

Dell EMC Ready Bundle for Red Hat OpenStack

Software Automated Deployment Guide

Version 6.0.1



Dell EMC Converged Platforms and Solutions

Contents

List of Tables.....	iii
Trademarks.....	4
Notes, Cautions, and Warnings.....	5
Chapter 1: Overview.....	6
Deployment Methodology.....	7
Intended Audience.....	7
Before You Begin.....	7
Prerequisites.....	7
Dependencies.....	8
Chapter 2: Red Hat Subscriptions.....	9
Red Hat Subscription Manager Pool IDs.....	10
Chapter 3: Automation Configuration Files.....	12
Downloading and Extracting Automation Files.....	13
Chapter 4: Preparing and Deploying the Solution Admin Host.....	14
Preparing the Solution Admin Host Deployment.....	15
Deploying the SAH Node.....	16
Presenting the Image to the RHEL OS Installation Process.....	16
Chapter 5: Deploying the Undercloud and the OpenStack Cluster.....	18
Deploying and Validating the Cluster.....	19
Appendix A: Example Files.....	21
Stamp-specific Initialization File Example.....	22
Stamp-specific Properties File Example.....	26
Appendix B: References.....	29
To Learn More.....	30

List of Tables

Table 1: Red Hat Subscription Entitlements.....10

Table 2: VLAN IDs.....15

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.

The Midokura® name and logo, as well as the MidoNet® name and logo, are registered trademarks of Midokura SARL.

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

Overview

Topics:

- [*Deployment Methodology*](#)
- [*Intended Audience*](#)
- [*Before You Begin*](#)
- [*Prerequisites*](#)
- [*Dependencies*](#)

This guide provides information necessary to deploy the Dell EMC Ready Bundle for Red Hat OpenStack, on Dell EMC PowerEdge R630 and Dell EMC PowerEdge R730xd servers with the Dell EMC PowerEdge H730 disk controller; and the network with Dell Networking S3048-ON and S4048-ON switches.

Deployment Methodology

To perform a deployment of the Dell EMC Ready Bundle for Red Hat OpenStack:

1. Use the [Dell EMC Ready Bundle for Red Hat OpenStack Hardware Deployment Guide](#).
2. Then, depending on the methodology that you prefer, use either the:
 - a. [Dell EMC Ready Bundle for Red Hat OpenStack Software Manual Deployment Guide](#), or
 - b. [Dell EMC Ready Bundle for Red Hat OpenStack Software Automated Deployment Guide](#) in order to perform an automated deployment using scripts and methods developed and validated by Dell EMC

Intended Audience

This guide assumes the reader is familiar with:

- OpenStack
- Dell EMC PowerEdge R630 and Dell EMC PowerEdge R730xd RAID and BIOS configuration
- Red Hat Enterprise Linux (RHEL)
- Red Hat OpenStack Platform (RHOSP) documentation
- Network Configuration
- The concepts and procedures in [Red Hat's Red Hat OpenStack Platform update/upgrade documentation](#)

Before You Begin

This guide assumes that you have racked the servers and networking hardware, and completed power and network cabling, as per the [Dell EMC Ready Bundle for Red Hat OpenStack Reference Architecture Guide](#).

The high-level steps required to install the Dell EMC Ready Bundle for Red Hat OpenStack using the automated installation procedures include:

1. Ensuring that your environment meets the [Prerequisites](#) on page 7
2. Determining the [Red Hat Subscription Manager Pool IDs](#) on page 10
3. [Downloading and Extracting Automation Files](#) on page 13
4. [Preparing the Solution Admin Host Deployment](#) on page 15
5. [Deploying the SAH Node](#) on page 16
6. [Deploying the Undercloud and the OpenStack Cluster](#) on page 18

Prerequisites

The following prerequisites must be satisfied before proceeding with an automated deployment of the Dell EMC Ready Bundle for Red Hat OpenStack:



Note: All nodes in the same roles must be of the same server models, with identical HDD, RAM, and NIC configurations. So, all Controller nodes must be identical to each other; all Compute nodes must be identical to each other; and so on. See the [Dell EMC Ready Bundle for Red Hat OpenStack Reference Architecture](#) for configuration options for each node role.

