete, documents rendered via some client applications make unpredictable changes to the actual spacing of the data elements, and lose fidelity to what a proper code or configuration setting should actually be to work properly. We see very impactful changes, for example, between the Firefox PDF display and the Adobe Acrobat Reader PDF display.

Copy and paste from this document only with full understanding of the necessary formatting changes that you'll have to make. We have made efforts to provide online verbatim copies of the essential data, as well as pointing the user to appropriate external documentation to achieve the proper formatting.

Intended Audience

This guide assumes the reader is familiar with:

- OpenStack

- Hadoop
- Dell EMC PowerEdge R730xd RAID and BIOS configuration
- Red Hat Enterprise Linux Server 7.3
- Red Hat OpenStack Platform (RHOSP) documentation
- Network Configuration

Prerequisites

The following prerequisites must be satisfied before proceeding:

1. Hardware racked and stacked as per the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Hardware Deployment Guide](#)
2. Switches wired and configured as per the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Architecture Guide](#)
3. Prior to deployment, the Infrastructure administrator must put all disks, in Compute nodes with local storage, into non-RAID mode (JBOD) in the BIOS, **except**:
 - The two disks that will host the operating system, which will be in RAID 1

Dependencies

For customers performing a self-installation, these files are available on request from Dell EMC. Please contact your account representative, or email openstack@dell.com for instructions. Dell EMC Ready Bundle for Red Hat OpenStack Platform automated and manual deployment dependencies include:

- *JetPack-manual-10.0.1.tgz* — Contains all manual deployment solution scripts
- *JetPack-automation-10.0.1.tgz* — Contains all automation deployment solution scripts

Related Documentation

Obtain the following documents for reference:

- [Dell EMC Ready Bundle for Red Hat OpenStack Platform Architecture Guide](#)
- [Dell EMC Ready Bundle for Red Hat OpenStack Platform Hardware Deployment Guide](#)
- [Dell EMC Ready Bundle for Red Hat OpenStack Platform Software Deployment Guide](#)
- [Dell EMC Ready Bundle for Cloudera Hadoop Architecture Guide](#)

Deployment Workflow

[Table 1: Deployment Workflow](#) on page 15 describes the basic deployment sequence:

Table 1: Deployment Workflow

Deployment Step	Information Reference
Cloud Administrator Tasks	
1. Configure RAID on the Compute Nodes	Configure RAID on Compute Nodes on page 20
2. Deploy the Dell EMC Ready Bundle for Red Hat OpenStack Platform Undercloud	Deploy the Undercloud on page 20

Deployment Step	Information Reference
3. Map the Cinder Volume Service on Compute Nodes	Map the Cinder Volume Service on page 21
4. Deploy the Dell EMC Ready Bundle for Red Hat OpenStack Platform Overcloud	Deploy the Overcloud on page 21
5. Create the Cinder type for LVM Storage	Create the Cinder Type for LVM Storage on page 22
6. Create a Hadoop Tenant (Project) and Admin User	Create a Hadoop Tenant and Administrator User on page 22
7. Create an RC file for the Hadoop Admin	Create a Resource Configuration File for the Hadoop Administrator on page 23
8. Set the Quota for the Hadoop Tenant	Set the Quota Values for the Hadoop Tenant on page 23
9. Create an External Network and Subnet	Create an External Network and Subnet on page 24
10. Create Host Aggregates for Compute Nodes	Create Host Aggregates for the Compute Nodes on page 24
11. Create Nova Instance Flavors	Create Nova Node Instance Flavors on page 25
Hadoop Administrator Tasks	
12. Create a Hadoop Image	Create a Hadoop Image on page 27
13. Upload the Hadoop Image into the OpenStack Image Service	Upload the Hadoop Image to the OpenStack Image Service on page 28
14. Register the Image with OpenStack Sahara	Register the Image with the Sahara Image Registry on page 28
15. Create a Nova Key Pair for OpenStack Sahara	Create a Nova Key Pair for Sahara on page 28
16. Create a Tenant Network, Subnet, and Router	Create a Tenant Network, Subnet, and Router on page 29
17. Create Node Group Templates	Create Node Group Templates on page 29
18. Create a Hadoop Cluster Template	Create a Hadoop Cluster Template on page 34
19. Create the Hadoop Cluster	Create the Hadoop Cluster on page 34
20. Configure Cloudera Hadoop	Configure Cloudera Hadoop on page 35

Chapter

2

High-level Architecture

Topics:

- [Hardware Configuration](#)
- [Network Architecture](#)

Follow the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Architecture Guide](#), version 10.0.1, for the hardware configuration and node choices. The Dell EMC Ready Bundle for Red Hat OpenStack Platform supports Dell EMC PowerEdge R730xd servers as Compute nodes.

Hardware Configuration

Table 2: Hardware Configuration on page 18 displays the recommended Compute node configuration for Hadoop workloads from Dell EMC Ready Bundle for Red Hat OpenStack Platform version 10.0.1.

Table 2: Hardware Configuration

Compute Nodes with Local Storage		
Server Model	Dell EMC PowerEdge R730xd	
CPU	Intel™ E5-2690 v4	
Memory	128GB, 2400 MT/s	
HDD	2x 300GB, 2.5", SAS, HDD	16x 1.1TB, 2.5", SAS, HDD
HDD Configuration	H730 RAID 1 with flex bay drives 24, 25	JBOD drives 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15
Read Policy	Adaptive Read Ahead	-
Write Policy	Write Back	-
Disk Cache Policy	Default	-

Network Architecture

Figure 1: Simple Network Schematic on page 18 displays a simplistic network schematic. Follow the Dell EMC Ready Bundle for Red Hat OpenStack Platform Architecture Guide for the detailed network architecture.

