



Red Hat OpenStack Platform 13

Advanced Overcloud Customization

Methods for configuring advanced features using Red Hat OpenStack Platform
director

Red Hat OpenStack Platform 13 Advanced Overcloud Customization

Methods for configuring advanced features using Red Hat OpenStack Platform director

OpenStack Team
rhos-docs@redhat.com

Legal Notice

Copyright © 2023 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

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.

All other trademarks are the property of their respective owners.

Abstract

This guide explains how to configure certain advanced features for a Red Hat OpenStack Platform enterprise environment using the Red Hat OpenStack Platform Director. This includes features such as network isolation, storage configuration, SSL communication, and general configuration methods.

Table of Contents

CHAPTER 1. INTRODUCTION	6
CHAPTER 2. UNDERSTANDING HEAT TEMPLATES	7
2.1. HEAT TEMPLATES	7
2.2. ENVIRONMENT FILES	8
2.3. CORE OVERCLOUD HEAT TEMPLATES	9
2.4. PLAN ENVIRONMENT METADATA	10
2.5. CAPABILITIES MAP	11
2.6. INCLUDING ENVIRONMENT FILES IN OVERCLOUD CREATION	12
2.7. USING CUSTOMIZED CORE HEAT TEMPLATES	13
2.8. JINJA2 RENDERING	16
CHAPTER 3. PARAMETERS	19
3.1. EXAMPLE 1: CONFIGURING THE TIMEZONE	19
3.2. EXAMPLE 2: ENABLING NETWORKING DISTRIBUTED VIRTUAL ROUTING (DVR)	19
3.3. EXAMPLE 3: CONFIGURING RABBITMQ FILE DESCRIPTOR LIMIT	19
3.4. EXAMPLE 4: ENABLING AND DISABLING PARAMETERS	20
3.5. EXAMPLE 5: ROLE-BASED PARAMETERS	20
3.6. IDENTIFYING PARAMETERS TO MODIFY	20
CHAPTER 4. CONFIGURATION HOOKS	23
4.1. FIRST BOOT: CUSTOMIZING FIRST BOOT CONFIGURATION	23
4.2. PRE-CONFIGURATION: CUSTOMIZING SPECIFIC OVERCLOUD ROLES	24
4.3. PRE-CONFIGURATION: CUSTOMIZING ALL OVERCLOUD ROLES	26
4.4. POST-CONFIGURATION: CUSTOMIZING ALL OVERCLOUD ROLES	28
4.5. PUPPET: CUSTOMIZING HIERADATA FOR ROLES	30
4.6. PUPPET: CUSTOMIZING HIERADATA FOR INDIVIDUAL NODES	31
4.7. PUPPET: APPLYING CUSTOM MANIFESTS	31
CHAPTER 5. OVERCLOUD REGISTRATION	33
5.1. REGISTERING THE OVERCLOUD WITH AN ENVIRONMENT FILE	33
5.2. EXAMPLE 1: REGISTERING TO THE CUSTOMER PORTAL	35
5.3. EXAMPLE 2: REGISTERING TO A RED HAT SATELLITE 6 SERVER	35
5.4. EXAMPLE 3: REGISTERING TO A RED HAT SATELLITE 5 SERVER	36
5.5. EXAMPLE 4: REGISTERING THROUGH A HTTP PROXY	37
5.6. ADVANCED REGISTRATION METHODS	37
CHAPTER 6. ANSIBLE-BASED OVERCLOUD REGISTRATION	39
6.1. RED HAT SUBSCRIPTION MANAGER (RHSM) COMPOSABLE SERVICE	39
6.2. RHSMVARS SUB-PARAMETERS	39
6.3. REGISTERING THE OVERCLOUD WITH THE RHSM COMPOSABLE SERVICE	40
6.4. APPLYING THE RHSM COMPOSABLE SERVICE TO DIFFERENT ROLES	41
6.5. REGISTERING THE OVERCLOUD TO RED HAT SATELLITE SERVER	42
6.6. SWITCHING TO THE RHSM COMPOSABLE SERVICE	43
6.7. RHEL-REGISTRATION TO RHSM MAPPINGS	43
6.8. DEPLOYING THE OVERCLOUD WITH THE RHSM COMPOSABLE SERVICE	44
CHAPTER 7. COMPOSABLE SERVICES AND CUSTOM ROLES	46
7.1. SUPPORTED ROLE ARCHITECTURE	46
7.2. ROLES	46
7.2.1. Examining the roles_data File	46
7.2.2. Creating a roles_data File	47
7.2.3. Supported Custom Roles	48

7.2.4. Examining Role Parameters	51
7.2.5. Creating a New Role	53
7.3. COMPOSABLE SERVICES	54
7.3.1. Guidelines and Limitations	54
7.3.2. Examining Composable Service Architecture	55
7.3.3. Adding and Removing Services from Roles	57
7.3.4. Enabling Disabled Services	58
7.3.5. Creating a Generic Node with No Services	58
CHAPTER 8. CONTAINERIZED SERVICES	60
8.1. CONTAINERIZED SERVICE ARCHITECTURE	60
8.2. CONTAINERIZED SERVICE PARAMETERS	60
8.3. MODIFYING OPENSTACK PLATFORM CONTAINERS	61
8.4. DEPLOYING A VENDOR PLUGIN	63
CHAPTER 9. BASIC NETWORK ISOLATION	64
9.1. NETWORK ISOLATION	64
9.2. MODIFYING ISOLATED NETWORK CONFIGURATION	65
9.3. NETWORK INTERFACE TEMPLATES	66
9.4. DEFAULT NETWORK INTERFACE TEMPLATES	67
9.5. ENABLING BASIC NETWORK ISOLATION	68
CHAPTER 10. CUSTOM COMPOSABLE NETWORKS	70
10.1. COMPOSABLE NETWORKS	70
10.2. ADDING A COMPOSABLE NETWORK	71
10.3. INCLUDING A COMPOSABLE NETWORK IN A ROLE	71
10.4. ASSIGNING OPENSTACK SERVICES TO COMPOSABLE NETWORKS	72
10.5. ENABLING CUSTOM COMPOSABLE NETWORKS	73
10.6. RENAMING THE DEFAULT NETWORKS	73
CHAPTER 11. CUSTOM NETWORK INTERFACE TEMPLATES	75
11.1. CUSTOM NETWORK ARCHITECTURE	75
11.2. RENDERING DEFAULT NETWORK INTERFACE TEMPLATES FOR CUSTOMIZATION	76
11.3. NETWORK INTERFACE ARCHITECTURE	76
11.4. NETWORK INTERFACE REFERENCE	77
11.5. EXAMPLE NETWORK INTERFACE LAYOUT	86
11.6. NETWORK INTERFACE TEMPLATE CONSIDERATIONS FOR CUSTOM NETWORKS	88
11.7. CUSTOM NETWORK ENVIRONMENT FILE	89
11.8. NETWORK ENVIRONMENT PARAMETERS	90
11.9. EXAMPLE CUSTOM NETWORK ENVIRONMENT FILE	93
11.10. ENABLING NETWORK ISOLATION WITH CUSTOM NICs	94
CHAPTER 12. ADDITIONAL NETWORK CONFIGURATION	95
12.1. CONFIGURING CUSTOM INTERFACES	95
12.2. CONFIGURING ROUTES AND DEFAULT ROUTES	96
12.3. CONFIGURING JUMBO FRAMES	97
12.4. CONFIGURING THE NATIVE VLAN ON A TRUNKED INTERFACE	98
12.5. INCREASING THE MAXIMUM NUMBER OF CONNECTIONS THAT NETFILTER TRACKS	99
CHAPTER 13. NETWORK INTERFACE BONDING	101
13.1. NETWORK INTERFACE BONDING AND LINK AGGREGATION CONTROL PROTOCOL (LACP)	101
13.2. OPEN VSWITCH BONDING OPTIONS	103
13.3. LINUX BONDING OPTIONS	104
13.4. OVS BONDING OPTIONS	104

CHAPTER 14. CONTROLLING NODE PLACEMENT	107
14.1. ASSIGNING SPECIFIC NODE IDS	107
14.2. ASSIGNING CUSTOM HOSTNAMES	108
14.3. ASSIGNING PREDICTABLE IPS	108
14.4. ASSIGNING PREDICTABLE VIRTUAL IPS	110
CHAPTER 15. ENABLING SSL/TLS ON OVERCLOUD PUBLIC ENDPOINTS	112
15.1. INITIALIZING THE SIGNING HOST	112
15.2. CREATING A CERTIFICATE AUTHORITY	112
15.3. ADDING THE CERTIFICATE AUTHORITY TO CLIENTS	112
15.4. CREATING AN SSL/TLS KEY	113
15.5. CREATING AN SSL/TLS CERTIFICATE SIGNING REQUEST	113
15.6. CREATING THE SSL/TLS CERTIFICATE	114
15.7. ENABLING SSL/TLS	114
15.8. INJECTING A ROOT CERTIFICATE	116
15.9. CONFIGURING DNS ENDPOINTS	116
15.10. ADDING ENVIRONMENT FILES DURING OVERCLOUD CREATION	117
15.11. UPDATING SSL/TLS CERTIFICATES	118
CHAPTER 16. ENABLING SSL/TLS ON INTERNAL AND PUBLIC ENDPOINTS WITH IDENTITY MANAGEMENT	119
16.1. ADD THE UNDERCLOUD TO THE CA	119
16.2. ADD THE UNDERCLOUD TO IDM	119
16.3. CONFIGURE OVERCLOUD DNS	121
16.4. CONFIGURE OVERCLOUD TO USE NOVAJOIN	121
CHAPTER 17. CONVERTING YOUR EXISTING DEPLOYMENT TO USE TLS	123
17.1. REQUIREMENTS	123
17.2. REVIEWING YOUR ENDPOINTS	123
17.3. APPLY WORKAROUND FOR KNOWN ISSUE	124
17.4. CONFIGURING ENDPOINTS TO USE TLS	124
17.4.1. Configuring undercloud integration for deployments using the same domain as IdM	125
17.4.2. Configuring overcloud integration for deployments that use the same domain as IdM, and retain the existing public endpoint certificates	126
17.4.3. Configuring overcloud integration for deployments that use the same domain as IdM, and replace the existing public endpoint certificates with an IdM generated certificate	127
17.4.4. Configuring undercloud integration for deployments that use an IdM subdomain	128
17.4.5. Configuring undercloud integration for deployments that use an IdM subdomain, and retain the existing public endpoint certificates	130
17.4.6. Configuring undercloud integration for deployments that use an IdM subdomain, and replace the existing public endpoint certificates with an IdM generated certificate	131
17.5. CHECKING TLS ENCRYPTION	132
CHAPTER 18. DEBUG MODES	134
CHAPTER 19. STORAGE CONFIGURATION	135
19.1. CONFIGURING NFS STORAGE	135
19.2. CONFIGURING CEPH STORAGE	137
19.3. USING AN EXTERNAL OBJECT STORAGE CLUSTER	137
19.4. CONFIGURING THE IMAGE IMPORT METHOD AND SHARED STAGING AREA	138
19.4.1. Creating and Deploying the glance-settings.yaml File	138
19.5. CONFIGURING CINDER BACK END FOR THE IMAGE SERVICE	139
19.6. CONFIGURING THE MAXIMUM NUMBER OF STORAGE DEVICES TO ATTACH TO ONE INSTANCE	140
19.7. CONFIGURING THIRD PARTY STORAGE	140

CHAPTER 20. SECURITY ENHANCEMENTS	142
20.1. MANAGING THE OVERCLOUD FIREWALL	142
20.2. CHANGING THE SIMPLE NETWORK MANAGEMENT PROTOCOL (SNMP) STRINGS	143
20.3. CHANGING THE SSL/TLS CIPHER AND RULES FOR HAPROXY	144
20.4. USING THE OPEN VSWITCH FIREWALL	145
20.5. USING SECURE ROOT USER ACCESS	146
CHAPTER 21. CONFIGURING NETWORK PLUGINS	147
21.1. FUJITSU CONVERGED FABRIC (C-FABRIC)	147
21.2. FUJITSU FOS SWITCH	147
CHAPTER 22. CONFIGURING IDENTITY	149
22.1. REGION NAME	149
CHAPTER 23. OTHER CONFIGURATIONS	150
23.1. CONFIGURING THE KERNEL ON OVERCLOUD NODES	150
23.2. CONFIGURING EXTERNAL LOAD BALANCING	150
23.3. CONFIGURING IPV6 NETWORKING	150

CHAPTER 1. INTRODUCTION

The Red Hat OpenStack Platform director provides a set of tools to provision and create a fully featured OpenStack environment, also known as the Overcloud. The [Director Installation and Usage Guide](#) covers the preparation and configuration of the Overcloud. However, a proper production-level Overcloud might require additional configuration, including:

- Basic network configuration to integrate the Overcloud into your existing network infrastructure.
- Network traffic isolation on separate VLANs for certain OpenStack network traffic types.
- SSL configuration to secure communication on public endpoints
- Storage options such as NFS, iSCSI, Red Hat Ceph Storage, and multiple third-party storage devices.
- Registration of nodes to the Red Hat Content Delivery Network or your internal Red Hat Satellite 5 or 6 server.
- Various system level options.
- Various OpenStack service options.

This guide provides instructions for augmenting your Overcloud through the director. At this point, the director has registered the nodes and configured the necessary services for Overcloud creation. Now you can customize your Overcloud using the methods in this guide.



NOTE

The examples in this guide are optional steps for configuring the Overcloud. These steps are only required to provide the Overcloud with additional functionality. Use only the steps that apply to the needs of your environment.

CHAPTER 2. UNDERSTANDING HEAT TEMPLATES

The custom configurations in this guide use Heat templates and environment files to define certain aspects of the OpenStack. This chapter provides a basic introduction to Heat templates so that you can understand the structure and format of these templates in the context of the Red Hat OpenStack Platform director.