- Hardware racked and wired per the [Dell EMC Ready Bundle for Red Hat OpenStack Reference Architecture](#)
- Hardware configured as per the [Dell EMC Ready Bundle for Red Hat OpenStack Hardware Deployment Guide](#)
- Hardware is powered off after the hardware is configured per the [Dell EMC Ready Bundle for Red Hat OpenStack Hardware Deployment Guide](#)
- Internet access to Red Hat's subscription manager service and repositories
- Valid Red Hat subscriptions

Dependencies

For customers performing a self-installation, these files are available upon request from Dell EMC. Please contact your account representative, or email openstack@dell.com for instructions.

[Dell EMC Ready Bundle for Red Hat OpenStack Software Automated Deployment Guide](#) dependencies include:

- [Downloading and Extracting Automation Files](#) on page 13
- The automated install also requires that you have a RHEL 7.2 ISO. It can be downloaded from the Red Hat Customer Portal here: https://access.redhat.com/downloads/content/69/ver=rhel---7/7.2/x86_64/product-software

Chapter 2

Red Hat Subscriptions

Topics:

- *Red Hat Subscription Manager
Pool IDs*

Once all prerequisites have been met, you must determine the appropriate Red Hat subscription entitlements for each cluster node.

Red Hat Subscription Manager Pool IDs

You must determine the pool ID to use for the Automation Control System, the Solution Admin Host (SAH) and each node in the cluster before proceeding with the installation. To determine the pool IDs, you must have an existing server that is registered to the Red Hat Hosted Services. This server must also be registered using the same credentials as the ones being used in this environment.

1. Once the server is correctly registered, execute the following command to see the available subscription pools.

```
# subscription-manager list --all --available
```

The command will output a list of available pools. Each section of information lists what the subscription provides, its pool ID, how many are available, the type of system it is for, as well as other information.

2. Determine the correct pool ID needed for this environment and take note of it.



Note: Pay close attention to the **System Type**. The System Type can be *Virtual* or *Physical*. If necessary you can use a physical license for a virtual node. However, you cannot use a virtual license for a physical node.

```
# subscription-manager list --all --available
```

[OUTPUT ABBREVIATED]

```
Subscription Name: Red Hat Cloud Infrastructure, Standard (8-sockets)
Provides:          Red Hat Beta
                  Red Hat OpenStack Beta
                  JBoss Enterprise Application Platform
                  Red Hat Software Collections (for RHEL Server)
                  Red Hat Enterprise Virtualization
                  Oracle Java (for RHEL Server)
                  Red Hat OpenStack
                  Red Hat Enterprise MRG Messaging
                  Red Hat Enterprise Linux Server
                  Red Hat Enterprise Linux High Availability (for RHEL
Server)
                  Red Hat Software Collections Beta (for RHEL Server)
                  Red Hat Enterprise Linux Load Balancer (for RHEL Server)
                  Red Hat CloudForms
SKU:               MCT2861
Pool ID:           aaaa111bbb222ccc333ddd444eee5556
Available:         7
Suggested:         1
Service Level:     Standard
Service Type:      L1-L3
Multi-Entitlement: No
Ends:              09/23/2015
System Type:       Physical
```

[OUTPUT ABBREVIATED]

The above output shows a subscription that contains the Red Hat OpenStack entitlement. The required entitlement types for each node are shown in [Table 1: Red Hat Subscription Entitlements](#) on page 10.

Table 1: Red Hat Subscription Entitlements

Node Role	Entitlement	System Type
Automation Control System	Red Hat Enterprise Linux Server	physical

Node Role	Entitlement	System Type
Solution Admin Host	Red Hat Enterprise Linux Server	physical
Director Node	Red Hat OpenStack	virtual
Red Hat Ceph Storage Admin Node	Red Hat Ceph Storage Calamari	physical (no virtual available at this time)
Controller Node	Red Hat OpenStack	physical
Compute Node	Red Hat OpenStack	physical
Storage Node	Red Hat Ceph Storage	physical

Chapter

3

Automation Configuration Files

Topics:

- [*Downloading and Extracting Automation Files*](#)

This chapter details obtaining the required configuration files.