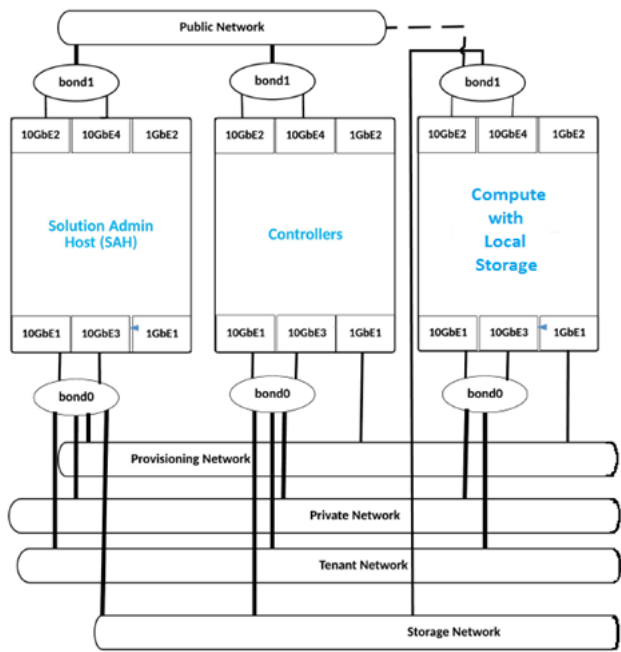


Figure 1: Simple Network Schematic

Chapter

3

Cloud Administrator Configurations

Topics:

- [*Configure RAID on Compute Nodes*](#)
- [*Deploy the Undercloud*](#)
- [*Map the Cinder Volume Service*](#)
- [*Deploy the Overcloud*](#)
- [*Create the Cinder Type for LVM Storage*](#)
- [*Create a Hadoop Tenant and Administrator User*](#)
- [*Create a Resource Configuration File for the Hadoop Administrator*](#)
- [*Set the Quota Values for the Hadoop Tenant*](#)
- [*Create an External Network and Subnet*](#)
- [*Create Host Aggregates for the Compute Nodes*](#)
- [*Create Nova Node Instance Flavors*](#)

The configurations described in this chapter are to be made by the Cloud Administrator, on the nodes specified in each step.



Note: These steps must be performed sequentially.

Configure RAID on Compute Nodes

Perform this procedure **before** deploying OpenStack.

To configure RAID on Compute Nodes with local storage:

1. On each Compute Node with local storage, start the RAID configuration utility by pressing **[CTRL + R]** during system boot-up.
2. Configure RAID 1 on the two disks in the flex bay. This virtual disk will host the OS for the Compute Node.
3. Put all other disks in the non-RAID/JBOD mode. These disks will serve the local storage on the Compute Nodes.

Deploy the Undercloud



Note: This procedure assumes an automated deployment. For a manual deployment, replace steps 1 and 2 with the Solution Admin Host deployment procedure as described in Chapter 3, *Software Manual Deployment*, of the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Software Deployment Guide](#). Then start here at step 3.

To deploy the Undercloud:

1. Kickstart the automated deployment as described in Chapter 2, *Software Automation Deployment*, of the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Software Deployment Guide Version 10.0.1](#).
2. After the Solution Admin Host (SAH) Node successfully deploys, initiate a SSH session to the SAH Node as the *root* user.
3. Move all files from the `~/JetPack/src/pilot/templates/big-data-configs` directory to `~/JetPack/src/pilot/templates`:

```
# cd ~/JetPack/src/pilot/templates/big-data-configs
# cp * ~/JetPack/src/pilot/templates/
```



Note: The updated templates contain the configurations for installation of OpenStack Sahara service, LVM backend and NUMA configurations. Please refer to Appendix A for more details on these files.

4. Open your hardware stamp's `.properties` file in a text editor.
5. Add the tag `skip_raid_config` : `"true"` for **all** Compute configurations. For example:

```
{
  "is_compute": "true",
  "idrac_ip": "192.168.110.15",
  "skip_raid_config": "true",
  "private_api_ip": "192.168.140.15",
  "storage_ip": "192.168.170.15",
  "tenant_tunnel_ip": "192.168.130.15"
},
```

6. Prepare your hardware stamp's `.ini` and `.properties` files ready as directed in the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Software Deployment Guide](#).

7. Start the Undercloud-only deployment by executing the following command, replacing `<path_to_settings.ini_file>` with the path to your `.ini` file:

```
# python ~/JetPack/src/deploy/osp_deployer/deployer.py -s
<path_to_settings.ini_file> -undercloud_only
```

Wait for the Undercloud deployment to complete. This may take up to 1 hour.

Map the Cinder Volume Service

To map the Cinder Volume Service on Compute Nodes:

1. After the Undercloud successfully deploys, initiate a SSH session to the Director Node as the `osp_admin` user.
2. On the Director Node, open the `~/pilot/templates/overcloud/roles_data.yaml` file in a text editor.
3. To configure and enable the OpenStack Cinder Volume Service on the Compute Nodes, add the following to the service list for the Compute Node's `ServicesDefault` parameter:

```
OS::TripleO::Services::BlockStorageCinderVolume
```

The resulting `Compute` Role section should look like this:

```
- name: Compute
  CountDefault: 1
  HostnameFormatDefault: '%stackname%-compute-%index%'
  ServicesDefault:
    - OS::TripleO::Services::BlockStorageCinderVolume
    - OS::TripleO::Services::CACerts
    - OS::TripleO::Services::CephClient
    - OS::TripleO::Services::CephExternal
    - OS::TripleO::Services::Timezone
    - OS::TripleO::Services::Ntp
    - OS::TripleO::Services::Snmp
    - OS::TripleO::Services::Sshd
    - OS::TripleO::Services::NovaCompute
    - OS::TripleO::Services::NovaLibvirt
    - OS::TripleO::Services::Kernel
    - OS::TripleO::Services::ComputeNeutronCorePlugin
    - OS::TripleO::Services::ComputeNeutronOvsAgent
    - OS::TripleO::Services::ComputeCeilometerAgent
    - OS::TripleO::Services::ComputeNeutronL3Agent
    - OS::TripleO::Services::ComputeNeutronMetadataAgent
    - OS::TripleO::Services::TripleoPackages
    - OS::TripleO::Services::TripleoFirewall
    - OS::TripleO::Services::NeutronSriovAgent
    - OS::TripleO::Services::OpenDaylightOvs
    - OS::TripleO::Services::SensuClient
    - OS::TripleO::Services::FluentdClient
    - OS::TripleO::Services::VipHosts
```

Deploy the Overcloud

To deploy the Overcloud:

1. Initiate a SSH session to the SAH Node as the *root* user.
2. Start the Overcloud-only deployment by executing the following command, replacing `<path_to_settings.ini_file>` with the path to your `.ini` file:

```
# cd ~/JetPack/src/deploy/osp_deployer
# python deployer.py -s <path_to_settings.ini_file> -overcloud_only
```



Note: Since the Undercloud VM is already deployed, the `-overcloud_only` parameter will ensure that only the Overcloud is deployed.