2.1. HEAT TEMPLATES

Red Hat OpenStack Platform (RHOSP) director uses Heat Orchestration Templates (HOT) as a template format for its overcloud deployment plan. Templates in HOT format are usually expressed in YAML format. The purpose of a template is to define and create a *stack*, which is a collection of resources that heat creates, and the configuration of the resources. Resources are objects in RHOSP and can include compute resources, network configuration, security groups, scaling rules, and custom resources.



NOTE

For RHOSP to use the heat template file as a custom template resource, the file extension must be either **.yaml** or **.template**.

Heat templates have three main sections:

Parameters

These are settings passed to heat to customize a stack. You can also use heat parameters to customize default values. These settings are defined in the **parameters** section of a template.

Resources

These are the specific objects to create and configure as part of a stack. Red Hat OpenStack Platform (RHOSP) contains a set of core resources that span across all components. These are defined in the **resources** section of a template.

Output

These are values passed from heat after the creation of the stack. You can access these values either through the heat API or client tools. These are defined in the **output** section of a template.

Here is an example of a basic heat template:

```
heat_template_version: 2013-05-23

description: > A very basic Heat template.

parameters:
  key_name:
    type: string
    default: lars
    description: Name of an existing key pair to use for the instance
  flavor:
    type: string
    description: Instance type for the instance to be created
    default: m1.small
  image:
    type: string
    default: cirros
```

```

description: ID or name of the image to use for the instance

resources:
  my_instance:
    type: OS::Nova::Server
    properties:
      name: My Cirros Instance
      image: { get_param: image }
      flavor: { get_param: flavor }
      key_name: { get_param: key_name }

output:
  instance_name:
    description: Get the instance's name
    value: { get_attr: [ my_instance, name ] }

```

This template uses the resource type **type: OS::Nova::Server** to create an instance called **my_instance** with a particular flavor, image, and key. The stack can return the value of **instance_name**, which is called **My Cirros Instance**.

When Heat processes a template it creates a stack for the template and a set of child stacks for resource templates. This creates a hierarchy of stacks that descend from the main stack you define with your template. You can view the stack hierarchy using this following command:

```
$ openstack stack list --nested
```

2.2. ENVIRONMENT FILES

An environment file is a special type of template that provides customization for your heat templates. This includes three key parts:

Resource Registry


This section defines custom resource names that are linked to other heat templates. This provides a method to create custom resources that do not exist within the core resource collection. These are defined in the **resource_registry** section of an environment file.

Parameters

These are common settings that you apply to the parameters of the top-level template. For example, if you have a template that deploys nested stacks, such as resource registry mappings, the parameters apply only to the top-level template and not templates for the nested resources. Parameters are defined in the **parameters** section of an environment file.

Parameter Defaults

These parameters modify the default values for parameters in all templates. For example, if you have a heat template that deploys nested stacks, such as resource registry mappings, the parameter defaults apply to all templates. The parameter defaults are defined in the **parameter_defaults** section of an environment file.



IMPORTANT

Use **parameter_defaults** instead of **parameters** when you create custom environment files for your overcloud. This is so that the parameters apply to all stack templates for the overcloud.

Example of a basic environment file:

```
resource_registry:
  OS::Nova::Server::MyServer: myserver.yaml

parameter_defaults:
  NetworkName: my_network

parameters:
  MyIP: 192.168.0.1
```

The environment file, **my_env.yaml**, might be included when creating a stack from a heat template, **my_template.yaml**. The **my_env.yaml** file creates a new resource type called **OS::Nova::Server::MyServer**. The **myserver.yaml** file is a heat template file that provides an implementation for this resource type that overrides any built-in ones. You can include the **OS::Nova::Server::MyServer** resource in your **my_template.yaml** file.

The **MyIP** applies a parameter only to the main heat template that deploys with this environment file. In this example, it only applies to the parameters in **my_template.yaml**.

The **NetworkName** applies to both the main heat template, **my_template.yaml**, and the templates that are associated with the resources that are included the main template, such as the **OS::Nova::Server::MyServer** resource and its **myserver.yaml** template in this example.

**NOTE**

For RHOSP to use the heat template file as a custom template resource, the file extension must be either **.yaml** or **.template**.

2.3. CORE OVERCLOUD HEAT TEMPLATES

The director contains a core heat template collection for the Overcloud. This collection is stored in **/usr/share/openstack-tripleo-heat-templates**.

The main files and directories in this template collection are:

overcloud.j2.yaml

This is the main template file that creates the overcloud environment. This file uses Jinja2 syntax to iterate over certain sections in the template to create custom roles. The Jinja2 formatting is rendered into YAML during the overcloud deployment process.

overcloud-resource-registry-puppet.j2.yaml

This is the main environment file that creates the overcloud environment. It provides a set of configurations for Puppet modules that are stored on the overcloud image. After director writes the overcloud image to each node, heat starts the Puppet configuration for each node by using the resources registered in this environment file. This file uses Jinja2 syntax to iterate over certain sections in the template to create custom roles. The Jinja2 formatting is rendered into YAML during the overcloud deployment process.

roles_data.yaml

This is a file that defines the roles in an overcloud and maps services to each role.

network_data.yaml

This is a file that defines the networks in an overcloud and their properties such as subnets, allocation pools, and VIP status. The default **network_data** file contains the default networks: External,

Internal Api, Storage, Storage Management, Tenant, and Management. You can create a custom **network_data** file and add it to your **openstack overcloud deploy** command with the **-n** option.

plan-environment.yaml

This is a file that defines the metadata for your overcloud plan. This includes the plan name, main template to use, and environment files to apply to the overcloud.

capabilities-map.yaml

This is a mapping of environment files for an overcloud plan. Use this file to describe and enable environment files on the director web UI. Custom environment files that are detected in the **environments** directory in an overcloud plan but are not defined in the **capabilities-map.yaml** are listed in the **Other** subtab of **2 Specify Deployment Configuration > Overall Settings** on the web UI.

environments

Contains additional heat environment files that you can use with your overcloud creation. These environment files enable extra functions for your resulting Red Hat OpenStack Platform (RHOSP) environment. For example, the directory contains an environment file to enable Cinder NetApp backend storage (**cinder-netapp-config.yaml**). Any environment files that are detected in this directory that are not defined in the **capabilities-map.yaml** file are listed in the **Other** subtab of **2 Specify Deployment Configuration > Overall Settings** in the director's web UI.

network

This is a set of heat templates to help create isolated networks and ports.

puppet

These are templates that are mostly driven by configuration with Puppet. The **overcloud-resource-registry-puppet.j2.yaml** environment file uses the files in this directory to drive the application of the Puppet configuration on each node.

puppet/services

This is a directory that contains heat templates for all services in the composable service architecture.

extraconfig

These are templates that enable extra functionality.

firstboot

Provides example **first_boot** scripts that director uses when it initially creates the nodes.

2.4. PLAN ENVIRONMENT METADATA

A plan environment metadata file allows you to define metadata about your overcloud plan. This information is used when importing and exporting your overcloud plan, plus used during the overcloud creation from your plan.

A plan environment metadata file includes the following parameters:

version

The version of the template.

name

The name of the overcloud plan and the container in OpenStack Object Storage (swift) used to store the plan files.

template

The core parent template to use for the overcloud deployment. This is most often **overcloud.yaml**, which is the rendered version of the **overcloud.yaml.j2** template.

environments

Defines a list of environment files to use. Specify the path of each environment file with the **path** sub-parameter.

parameter_defaults

A set of parameters to use in your overcloud. This functions in the same way as the **parameter_defaults** section in a standard environment file.

passwords

A set of parameters to use for overcloud passwords. This functions in the same way as the **parameter_defaults** section in a standard environment file. Normally, the director automatically populates this section with randomly generated passwords.

workflow_parameters

Allows you to provide a set of parameters to OpenStack Workflow (mistral) namespaces. You can use this to calculate and automatically generate certain overcloud parameters.

The following is an example of the syntax of a plan environment file:

```
version: 1.0
name: myovercloud
description: 'My Overcloud Plan'
template: overcloud.yaml
environments:
- path: overcloud-resource-registry-puppet.yaml
- path: environments/docker.yaml
- path: environments/docker-ha.yaml
- path: environments/containers-default-parameters.yaml
- path: user-environment.yaml
parameter_defaults:
  ControllerCount: 1
  ComputeCount: 1
  OvercloudComputeFlavor: compute
  OvercloudControllerFlavor: control
workflow_parameters:
  tripleo.derive_params.v1.derive_parameters:
    num_phy_cores_per_numa_node_for_pmd: 2
```

You can include the plan environment metadata file with the **openstack overcloud deploy** command using the **-p** option. For example:

```
(undercloud) $ openstack overcloud deploy --templates \
-p /my-plan-environment.yaml \
[OTHER OPTIONS]
```

You can also view plan metadata for an existing overcloud plan using the following command:

```
(undercloud) $ openstack object save overcloud plan-environment.yaml --file -
```

2.5. CAPABILITIES MAP

The capabilities map provides a mapping of environment files in your plan and their dependencies. Use this file to describe and enable environment files through the director's web UI. Custom environment files detected in an overcloud plan but not listed in the **capabilities-map.yaml** are listed in the **Other** subtab of **2 Specify Deployment Configuration > Overall Settings** on the web UI.

The default file is located at `/usr/share/openstack-tripleo-heat-templates/capabilities-map.yaml`.

The following is an example of the syntax for a capabilities map:

```
topics: 1
- title: My Parent Section
  description: This contains a main section for different environment files
  environment_groups: 2
  - name: my-environment-group
    title: My Environment Group
    description: A list of environment files grouped together
    environments: 3
    - file: environment_file_1.yaml
      title: Environment File 1
      description: Enables environment file 1
      requires: 4
      - dependent_environment_file.yaml
    - file: environment_file_2.yaml
      title: Environment File 2
      description: Enables environment file 2
      requires: 5
      - dependent_environment_file.yaml
    - file: dependent_environment_file.yaml
      title: Dependent Environment File
      description: Enables the dependent environment file
```

1 The **topics** parameter contains a list of sections in the UI's deployment configuration. Each topic is displayed as a single screen of environment options and contains multiple environment groups, which you define with the **environment_groups** parameter. Each topic can have a plain text **title** and **description**.

2 The **environment_groups** parameter lists groupings of environment files in the UI's deployment configuration. For example, on a storage topic, you might have an environment group for Ceph-related environment files. Each environment group can have a plain text **title** and **description**.

3 The **environments** parameter shows all environment files that belong to an environment group. The **file** parameter is the location of the environment file. Each environment entry can have a plain text **title** and **description**.

4 5 The **requires** parameter is a list of dependencies for an environment file. In this example, both **environment_file_1.yaml** and **environment_file_2.yaml** require you to enable **dependent_environment_file.yaml** too.



NOTE

Red Hat OpenStack Platform uses this file to add access to features to the director UI. It is recommended not to modify this file as newer versions of Red Hat OpenStack Platform might override this file.

2.6. INCLUDING ENVIRONMENT FILES IN OVERCLOUD CREATION

The deployment command (**openstack overcloud deploy**) uses the **-e** option to include an environment file to customize your Overcloud. You can include as many environment files as necessary.

However, the order of the environment files is important as the parameters and resources defined in subsequent environment files take precedence. For example, you might have two environment files:

environment-file-1.yaml

```
resource_registry:
  OS::TripleO::NodeExtraConfigPost: /home/stack/templates/template-1.yaml

parameter_defaults:
  RabbitFDLimit: 65536
  TimeZone: 'Japan'
```

environment-file-2.yaml

```
resource_registry:
  OS::TripleO::NodeExtraConfigPost: /home/stack/templates/template-2.yaml

parameter_defaults:
  TimeZone: 'Hongkong'
```

Then deploy with both environment files included:

```
$ openstack overcloud deploy --templates -e environment-file-1.yaml -e environment-file-2.yaml
```

In this example, both environment files contain a common resource type (**OS::TripleO::NodeExtraConfigPost**) and a common parameter (**TimeZone**). The **openstack overcloud deploy** command runs through the following process:

1. Loads the default configuration from the core Heat template collection as per the **--template** option.
2. Applies the configuration from **environment-file-1.yaml**, which overrides any common settings from the default configuration.
3. Applies the configuration from **environment-file-2.yaml**, which overrides any common settings from the default configuration and **environment-file-1.yaml**.

This results in the following changes to the default configuration of the Overcloud:

- **OS::TripleO::NodeExtraConfigPost** resource is set to **/home/stack/templates/template-2.yaml** as per **environment-file-2.yaml**.
- **TimeZone** parameter is set to **Hongkong** as per **environment-file-2.yaml**.
- **RabbitFDLimit** parameter is set to **65536** as per **environment-file-1.yaml**. **environment-file-2.yaml** does not change this value.

This provides a method for defining custom configuration to the your Overcloud without values from multiple environment files conflicting.

2.7. USING CUSTOMIZED CORE HEAT TEMPLATES

When creating the overcloud, the director uses a core set of Heat templates located in **/usr/share/openstack-tripleo-heat-templates**. If you want to customize this core template collection,

use a Git workflow to track changes and merge updates. Use the following git processes to help manage your custom template collection:

Initializing a Custom Template Collection

Use the following procedure to create an initial Git repository containing the Heat template collection:

1. Copy the template collection to the **stack** users directory. This example copies the collection to the **~/templates** directory:

```
$ cd ~/templates
$ cp -r /usr/share/openstack-tripleo-heat-templates .
```

2. Change to the custom template directory and initialize a Git repository:

```
$ cd openstack-tripleo-heat-templates
$ git init .
```

3. Configure your Git user name and email address:

```
$ git config --global user.name "<USER_NAME>"
$ git config --global user.email "<EMAIL_ADDRESS>"
```

Replace **<USER_NAME>** with the user name that you want to use. Replace **<EMAIL_ADDRESS>** with your email address.