Downloading and Extracting Automation Files

The following procedure installs the required configuration files and scripts

1. Log into your RHEL 7.2 system as user *root*.
2. Create the */root/JetStream* directory.
3. Download the *automation-JS-6.0.1.tar.gz* file to the */root/JetStream* directory.
4. Change the working directory to */root/JetStream*.
5. Extract the tar file contents:

```
# tar -xvf automation-JS-6.0.1.tar.gz
```

6. Download or copy the ISO of the Red Hat Enterprise Linux Server 7 installation DVD to *~/JetStream/rhel72.iso*.

Chapter

4

Preparing and Deploying the Solution Admin Host

Topics:

- [*Preparing the Solution Admin Host Deployment*](#)
- [*Deploying the SAH Node*](#)

This topic describes preparing for, and performing, the Solution Admin Host (SAH) deployment.

Preparing the Solution Admin Host Deployment

1. Log in as your RHEL 7.2 system as the *root* user.
2. Change the working directory to `/root/JetStream/deploy-auto/osp_deployer/settings/`.
3. Copy the `settings_sample.ini` and `sample.properties` files to stamp-specific files (e.g., `acme.ini` and `acme.properties`), and place them in the `~/JetStream` directory:

```
# cp ~/deploy-auto/osp_deployer/settings/sample.properties ~/JetStream/
acme.properties
# cp ~/deploy-auto/osp_deployer/settings/settings_sample.ini ~/JetStream/
acme.ini
```

4. Edit your hardware stamp's `.ini` and `.properties` files to match your hardware stamp documentation (i.e., a Solution Workbook). Use a text editor of your choice; our example uses `vi`:

```
# vi ~/JetStream/acme.ini
```

5. Change the values in your stamp-specific `.ini` file to match your specific environment. You must supply a value for each `CHANGEME` token in the file. In addition, the IP addresses and the Subscription Manager Pool IDs must be changed to match your deployment. Each section will have a brief description of the attributes.

See [Stamp-specific Initialization File Example](#) on page 22 for an example stamp-specific `.ini` file.



Note: Instance HA and Ephemeral storage selections must both be performed at installation time, as Dell EMC does not support changing these settings post-installation. See the [Technical Guide - Using Instance High Availability in the Dell EMC Ready Bundle for Red Hat OpenStack](#) for more information.

6. Edit the stamp-specific `.properties` file:

```
# vi ~/JetStream/acme.properties
```

7. Change the values in your `.properties` file to match your specific environment. You must supply a value for IP addresses, host names, passwords, interfaces, and storage OSDs/journals.

The examples in this file are based on the [Dell EMC Ready Bundle for Red Hat OpenStack Reference Architecture Guide](#), and the installation scripts rely on the VLAN IDs as specified in this file. For example, the Private API VLAN ID is 140. So, all addresses on the Private API network must have 140 as the third octet (e.g., 192.168.140.114). [Table 2: VLAN IDs](#) on page 15 below lists the VLAN IDs.

See [Stamp-specific Properties File Example](#) on page 26 for an example stamp-specific `.properties` file.



Caution: Deviating from the VLAN IDs listed below will cause the installation to fail.

Table 2: VLAN IDs

VLAN ID	Name
110	Management/Out of Band (OOB) Network (iDRAC)
120	Provisioning Network
140	Private API Network
170	Storage Network
180	Storage Clustering Network
190	Public API Network

VLAN ID	Name
191	External Tenant Network (Used for floating IP addresses)
201-250	Internal Tenant Network

The *anaconda_ip* is used for the initial installation of the SAH node, and requires an address that can access the Internet to obtain Red Hat software. The *anaconda_iface* must be a dedicated interface that is only used for this purpose, and is not used in any other part of the configuration.

8. Update your `python` path:

```
# export PYTHONPATH=/usr/bin/python:/lib/python2.7:/lib/python2.7/site-
packages:~/JetStream/deploy-auto
```

9. You can install the SAH node using either of the following methods:

a. Using a physical USB key:

- a. Plug your USB key into your RHEL 7.2 system.
- b. Run the setup script to prepare your USB key, passing in the USB device ID (*/dev/sdb* in the example below).