3. Wait for the Overcloud deployment to complete. This may take up to around 60 minutes.

Create the Cinder Type for LVM Storage

To create the Cinder type for Logical Volume Management (LVM) Storage:

1. After the Overcloud successfully deploys, initiate a SSH session to the Director Node as the *osp_admin* user.
2. Source the Overcloud admin credentials, replacing the `<stack-rc>` with your Overcloud stack RC file:

```
$ source <stack-rc>
```

3. Create a Cinder type for LVM storage:

```
$ cinder type-create LVM
```

4. Set the `volume_backend_name` so that the volume creation requests can be directed to the correct nodes:

```
$ cinder type-key LVM set volume_backend_name=lvm_backend
```

Create a Hadoop Tenant and Administrator User

The Hadoop cluster should be deployed by the Hadoop administrator in a separate tenant.

To create a Hadoop tenant and admin user:

1. If not currently open, initiate a SSH session to the Director Node as the *osp_admin* user.
2. Source the Overcloud admin credentials, replacing the `<stack-rc>` with your Overcloud stack RC file:

```
$ source <stack-rc>
```

3. Create a hadoop project, or tenant:

```
$ openstack project create hadoop
```

4. Create a *hadoop_admin* user, replacing `<PASSWORD>` with a suitable password:

```
$ openstack user create hadoop_admin --password <PASSWORD>
```

5. Assign the *hadoop_admin* user the Member role:

```
$ openstack role add --project hadoop --user hadoop_admin Member
```

Create a Resource Configuration File for the Hadoop Administrator

This Resource Configuration (RC) file will be used by the Hadoop administrator to perform administrative tasks in the `hadoop` tenant.

To Create an RC file for the Hadoop administrator:

1. If not currently open, initiate a SSH session to the Director Node as the `osp_admin` user.
2. Make a copy of your Overcloud admin RC file, with the name `hadoop_adminrc`:

```
$ cp ~/<Overcloud-name-rc-file> ~/hadoop_adminrc
```

3. Open the `hadoop_adminrc` file in a text editor:

```
$ vi ~/hadoop_adminrc
```

The file looks similar to this example:

```
export OS_NO_CACHE=True
export OS_CLOUDNAME=<stack-name>
export OS_AUTH_URL=http://100.82.37.60:5000/v2.0
export NOVA_VERSION=1.1
export COMPUTE_API_VERSION=1.1
export OS_USERNAME=hadoop_admin
export OS_PASSWORD=<PASSWORD>
export no_proxy=,100.82.37.60,192.168.120.121
export OS_PROJECT_NAME=hadoop
export PYTHONWARNINGS="ignore:Certificate has no, ignore:A true SSLContext
object is not available"
```

4. Replace the following information in the file:
 - `OS_CLOUDNAME` — Replace with your Overcloud stack name
 - `OS_USERNAME` — Replace with the `hadoop_admin` user name
 - `OS_PASSWORD` — Replace with the `hadoop_admin` user password
 - `OS_PROJECT_NAME` — Replace with the `hadoop` tenant name



Note: Leave all other fields unchanged.

Set the Quota Values for the Hadoop Tenant

The default quota does not provide enough resources, and must be modified for the Hadoop tenant.

To set the quota values for the Hadoop Tenant:

1. If not currently open, initiate a SSH session to the Director Node as the `osp_admin` user.
2. Source the Overcloud admin credentials, replacing the `<stack-rc>` with your Overcloud stack RC file:

```
$ source <stack-rc>
```

3. Update the quota values of Hadoop tenant by executing the following commands:

```
$ openstack quota set hadoop --gigabytes 50176
$ openstack quota set hadoop --volumes 50
$ openstack quota set hadoop --ram 270336
$ openstack quota set hadoop --cores 136
```

```
$ openstack quota set hadoop --instances 14
$ openstack quota set hadoop --secgroup-rules 200
```

Create an External Network and Subnet

The Hadoop tenant requires an external network and subnet.

To create an external network and subnet:

1. If not currently open, initiate a SSH session to the Director Node as the *osp_admin* user.
2. Source the Overcloud admin credentials, replacing the `<stack-rc>` with your Overcloud stack RC file:

```
$ source <stack-rc>
```

3. Execute the following commands, adjusting the `segmentation_id`, `allocation-pool`, `dns-nameserver`, `gateway`, and `network` address values to suit your environment:

```
$ neutron net-create public --router:external \
  --provider:network_type vlan --provider:physical_network physext \
  --provider:segmentation_id 372
$ neutron subnet-create --name external_sub --allocation-pool \
  start=100.82.37.70,end=100.82.37.120 --dns-nameserver 100.82.32.10 \
  --gateway 100.82.37.65 public 100.82.37.64/26
```



Note: Ensure that the external network/floating IP address pool has Internet access available with a working DNS server assigned to it.

Create Host Aggregates for the Compute Nodes

To create a separate host aggregate for each Compute Node:

1. If not currently open, initiate a SSH session to the Director Node as the *osp_admin* user.
2. Source the Overcloud admin credentials, replacing the `<stack-rc>` with your Overcloud stack RC file:

```
$ source <stack-rc>
```

3. Retrieve the list of known hosts:

```
$ openstack host list
```

4. Execute the following commands, copying and replacing `<fqdn_of_compute_node_x>` with the respective FQDNs as retrieved in step 3 on page 24:

```
$ openstack aggregate create --zone <fqdn_of_compute_node_0> nova0
$ openstack aggregate add host nova0 <fqdn_of_compute_node_0>

$ openstack aggregate create --zone <fqdn_of_compute_node_1> nova1
$ openstack aggregate add host nova1 <fqdn_of_compute_node_1>

$ openstack aggregate create --zone <fqdn_of_compute_node_2> nova2
$ openstack aggregate add host nova2 <fqdn_of_compute_node_2>
```


Create Nova Node Instance Flavors

This procedure creates Nova flavors for the Name Node, Data Node, and Cloudera Manager instances; with the resources listed in [Table 3: Nova Node Flavor Resources](#) on page 25.