4. Stage all templates for the initial commit:

```
$ git add *
```

5. Create an initial commit:

```
$ git commit -m "Initial creation of custom core heat templates"
```

This creates an initial **master** branch containing the latest core template collection. Use this branch as the basis for your custom branch and merge new template versions to this branch.

Creating a Custom Branch and Committing Changes

Use a custom branch to store your changes to the core template collection. Use the following procedure to create a **my-customizations** branch and add customizations to it:

1. Create the **my-customizations** branch and switch to it:

```
$ git checkout -b my-customizations
```

2. Edit the files in the custom branch.

3. Stage the changes in git:

```
$ git add [edited files]
```

4. Commit the changes to the custom branch:

```
$ git commit -m "[Commit message for custom changes]"
```

This adds your changes as commits to the **my-customizations** branch. When the **master** branch updates, you can rebase **my-customizations** off **master**, which causes git to add these commits on to the updated template collection. This helps track your customizations and replay them on future template updates.

Updating the Custom Template Collection:

When updating the undercloud, the **openstack-tripleo-heat-templates** package might also update. When this occurs, use the following procedure to update your custom template collection:

1. Save the **openstack-tripleo-heat-templates** package version as an environment variable:

```
$ export PACKAGE=$(rpm -qv openstack-tripleo-heat-templates)
```

2. Change to your template collection directory and create a new branch for the updated templates:

```
$ cd ~/templates/openstack-tripleo-heat-templates
$ git checkout -b $PACKAGE
```

3. Remove all files in the branch and replace them with the new versions:

```
$ git rm -rf *
$ cp -r /usr/share/openstack-tripleo-heat-templates/* .
```

4. Add all templates for the initial commit:

```
$ git add *
```

5. Create a commit for the package update:

```
$ git commit -m "Updates for $PACKAGE"
```

6. Merge the branch into master. If using a Git management system (such as GitLab) use the management workflow. If using git locally, merge by switching to the **master** branch and run the **git merge** command:

```
$ git checkout master
$ git merge $PACKAGE
```

The **master** branch now contains the latest version of the core template collection. You can now rebase the **my-customization** branch from this updated collection.

Rebasing the Custom Branch

Use the following procedure to update the **my-customization** branch,;

1. Change to the **my-customizations** branch:

```
$ git checkout my-customizations
```

2. Rebase the branch off **master**:

```
$ git rebase master
```

This updates the **my-customizations** branch and replays the custom commits made to this branch.

If git reports any conflicts during the rebase, use this procedure:

1. Check which files contain the conflicts:

```
$ git status
```

2. Resolve the conflicts of the template files identified.
3. Add the resolved files

```
$ git add [resolved files]
```

4. Continue the rebase:

```
$ git rebase --continue
```

Deploying Custom Templates

Use the following procedure to deploy the custom template collection:

1. Make sure you have switched to the **my-customization** branch:

```
git checkout my-customizations
```


2. Run the **openstack overcloud deploy** command with the **--templates** option to specify your local template directory:

```
$ openstack overcloud deploy --templates /home/stack/templates/openstack-tripleo-heat-templates [OTHER OPTIONS]
```



NOTE

The director uses the default template directory (**/usr/share/openstack-tripleo-heat-templates**) if you specify the **--templates** option without a directory.



IMPORTANT

Red Hat recommends using the methods in [Chapter 4, Configuration Hooks](#) instead of modifying the heat template collection.

2.8. JINJA2 RENDERING

The core Heat templates in **/usr/share/openstack-tripleo-heat-templates** contains a number of files ending with a **j2.yaml** extension. These files contain Jinja2 template syntax and the director renders these files to their static Heat template equivalents ending in **.yaml**. For example, the main **overcloud.j2.yaml** file renders into **overcloud.yaml**. The director uses the resulting **overcloud.yaml** file.

The Jinja2-enabled Heat templates use Jinja2 syntax to create parameters and resources for iterative values. For example, the **overcloud.j2.yaml** file contains the following snippet:

```
parameters:
...
{% for role in roles %}
...
  {{role.name}}Count:
    description: Number of {{role.name}} nodes to deploy
    type: number
    default: {{role.CountDefault|default(0)}}
...
{% endfor %}
```

When the director renders the Jinja2 syntax, the director iterates over the roles defined in the **roles_data.yaml** file and populates the **{{role.name}}Count** parameter with the name of the role. The default **roles_data.yaml** file contains five roles and results in the the following parameters from our example:

- **ControllerCount**
- **ComputeCount**
- **BlockStorageCount**
- **ObjectStorageCount**
- **CephStorageCount**

A example rendered version of the parameter looks like this:

```
parameters:
...
ControllerCount:
  description: Number of Controller nodes to deploy
  type: number
  default: 1
...
```

The director only renders Jinja2-enabled templates and environment files within the directory of your core Heat templates. The following use cases demonstrate the correct method to render the Jinja2 templates.

Use case 1: Default core templates

Template directory: `/usr/share/openstack-tripleo-heat-templates/`

Environment file: `/usr/share/openstack-tripleo-heat-templates/environments/network-isolation.j2.yaml`

The director uses the default core template location (`--templates`). The director renders the **network-isolation.j2.yaml** file into **network-isolation.yaml**. When running the **openstack overcloud deploy** command, use the **-e** option to include the name of rendered **network-isolation.yaml** file.

```
$ openstack overcloud deploy --templates \  
-e /usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml  
...
```

Use case 2: Custom core templates

Template directory: **/home/stack/tripleo-heat-templates**

Environment file: **/home/stack/tripleo-heat-templates/environments/network-isolation.j2.yaml**

The director uses a custom core template location (**--templates /home/stack/tripleo-heat-templates**). The director renders the **network-isolation.j2.yaml** file within the custom core templates into **network-isolation.yaml**. When running the **openstack overcloud deploy** command, use the **-e** option to include the name of rendered **network-isolation.yaml** file.

```
$ openstack overcloud deploy --templates /home/stack/tripleo-heat-templates \  
-e /home/stack/tripleo-heat-templates/environments/network-isolation.yaml  
...
```

Use case 3: Incorrect usage

Template directory: **/usr/share/openstack-tripleo-heat-templates/**

Environment file: **/home/stack/tripleo-heat-templates/environments/network-isolation.j2.yaml**

This director uses a custom core template location (**--templates /home/stack/tripleo-heat-templates**). However, the chosen **network-isolation.j2.yaml** is not located within the custom core templates, so it *will not* render into **network-isolation.yaml**. This causes the deployment to fail.

CHAPTER 3. PARAMETERS

Each Heat template in the director's template collection contains a **parameters** section. This section defines all parameters specific to a particular overcloud service. This includes the following:

- **overcloud.j2.yaml** - Default base parameters
- **roles_data.yaml** - Default parameters for composable roles
- **puppet/services/*.yaml** - Default parameters for specific services

You can modify the values for these parameters using the following method:

1. Create an environment file for your custom parameters.
2. Include your custom parameters in the **parameter_defaults** section of the environment file.
3. Include the environment file with the **openstack overcloud deploy** command.

The next few sections contain examples to demonstrate how to configure specific parameters for services in the **puppet/services** directory.

3.1. EXAMPLE 1: CONFIGURING THE TIMEZONE

The Heat template for setting the timezone (**puppet/services/time/timezone.yaml**) contains a **TimeZone** parameter. If you leave the **TimeZone** parameter blank, the overcloud sets the time to **UTC** as a default.

To obtain lists of timezones run the **timedatectl list-timezones** command. The following example command retrieves the timezones for Asia:

```
$ sudo timedatectl list-timezones|grep "Asia"
```

After you identify your timezone, set the *TimeZone* parameter in an environment file. The following example environment file sets the value of *TimeZone* to *Asia/Tokyo*:

```
parameter_defaults:
  TimeZone: 'Asia/Tokyo'
```

3.2. EXAMPLE 2: ENABLING NETWORKING DISTRIBUTED VIRTUAL ROUTING (DVR)

The Heat template for the OpenStack Networking (neutron) API (**puppet/services/neutron-api.yaml**) contains a parameter to enable and disable Distributed Virtual Routing (DVR). The default for the parameter is **false**. However, you can enable it using the following in an environment file:

```
parameter_defaults:
  NeutronEnableDVR: true
```

3.3. EXAMPLE 3: CONFIGURING RABBITMQ FILE DESCRIPTOR LIMIT

For certain configurations, you might need to increase the file descriptor limit for the RabbitMQ server. The `puppet/services/rabbitmq.yaml` Heat template allows you to set the `RabbitFDLimit` parameter to the limit you require. Add the following to an environment file.

```
parameter_defaults:
  RabbitFDLimit: 65536
```

3.4. EXAMPLE 4: ENABLING AND DISABLING PARAMETERS

In some case, you might need to initially set a parameters during a deployment, then disable the parameter for a future deployment operation, such as updates or scaling operations. For example, to include a custom RPM during the overcloud creation, you would include the following:

```
parameter_defaults:
  DeployArtifactURLs: ["http://www.example.com/myfile.rpm"]
```

If you need to disable this parameter from a future deployment, it is not enough to remove the parameter. Instead, you set the parameter to an empty value:

```
parameter_defaults:
  DeployArtifactURLs: []
```

This ensures the parameter is no longer set for subsequent deployments operations.

3.5. EXAMPLE 5: ROLE-BASED PARAMETERS

Use the `[ROLE]Parameters` parameters, replacing `[ROLE]` with a composable role, to set parameters for a specific role.

For example, director configures `logrotate` on both Controller and Compute nodes. To set a different different `logrotate` parameters for Controller and Compute nodes, create an environment file that contains both the 'ControllerParameters' and 'ComputeParameters' parameter and set the `logrotate` parameter for each specific role:

```
parameter_defaults:
  ControllerParameters:
    LogrotateMaxsize: 10M
    LogrotatePurgeAfterDays: 30
  ComputeParameters:
    LogrotateMaxsize: 20M
    LogrotatePurgeAfterDays: 15
```

3.6. IDENTIFYING PARAMETERS TO MODIFY

Red Hat OpenStack Platform director provides many parameters for configuration. In some cases, you might experience difficulty identifying a certain option to configure and the corresponding director parameter. If there is an option you want to configure through the director, use the following workflow to identify and map the option to a specific overcloud parameter:

1. Identify the option you aim to configure. Make a note of the service that uses the option.
2. Check the corresponding Puppet module for this option. The Puppet modules for Red Hat OpenStack Platform are located under `/etc/puppet/modules` on the director node. Each

module corresponds to a particular service. For example, the **keystone** module corresponds to the OpenStack Identity (keystone).

- If the Puppet module contains a variable that controls the chosen option, move to the next step.
 - If the Puppet module does not contain a variable that controls the chosen option, then no hieradata exists for this option. If possible, you can set the option manually after the overcloud completes deployment.
3. Check the director's core Heat template collection for the Puppet variable in the form of hieradata. The templates in **puppet/services/*** usually correspond to the Puppet modules of the same services. For example, the **puppet/services/keystone.yaml** template provides hieradata to the **keystone** module.
 - If the Heat template sets hieradata for the Puppet variable, the template should also disclose the director-based parameter to modify.
 - If the Heat template does not set hieradata for the Puppet variable, use the configuration hooks to pass the hieradata using an environment file. See [Section 4.5, "Puppet: Customizing Hieradata for Roles"](#) for more information on customizing hieradata.



IMPORTANT

Do not define multiple instances of the same custom hieradata hashes. Multiple instances of the same custom hieradata can cause conflicts during Puppet runs and result in unexpected values set for configuration options.

Workflow Example

You might aim to change the notification format for OpenStack Identity (keystone). Using the workflow, you would:

1. Identify the OpenStack parameter to configure (**notification_format**).
2. Search the **keystone** Puppet module for the **notification_format** setting. For example:

```
$ grep notification_format /etc/puppet/modules/keystone/manifests/*
```

In this case, the **keystone** module manages this option using the **keystone::notification_format** variable.

3. Search the **keystone** service template for this variable. For example:

```
$ grep "keystone::notification_format" /usr/share/openstack-tripleo-heat-templates/puppet/services/keystone.yaml
```

The output shows the director using the **KeystoneNotificationFormat** parameter to set the **keystone::notification_format** hieradata.

The following table shows the eventual mapping:

Director Parameter	Puppet Hieradata	OpenStack Identity (keystone) option
KeystoneNotificationFormat	keystone::notification_format	notification_format

This means setting the **KeystoneNotificationFormat** in an overcloud's environment file would set the **notification_format** option in the **keystone.conf** file during the overcloud's configuration.

CHAPTER 4. CONFIGURATION HOOKS

The configuration hooks provide a method to inject your own configuration functions into the OpenStack deployment process. This includes hooks for injecting custom configuration before and after the main OpenStack services configuration and hook for modifying and including Puppet-based configuration.

4.1. FIRST BOOT: CUSTOMIZING FIRST BOOT CONFIGURATION

The director provides a mechanism to perform configuration on all nodes upon the initial creation of the OpenStack. The director achieves this through **cloud-init**, which you can call using the **OS::TripleO::NodeUserData** resource type.

In this example, you will update the nameserver with a custom IP address on all nodes. You must first create a basic heat template (**/home/stack/templates/nameserver.yaml**) that runs a script to append each node's **resolv.conf** with a specific nameserver. You can use the **OS::TripleO::MultipartMime** resource type to send the configuration script.

```
heat_template_version: 2014-10-16

description: >
  Extra hostname configuration

resources:
  userdata:
    type: OS::Heat::MultipartMime
    properties:
      parts:
        - config: {get_resource: nameserver_config}

  nameserver_config:
    type: OS::Heat::SoftwareConfig
    properties:
      config: |
        #!/bin/bash
        echo "nameserver 192.168.1.1" >> /etc/resolv.conf

outputs:
  OS::stack_id:
    value: {get_resource: userdata}
```

Next, create an environment file (**/home/stack/templates/firstboot.yaml**) that registers your heat template as the **OS::TripleO::NodeUserData** resource type.

```
resource_registry:
  OS::TripleO::NodeUserData: /home/stack/templates/nameserver.yaml
```

To add the first boot configuration, add the environment file to the stack along with your other environment files when first creating the OpenStack. For example:

```
$ openstack overcloud deploy --templates \
...
-e /home/stack/templates/firstboot.yaml \
...
```

The **-e** applies the environment file to the Overcloud stack.