Note: Use full paths.

```
# cd ~/JetStream/deploy-auto/setup
# python setup_usb_idrac.py -s /root/JetStream/acme.ini -usb_key /
dev/sdb
```

b. Using an iDRAC virtual media image file. This requires your RHEL 7.2 system to have access to the iDRAC consoles to attach the image.

- a. Run the setup script to generate an image file that can later be attached to the SAH node.



Note: Use full paths.

```
# cd ~/JetStream/deploy-auto/setup
# python setup_usb_idrac.py -s /root/JetStream/acme.ini -
idrac_vmedia_img
```

- b. The output will be an image file generated in *~/osp_ks.img*.

Deploying the SAH Node

You can deploy the SAH node by one of two methods:

- Using a physical USB key generated above, plugged into the SAH node
- Using an iDRAC virtual media image generated above, made available using the *Map the image as Removable Media* option on the iDRAC.

Proceed to [Presenting the Image to the RHEL OS Installation Process](#) on page 16.

Presenting the Image to the RHEL OS Installation Process

1. Attach the RHEL 7.2 ISO as a virtual CD/DVD using the VirtualMedia-> MapVirtual CD/DVD option
2. Set the SAH node to boot from the virtual CD/DVD using the Next Boot-> Virtual CD/DVD/ISO option
3. Boot the SAH node.
 - a. At the installation menu, select the **Install** option. *Do not press the [Enter] key.*
 - b. Press the **Tab** key.

- c. Move the cursor to the end of the line that begins with `vmlinux`.
- d. Append the following to the end of the line:



Note: The device `sdb` can change, depending upon the quantity of disks being presented to the installation environment. These instructions assume that a single disk is presented. If otherwise, adjust accordingly.

```
ks=hd:sdb:/JetStream/osp-sah.ks
```

4. Press the **[Enter]** key to start the installation.



Note: It may take a few minutes before progress is seen on the screen. Press the **[ESC]** key at the memory check to speed up the process.

Chapter

5

Deploying the Undercloud and the OpenStack Cluster

Topics:

- [*Deploying and Validating the Cluster*](#)

Now that the SAH node is installed you can deploy and validate the rest of the Dell EMC Ready Bundle for Red Hat OpenStack nodes.

Deploying and Validating the Cluster

To deploy and validate the rest of the cluster:

1. Log in through the iDRAC console as *root*, or *ssh* into the SAH node.
2. Mount the USB media:

```
# mount /dev/sdb /mnt
```

3. Copy all the files locally:

```
# cp -rf /mnt/JetStream /root
```

4. Start a *tmux* session to avoid losing progress if the connection drops:

```
# tmux
```

5. Run the deployment, passing in your settings files. There are also some post-deployment validation options in the `[Bastion Settings]` group of the stamp-specific initialization file you should consider prior to deployment:

- a. *run_sanity* - If set to *true* the *sanity_test.sh* script will be executed that will verify the basic functionality of your overcloud deployment.
- b. *run_tempest* - If set to *true* the Tempest integration test suite will be executed against your overcloud deployment.



Note: Tempest requires that the sanity test must be run first so *run_sanity*, above, must also be set to *true*. For more information about Tempest please refer to the [Dell EMC Ready Bundle for Red Hat OpenStack Overcloud Validation Guide](#).

- c. *tempest_smoke_only* - If *run_tempest*, above, is set to *true* this option, which is set to *true* by default, will cause Tempest to run only a small subset of the test suite, where the tests are tagged as *"smoke"*. If set to *false* the entire Tempest suite will be run, which can take an hour or more to complete.

```
# cd /root/JetStream/deploy-auto/osp_deployer
# python deployer.py -s /root/JetStream/acme.ini
```

6. For installation details, execute a *tail* command on the */deployer.log.xxx* file on the SAH node. For example:

```
# tail -f /auto_results/deployer.log.2016.04.26-09.09
```

7. If issues are discovered during the installation process:

- a. Identify the issue in the *deployer.log*
- b. Address the issue.
- c. Rerun the *python deployer.py* command above.