Table 3: Nova Node Flavor Resources

Role	Flavor Name	RAM	vCPUs	Ephemeral Storage	Properties
Cloudera Manager Instance	m1.clouderamanager	8 GB	8	40 GB	hw:cpu_policy=dedicated, hw:numa_nodes=2
Name Node Instance	m1.namenode	16 GB	8	80 GB	hw:cpu_policy=dedicated, hw:numa_nodes=2
Data Node Instance	m1.datanode	20 GB	10	100 GB	hw:cpu_policy=dedicated, hw:numa_nodes=1

1. If not currently open, initiate a SSH session to the Director Node as the `osp_admin` user.
2. Source the Overcloud admin credentials, replacing the `<stack-rc>` with your Overcloud stack RC file:

```
$ source <stack-rc>
```

3. Create the Nova flavors by executing the following commands:

```
$ openstack flavor create m1.clouderamanager --ram 8192 \
--vcpus 8 --ephemeral 200 --public
$ openstack flavor set m1.clouderamanager \
--property hw:cpu_policy=dedicated
$ openstack flavor set m1.clouderamanager --property hw:numa_nodes=2

$ openstack flavor create m1.namenode --ram 16384 --vcpus 8 \
--ephemeral 200 --public
$ openstack flavor set m1.namenode --property hw:cpu_policy=dedicated
$ openstack flavor set m1.namenode --property hw:numa_nodes=2

$ openstack flavor create m1.datanode --ram 20480 --vcpus 10 \
--ephemeral 100 --public
$ openstack flavor set m1.datanode --property hw:cpu_policy=dedicated
$ openstack flavor set m1.datanode --property hw:numa_nodes=1
```

Chapter

4

Hadoop Administrator Configurations

Topics:

- [Deploy Hadoop Clusters with OpenStack Sahara](#)
- [Configure Cloudera Hadoop](#)

The configurations described in this chapter are to be made by the Hadoop Administrator. The following steps are specified for Cloudera Hadoop deployment and analogous steps can be done for other Hadoop distributions.



Note: These steps must be performed sequentially from a RHEL or CentOS Machine. See [Configuring a Node or VM for Hadoop Admin Use](#) on page 42 for more details.

Deploy Hadoop Clusters with OpenStack Sahara

Dell EMC tested and validated this procedure using CDH 5.7.0, and with RHEL 7.3 as the base VM operating system. If you prefer a different Cloudera version, replace the CDH 5.7.0 image with your preferred version image.

Perform the following procedures, in the order listed, to deploy Hadoop clusters using OpenStack Sahara:

1. [Create a Hadoop Image](#) on page 27
2. [Upload the Hadoop Image to the OpenStack Image Service](#) on page 28
3. [Register the Image with the Sahara Image Registry](#) on page 28
4. [Create a Nova Key Pair for Sahara](#) on page 28
5. [Create a Tenant Network, Subnet, and Router](#) on page 29
6. [Create Node Group Templates](#) on page 29
7. [Create a Hadoop Cluster Template](#) on page 34
8. [Create the Hadoop Cluster](#) on page 34

Create a Hadoop Image

To create Hadoop image based upon Cloudera 5.7.0:

1. If not currently open, initiate a SSH session to your VM.
2. Download a Red Hat Enterprise Linux Server 7.3 KVM guest image from https://access.redhat.com/downloads/content/69/ver=/rhel---7/7.3/x86_64/product-software.
 Note: Requires a Red Hat subscription.
 - a) Look for the *KVM Guest Image* in the list.
 - b) Right-click on the **Download Now** button, and copy the link address.
3. On the console of your VM, execute the following command replacing <paste_copied_link_here> with the link copied in the previous step.

```
$ wget <paste_copied_link_here> -O rhel-guest-image-7.3-35.x86_64.qcow2
```

- a) Wait until the download completes.
4. Configure your VM to use your Red Hat subscription details.



Note: This allows your VM to download the tools necessary for creating a Hadoop image from the downloaded Red Hat Enterprise Linux image.

```
$ export REG_USER=USERNAME
$ export REG_PASSWORD=PASSWORD
$ export REG_POOL_ID=POOLID
```

Where:

- USERNAME and PASSWORD are your Red Hat subscription credentials.
- POOLID is your Red Hat subscription's Pool ID.

- a) To determine your subscription's Pool ID, execute following command:

```
$ sudo subscription-manager list --available
```

This command displays important details about your subscription, including the Pool ID value.

5. Install the Sahara Image Elements package by executing the following command:

```
$ sudo pip install sahara-image-elements==5.0.1
```

6. Navigate to the directory where your downloaded Red Hat Enterprise Linux image resides.
7. Create the Hadoop RHEL 7.3-based image with Cloudera 5.7 packages by executing the following command:

```
$ sahara-image-create -p cloudera -v 5.7
```

A RHEL 7.3-based image with Cloudera 5.7 packages is created and saved in the current working directory, as `rhel7_sahara_cloudera_5.7.0.qcow2`.

Upload the Hadoop Image to the OpenStack Image Service

To upload the image created in [Create a Hadoop Image](#) on page 27 to the OpenStack Image Service:

1. If not currently open, initiate a SSH session to your VM.
2. Source the `hadoop_adminrc` credentials:

```
$ source ~/hadoop_adminrc
```

3. Upload the image by executing the following command:

```
$ openstack image create --disk-format qcow2 --container-format bare \
  --file rhel7_sahara_cloudera_5.7.0.qcow2 sahara-rhel7-cdh5-7
```

Register the Image with the Sahara Image Registry

To register the image:

1. If not currently open, initiate a SSH session to your VM.
2. Source the `hadoop_adminrc` credentials:

```
$ source ~/hadoop_adminrc
```

3. Register the image with the Sahara Image Registry, with the username `cloud-user`:

```
$ openstack dataprocessing image register sahara-rhel7-cdh5-7 \
  --username cloud-user
```

4. Tag the image to inform Sahara about the plugin, and the version with which it will be used:

```
$ openstack dataprocessing image tags add sahara-rhel7-cdh5-7 \
  --tags cdh 5.7.0
```

Create a Nova Key Pair for Sahara

OpenStack Sahara requires a key pair to access and configure the Hadoop cluster.