This adds the configuration to all nodes when they are first created and boot for the first time. Subsequent inclusions of these templates, such as updating the Overcloud stack, does not run these scripts.



IMPORTANT

You can only register the **OS::TripleO::NodeUserData** to one heat template. Subsequent usage overrides the heat template to use.

4.2. PRE-CONFIGURATION: CUSTOMIZING SPECIFIC OVERCLOUD ROLES



IMPORTANT

Previous versions of this document used the **OS::TripleO::Tasks::*PreConfig** resources to provide pre-configuration hooks on a per role basis. The director's Heat template collection requires dedicated use of these hooks, which means you should not use them for custom use. Instead, use the **OS::TripleO::*ExtraConfigPre** hooks outlined below.

The Overcloud uses Puppet for the core configuration of OpenStack components. The director provides a set of hooks to provide custom configuration for specific node roles after the first boot completes and before the core configuration begins. These hooks include:

OS::TripleO::ControllerExtraConfigPre

Additional configuration applied to Controller nodes before the core Puppet configuration.

OS::TripleO::ComputeExtraConfigPre

Additional configuration applied to Compute nodes before the core Puppet configuration.

OS::TripleO::CephStorageExtraConfigPre

Additional configuration applied to Ceph Storage nodes before the core Puppet configuration.

OS::TripleO::ObjectStorageExtraConfigPre

Additional configuration applied to Object Storage nodes before the core Puppet configuration.

OS::TripleO::BlockStorageExtraConfigPre

Additional configuration applied to Block Storage nodes before the core Puppet configuration.

OS::TripleO::[ROLE]ExtraConfigPre

Additional configuration applied to custom nodes before the core Puppet configuration. Replace **[ROLE]** with the composable role name.

In this example, you first create a basic heat template (**/home/stack/templates/nameserver.yaml**) that runs a script to write to a node's **resolv.conf** with a variable nameserver.

```
heat_template_version: 2014-10-16
```

```
description: >
  Extra hostname configuration
```

```
parameters:
  server:
    type: string
```

```

nameserver_ip:
  type: string
DeployIdentifier:
  type: string

resources:
  CustomExtraConfigPre:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script
      config:
        str_replace:
          template: |
            #!/bin/sh
            echo "nameserver _NAMESERVER_IP_" > /etc/resolv.conf
        params:
          _NAMESERVER_IP_: {get_param: nameserver_ip}

  CustomExtraDeploymentPre:
    type: OS::Heat::SoftwareDeployment
    properties:
      server: {get_param: server}
      config: {get_resource: CustomExtraConfigPre}
      actions: ['CREATE','UPDATE']
      input_values:
        deploy_identifier: {get_param: DeployIdentifier}

outputs:
  deploy_stdout:
    description: Deployment reference, used to trigger pre-deploy on changes
    value: {get_attr: [CustomExtraDeploymentPre, deploy_stdout]}

```

In this example, the **resources** section contains the following:

CustomExtraConfigPre

This defines a software configuration. In this example, we define a Bash **script** and Heat replaces **_NAMESERVER_IP_** with the value stored in the **nameserver_ip** parameter.

CustomExtraDeploymentPre

This executes a software configuration, which is the software configuration from the **CustomExtraConfigPre** resource. Note the following:

- The **config** parameter makes a reference to the **CustomExtraConfigPre** resource so Heat knows what configuration to apply.
- The **server** parameter retrieves a map of the Overcloud nodes. This parameter is provided by the parent template and is mandatory in templates for this hook.
- The **actions** parameter defines when to apply the configuration. In this case, we only apply the configuration when the Overcloud is created or updated. Possible actions include **CREATE, UPDATE, DELETE, SUSPEND, and RESUME**.
- **input_values** contains a parameter called **deploy_identifier**, which stores the **DeployIdentifier** from the parent template. This parameter provides a timestamp to the resource for each deployment update. This ensures the resource reapplies on subsequent overcloud updates.

Next, create an environment file (`/home/stack/templates/pre_config.yaml`) that registers your heat template to the role-based resource type. For example, to apply only to Controller nodes, use the **ControllerExtraConfigPre** hook:

```
resource_registry:
  OS::TripleO::ControllerExtraConfigPre: /home/stack/templates/nameserver.yaml

parameter_defaults:
  nameserver_ip: 192.168.1.1
```

To apply the configuration, add the environment file to the stack along with your other environment files when creating or updating the Overcloud. For example:

```
$ openstack overcloud deploy --templates \
...
-e /home/stack/templates/pre_config.yaml \
...
```

This applies the configuration to all Controller nodes before the core configuration begins on either the initial Overcloud creation or subsequent updates.



IMPORTANT

You can only register each resource to only one Heat template per hook. Subsequent usage overrides the Heat template to use.

4.3. PRE-CONFIGURATION: CUSTOMIZING ALL OVERCLOUD ROLES

The Overcloud uses Puppet for the core configuration of OpenStack components. The director provides a hook to configure all node types after the first boot completes and before the core configuration begins:

OS::TripleO::NodeExtraConfig

Additional configuration applied to all nodes roles before the core Puppet configuration.

In this example, you first create a basic heat template (`/home/stack/templates/nameserver.yaml`) that runs a script to append each node's **resolv.conf** with a variable nameserver.

```
heat_template_version: 2014-10-16

description: >
  Extra hostname configuration

parameters:
  server:
    type: string
  nameserver_ip:
    type: string
  DeployIdentifier:
    type: string

resources:
  CustomExtraConfigPre:
    type: OS::Heat::SoftwareConfig
```

```

properties:
  group: script
  config:
    str_replace:
      template: |
        #!/bin/sh
        echo "nameserver _NAMESERVER_IP_" >> /etc/resolv.conf
    params:
      _NAMESERVER_IP_: {get_param: nameserver_ip}

```

```

CustomExtraDeploymentPre:
  type: OS::Heat::SoftwareDeployment
  properties:
    server: {get_param: server}
    config: {get_resource: CustomExtraConfigPre}
    actions: ['CREATE','UPDATE']
    input_values:
      deploy_identifier: {get_param: DeployIdentifier}

```

```

outputs:
  deploy_stdout:
    description: Deployment reference, used to trigger pre-deploy on changes
    value: {get_attr: [CustomExtraDeploymentPre, deploy_stdout]}

```

In this example, the **resources** section contains the following:

CustomExtraConfigPre

This defines a software configuration. In this example, we define a Bash **script** and Heat replaces **_NAMESERVER_IP_** with the value stored in the **nameserver_ip** parameter.

CustomExtraDeploymentPre

This executes a software configuration, which is the software configuration from the **CustomExtraConfigPre** resource. Note the following:

- The **config** parameter makes a reference to the **CustomExtraConfigPre** resource so Heat knows what configuration to apply.
- The **server** parameter retrieves a map of the Overcloud nodes. This parameter is provided by the parent template and is mandatory in templates for this hook.
- The **actions** parameter defines when to apply the configuration. In this case, we only apply the configuration when the Overcloud is created or updated. Possible actions include **CREATE, UPDATE, DELETE, SUSPEND, and RESUME**.
- The **input_values** parameter contains a sub-parameter called **deploy_identifier**, which stores the **DeployIdentifier** from the parent template. This parameter provides a timestamp to the resource for each deployment update. This ensures the resource reapplies on subsequent overcloud updates.

Next, create an environment file (**/home/stack/templates/pre_config.yaml**) that registers your heat template as the **OS::TripleO::NodeExtraConfig** resource type.

```

resource_registry:
  OS::TripleO::NodeExtraConfig: /home/stack/templates/nameserver.yaml

```

```
parameter_defaults:
  nameserver_ip: 192.168.1.1
```

To apply the configuration, add the environment file to the stack along with your other environment files when creating or updating the Overcloud. For example:

```
$ openstack overcloud deploy --templates \
...
-e /home/stack/templates/pre_config.yaml \
...
```

This applies the configuration to all nodes before the core configuration begins on either the initial Overcloud creation or subsequent updates.



IMPORTANT

You can only register the **OS::TripleO::NodeExtraConfig** to only one Heat template. Subsequent usage overrides the Heat template to use.

4.4. POST-CONFIGURATION: CUSTOMIZING ALL OVERCLOUD ROLES



IMPORTANT

Previous versions of this document used the **OS::TripleO::Tasks::*PostConfig** resources to provide post-configuration hooks on a per role basis. The director's Heat template collection requires dedicated use of these hooks, which means you should not use them for custom use. Instead, use the **OS::TripleO::NodeExtraConfigPost** hook outlined below.

A situation might occur where you have completed the creation of your Overcloud but want to add additional configuration to all roles, either on initial creation or on a subsequent update of the Overcloud. In this case, you use the following post-configuration hook:

OS::TripleO::NodeExtraConfigPost

Additional configuration applied to all nodes roles after the core Puppet configuration.

In this example, you first create a basic heat template (`/home/stack/templates/nameserver.yaml`) that runs a script to append each node's **resolv.conf** with a variable nameserver.

```
description: >
  Extra hostname configuration

parameters:
  servers:
    type: json
  nameserver_ip:
    type: string
  DeployIdentifier:
    type: string

resources:
  CustomExtraConfig:
    type: OS::Heat::SoftwareConfig
```



```

properties:
  group: script
  config:
    str_replace:
      template: |
        #!/bin/sh
        echo "nameserver _NAMESERVER_IP_" >> /etc/resolv.conf
    params:
      _NAMESERVER_IP_: {get_param: nameserver_ip}

```

```

CustomExtraDeployments:
  type: OS::Heat::SoftwareDeploymentGroup
  properties:
    servers: {get_param: servers}
    config: {get_resource: CustomExtraConfig}
    actions: ['CREATE','UPDATE']
    input_values:
      deploy_identifier: {get_param: DeployIdentifier}

```

In this example, the **resources** section contains the following:

CustomExtraConfig

This defines a software configuration. In this example, we define a Bash **script** and Heat replaces **_NAMESERVER_IP_** with the value stored in the **nameserver_ip** parameter.

CustomExtraDeployments

This executes a software configuration, which is the software configuration from the **CustomExtraConfig** resource. Note the following:

- The **config** parameter makes a reference to the **CustomExtraConfig** resource so Heat knows what configuration to apply.
- The **servers** parameter retrieves a map of the OpenStack nodes. This parameter is provided by the parent template and is mandatory in templates for this hook.
- The **actions** parameter defines when to apply the configuration. In this case, we only apply the configuration when the OpenStack is created. Possible actions include **CREATE**, **UPDATE**, **DELETE**, **SUSPEND**, and **RESUME**.
- **input_values** contains a parameter called **deploy_identifier**, which stores the **DeployIdentifier** from the parent template. This parameter provides a timestamp to the resource for each deployment update. This ensures the resource reapplies on subsequent openstack updates.

Next, create an environment file (**/home/stack/templates/post_config.yaml**) that registers your heat template as the **OS::TripleO::NodeExtraConfigPost** resource type.

```

resource_registry:
  OS::TripleO::NodeExtraConfigPost: /home/stack/templates/nameserver.yaml

parameter_defaults:
  nameserver_ip: 192.168.1.1

```

To apply the configuration, add the environment file to the stack along with your other environment files when creating or updating the OpenStack. For example:

```
$ openstack overcloud deploy --templates \
...
-e /home/stack/templates/post_config.yaml \
...
```

This applies the configuration to all nodes after the core configuration completes on either initial Overcloud creation or subsequent updates.



IMPORTANT

You can only register the **OS::TripleO::NodeExtraConfigPost** to only one Heat template. Subsequent usage overrides the Heat template to use.

4.5. PUPPET: CUSTOMIZING HIERADATA FOR ROLES

The Heat template collection contains a set of parameters to pass extra configuration to certain node types. These parameters save the configuration as hieradata for the node's Puppet configuration. These parameters are:

ControllerExtraConfig

Configuration to add to all Controller nodes.

ComputeExtraConfig

Configuration to add to all Compute nodes.

BlockStorageExtraConfig

Configuration to add to all Block Storage nodes.

ObjectStorageExtraConfig

Configuration to add to all Object Storage nodes

CephStorageExtraConfig

Configuration to add to all Ceph Storage nodes

[ROLE]ExtraConfig

Configuration to add to a composable role. Replace **[ROLE]** with the composable role name.


ExtraConfig

Configuration to add to all nodes.

To add extra configuration to the post-deployment configuration process, create an environment file that contains these parameters in the **parameter_defaults** section. For example, to increase the reserved memory for Compute hosts to 1024 MB and set the VNC keymap to Japanese:

```
parameter_defaults:
  ComputeExtraConfig:
    nova::compute::reserved_host_memory: 1024
    nova::compute::vnc_keymap: ja
```

Include this environment file when running **openstack overcloud deploy**.



IMPORTANT

You can only define each parameter once. Subsequent usage overrides previous values.