8. If the installation is successful, the *deployment_summary.log* file will display some useful information for accessing the Dell EMC Ready Bundle for Red Hat OpenStack.

```
# cd auto_results
# cat deployment_summary.log
```

The output will appear similar to this:

```
=====
### nodes ip information ###
### Controllers ###
```

```

mercury-controller-0 :
- provisioning ip : 192.168.120.134
- nova private ip : 192.168.140.22
- nova public ip  : 192.168.190.32
- storage ip      : 192.168.170.21
mercury-controller-1 :
- provisioning ip : 192.168.120.128
- nova private ip : 192.168.140.26
- nova public ip  : 192.168.190.34
- storage ip      : 192.168.170.26
mercury-controller-2 :
- provisioning ip : 192.168.120.126
- nova private ip : 192.168.140.23
- nova public ip  : 192.168.190.33
- storage ip      : 192.168.170.22
### Compute ###
mercury-compute-0 :
- provisioning ip : 192.168.120.133
- nova private ip : 192.168.140.25
- storage ip      : 192.168.170.25
mercury-compute-1 :
- provisioning ip : 192.168.120.132
- nova private ip : 192.168.140.24
- storage ip      : 192.168.170.24
### Storage ###
mercury-cephstorage-0 :
- provisioning ip : 192.168.120.131
- storage cluster ip : 192.168.180.20
- storage ip      : 192.168.170.23
mercury-cephstorage-1 :
- provisioning ip : 192.168.120.127
- storage cluster ip : 192.168.180.21
- storage ip      : 192.168.170.27
=====
OverCloud Horizon      : http://192.168.190.31:5000/v2.0

OverCloud admin password : GbXkxG99KtxHtbmTVzMK9QnUv

```

Appendix

A

Example Files

Topics:

- [*Stamp-specific Initialization File Example*](#)
- [*Stamp-specific Properties File Example*](#)

This appendix contains examples of files that you will edit and use in the deployment.

Stamp-specific Initialization File Example

```
#####
#
#Copy & rename this file for your own stamp.
#Review ALL settings below, pay particular attention to paths/ip's etc..#
#
#####

[Cluster Settings]
# Only developers should set to false.
enable_version_locking=true

# This pathname must be the full path to the properties file which
# describes the cluster. You should copy *this* sample settings file
# (settings_sample.ini) and the sample properties file
# (sample.properties) to another directory, and customize them for your
# cluster. Then use the path to your customized properties file here.
cluster_nodes_configuration_file=/root/JetStream/acme.properties

# Domain name for the cluster (i.e., mycluster.lab)
domain=domain.net

# DRAC credentials with IPMI privilege for the nodes
ipmi_user=root
ipmi_password=xxxxxxx

# User for the undercloud/overcloud installation
director_install_user=osp_admin
director_install_user_password=xxxxxxx

# Name of the Overcloud.
# The nodes hostnames will be prepended with the given name and a dash
overcloud_name=overcloud
use_internal_repo=false
# Semi-colon ( ; ) separated list of internal repos to use, if needed.
# Typically
# not used.
internal_repos_locations=CHANGEME_INTERNAL_REPO_URL

# Subscription Manager account info for registering Red Hat subscriptions
subscription_manager_user=xxxxxxx
subscription_manager_password=xxxxxxxxxxxxxxxxx

# The following pool IDs provide different collections of repositories.
# Each is labeled with possible subscription names.

# Red Hat Enterprise Linux Self-Supported Business Partner NFR
subscription_manager_pool_sah=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx44f5

# Red Hat Enterprise Linux OpenStack Platform Self-Supported Business
# Partner NFR
subscription_manager_pool_vm_rhel=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx454a

# Red Hat Ceph Storage Business Partner Self-Supported NFR, (Physical Node)
subscription_manager_vm_ceph=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx7826

subscription_check_retries=20
ntp_servers=0.centos.pool.ntp.org
time_zone=America/Chicago
```

```

# Use static ip addresses for the overcloud nodes if set to true (ips need to
# be defined in the .properties)
# Use dhcp if set to false (ips not required in the .properties)
overcloud_static_ips=true

# External network details
external_netmask=255.255.255.0
external_gateway=10.7.101.1

# Nova public network details
public_api_network=192.168.190.0/24
public_api_vlanid=190
public_api_gateway=192.168.190.1
public_api_netmask=255.255.255.0
public_api_allocation_pool_start=192.168.190.121
public_api_allocation_pool_end=192.168.190.250

# Private API network details
private_api_network=192.168.140.0/24
private_api_vlanid=140
private_api_netmask=255.255.255.0
private_api_allocation_pool_start=192.168.140.121
private_api_allocation_pool_end=192.168.140.250

# Storage network details
storage_network=192.168.170.0/24
storage_vlanid=170
storage_netmask=255.255.255.0
storage_allocation_pool_start=192.168.170.125
storage_allocation_pool_end=192.168.170.250

# Provisioning network details
provisioning_network=192.168.120.0/24
provisioning_vlanid=120
provisioning_netmask=255.255.255.0
provisioning_gateway=192.168.120.1
provisioning_net_dhcp_start=192.168.120.121
provisioning_net_dhcp_end=192.168.120.250
discovery_ip_range=192.168.120.21,192.168.120.120

# Storage cluster network details
storage_cluster_network=192.168.180.0/24
storage_cluster_vlanid=180
storage_cluster_allocation_pool_start=192.168.180.121
storage_cluster_allocation_pool_end=192.168.180.250

# Management network details
management_network=192.168.110.0/24
managment_vlanid=110
managment_netmask=255.255.255.0
name_server=8.8.8.8

# Tenant network details
# Not used unless you wish to configure Generic Routing Encapsulation (GRE)
# networks.
tenant_network=192.168.130.0/24
tenant_network_allocation_pool_start=192.168.130.121
tenant_network_allocation_pool_end=192.168.130.250

# Nova Private network details
tenant_vlan_range=201:220