To create a Nova key pair for Sahara:

1. If not currently open, initiate a SSH session to your VM.
2. Source the `hadoop_adminrc` credentials:

```
$ source ~/hadoop_adminrc
```

3. Execute the following command:

```
$ openstack keypair create saharakey > saharapub
```

This command generates a key pair, and writes the key in the `saharapub` file. The `hadoop_admin` user can access the Hadoop cluster using this key.

Create a Tenant Network, Subnet, and Router

To create a network, subnet, and router for the Hadoop tenant:

1. If not currently open, initiate a SSH session to your VM.
2. Source the `hadoop_adminrc` credentials:

```
$ source ~/hadoop_adminrc
```

3. Create a tenant network, `tenant-net`, by executing the following command:

```
$ openstack network create tenant-net
```

4. Create a subnet, `tenant-subnet`, and assign it to the tenant network, adjusting the `subnet-range` and `dns-nameserver` values to suit your environment:

```
$ openstack subnet create tenant-subnet --network tenant-net \
  --subnet-range 192.168.201.0/24 --dhcp \
  --dns-nameserver <100.82.32.10>
```



Note: Ensure that the tenant network has a working DNS server assigned to it.

5. Create a tenant router, `tenant-router`:

```
$ openstack router create tenant-router
```

6. Attach an interface from the `tenant-network` to the `tenant-router`:

```
$ openstack router add subnet tenant-router tenant-subnet
```

7. Attach the external network to the `tenant-router`:

```
$ neutron router-gateway-set tenant-router public
```

Create Node Group Templates

The following procedures describe creating Node Group Templates:

- [Create a Name Node Group Template](#) on page 29
- [Create a Cloudera Manager Group Template](#) on page 30
- [Create Data Node Group Templates](#) on page 31

Create a Name Node Group Template

In this task you will create a Name Node group template with the following configurations:

Table 4: NameNode Group Template Configurations

Configuration	Value
Plugin	Cloudera Plugin
Hadoop Version	5.7.0

Configuration	Value
Name	NameNode
Flavor	m1.namenode
Availability Zone	<FQDN of Compute-0>
Storage Location	Cinder Volume
Volumes per node	1
Volumes Size (GB)	100
Volumes Type	LVM
Volumes Availability Zone	<FQDN of Compute-0>
Floating IP Pool	public
Auto Security Group	True
Node Processes	YARN_JOBHISTORY OOZIE_SERVER YARN_RESOURCEMANAGER HDFS_SECONDARYNAMENODE HDFS_NAMENODE

1. If not currently open, initiate a SSH session to your VM.
2. Source the *hadoop_adminrc* credentials:

```
$ source ~/hadoop_adminrc
```

3. Execute the following command on the console of your VM:

```
$ openstack dataprocessing node group template create --name NameNode \
--plugin cdh --plugin-version 5.7.0 \
--processes YARN_JOBHISTORY OOZIE_SERVER \
YARN_RESOURCEMANAGER HDFS_SECONDARYNAMENODE HDFS_NAMENODE \
--flavor m1.namenode --auto-security-group \
--availability-zone <FQDN of Compute-0> \
--floating-ip-pool public --volumes-per-node 1 \
--volumes-size 100 --volumes-type LVM \
--volumes-availability-zone <FQDN of Compute-0> \
--autoconfig
```

Create a Cloudera Manager Group Template

In this task you will create a Cloudera Manager group template with the following configurations:

Table 5: Cloudera Manager Group Template Configuration

Configuration	Value
Plugin	Cloudera Plugin
Hadoop Version	5.7.0
Name	ClouderaManager
Flavor	m1.clouderamanager

Configuration	Value
Availability Zone	<FQDN of Compute-1>
Storage Location	Cinder Volume
Volumes per node	1
Volumes Size (GB)	100
Volumes Type	LVM
Volumes Availability Zone	<FQDN of Compute-1>
Floating IP Pool	public
Auto Security Group	True
Node Processes	CLOUDERA_MANAGER

1. If not currently open, initiate a SSH session to your VM.
2. Source the `hadoop_adminrc` credentials:

```
$ source ~/hadoop_adminrc
```

3. Execute the following command on the console of your VM:

```
$ openstack dataprocessing node group template create \
--name ClouderaManager --plugin cdh --plugin-version 5.7.0 \
--processes CLOUDERA_MANAGER --flavor m1.clouderamanager \
--auto-security-group \
--availability-zone <FQDN of Compute-1> \
--floating-ip-pool public --volumes-per-node 1 \
--volumes-size 100 --volumes-type LVM \
--volumes-availability-zone <FQDN of Compute-1> \
--autoconfig
```

Create Data Node Group Templates

The following procedures describe creating three (3) separate DataNode group templates, one for each Availability Zone:

- [Create the Data Node 0 Group Template](#) on page 31
- [Create the Data Node 1 Group Template](#) on page 32
- [Create the Data Node 2 Group Template](#) on page 33

Create the Data Node 0 Group Template

In this task you will create a Data Node group template with the following configurations:

Table 6: Data Node 0 Group Template Configuration

Configuration	Value
Plugin	Cloudera Plugin
Hadoop Version	5.7.0
Name	DataNode0
Flavor	m1.datanode
Availability Zone	<FQDN of Compute-0>
Storage Location	Cinder Volume

Configuration	Value
Volumes per node	4
Volumes Size (GB)	1024
Volumes Type	LVM
Volumes Availability Zone	<FQDN of Compute-0>
Floating IP Pool	public
Auto Security Group	True
Node Processes	HDFS_DATANODE YARN_NODEMANAGER

1. If not currently open, initiate a SSH session to your VM.
2. Source the *hadoop_adminrc* credentials:

```
$ source ~/hadoop_adminrc
```

3. Execute the following command on the console of your VM, replacing <FQDN of Compute-0> to suit your environment:

```
$ openstack dataprocessing node group template create --name DataNode0 \
--plugin cdh --plugin-version 5.7.0 \
--processes HDFS_DATANODE YARN_NODEMANAGER --flavor m1.datanode \
--auto-security-group --availability-zone <FQDN of Compute-0> \
--floating-ip-pool public --volumes-per-node 4 \
--volumes-size 1024 --volumes-type LVM \
--volumes-availability-zone <FQDN of Compute-0> \
--autoconfig
```