4.6. PUPPET: CUSTOMIZING HIERADATA FOR INDIVIDUAL NODES

You can set Puppet hieradata for individual nodes using the Heat template collection. To accomplish this, you need to acquire the system UUID saved as part of the introspection data for a node:

```
$ openstack baremetal introspection data save 9dcc87ae-4c6d-4ede-81a5-9b20d7dc4a14 | jq
.extra.system.product.uuid
```

This outputs a system UUID. For example:

```
"F5055C6C-477F-47FB-AFE5-95C6928C407F"
```

Use this system UUID in an environment file that defines node-specific hieradata and registers the **per_node.yaml** template to a pre-configuration hook. For example:

```
resource_registry:
  OS::TripleO::ComputeExtraConfigPre: /usr/share/openstack-tripleo-heat-
  templates/puppet/extraconfig/pre_deploy/per_node.yaml
parameter_defaults:
  NodeDataLookup: '{"F5055C6C-477F-47FB-AFE5-95C6928C407F":
{"nova::compute::vcpu_pin_set": [ "2", "3" ]}}'
```

Include this environment file when running **openstack overcloud deploy**.

The **per_node.yaml** template generates a set of hieradata files on nodes that correspond to each system UUID and contains the hieradata you defined. If a UUID is not defined, the resulting hieradata file is empty. In the previous example, the **per_node.yaml** template runs on all Compute nodes (as per the **OS::TripleO::ComputeExtraConfigPre** hook), but only the Compute node with system UUID **F5055C6C-477F-47FB-AFE5-95C6928C407F** receives hieradata.

This provides a method of tailoring each node to specific requirements.

For more information about NodeDataLookup, see [Configuring Ceph Storage Cluster Setting](#) in the *Deploying an Overcloud with Containerized Red Hat Ceph* guide.

4.7. PUPPET: APPLYING CUSTOM MANIFESTS

In certain circumstances, you might need to install and configure some additional components to your Overcloud nodes. You can achieve this with a custom Puppet manifest that applies to nodes on after the main configuration completes. As a basic example, you might intend to install **motd** to each node. The process for accomplishing is to first create a Heat template (**/home/stack/templates/custom_puppet_config.yaml**) that launches Puppet configuration.

```
heat_template_version: 2014-10-16

description: >
  Run Puppet extra configuration to set new MOTD

parameters:
  servers:
    type: json

resources:
  ExtraPuppetConfig:
```

```

type: OS::Heat::SoftwareConfig
properties:
  config: {get_file: motd.pp}
  group: puppet
  options:
    enable_hiera: True
    enable_factor: False

```

```

ExtraPuppetDeployments:
type: OS::Heat::SoftwareDeploymentGroup
properties:
  config: {get_resource: ExtraPuppetConfig}
  servers: {get_param: servers}

```

This includes the **/home/stack/templates/motd.pp** within the template and passes it to nodes for configuration. The **motd.pp** file itself contains the Puppet classes to install and configure **motd**.

Next, create an environment file (**/home/stack/templates/puppet_post_config.yaml**) that registers your heat template as the **OS::TripleO::NodeExtraConfigPost:** resource type.

```

resource_registry:
  OS::TripleO::NodeExtraConfigPost: /home/stack/templates/custom_puppet_config.yaml

```

And finally include this environment file along with your other environment files when creating or updating the Overcloud stack:

```

$ openstack overcloud deploy --templates \
...
-e /home/stack/templates/puppet_post_config.yaml \
...

```

This applies the configuration from **motd.pp** to all nodes in the Overcloud.



IMPORTANT

Do not define multiple instances of the same custom hieradata hashes. Multiple instances of the same custom hieradata can cause conflicts during Puppet runs and result in unexpected values set for configuration options.

CHAPTER 5. OVERCLOUD REGISTRATION

The Overcloud provides a method to register nodes to either the Red Hat Content Delivery Network, Red Hat Satellite Server 5, or Red Hat Satellite Server 6.

5.1. REGISTERING THE OVERCLOUD WITH AN ENVIRONMENT FILE

Copy the registration files from the Heat template collection:

```
$ cp -r /usr/share/openstack-tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration
~/templates/.
```

Edit the `~/templates/rhel-registration/environment-rhel-registration.yaml` file and change the values of the parameters that apply to your registration method and details.

General Parameters

`rhel_reg_method`

Choose the registration method. Either **portal**, **satellite**, or **disable**.

`rhel_reg_type`

The type of unit to register. Leave blank to register as a **system**

`rhel_reg_auto_attach`

Automatically attach compatible subscriptions to this system. Set to **true** to enable. To disable this feature, leave this parameter blank.

`rhel_reg_service_level`

The service level that you want to use for auto attachment.

`rhel_reg_release`

Use this parameter to set a release version for auto attachment. Leave blank to use the default from Red Hat Subscription Manager.

`rhel_reg_pool_id`

The subscription pool ID that you want to use. Use this if not auto-attaching subscriptions. To locate this ID, run **sudo subscription-manager list --available --all --matches="*OpenStack*"** from the undercloud node, and use the resulting **Pool ID** value.

`rhel_reg_sat_url`

The base URL of the Satellite Server that you want to register the Overcloud nodes with. Use the Satellite Server HTTP URL and not the HTTPS URL for this parameter. For example, use <http://satellite.example.com> and not <https://satellite.example.com>. The Overcloud creation process uses this URL to determine whether you are using Red Hat Satellite Server 5 or Red Hat Satellite Server 6. When using Red Hat Satellite Server 6, the Overcloud obtains the **katello-ca-consumer-latest.noarch.rpm** file, registers with **subscription-manager**, and installs **katello-agent**. When using Red Hat Satellite Server 5, the Overcloud obtains the **RHN-ORG-TRUSTED-SSL-CERT** file and registers with **rhnreg_ks**.

`rhel_reg_server_url`

The hostname of the subscription service that you want to use. The default is for Customer Portal Subscription Management, `subscription.rhn.redhat.com`. If this option is not used, the system is registered with Customer Portal Subscription Management. The subscription server URL uses the form of <https://hostname:port/prefix>.

`rhel_reg_base_url`

The hostname of the content delivery server that you want to use to receive updates. The default is <https://cdn.redhat.com>. Since Satellite 6 hosts its own content, the URL must be used for systems registered with Satellite 6. The base URL for content uses the form of <https://hostname:port/prefix>.

rhel_reg_org

The organization that you want to use for registration. To locate this ID, run **sudo subscription-manager orgs** from the undercloud node. Enter your Red Hat credentials when prompted, and use the resulting **Key** value.

rhel_reg_environment

The environment that you want to use within the chosen organization.

rhel_reg_repos

A comma-separated list of repositories to enable.

rhel_reg_activation_key

The activation key that you want to use for registration. When using an activation key for registration, you must also specify the organization that you want to use for registration.

rhel_reg_user; rhel_reg_password

The username and password that you want to use for registration. If possible, use activation keys for registration.

rhel_reg_machine_name

The machine name that you want to use for registration. Leave this blank if you want to use the hostname of the node.

rhel_reg_force

Set to **true** to force your registration options, for example, when re-registering nodes.

rhel_reg_sat_repo

The repository that contains Red Hat Satellite 6 Server management tools, such as **katello-agent**. Ensure that the repository name corresponds to your Satellite Server version and that the repository is synchronized on Satellite Server. For example, **rhel-7-server-satellite-tools-6.2-rpms** corresponds to Red Hat Satellite 6.2.

Upgrade Parameters

UpdateOnRHELRegistration

If set to **True**, this triggers an update of the overcloud packages after registration completes. Set to **False** by default.

HTTP Proxy Parameters

rhel_reg_http_proxy_host

The hostname for the HTTP proxy. For example: **proxy.example.com**.

rhel_reg_http_proxy_port


The port for HTTP proxy communication. For example: **8080**.

rhel_reg_http_proxy_username

The username to access the HTTP proxy.

rhel_reg_http_proxy_password


The password to access the HTTP proxy.

**IMPORTANT**

If using a proxy server, ensure all overcloud nodes have a route to the host defined in the **rhel_reg_http_proxy_host** parameter. Without a route to this host, **subscription-manager** will time out and cause deployment failure.

The deployment command (**openstack overcloud deploy**) uses the **-e** option to add environment files. Add both **~/templates/rhel-registration/environment-rhel-registration.yaml** and **~/templates/rhel-registration/rhel-registration-resource-registry.yaml**. For example:

```
$ openstack overcloud deploy --templates [...] -e /home/stack/templates/rhel-
registration/environment-rhel-registration.yaml -e /home/stack/templates/rhel-
registration/rhel-registration-resource-registry.yaml
```

**IMPORTANT**

Registration is set as the **OS::TripleO::NodeExtraConfig** Heat resource. This means you can only use this resource for registration. See [Section 4.2, “Pre-Configuration: Customizing Specific Overcloud Roles”](#) for more information.

5.2. EXAMPLE 1: REGISTERING TO THE CUSTOMER PORTAL

The following registers the overcloud nodes to the Red Hat Customer Portal using the **my-openstack** activation key and subscribes to pool **1a85f9223e3d5e43013e3d6e8ff506fd**.

```
parameter_defaults:
  rhel_reg_auto_attach: ""
  rhel_reg_activation_key: "my-openstack"
  rhel_reg_org: "1234567"
  rhel_reg_pool_id: "1a85f9223e3d5e43013e3d6e8ff506fd"
  rhel_reg_repos: "rhel-7-server-rpms,rhel-7-server-extras-rpms,rhel-7-server-rh-common-rpms,rhel-
ha-for-rhel-7-server-rpms,rhel-7-server-openstack-13-rpms,rhel-7-server-rhceph-3-osd-rpms,rhel-7-
server-rhceph-3-mon-rpms,rhel-7-server-rhceph-3-tools-rpms"
  rhel_reg_method: "portal"
  rhel_reg_sat_repo: ""
  rhel_reg_base_url: ""
  rhel_reg_environment: ""
  rhel_reg_force: ""
  rhel_reg_machine_name: ""
  rhel_reg_password: ""
  rhel_reg_release: ""
  rhel_reg_sat_url: ""
  rhel_reg_server_url: ""
  rhel_reg_service_level: ""
  rhel_reg_user: ""
  rhel_reg_type: ""
  rhel_reg_http_proxy_host: ""
  rhel_reg_http_proxy_port: ""
  rhel_reg_http_proxy_username: ""
  rhel_reg_http_proxy_password: ""
```

5.3. EXAMPLE 2: REGISTERING TO A RED HAT SATELLITE 6 SERVER

The following registers the overcloud nodes to a Red Hat Satellite 6 Server at `sat6.example.com` and uses the **my-openstack** activation key to subscribe to pool `1a85f9223e3d5e43013e3d6e8ff506fd`. In this situation, the activation key also provides the repositories to enable.

```
parameter_defaults:
  rhel_reg_activation_key: "my-openstack"
  rhel_reg_org: "1"
  rhel_reg_pool_id: "1a85f9223e3d5e43013e3d6e8ff506fd"
  rhel_reg_method: "satellite"
  rhel_reg_sat_url: "http://sat6.example.com"
  rhel_reg_sat_repo: "rhel-7-server-satellite-tools-6.2-rpms"
  rhel_reg_repos: ""
  rhel_reg_auto_attach: ""
  rhel_reg_base_url: ""
  rhel_reg_environment: ""
  rhel_reg_force: ""
  rhel_reg_machine_name: ""
  rhel_reg_password: ""
  rhel_reg_release: ""
  rhel_reg_server_url: ""
  rhel_reg_service_level: ""
  rhel_reg_user: ""
  rhel_reg_type: ""
  rhel_reg_http_proxy_host: ""
  rhel_reg_http_proxy_port: ""
  rhel_reg_http_proxy_username: ""
  rhel_reg_http_proxy_password: ""
```

5.4. EXAMPLE 3: REGISTERING TO A RED HAT SATELLITE 5 SERVER

The following registers the overcloud nodes to a Red Hat Satellite 5 Server at `sat5.example.com`, uses the **my-openstack** activation key, and automatically attaches subscriptions. In this situation, the activation key also provides the repositories to enable.

```
parameter_defaults:
  rhel_reg_auto_attach: ""
  rhel_reg_activation_key: "my-openstack"
  rhel_reg_org: "1"
  rhel_reg_method: "satellite"
  rhel_reg_sat_url: "http://sat5.example.com"
  rhel_reg_repos: ""
  rhel_reg_base_url: ""
  rhel_reg_environment: ""
  rhel_reg_force: ""
  rhel_reg_machine_name: ""
  rhel_reg_password: ""
  rhel_reg_pool_id: ""
  rhel_reg_release: ""
  rhel_reg_server_url: ""
  rhel_reg_service_level: ""
  rhel_reg_user: ""
  rhel_reg_type: ""
  rhel_reg_sat_repo: ""
  rhel_reg_http_proxy_host: ""
```



```

rhel_reg_http_proxy_port: ""
rhel_reg_http_proxy_username: ""
rhel_reg_http_proxy_password: ""

```

5.5. EXAMPLE 4: REGISTERING THROUGH A HTTP PROXY

The following sample parameters set the HTTP proxy settings for your desired registration method:

```

parameter_defaults:
  ...
  rhel_reg_http_proxy_host: "proxy.example.com"
  rhel_reg_http_proxy_port: "8080"
  rhel_reg_http_proxy_username: "proxyuser"
  rhel_reg_http_proxy_password: "p@55w0rd!"
  ...

```

5.6. ADVANCED REGISTRATION METHODS

In some situations, you might aim to register different roles to different subscription types. For example, you might aim to only subscribe Controller nodes to an OpenStack Platform subscription and Ceph Storage nodes to a Ceph Storage subscription. This section provides some advanced registration methods to help with assigning separate subscriptions to different roles.