```

```

# Use static VIPs ip addresses for the overcloud
use_static_vips=true

# The following VIP settings apply if the above use_static_vips is enabled.

# VIP for the redis service on the Private API api network
# Note that this IP must lie OUTSIDE the private_api_allocation_pool_start/
end range
redis_vip=192.168.140.251

# VIP for the provisioning network
# Note that this IP must lie INSIDE the provisioning_net_dhcp_start/end
range
# but cannot be the first IP in that range
provisioning_vip=192.168.120.250

# VIP for the Private API network
# Note that this IP must lie INSIDE the private_api_allocation_pool_start/
end range
private_api_vip=192.168.140.250

# VIP for the Public API network
# Note that this IP must lie INSIDE the public_api_allocation_pool_start/end
range
public_api_vip=192.168.190.250

# VIP for the Storage network
# Note that this IP must lie INSIDE the storage_allocation_pool_start/end
range
storage_vip=192.168.170.250

# VIP for the Storage cluster network
# The Storage Clustering network is not connected to the controller nodes,
# so the VIP for this network must be mapped to the provisioning network
# Note that this IP must lie INSIDE the provisioning_net_dhcp_start/range
but
# cannot be the first IP in that range
storage_cluster_vip=192.168.120.249

# iDRAC doesn't always play nice due to a bug. Workaround built in the
deployer already, but occasionally fails.
# Option below is to use a custom instack.json (i.e., not using idracula)
use_custom_instack_json=false
custom_instack_json=n/a

# Bonding options configuration by node type
controller_bond_opts=802.3ad miimon=100
compute_bond_opts=802.3ad miimon=100
storage_bond_opts=802.3ad miimon=100

# Overcloud deployment timeout value - default is 120mns, but can be tweaked
here if required.
overcloud_deploy_timeout=120

# Interfaces per node type
controller_bond0_interfaces=em1 plp1
controller_bond1_interfaces=em2 plp2
controller_provisioning_interface=em3

compute_bond0_interfaces=em1 plp1