Create the Data Node 1 Group Template

In this task you will create a Data Node group template with the following configurations:

Table 7: Data Node 1 Group Template Configuration

Configuration	Value
Plugin	Cloudera Plugin
Hadoop Version	5.7.0
Name	DataNode1
Flavor	m1.datanode
Availability Zone	<FQDN of Compute-1>
Storage Location	Cinder Volume
Volumes per node	4
Volumes Size (GB)	1024
Volumes Type	LVM
Volumes Availability Zone	<FQDN of Compute-1>
Floating IP Pool	public

Configuration	Value
Auto Security Group	True
Node Processes	HDFS_DATANODE YARN_NODEMANAGER

1. If not currently open, initiate a SSH session to your VM.
2. Source the `hadoop_adminrc` credentials:

```
$ source ~/hadoop_adminrc
```

3. Execute the following command on the console of your VM, replacing <FQDN of Compute-1> to suit your environment:

```
$ openstack dataprocessing node group template create --name DataNode1 \
--plugin cdh --plugin-version 5.7.0 \
--processes HDFS_DATANODE YARN_NODEMANAGER --flavor m1.datanode \
--auto-security-group --availability-zone <FQDN of Compute-1> \
--floating-ip-pool public --volumes-per-node 4 \
--volumes-size 1024 --volumes-type LVM \
--volumes-availability-zone <FQDN of Compute-1> \
--autoconfig
```

Create the Data Node 2 Group Template

In this task you will create a Data Node group template with the following configurations:

Table 8: Data Node 2 Group Template Configuration

Configuration	Value
Plugin	Cloudera Plugin
Hadoop Version	5.7.0
Name	DataNode2
Flavor	m1.datanode
Availability Zone	<FQDN of Compute-2>
Storage Location	Cinder Volume
Volumes per node	4
Volumes Size (GB)	1024
Volumes Type	LVM
Volumes Availability Zone	<FQDN of Compute-2>
Floating IP Pool	public
Auto Security Group	True
Node Processes	HDFS_DATANODE YARN_NODEMANAGER

1. If not currently open, initiate a SSH session to your VM.

2. Source the *hadoop_adminrc* credentials:

```
$ source ~/hadoop_adminrc
```

3. Execute the following command on the console of your VM, replacing <FQDN of Compute-2> to suit your environment:

```
$ openstack dataprocessing node group template create --name DataNode2 \
--plugin cdh --plugin-version 5.7.0 \
--processes HDFS_DATANODE YARN_NODEMANAGER --flavor m1.datanode \
--auto-security-group --availability-zone <FQDN of Compute-2> \
--floating-ip-pool public --volumes-per-node 4 \
--volumes-size 1024 --volumes-type LVM \
--volumes-availability-zone <FQDN of Compute-2> \
--autoconfig
```

Create a Hadoop Cluster Template

To create the Cluster template:

1. If not currently open, initiate a SSH session to your VM.
2. Source the *hadoop_adminrc* credentials:

```
$ source ~/hadoop_adminrc
```

3. Create a Cluster template with the configurations listed in:
 - [Table 9: Cluster Template Configurations](#) on page 34

Table 9: Cluster Template Configurations

Parameter	Value	
Plugin	Cloudera Plugin	
Hadoop Version	5.7.0	
Name	Bravo	
Node Groups	Name	Quantity
	ClouderaManager	1
	NameNode	1
	DataNode0	4
	DataNode1	4
	DataNode2	4

4. Execute the following command on the console of your VM:

```
$ openstack dataprocessing cluster template create --name Bravo \
--nodegroups ClouderaManager:1 NameNode:1 DataNode0:4 \
DataNode1:4 DataNode2:4 --autoconfig
```

Create the Hadoop Cluster

To create the Hadoop Cluster:

1. If not currently open, initiate a SSH session to your VM.

2. Source the `hadoop_adminrc` credentials:

```
$ source ~/hadoop_adminrc
```

3. Create a Cluster template with the parameters listed in:

- [Table 10: Hadoop Cluster Parameters](#) on page 35

Table 10: Hadoop Cluster Parameters

Parameter	Value
Plugin	Cloudera Plugin
Hadoop Version	5.7.0
Name	ClouderaJS
Cluster Template	Bravo
Base Image	Name of Image (default = sahara-rhel7-cdh5-7)
Keypair	saharakey
Management Network	Name of Tenant Network

4. Execute the following command on the console of your VM:

```
$ openstack dataprocessing cluster create --name ClouderaJS \
  --clustertemplate Bravo --image sahara-rhel7-cdh5-7 \
  --user-keypair saharakey --neutron-network tenant-net
```

5. Wait for the cluster deployment to complete. This may take around 25 minutes.

During the deployment, the instances in the cluster get dedicated vCPU pins from the physical hosts and volumes get created and attached to those instances automatically under the covers. When the cluster is successfully deployed you can find the Cloudera Manager login credentials in the **General Info** tab, which can be accessed by clicking on the name of the deployed cluster.

Configure Cloudera Hadoop

To configure the Hadoop cluster:

1. Log into the Cloudera Manager GUI with the credentials found in [Create the Hadoop Cluster](#) on page 34.
2. Configure and tune the Cloudera Hadoop cluster with the values presented in [Table 11: Hadoop Cluster Configuration Values](#) on page 35.

Table 11: Hadoop Cluster Configuration Values

S. No.	Configuration Name	Value
1	dfs.replication	3
2	dfs.blocksize	512 MB
3	mapreduce.map.cpu.vcores	1
4	mapreduce.map.memory.mb	1024 MB

S. No.	Configuration Name	Value
5	mapreduce.reduce.cpu.vcores	1
6	mapreduce.reduce.memory.mb	2048 MB
7	yarn.scheduler.minimum-allocation-vcores	1
8	yarn.scheduler.minimum-allocation-mb	1024 MB
9	yarn.scheduler.increment-allocation-vcores	1
10	yarn.scheduler.increment-allocation-mb	512 MB
11	yarn.nodemanager.resource.memory-mb	20 GB
12	yarn.nodemanager.resource.cpu-vcores	10
13	yarn.scheduler.maximum-allocation-mb	20 GB
14	yarn.scheduler.maximum-allocation-vcores	10
15	yarn.app.mapreduce.am.resource.mb	2048 MB
16	mapreduce.map.sort.spill.percent	0.8
17	mapreduce.task.io.sort.mb	256 MB
18	mapreduce.job.reduce.slowstart.completedmaps	0.8
19	dfs.datanode.handler.count	30
20	dfs.namenode.handler.count	49

The cluster is now ready to run Hadoop Applications with near bare-metal performance.