Configuration Hooks

One method is to write role-specific scripts and include them with a role-specific hook. For example, the following snippet could be added to the **OS::TripleO::ControllerExtraConfigPre** resource's template, which ensures only the Controller nodes receive these subscription details.

```

ControllerRegistrationConfig:
  type: OS::Heat::SoftwareConfig
  properties:
    group: script
    config: |
      #!/bin/sh
      sudo subscription-manager register --org 1234567 \
        --activationkey "my-openstack"
      sudo subscription-manager attach --pool 1a85f9223e3d5e43013e3d6e8ff506fd
      sudo subscription-manager repos --enable rhel-7-server-rpms \
        --enable rhel-7-server-extras-rpms \
        --enable rhel-7-server-rh-common-rpms \
        --enable rhel-ha-for-rhel-7-server-rpms \
        --enable rhel-7-server-openstack-13-rpms \
        --enable rhel-7-server-rhceph-3-mon-rpms

ControllerRegistrationDeployment:
  type: OS::Heat::SoftwareDeployment
  properties:
    server: {get_param: server}
    config: {get_resource: ControllerRegistrationConfig}
    actions: ['CREATE','UPDATE']
    input_values:
      deploy_identifier: {get_param: DeployIdentifier}

```

The script uses a set of **subscription-manager** commands to register the system, attach the subscription, and enable the required repositories.

For more information about hooks, see [Chapter 4, Configuration Hooks](#).

Ansible-Based Configuration

You can perform Ansible-based registration on specific roles using the director's dynamic inventory script. For example, you might aim to register Controller nodes using the following play:

```
---
- name: Register Controller nodes
  hosts: Controller
  become: yes
  vars:
    repos:
      - rhel-7-server-rpms
      - rhel-7-server-extras-rpms
      - rhel-7-server-rh-common-rpms
      - rhel-ha-for-rhel-7-server-rpms
      - rhel-7-server-openstack-13-rpms
      - rhel-7-server-rhceph-3-mon-rpms
  tasks:
    - name: Register system
      redhat_subscription:
        activationkey: my-openstack
        org_id: 1234567
        pool_ids: 1a85f9223e3d5e43013e3d6e8ff506fd
    - name: Disable all repos
      command: "subscription-manager repos --disable *"
    - name: Enable Controller node repos
      command: "subscription-manager repos --enable {{ item }}"
      with_items: "{{ repos }}"
```

This play contains three tasks: - Register the node using an activation key - Disable any auto-enabled repositories - Enable only the repositories relevant to the Controller node. The repositories are listed with the **repos** variable.

After deploying the overcloud, you can run the following command so that Ansible executes the playbook (**ansible-osp-registration.yml**) against your overcloud:

```
$ ansible-playbook -i /usr/bin/tripleo-ansible-inventory ansible-osp-registration.yml
```

This command does the following: - Runs the dynamic inventory script to get a list of host and their groups - Applies the playbook tasks to the nodes in the group defined in the playbook's **hosts** parameter, which in this case is the **Controller** group.

For more information on the running Ansible automation on your overcloud, see ["Running Ansible Automation"](#) in the *Director Installation and Usage* guide.

CHAPTER 6. ANSIBLE-BASED OVERCLOUD REGISTRATION



IMPORTANT

This feature is available in this release as a *Technology Preview*, and therefore is not fully supported by Red Hat. It should only be used for testing, and should not be deployed in a production environment. For more information about Technology Preview features, see [Scope of Coverage Details](#).

As an alternative to the **rhel-registration** method from [Chapter 5, Overcloud Registration](#), the director can use an Ansible-based method to register overcloud nodes to the Red Hat Customer Portal or a Red Hat Satellite 6 server. This method relies on enabling Ansible-based configuration (**config-download**) in the overcloud.

6.1. RED HAT SUBSCRIPTION MANAGER (RHSM) COMPOSABLE SERVICE

The **rhsm** composable service provides a method to register overcloud nodes through Ansible. Each role in the default **roles_data** file contains a **OS::TripleO::Services::Rhsm** resource, which is disabled by default. To enable the service, you register the resource to the **rhsm** composable service file. For example:

```
resource_registry:
  OS::TripleO::Services::Rhsm: /usr/share/openstack-tripleo-heat-
    templates/extraconfig/services/rhsm.yaml
```

The **rhsm** composable service accepts a **RhsmVars** parameter, which allows you to define multiple sub-parameters relevant to your registration. For example:

```
parameter_defaults:
  RhsmVars:
    rhsm_repos:
      - rhel-7-server-rpms
      - rhel-7-server-extras-rpms
      - rhel-7-server-rh-common-rpms
      - rhel-ha-for-rhel-7-server-rpms
      - rhel-7-server-openstack-13-rpms
      - rhel-7-server-rhceph-3-osd-rpms
      - rhel-7-server-rhceph-3-mon-rpms
      - rhel-7-server-rhceph-3-tools-rpms
    rhsm_activation_key: "my-openstack"
    rhsm_org_id: "1234567"
```

You can also use the **RhsmVars** parameter in combination with role-specific parameters (e.g. **ControllerParameters**) to provide flexibility when enabling specific repositories for different nodes types.

The next section is a list of sub-parameters available to use with the **RhsmVars** parameter for use with the **rhsm** composable service.

6.2. RHSMVARS SUB-PARAMETERS

rhsm	Description
rhsm_method	Choose the registration method. Either portal , satellite , or disable .
rhsm_org_id	The organization to use for registration. To locate this ID, run sudo subscription-manager orgs from the undercloud node. Enter your Red Hat credentials when prompted, and use the resulting Key value.
rhsm_pool_ids	The subscription pool ID to use. Use this if not auto-attaching subscriptions. To locate this ID, run sudo subscription-manager list --available --all --matches="*OpenStack*" from the undercloud node, and use the resulting Pool ID value.
rhsm_activation_key	The activation key to use for registration.
rhsm_autosubscribe	Automatically attach compatible subscriptions to this system. Set to true to enable.
rhsm_satellite_url	The base URL of the Satellite server to register Overcloud nodes.
rhsm_repos	A list of repositories to enable.
rhsm_username	The username for registration. If possible, use activation keys for registration.
rhsm_password	The password for registration. If possible, use activation keys for registration.
rhsm_rhsm_proxy_host_name	The hostname for the HTTP proxy. For example: proxy.example.com .
rhsm_rhsm_proxy_port	The port for HTTP proxy communication. For example: 8080 .
rhsm_rhsm_proxy_user	The username to access the HTTP proxy.
rhsm_rhsm_proxy_password	The password to access the HTTP proxy.

Now that you have an understanding of how the **rhsm** composable service works and how to configure it, you can use the following procedures to configure your own registration details.

6.3. REGISTERING THE OVERCLOUD WITH THE RHSM COMPOSABLE SERVICE

Use the following procedure to create an environment file that enables and configures the **rhsm** composable service. The director uses this environment file to register and subscribe your nodes.

Procedure

1. Create an environment file (**templates/rhsm.yml**) to store the configuration.

2. Include your configuration in the environment file. For example:

```
resource_registry:
  OS::TripleO::Services::Rhsm: /usr/share/openstack-tripleo-heat-
  templates/extraconfig/services/rhsm.yaml
parameter_defaults:
  RhsmVars:
    rhsm_repos:
      - rhel-7-server-rpms
      - rhel-7-server-extras-rpms
      - rhel-7-server-rh-common-rpms
      - rhel-ha-for-rhel-7-server-rpms
      - rhel-7-server-openstack-13-rpms
      - rhel-7-server-rhceph-3-osd-rpms
      - rhel-7-server-rhceph-3-mon-rpms
      - rhel-7-server-rhceph-3-tools-rpms
    rhsm_activation_key: "my-openstack"
    rhsm_org_id: "1234567"
    rhsm_pool_ids: "1a85f9223e3d5e43013e3d6e8ff506fd"
    rhsm_method: "portal"
```

The **resource_registry** associates the **rhsm** composable service with the **OS::TripleO::Services::Rhsm** resource, which is available on each role.

The **RhsmVars** variable passes parameters to Ansible for configuring your Red Hat registration.

3. Save the environment file

You can also provide registration details to specific overcloud roles. The next section provides an example of this.

6.4. APPLYING THE RHSM COMPOSABLE SERVICE TO DIFFERENT ROLES

You can apply the **rhsm** composable service on a per-role basis. For example, you can apply one set of configuration to Controller nodes and a different set of configuration to Compute nodes.

Procedure

1. Create an environment file (**templates/rhsm.yml**) to store the configuration.
2. Include your configuration in the environment file. For example:

```
resource_registry:
  OS::TripleO::Services::Rhsm: /usr/share/openstack-tripleo-heat-
  templates/extraconfig/services/rhsm.yaml
parameter_defaults:
  ControllerParameters:
    RhsmVars:
      rhsm_repos:
        - rhel-7-server-rpms
        - rhel-7-server-extras-rpms
        - rhel-7-server-rh-common-rpms
        - rhel-ha-for-rhel-7-server-rpms
        - rhel-7-server-openstack-13-rpms
```

```

- rhel-7-server-rhceph-3-osd-rpms
- rhel-7-server-rhceph-3-mon-rpms
- rhel-7-server-rhceph-3-tools-rpms
rhsm_activation_key: "my-openstack"
rhsm_org_id: "1234567"
rhsm_pool_ids: "1a85f9223e3d5e43013e3d6e8ff506fd"
rhsm_method: "portal"
ComputeParameters:
  RhsmVars:
    rhsm_repos:
      - rhel-7-server-rpms
      - rhel-7-server-extras-rpms
      - rhel-7-server-rh-common-rpms
      - rhel-ha-for-rhel-7-server-rpms
      - rhel-7-server-openstack-13-rpms
      - rhel-7-server-rhceph-3-tools-rpms
    rhsm_activation_key: "my-openstack"
    rhsm_org_id: "1234567"
    rhsm_pool_ids: "1a85f9223e3d5e43013e3d6e8ff506fd"
    rhsm_method: "portal"

```

The **resource_registry** associates the **rhsm** composable service with the **OS::TripleO::Services::Rhsm** resource, which is available on each role.

Both **ControllerParameters** and **ComputeParameters** use their own **RhsmVars** parameter to pass subscription details to their respective roles.

3. Save the environment file

6.5. REGISTERING THE OVERCLOUD TO RED HAT SATELLITE SERVER

Create an environment file that enables and configures the **rhsm** composable service to register nodes to Red Hat Satellite instead of the Red Hat Customer Portal.

Procedure

1. Create an environment file named **templates/rhsm.yml** to store the configuration.
2. Include your configuration in the environment file. For example:

```

resource_registry:
  OS::TripleO::Services::Rhsm: /usr/share/openstack-tripleo-heat-
  templates/extraconfig/services/rhsm.yaml
parameter_defaults:
  RhsmVars:
    rhsm_activation_key: "myactivationkey"
    rhsm_method: "satellite"
    rhsm_org_id: "ACME"
    rhsm_server_hostname: "satellite.example.com"
    rhsm_baseurl: "https://satellite.example.com/pulp/repos"
    rhsm_release: 7.9

```

The **resource_registry** associates the **rhsm** composable service with the **OS::TripleO::Services::Rhsm** resource, which is available on each role.

The **RhsmVars** variable passes parameters to Ansible for configuring your Red Hat registration.

3. Save the environment file.

These procedures enable and configure **rhsm** on the overcloud. However, if you are using the **rhel-registration** method from [Chapter 5, Overcloud Registration](#), you must disable it to switch to the Ansible-based method. Use the following procedure to switch from the **rhel-registration** method to the Ansible-based method.

6.6. SWITCHING TO THE RHSM COMPOSABLE SERVICE

The **rhel-registration** method which runs a bash script to handle the overcloud registration. The scripts and environment files for this method are located in the core Heat template collection at `/usr/share/openstack-tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration/`.

This procedure shows how to switch from the **rhel-registration** method to the **rhsm** composable service.

Procedure

1. Exclude the **rhel-registration** environment files from future deployments operations. In most cases, this will be the following files:
 - **rhel-registration/environment-rhel-registration.yaml**
 - **rhel-registration/rhel-registration-resource-registry.yaml**
2. Add the environment file for **rhsm** composable service parameters to future deployment operations.

This method replaces the **rhel-registration** parameters with the **rhsm** service parameters and changes the Heat resource that enables the service from:

```
resource_registry:
  OS::TripleO::NodeExtraConfig: rhel-registration.yaml
```

To:

```
resource_registry:
  OS::TripleO::Services::Rhsm: /usr/share/openstack-tripleo-heat-
  templates/extraconfig/services/rhsm.yaml
```

To help transition your details from the **rhel-registration** method to the **rhsm** method, use the following table to map the your parameters and their values.

6.7. RHEL-REGISTRATION TO RHSM MAPPINGS

rhel-registration	rhsm / RhsmVars
rhel_reg_method	rhsm_method
rhel_reg_org	rhsm_org_id

rhel-registration	rhsm / RhsmVars
rhel_reg_pool_id	rhsm_pool_ids
rhel_reg_activation_key	rhsm_activation_key
rhel_reg_auto_attach	rhsm_autosubscribe
rhel_reg_sat_url	rhsm_satellite_url
rhel_reg_repos	rhsm_repos
rhel_reg_user	rhsm_username
rhel_reg_password	rhsm_password
rhel_reg_http_proxy_host	rhsm_rhsm_proxy_hostname
rhel_reg_http_proxy_port	rhsm_rhsm_proxy_port
rhel_reg_http_proxy_username	rhsm_rhsm_proxy_user
rhel_reg_http_proxy_password	rhsm_rhsm_proxy_password

Now that you have configured the environment file for the **rhsm** service, you can include it with your next overcloud deployment operation.

6.8. DEPLOYING THE OVERCLOUD WITH THE RHSM COMPOSABLE SERVICE

This process shows how to apply your **rhsm** configuration to the overcloud.

Procedure

1. When running the **openstack overcloud deploy** command, include the **config-download** option and environment file and the **rhsm.yml** environment file:

```
openstack overcloud deploy \
  <other cli args> \
  -e /usr/share/openstack-tripleo-heat-templates/environments/config-download-
environment.yaml \
  --config-download \
  -e ~/templates/rhsm.yaml
```

This enables the Ansible configuration of the overcloud and the Ansible-based registration.