```



```

compute_bond1_interfaces=em2 plp2
compute_provisioning_interface=em3

storage_bond0_interfaces=em1 p2p1
storage_bond1_interfaces=em2 p2p2
storage_provisioning_interface=em3

#Default driver is DRAC.
use_ipmi_driver=true

# EQLX backend settings, if applicable
enable_eqlx_backend = false
eqlx_backend_name=CHANGEME_EQLX_GROUP_NAME
eqlx_san_ip=CHANGEME_SAN_IP
eqlx_san_login=CHANGEME_SAN_USERNAME
eqlx_san_password=CHANGEME_SAN_PASSWORD
eqlx_ch_login=CHANGEME_CHAP_USERNAME
eqlx_ch_pass=CHANGEME_CHAP_PASSWORD
eqlx_group_n=CHANGEME_EQLX_GROUP_NAME
eqlx_thin_provisioning=true
eqlx_pool=default
eqlx_use_chap=false

# Compellent parameters. See the Deployment Guide for description of the
  parameters.
enable_dellsc_backend=false
dellsc_backend_name=CHANGEME
dellsc_api_port=3033
dellsc_iscsi_ip_address=CHANGEME
dellsc_iscsi_port=3260
dellsc_san_ip=CHANGEME
dellsc_san_login=CHANGEME
dellsc_san_password=CHANGEME
dellsc_ssn=CHANGEME
dellsc_server_folder=cmpl_iscsi_servers
dellsc_volume_folder=cmpl_iscsi_volumes

# Set to true to enable Nova usage of Ceph for ephemeral storage.
# If set to false, Nova uses the storage local to the compute.
enable_rbd_backend=true

# Set to true to enable fencing
# Please refer to the following document for more details on fencing :
# Dell EMC Ready Bundle for Red Hat OpenStack Deployment Guide,
  Version 6.0.1
enable_fencing=false

# Set to true to enable instance HA
# Note : fencing must also be enabled (setting above)
# Please refer to following technical document for more details on Instance
  HA :
# Dell EMC Ready Bundle for Red Hat OpenStack Deployment Guide,
  Version 6.0.1
enable_instance_ha=false

# The list of RHSM repositories to enable to access the product,
# the list separator is ','.
rhsm_repos=

[Bastion Settings]
cloud_repo_dir=/root/JetStream/cloud_repo
rh172_iso=/root/JetStream/rhel72.iso
# If you want the sanity script to run on deployment completion (Appendix C,
  etc.), you may do so.

```

```

run_sanity=false
# If you want to run Tempest post-deployment, you may do so. The sanity
# script must also run to create networks for Tempest.
# For more on Tempest please see the Overcloud Validation Guide
run_tempest=false
tempest_smoke_only=true
# RDO cloud images
# Available to download @ https://access.redhat.com/downloads/content/191/
# ver=8/rhel---7/8/x86_64/product-software
discovery_ram_disk_image=/pathto/discovery-ramdisk-7.1.0-39.tar
overcloud_image=/pathto/overcloud-full-7.1.0-39.tar
# if option below is enabled, images will be pulled from the cdn (and the
# above x2 settings ignored)
pull_images_from_cdn=true

```

Stamp-specific Properties File Example

```

[
  {
    "is_sah": "true",
    "hostname": "sah",
    "idrac_ip": "192.168.110.20",
    "root_password": "xxxxxxxxxx",

    "anaconda_ip": "10.7.101.134",
    "anaconda_iface": "em4",

    "external_bond": "bond1",
    "external_slaves": "em2 plp2",
    "external_ip": "10.7.101.12",

    "private_bond": "bond0",
    "private_slaves": "em1 plp1",
    "provisioning_ip": "192.168.120.12",
    "storage_ip": "192.168.170.12",
    "public_api_ip": "192.168.190.12",
    "private_api_ip": "192.168.140.12",
    "managment_ip": "192.168.110.12"
  },
  {
    "is_director": "true",
    "hostname": "director",
    "root_password": "xxxxxxxxxx",

    "external_ip": "10.7.101.13",
    "provisioning_ip": "192.168.120.13",

    "managment_ip": "192.168.110.13",
    "public_api_ip": "192.168.190.13",

    "private_api_ip": "192.168.140.13"
  },
  {
    "is_ceph": "true",
    "hostname": "ceph",
    "root_password": "xxxxxxxxxx",

    "external_ip": "10.7.101.14",
    "storage_ip": "192.168.170.14"
  },
]