Appendix

A

Configuration Files Notes

Topics:

- [*dell_environment.yaml*](#)
- [*first_boot.yaml*](#)
- [*post_install.yaml*](#)

This appendix explains important sections of the templates that are required to install and configure the Dell EMC Ready Bundle for Red Hat OpenStack Platform for Hadoop.

dell_environment.yaml

Take note of the following sections of the `dell_environment.yaml` configuration file.

```
resource_registry:
  OS::TripleO::NodeUserData: ./first-boot.yaml
  OS::TripleO::NodeExtraConfigPost: ./post-install.yaml1
  OS::TripleO::Services::SaharaApi: ./overcloud/puppet/services/sahara-
    api.yaml2
  OS::TripleO::Services::SaharaEngine: ./overcloud/puppet/services/sahara-
    engine.yaml

cinder::config::cinder_config:3
  #LVM
  lvm_backend/volume_backend_name:
    value: lvm_backend
  lvm_backend/volume_driver:
    value: cinder.volume.drivers.lvm.LVMVolumeDriver
  lvm_backend/volume_group:
    value: cinder-volumes
  lvm_backend/iscsi_protocol:
    value: iscsi
  lvm_backend/iscsi_helper:
    value: lioadm
  lvm_backend/volume_clear:
    value: none
  lvm_backend/volume_clear_size:
    value: 50
  #LVM-END
  cinder_user_enabled_backends: ['lvm_backend']
  CinderEnableRbdBackend: False
  ControllerExtraConfig:
    cinder::config::cinder_config:
      #LVM
      lvm_backend/volume_backend_name:
        value: lvm_backend
      lvm_backend/volume_driver:
        value: cinder.volume.drivers.lvm.LVMVolumeDriver
      #LVM-END
      DEFAULT/scheduler_default_filters:
        value:
          'AvailabilityZoneFilter,CapacityFilter,CapabilitiesFilter,InstanceLocality'
      cinder_user_enabled_backends: ['lvm_backend']
      CinderEnableRbdBackend: True
```

¹ Including the `post-install.yaml` file, which contains the custom configurations that must be pushed to the nodes after Overcloud deployment.

² The OpenStack Sahara API and Engine Services are being linked to their respective Heat templates in the `/usr/share/openstack-tripleo-heat-templates/puppet/services` directory. They are disabled by default. This ensures that the Sahara API and Engine services become enabled and configured on the Controller Nodes during Overcloud deployment.

³ LVM storage backend configurations for Cinder Block Storage are being specified, including Volume Backend Name, Volume Driver, and Volume Group Name. These configurations are pushed to the `/etc/cinder/cinder.conf` file on the Compute and Controller nodes.

first_boot.yaml

Take note of the following section of the `first_boot.yaml` configuration file.

```
resources:
  userdata:
    type: OS::Heat::MultipartMime
    properties:
      parts:
        - config: {get_resource: wipe_disk}
        - config: {get_resource: numa_config_grub}

  numa_config_grub:4
    type: OS::Heat::SoftwareConfig
    properties:
      config: |
        #!/bin/bash
        if [[ $HOSTNAME =~ "compute" ]]; then
        {
          # Destroy all previous logical volumes and volume groups
          vgs --noheadings -o vg_name | awk '{print $1}' | while read lvg;
          do
            cmd="lvremove -f $lvg"
            echo $cmd
            $cmd

            cmd="vgremove -f $lvg"
            echo $cmd
            $cmd
          done

          # Create "cinder-volume" volume group and add all disks to this
          volume group
          let count=1
          echo $count
          lsblk -dno NAME,TYPE | \
          while read disk type; do
            # Skip if the device type isn't "disk" or if it's mounted
            [ "${type}" == "disk" ] || continue
            device="/dev/${disk}"
            if grep -q ^${device}[1-9] /proc/mounts; then
              echo "Skipping ${device} because it's mounted"
              continue
            fi
            if [ $count -eq 1 ]; then
              cmd="vgcreate cinder-volumes $device"
              echo $cmd
              $cmd
            else
              # Add all disks to cinder-volumes volume group
              cmd="vgextend cinder-volumes $device"
```

⁴ Including a bash script that is executed on the Compute Nodes upon first boot. This script destroys all the previous logical volumes and volume groups that might exist on the Compute Nodes; creates a "cinder-volume" volume group, and adds all data disks to this volume group. This script also performs the Grub configurations required for NUMA enablement. It calculates the available vCPUs on the host, isolates 8 sibling vCPUs for the physical host, and dedicates the rest to be used for Instances.

```

        echo $cmd
        $cmd
    fi
    let count++
done

# Grub Configurations for NUMA
# echo "I'm a Compute Node"
set -x
let num_cpus=$(lscpu | grep CPU\s\): | grep -v NUMA | awk
'{print $2}')
let i=$((num_cpus/2))
host_vcpu_list="0,1,2,3,$i,$((i+1)),$((i+2)),$((i+3))"
for vcpu in `seq 4 $((i-1))` `seq $((i+4)) $((num_cpus-1))`; do
    isolcpu+=${vcpu},
done
echo host_vcpu_list=$host_vcpu_list
isolcpu=$(echo $isolcpu | sed 's/,,$//')
echo isolcpu=$isolcpu
sed '/GRUB_CMDLINE_LINUX/cGRUB_CMDLINE_LINUX="console=tty0
console=ttyS0,115200n8 crashkernel=auto rhgb quiet isolcpus=\'$isolcpu\'" -
i /etc/default/grub
grub2-mkconfig -o /boot/grub2/grub.cfg
shutdown -r now
} > /root/compute-first-boot.log 2>&1