2. Wait until the overcloud deployment completes.

3. Check the subscription details on your overcloud nodes. For example, log into a Controller node and run the following commands:

```
$ sudo subscription-manager status  
$ sudo subscription-manager list --consumed
```

CHAPTER 7. COMPOSABLE SERVICES AND CUSTOM ROLES

The Overcloud usually consists of nodes in predefined roles such as Controller nodes, Compute nodes, and different storage node types. Each of these default roles contains a set of services defined in the core Heat template collection on the director node. However, the architecture of the core Heat templates provide methods to:

- Create custom roles
- Add and remove services from each role

This allows the possibility to create different combinations of services on different roles. This chapter explores the architecture of custom roles, composable services, and methods for using them.

7.1. SUPPORTED ROLE ARCHITECTURE

The following architectures are available when using custom roles and composable services:

Architecture 1 - Default Architecture

Uses the default **roles_data** files. All controller services are contained within one Controller role.

Architecture 2 - Supported Standalone Roles

Use the predefined files in `/usr/share/openstack-tripleo-heat-templates/roles` to generate a custom **roles_data** file`. See [Section 7.2.3, "Supported Custom Roles"](#).

Architecture 3 - Custom Composable Services

Create your own **roles** and use them to generate a custom **roles_data** file. Note that only a limited number of composable service combinations have been tested and verified and Red Hat cannot support all composable service combinations.

7.2. ROLES

7.2.1. Examining the roles_data File

The Overcloud creation process defines its roles using a **roles_data** file. The **roles_data** file contains a YAML-formatted list of the roles. The following is a shortened example of the **roles_data** syntax:

```
- name: Controller
  description: |
    Controller role that has all the controller services loaded and handles
    Database, Messaging and Network functions.
  ServicesDefault:
    - OS::TripleO::Services::AuditD
    - OS::TripleO::Services::CACerts
    - OS::TripleO::Services::CephClient
    ...
- name: Compute
  description: |
    Basic Compute Node role
  ServicesDefault:
    - OS::TripleO::Services::AuditD
    - OS::TripleO::Services::CACerts
    - OS::TripleO::Services::CephClient
    ...
```

The core Heat template collection contains a default **roles_data** file located at `/usr/share/openstack-tripleo-heat-templates/roles_data.yaml`. The default file defines the following role types:

- **Controller**
- **Compute**
- **BlockStorage**
- **ObjectStorage**
- **CephStorage**.

The **openstack overcloud deploy** command includes this file during deployment. You can override this file with a custom **roles_data** file using the **-r** argument. For example:

```
$ openstack overcloud deploy --templates -r ~/templates/roles_data-custom.yaml
```

7.2.2. Creating a roles_data File

Although you can manually create a custom **roles_data** file, you can also automatically generate the file using individual role templates. The **director** provides a several commands to manage role templates and automatically generate a custom **roles_data** file.

To list the default role templates, use the **openstack overcloud role list** command:

```
$ openstack overcloud role list
BlockStorage
CephStorage
Compute
ComputeHCI
ComputeOvsDpdk
Controller
...
```

To see the role's YAML definition, use the **openstack overcloud role show** command:

```
$ openstack overcloud role show Compute
```

To generate a custom **roles_data** file, use the **openstack overcloud roles generate** command to join multiple predefined roles into a single file. For example, the following command joins the **Controller**, **Compute**, and **Networker** roles into a single file:

```
$ openstack overcloud roles generate -o ~/roles_data.yaml Controller Compute Networker
```

The **-o** defines the name of the file to create.

This creates a custom **roles_data** file. However, the previous example uses the **Controller** and **Networker** roles, which both contain the same networking agents. This means the networking services scale from **Controller** to the **Networker** role. The overcloud balance the load for networking services between the **Controller** and **Networker** nodes.

To make this **Networker** role standalone, you can create your own custom **Controller** role, as well as any other role needed. This allows you to easily generate a **roles_data** file from your own custom roles.

Copy the directory from the core Heat template collection to the **stack** user's home directory:

```
$ cp -r /usr/share/openstack-tripleo-heat-templates/roles ~/.
```

Add or modify the custom role files in this directory. Use the **--roles-path** option with any of the aforementioned role sub-commands to use this directory as the source for your custom roles. For example:

```
$ openstack overcloud roles generate -o my_roles_data.yaml \
  --roles-path ~/roles \
  Controller Compute Networker
```

This generates a single **my_roles_data.yaml** file from the individual roles in the **~/roles** directory.



NOTE

The default roles collection also contains the **ControllerOpenStack** role, which does not include services for **Networker**, **Messaging**, and **Database** roles. You can use the **ControllerOpenStack** combined with with the standalone **Networker**, **Messaging**, and **Database** roles.

7.2.3. Supported Custom Roles

The following table describes of all supported roles available in **/usr/share/openstack-tripleo-heat-templates/roles**.

Role	Description	File
BlockStorage	OpenStack Block Storage (cinder) node.	BlockStorage.yaml
CephAll	Full standalone Ceph Storage node. Includes OSD, MON, Object Gateway (RGW), Object Operations (MDS), Manager (MGR), and RBD Mirroring.	CephAll.yaml
CephFile	Standalone scale-out Ceph Storage file role. Includes OSD and Object Operations (MDS).	CephFile.yaml
CephObject	Standalone scale-out Ceph Storage object role. Includes OSD and Object Gateway (RGW).	CephObject.yaml
CephStorage	Ceph Storage OSD node role.	CephStorage.yaml
ComputeAlt	Alternate Compute node role.	ComputeAlt.yaml
ComputeDVR	DVR enabled Compute node role.	ComputeDVR.yaml
ComputeHCI	Compute node with hyper-converged infrastructure. Includes Compute and Ceph OSD services.	ComputeHCI.yaml

Role	Description	File
ComputeInstanceHA	Compute Instance HA node role. Use in conjunction with the environments/compute-instanceha.yaml environment file.	ComputeInstanceHA.yaml
ComputeLiquidio	Compute node with Cavium Liquidio Smart NIC.	ComputeLiquidio.yaml
ComputeOvsDpdkRT	Compute OVS DPDK RealTime role.	ComputeOvsDpdkRT.yaml
ComputeOvsDpdk	Compute OVS DPDK role.	ComputeOvsDpdk.yaml
ComputePPC64LE	Compute role for ppc64le servers.	ComputePPC64LE.yaml
ComputeRealTime	Compute role optimized for real-time behaviour. When using this role, it is mandatory that an overcloud-realtime-compute image is available and the role specific parameters IsolCpusList and NovaVcpuPinSet are set accordingly to the hardware of the real-time compute nodes.	ComputeRealTime.yaml
ComputeSriovRT	Compute SR-IOV RealTime role.	ComputeSriovRT.yaml
ComputeSriov	Compute SR-IOV role.	ComputeSriov.yaml
Compute	Standard Compute node role.	Compute.yaml
ControllerAllNovaStandalone	Controller role that does not contain the database, messaging, networking, and OpenStack Compute (nova) control components. Use in combination with the Database, Messaging, Networker , and Novacontrol roles.	ControllerAllNovaStandalone.yaml
ControllerNoCeph	Controller role with core Controller services loaded but no Ceph Storage (MON) components. This role handles database, messaging, and network functions but not any Ceph Storage functions.	ControllerNoCeph.yaml
ControllerNovaStand alone	Controller role that does not contain the OpenStack Compute (nova) control component. Use in combination with the Novacontrol role.	ControllerNovaStand alone.yaml

Role	Description	File
ControllerOpenstack	Controller role that does not contain the database, messaging, and networking components. Use in combination with the Database , Messaging , and Networker roles.	ControllerOpenstack.yaml
ControllerStorageNfs	Controller role with all core services loaded and uses Ceph NFS. This role handles database, messaging, and network functions.	ControllerStorageNfs.yaml
Controller	Controller role with all core services loaded. This role handles database, messaging, and network functions.	Controller.yaml
Database	Standalone database role. Database managed as a Galera cluster using Pacemaker.	Database.yaml
HciCephAll	Compute node with hyper-converged infrastructure and all Ceph Storage services. Includes OSD, MON, Object Gateway (RGW), Object Operations (MDS), Manager (MGR), and RBD Mirroring.	HciCephAll.yaml
HciCephFile	Compute node with hyper-converged infrastructure and Ceph Storage file services. Includes OSD and Object Operations (MDS).	HciCephFile.yaml
HciCephMon	Compute node with hyper-converged infrastructure and Ceph Storage block services. Includes OSD, MON, and Manager.	HciCephMon.yaml
HciCephObject	Compute node with hyper-converged infrastructure and Ceph Storage object services. Includes OSD and Object Gateway (RGW).	HciCephObject.yaml
IronicConductor	Ironic Conductor node role.	IronicConductor.yaml
Messaging	Standalone messaging role. RabbitMQ managed with Pacemaker.	Messaging.yaml
Networker	Standalone networking role. Runs OpenStack networking (neutron) agents on their own. If your deployment uses the ML2/OVN mechanism driver, see additional steps in Deploying a Custom Role with ML2/OVN .	Networker.yaml
Novacontrol	Standalone nova-control role to run OpenStack Compute (nova) control agents on their own.	Novacontrol.yaml
ObjectStorage	Swift Object Storage node role.	ObjectStorage.yaml

Role	Description	File
Telemetry	Telemetry role with all the metrics and alarming services.	Telemetry.yaml

7.2.4. Examining Role Parameters

Each role uses the following parameters:

name

(Mandatory) The name of the role, which is a plain text name with no spaces or special characters. Check that the chosen name does not cause conflicts with other resources. For example, use **Networker** as a name instead of **Network**.

description

(Optional) A plain text description for the role.

tags

(Optional) A YAML list of tags that define role properties. Use this parameter to define the primary role with both the **controller** and **primary** tags together:

```
- name: Controller
  ...
  tags:
    - primary
    - controller
  ...
```



IMPORTANT

If you do not tag the primary role, the first role defined becomes the primary role. Ensure this role is the Controller role.

networks

A YAML list of networks to configure on the role:

```
networks:
  - External
  - InternalApi
  - Storage
  - StorageMgmt
  - Tenant
```

Default networks include **External**, **InternalApi**, **Storage**, **StorageMgmt**, **Tenant**, and **Management**.

CountDefault

(Optional) Defines the default number of nodes to deploy for this role.

HostnameFormatDefault

(Optional) Defines the default hostname format for the role. The default naming convention uses the following format:

```
[STACK NAME]-[ROLE NAME]-[NODE ID]
```

For example, the default Controller nodes are named:

```
overcloud-controller-0
overcloud-controller-1
overcloud-controller-2
...
```

disable_constraints

(Optional) Defines whether to disable OpenStack Compute (nova) and OpenStack Image Storage (glance) constraints when deploying with the director. Used when deploying an overcloud with pre-provisioned nodes. For more information, see "[Configuring a Basic Overcloud using Pre-Provisioned Nodes](#)" in the *Director Installation and Usage Guide*.

disable_upgrade_deployment

(Optional) Defines whether to disable upgrades for a specific role. This provides a method to upgrade individual nodes in a role and ensure availability of services. For example, the Compute and Swift Storage roles use this parameter.

update_serial

(Optional) Defines how many nodes to update simultaneously during the OpenStack update options. In the default **roles_data.yaml** file:

- The default is **1** for Controller, Object Storage, and Ceph Storage nodes.
- The default is **25** for Compute and Block Storage nodes.

If you omit this parameter from a custom role, the default is **1**.

ServicesDefault

(Optional) Defines the default list of services to include on the node. See [Section 7.3.2, "Examining Composable Service Architecture"](#) for more information.

These parameters provide a means to create new roles and also define which services to include.

The **openstack overcloud deploy** command integrates the parameters from the **roles_data** file into some of the Jinja2-based templates. For example, at certain points, the **overcloud.j2.yaml** Heat template iterates over the list of roles from **roles_data.yaml** and creates parameters and resources specific to each respective role.

The resource definition for each role in the **overcloud.j2.yaml** Heat template appears as the following snippet:

```
{{role.name}}:
  type: OS::Heat::ResourceGroup
  depends_on: Networks
  properties:
    count: {get_param: {{role.name}}Count}
    removal_policies: {get_param: {{role.name}}RemovalPolicies}
  resource_def:
    type: OS::TripleO::{{role.name}}
    properties:
      CloudDomain: {get_param: CloudDomain}
```



```
ServiceNetMap: {get_attr: [ServiceNetMap, service_net_map]}
EndpointMap: {get_attr: [EndpointMap, endpoint_map]}
```

...

This snippet shows how the Jinja2-based template incorporates the `{{role.name}}` variable to define the name of each role as a `OS::Heat::ResourceGroup` resource. This in turn uses each `name` parameter from the `roles_data` file to name each respective `OS::Heat::ResourceGroup` resource.

7.2.5. Creating a New Role

In this example, the aim is to create a new **Horizon** role to host the OpenStack Dashboard (**horizon**) only. In this situation, you create a custom `roles` directory that includes the new role information.

Create a custom copy of the default `roles` directory:

```
$ cp -r /usr/share/openstack-tripleo-heat-templates/roles ~/.
```

Create a new file called `~/roles/Horizon.yaml` and create a new **Horizon** role containing base and core OpenStack Dashboard services. For example:

```
- name: Horizon
  CountDefault: 1
  HostnameFormatDefault: '%stackname%-horizon-%index%'
  ServicesDefault:
    - OS::TripleO::Services::CACerts
    - OS::TripleO::Services::Kernel
    - OS::TripleO::Services::Ntp
    - OS::TripleO::Services::Snmp
    - OS::TripleO::Services::Sshd
    - OS::TripleO::Services::Timezone
    - OS::TripleO::Services::TripleoPackages
    - OS::TripleO::Services::TripleoFirewall
    - OS::TripleO::Services::SensuClient
    - OS::TripleO::Services::Fluentd
    - OS::TripleO::Services::AuditD
    - OS::TripleO::Services::Collectd
    - OS::TripleO::Services::MySQLClient
    - OS::TripleO::Services::Apache
    - OS::TripleO::Services::Horizon
```

It is also a good idea to set the `CountDefault` to `1` so that a default Overcloud always includes the **Horizon** node.