```

```

{
  "is_controller": "true",
  "idrac_ip": "192.168.110.21",

  "public_api_ip": "192.168.190.21",
  "private_api_ip": "192.168.140.21",
  "storage_ip": "192.168.170.21",
  "storage_cluster_ip": "192.168.180.21",
  "tenant_ip": "192.168.130.21"
},
{
  "is_controller": "true",
  "idrac_ip": "192.168.110.22",

  "public_api_ip": "192.168.190.22",
  "private_api_ip": "192.168.140.22",
  "storage_ip": "192.168.170.22",
  "storage_cluster_ip": "192.168.180.22",
  "tenant_ip": "192.168.130.22"
},
{
  "is_controller": "true",
  "idrac_ip": "192.168.110.23",

  "public_api_ip": "192.168.190.23",
  "private_api_ip": "192.168.140.23",
  "storage_ip": "192.168.170.23",
  "storage_cluster_ip": "192.168.180.23",
  "tenant_ip": "192.168.130.23"
},
{
  "is_compute": "true",
  "idrac_ip": "192.168.110.31",

  "private_api_ip": "192.168.140.31",
  "storage_ip": "192.168.170.31",
  "tenant_ip": "192.168.130.31"
},
{
  "is_compute": "true",
  "idrac_ip": "192.168.110.32",

  "private_api_ip": "192.168.140.32",
  "storage_ip": "192.168.170.32",
  "tenant_ip": "192.168.130.32"
},
{
  "is_compute": "true",
  "idrac_ip": "192.168.110.33",

  "private_api_ip": "192.168.140.33",
  "storage_ip": "192.168.170.33",
  "tenant_ip": "192.168.130.33"
},
{
  "is_ceph_storage": "true",
  "idrac_ip": "192.168.110.76",

  "storage_ip": "192.168.170.76",
  "storage_cluster_ip": "192.168.180.76",

```

```

        "osd_disks":[
            "":"/dev/sdd:/dev/sda",
            "":"/dev/sde:/dev/sda",
            "":"/dev/sdf:/dev/sda",
            "":"/dev/sgd:/dev/sda",
            "":"/dev/sdh:/dev/sdb",
            "":"/dev/sdi:/dev/sdb",
            "":"/dev/sdj:/dev/sdb",
            "":"/dev/sdk:/dev/sdb",
            "":"/dev/sdl:/dev/sdc",
            "":"/dev/sdm:/dev/sdc",
            "":"/dev/sdn:/dev/sdc",
            "":"/dev/sdo:/dev/sdc"
        ]
    },
    {
        "is_ceph_storage": "true",
        "idrac_ip": "192.168.110.77",

        "storage_ip": "192.168.170.77",
        "storage_cluster_ip": "192.168.180.77",

        "osd_disks":[
            "":"/dev/sdd:/dev/sda",
            "":"/dev/sde:/dev/sda",
            "":"/dev/sdf:/dev/sda",
            "":"/dev/sgd:/dev/sda",
            "":"/dev/sdh:/dev/sdb",
            "":"/dev/sdi:/dev/sdb",
            "":"/dev/sdj:/dev/sdb",
            "":"/dev/sdk:/dev/sdb",
            "":"/dev/sdl:/dev/sdc",
            "":"/dev/sdm:/dev/sdc",
            "":"/dev/sdn:/dev/sdc",
            "":"/dev/sdo:/dev/sdc"
        ]
    },
    {
        "is_ceph_storage": "true",
        "idrac_ip": "192.168.110.78",

        "storage_ip": "192.168.170.78",
        "storage_cluster_ip": "192.168.180.78",

        "osd_disks":[
            "":"/dev/sdd:/dev/sda",
            "":"/dev/sde:/dev/sda",
            "":"/dev/sdf:/dev/sda",
            "":"/dev/sgd:/dev/sda",
            "":"/dev/sdh:/dev/sdb",
            "":"/dev/sdi:/dev/sdb",
            "":"/dev/sdj:/dev/sdb",
            "":"/dev/sdk:/dev/sdb",
            "":"/dev/sdl:/dev/sdc",
            "":"/dev/sdm:/dev/sdc",
            "":"/dev/sdn:/dev/sdc",
            "":"/dev/sdo:/dev/sdc"
        ]
    }
]

```

Appendix

B

References

Topics:

- [To Learn More](#)

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 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.