```

post_install.yaml

Take note of the following section of the `post_install.yaml` configuration file.

```

config: |
    #!/bin/bash5

    if [[ $HOSTNAME =~ "compute" ]]; then
    {
        set -x
        # echo "I'm a Compute Node"
        let num_cpus=$(lscpu | grep CPU\s\): | grep -v NUMA | awk
        '{print $2}')
        let i=$((num_cpus/2))
        host_vcpu_list="0,1,2,3,$i,$((i+1)),$((i+2)),$((i+3))"
        for vcpu in `seq 4 $((i-1))` `seq $((i+4)) $((num_cpus-1))`; do
            isolcpu+=${vcpu},
        done
        echo host_vcpu_list=$host_vcpu_list
        isolcpu=$(echo $isolcpu | sed 's/,,$//')
    }

```

⁵ Including a bash script that is executed on the Compute and Controller Nodes once the Openstack is deployed. This script performs the remaining configurations in the `/etc/nova/nova.conf` file required for enabling CPU Pinning with NUMA. On the Compute nodes, it calculates the available vCPUs, isolates 8 sibling vCPUs to be used by the physical host, dedicates the rest to be used by Nova, and restarts the `openstack-nova-compute` service for the changes to take effect. This script also configures `/etc/cinder/cinder.conf`, updates the storage availability zone, and restarts the `openstack-cinder-volume` service for the changes to take effect. On the Controller nodes, this script enables the `NUMATopologyFilter` along with the other filters which are required to use CPU Pinning with NUMA.


```

        echo isolcpu=$isolcpu
        crudini --set /etc/nova/nova.conf DEFAULT vcpu_pin_set $isolcpu
        systemctl restart openstack-nova-compute
        crudini --set /etc/cinder/cinder.conf DEFAULT
storage_availability_zone $HOSTNAME
        systemctl restart openstack-cinder-volume
    } > /root/enable_numa.txt 2>&1

    elif [[ $HOSTNAME =~ "controller" ]]; then
    {
        set -x
        # echo "I'm a Controller"
        crudini --set /etc/nova/nova.conf DEFAULT
scheduler_default_filters
RetryFilter,AvailabilityZoneFilter,RamFilter,ComputeFilter,\
        ComputeCapabilitiesFilter,ImagePropertiesFilter,CoreFilter,\
        NUMATopologyFilter,AggregateInstanceExtraSpecsFilter \
        systemctl restart openstack-nova-scheduler
    } > /root/enable_numa_filter.txt 2>&1

```

Appendix

B

Configuring a Node or VM for Hadoop Admin Use

Topics:

- [*Configure a Node or VM for Hadoop Admin Use*](#)

This appendix describes the requirements and configurations for setting up a Node or VM to be used by the OpenStack Cloud Administrator user or Hadoop Admin user to interact with the OpenStack Cloud.

Configure a Node or VM for Hadoop Admin Use

Hadoop Admin is essentially an OpenStack User, and requires a node to use the OpenStack Cloud APIs in order to perform the tasks specified in [Hadoop Administrator Configurations](#) on page 26.

System Requirements:

- **OS:** Red Hat Enterprise Linux Server 7.3 or CentOS 7.3
 - A valid Red Hat Subscription is required.
- **vCPUs:** — Minimum 2
- **RAM:** — Minimum 4 GB
- **Storage:** — Minimum 20 GB
- **Network:** — Must provide access to the OpenStack Cloud's Public Network.

Software Requirements:

The following packages are required to be installed on the VM:

- python-devel
- gcc
- python-setuptools
- pip
- python-openstackclient

To install these packages:

1. Execute the following commands:

```
$ sudo yum install -y python-setuptools python-devel gcc
$ sudo easy_install pip
$ sudo pip install python-openstackclient==3.2.1
```

2. Ask the Cloud Admin to provide the "hadoop_adminrc" file created in [Create a Resource Configuration File for the Hadoop Administrator](#) on page 23.
3. Place the file in the home directory.
4. Load the credentials by executing the following command:

```
$ source hadoop_adminrc
```

Appendix

C

Getting Help

Topics:

- [Technical References](#)
- [Contacting Dell EMC](#)
- [References](#)

This appendix details contact and reference information for the Dell EMC Ready Bundle for Red Hat OpenStack Platform.

Technical References

See the following for additional technical information:

- [Performance and Benchmarking References](#) on page 45
- [Cloudera 5.7.0 Documentation](#) on page 45
- [Red Hat Documentation](#) on page 45

Performance and Benchmarking References

“TPCx-HS on the Cloud!”, Performance Evaluation and Benchmarking. Traditional - Big Data - Internet of Things, 8th TPC Technology Conference, TPCTC 2016, New Delhi, India, September 5-9, 2016, Revised Selected Papers.

Please see: <https://link.springer.com/book/10.1007%2F978-3-319-54334-5>.

Cloudera 5.7.0 Documentation

For the latest Cloudera Distribution for Apache Hadoop documentation, please see:

<http://www.cloudera.com/documentation/enterprise/latest.html>

Apache Hadoop Documentation

For the latest Apache Hadoop documentation, please see:

<http://hadoop.apache.org/>

Red Hat Documentation

For Red Hat OpenStack Platform installation and deployment information, please see:

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/System_Administrators_Guide/index.html

Contacting Dell EMC

For customers in the United States, call 800-WWW-DELL (800-999-3355).



Note: If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell EMC product catalog.

Dell EMC provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell EMC for sales, technical support, or customer service issues:

1. Visit dell.com/support.
2. Click your country/region at the bottom of the page. For a full listing of country/region, click **All**.
3. Click **All Support** from the **Support** menu.
4. Select the appropriate service or support link based on your need.
5. Choose the method of contacting Dell EMC that is convenient for you.

References

Additional information can be obtained at <http://www.dell.com/en-us/work/learn/openstack-cloud> or by e-mailing openstack@dell.com.

If you need additional services or implementation help, please contact your Dell EMC sales representative.

To Learn More

For more information on the Dell EMC Ready Bundle for Red Hat OpenStack Platform visit <http://www.dell.com/learn/us/en/04/solutions/red-hat-openstack>.

Copyright © 2014-2017 Dell Inc. or its subsidiaries. All rights reserved. Trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Specifications are correct at date of publication but are subject to availability or change without notice at any time. Dell EMC and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell EMC's Terms and Conditions of Sales and Service apply and are available on request. Dell EMC service offerings do not affect consumer's statutory rights.

Dell EMC, the DELL EMC logo, the DELL EMC badge, and PowerEdge are trademarks of Dell Inc.