If scaling the services in an existing overcloud, keep the existing services on the **Controller** role. If creating a new overcloud and you want the OpenStack Dashboard to remain on the standalone role, remove the OpenStack Dashboard components from the **Controller** role definition:

```
- name: Controller
  CountDefault: 1
  ServicesDefault:
    ...
    - OS::TripleO::Services::GnocchiMetricd
    - OS::TripleO::Services::GnocchiStatsd
    - OS::TripleO::Services::HAproxy
    - OS::TripleO::Services::HeatApi
```

```

- OS::TripleO::Services::HeatApiCfn
- OS::TripleO::Services::HeatApiCloudwatch
- OS::TripleO::Services::HeatEngine
# - OS::TripleO::Services::Horizon          # Remove this service
- OS::TripleO::Services::IronicApi
- OS::TripleO::Services::IronicConductor
- OS::TripleO::Services::Isccsid
- OS::TripleO::Services::Keepalived
...

```

Generate the new **roles_data** file using the **roles** directory as the source:

```

$ openstack overcloud roles generate -o roles_data-horizon.yaml \
--roles-path ~/roles \
Controller Compute Horizon

```

You might need to define a new flavor for this role so that you can tag specific nodes. For this example, use the following commands to create a **horizon** flavor:

```

$ openstack flavor create --id auto --ram 6144 --disk 40 --vcpus 4 horizon
$ openstack flavor set --property "cpu_arch"="x86_64" --property "capabilities:boot_option"="local" --
property "capabilities:profile"="horizon" horizon
$ openstack flavor set --property resources:VCPU=0 --property resources:MEMORY_MB=0 --
property resources:DISK_GB=0 --property resources:CUSTOM_BAREMETAL=1 horizon

```

Tag nodes into the new flavor using the following command:

```

$ openstack baremetal node set --property capabilities='profile:horizon,boot_option:local' 58c3d07e-
24f2-48a7-bbb6-6843f0e8ee13

```

Define the Horizon node count and flavor using the following environment file snippet:

```

parameter_defaults:
  OvercloudHorizonFlavor: horizon
  HorizonCount: 1

```

Include the new **roles_data** file and environment file when running the **openstack overcloud deploy** command. For example:

```

$ openstack overcloud deploy --templates -r ~/templates/roles_data-horizon.yaml -e
~/templates/node-count-flavor.yaml

```

When the deployment completes, this creates a three-node Overcloud consisting of one Controller node, one Compute node, and one Networker node. To view the Overcloud's list of nodes, run the following command:

```

$ openstack server list

```

7.3. COMPOSABLE SERVICES

7.3.1. Guidelines and Limitations

Note the following guidelines and limitations for the composable node architecture.

For services not managed by Pacemaker:

- You can assign services to standalone custom roles.
- You can create additional custom roles after the initial deployment and deploy them to scale existing services.

For services managed by Pacemaker:

- You can assign Pacemaker managed services to standalone custom roles.
- Pacemaker has a 16 node limit. If you assign the Pacemaker service (**OS::TripleO::Services::Pacemaker**) to 16 nodes, any subsequent nodes must use the Pacemaker Remote service (**OS::TripleO::Services::PacemakerRemote**) instead. You cannot have the Pacemaker service and Pacemaker Remote service on the same role.
- Do not include the Pacemaker service (**OS::TripleO::Services::Pacemaker**) on roles that do not contain Pacemaker managed services.
- You cannot scale up or scale down a custom role that contains **OS::TripleO::Services::Pacemaker** or **OS::TripleO::Services::PacemakerRemote** services.

General Limitations:

- You cannot change custom roles and composable services during the a major version upgrade.
- You cannot modify the list of services for any role after deploying an Overcloud. Modifying the service lists after Overcloud deployment can cause deployment errors and leave orphaned services on nodes.

7.3.2. Examining Composable Service Architecture

The core heat template collection contains two sets of composable service templates:

- **puppet/services** contains the base templates for configuring composable services.
- **docker/services** contains the containerized templates for key OpenStack Platform services. These templates act as augmentations for some of the base templates and reference back to the base templates.

Each template contains a description that identifies its purpose. For example, the **ntp.yaml** service template contains the following description:

```
description: >
  NTP service deployment using puppet, this YAML file
  creates the interface between the HOT template
  and the puppet manifest that actually installs
  and configure NTP.
```

These service templates are registered as resources specific to a RHOSP deployment. This means you can call each resource using a unique heat resource namespace defined in the **overcloud-resource-registry-puppet.j2.yaml** file. All services use the **OS::TripleO::Services** namespace for their resource type.

Some resources use the base composable service templates directly:

```
resource_registry:
  ...
  OS::TripleO::Services::Ntp: puppet/services/time/ntp.yaml
  ...
```

However, core services require containers and as such use the containerized service templates. For example, the **keystone** containerized service uses the following:

```
resource_registry:
  ...
  OS::TripleO::Services::Keystone: docker/services/keystone.yaml
  ...
```

These containerized templates usually reference back to the base templates in order to include Puppet configuration. For example, the **docker/services/keystone.yaml** template stores the output of the base template in the **KeystoneBase** parameter:

```
KeystoneBase:
  type: ../../puppet/services/keystone.yaml
```

The containerized template can then incorporate functions and data from the base template.

The **overcloud.j2.yaml** heat template includes a section of Jinja2-based code to define a service list for each custom role in the **roles_data.yaml** file:

```
{{role.name}}Services:
  description: A list of service resources (configured in the Heat
    resource_registry) which represent nested stacks
    for each service that should get installed on the {{role.name}} role.
  type: comma_delimited_list
  default: {{role.ServicesDefault|default([])}}
```

For the default roles, this creates the following service list parameters: **ControllerServices**, **ComputeServices**, **BlockStorageServices**, **ObjectStorageServices**, and **CephStorageServices**.

You define the default services for each custom role in the **roles_data.yaml** file. For example, the default Controller role contains the following content:

```
- name: Controller
  CountDefault: 1
  ServicesDefault:
    - OS::TripleO::Services::CACerts
    - OS::TripleO::Services::CephMon
    - OS::TripleO::Services::CephExternal
    - OS::TripleO::Services::CephRgw
    - OS::TripleO::Services::CinderApi
    - OS::TripleO::Services::CinderBackup
    - OS::TripleO::Services::CinderScheduler
    - OS::TripleO::Services::CinderVolume
    - OS::TripleO::Services::Core
    - OS::TripleO::Services::Kernel
    - OS::TripleO::Services::Keystone
```

```

- OS::TripleO::Services::GlanceApi
- OS::TripleO::Services::GlanceRegistry
...

```

These services are then defined as the default list for the **ControllerServices** parameter.

You can also use an environment file to override the default list for the service parameters. For example, you can define **ControllerServices** as a **parameter_default** in an environment file to override the services list from the **roles_data.yaml** file.

7.3.3. Adding and Removing Services from Roles

The basic method of adding or removing services involves creating a copy of the default service list for a node role and then adding or removing services. For example, you might aim to remove OpenStack Orchestration (**heat**) from the Controller nodes. In this situation, create a custom copy of the default **roles** directory:

```
$ cp -r /usr/share/openstack-tripleo-heat-templates/roles ~/.
```

Edit the **~/roles/Controller.yaml** file and modify the service list for the **ServicesDefault** parameter. Scroll to the OpenStack Orchestration services and remove them:

```

- OS::TripleO::Services::GlanceApi
- OS::TripleO::Services::GlanceRegistry
- OS::TripleO::Services::HeatApi          # Remove this service
- OS::TripleO::Services::HeatApiCfn      # Remove this service
- OS::TripleO::Services::HeatApiCloudwatch # Remove this service
- OS::TripleO::Services::HeatEngine      # Remove this service
- OS::TripleO::Services::MySQL
- OS::TripleO::Services::NeutronDhcpAgent

```


Generate the new **roles_data** file. For example:

```
$ openstack overcloud roles generate -o roles_data-no_heat.yaml \
  --roles-path ~/roles \
  Controller Compute Networker
```

Include this new **roles_data** file when running the **openstack overcloud deploy** command. For example:

```
$ openstack overcloud deploy --templates -r ~/templates/roles_data-no_heat.yaml
```

This deploys an Overcloud without OpenStack Orchestration services installed on the Controller nodes.

**NOTE**

You can also disable services in the **roles_data** file using a custom environment file. Redirect the services to disable to the **OS::Heat::None** resource. For example:

```
resource_registry:
  OS::TripleO::Services::HeatApi: OS::Heat::None
  OS::TripleO::Services::HeatApiCfn: OS::Heat::None
  OS::TripleO::Services::HeatApiCloudwatch: OS::Heat::None
  OS::TripleO::Services::HeatEngine: OS::Heat::None
```

7.3.4. Enabling Disabled Services

Some services are disabled by default. These services are registered as null operations (**OS::Heat::None**) in the **overcloud-resource-registry-puppet.j2.yaml** file. For example, the Block Storage backup service (**cinder-backup**) is disabled:

```
OS::TripleO::Services::CinderBackup: OS::Heat::None
```

To enable this service, include an environment file that links the resource to its respective Heat templates in the **puppet/services** directory. Some services have predefined environment files in the **environments** directory. For example, the Block Storage backup service uses the **environments/cinder-backup.yaml** file, which contains the following:

```
resource_registry:
  OS::TripleO::Services::CinderBackup: ../docker/services/pacemaker/cinder-backup.yaml
  ...
```

This overrides the default null operation resource and enables the service. Include this environment file when running the **openstack overcloud deploy** command.

```
$ openstack overcloud deploy --templates -e /usr/share/openstack-tripleo-heat-templates/environments/cinder-backup.yaml
```

TIP

For another example of how to enable disabled services, see [Installation](#) in the *OpenStack Data Processing* guide. This section contains instructions on how to enable the OpenStack Data Processing service (**sahara**) on the overcloud.

7.3.5. Creating a Generic Node with No Services

Red Hat OpenStack Platform provides the ability to create generic Red Hat Enterprise Linux 7 nodes without any OpenStack services configured. This is useful in situations where you need to host software outside of the core Red Hat OpenStack Platform environment. For example, OpenStack Platform provides integration with monitoring tools such as Kibana and Sensu (see [Monitoring Tools Configuration Guide](#)). While Red Hat does not provide support for the monitoring tools themselves, the director can create a generic Red Hat Enterprise Linux 7 node to host these tools.



NOTE

The generic node still uses the base **overcloud-full** image rather than a base Red Hat Enterprise Linux 7 image. This means the node has some Red Hat OpenStack Platform software installed but not enabled or configured.

Creating a generic node requires a new role without a **ServicesDefault** list:

```
- name: Generic
```

Include the role in your custom **roles_data** file (**roles_data_with_generic.yaml**). Make sure to keep the existing **Controller** and **Compute** roles.

You can also include an environment file (**generic-node-params.yaml**) to specify how many generic Red Hat Enterprise Linux 7 nodes you require and the flavor when selecting nodes to provision. For example:

```
parameter_defaults:  
  OvercloudGenericFlavor: baremetal  
  GenericCount: 1
```

Include both the roles file and the environment file when running the **openstack overcloud deploy** command. For example:

```
$ openstack overcloud deploy --templates -r ~/templates/roles_data_with_generic.yaml -e  
~/templates/generic-node-params.yaml
```

This deploys a three-node environment with one Controller node, one Compute node, and one generic Red Hat Enterprise Linux 7 node.

CHAPTER 8. CONTAINERIZED SERVICES

The director installs the core OpenStack Platform services as containers on the overcloud. This section provides some background information on how containerized services work.

8.1. CONTAINERIZED SERVICE ARCHITECTURE

The director installs the core OpenStack Platform services as containers on the overcloud. The templates for the containerized services are located in the `/usr/share/openstack-tripleo-heat-templates/docker/services/`. These templates reference their respective composable service templates. For example, the OpenStack Identity (keystone) containerized service template (`docker/services/keystone.yaml`) includes the following resource:

```
KeystoneBase:
  type: ../../puppet/services/keystone.yaml
  properties:
    EndpointMap: {get_param: EndpointMap}
    ServiceData: {get_param: ServiceData}
    ServiceNetMap: {get_param: ServiceNetMap}
    DefaultPasswords: {get_param: DefaultPasswords}
    RoleName: {get_param: RoleName}
    RoleParameters: {get_param: RoleParameters}
```

The **type** refers to the respective OpenStack Identity (keystone) composable service and pulls the **outputs** data from that template. The containerized service merges this data with its own container-specific data.

All nodes using containerized services must enable the **OS::TripleO::Services::Docker** service. When creating a `roles_data.yaml` file for your custom roles configuration, include the the **OS::TripleO::Services::Docker** service with the base composable services, as the containerized services. For example, the **Keystone** role uses the following role definition:

```
- name: Keystone
  ServicesDefault:
    - OS::TripleO::Services::CACerts
    - OS::TripleO::Services::Kernel
    - OS::TripleO::Services::Ntp
    - OS::TripleO::Services::Snmp
    - OS::TripleO::Services::Sshd
    - OS::TripleO::Services::Timezone
    - OS::TripleO::Services::TripleoPackages
    - OS::TripleO::Services::TripleoFirewall
    - OS::TripleO::Services::SensuClient
    - OS::TripleO::Services::Fluentd
    - OS::TripleO::Services::AuditD
    - OS::TripleO::Services::Collectd
    - OS::TripleO::Services::MySQLClient
    - OS::TripleO::Services::Docker
    - OS::TripleO::Services::Keystone
```

8.2. CONTAINERIZED SERVICE PARAMETERS

Each containerized service template contains an **outputs** section that defines a data set passed to the director's OpenStack Orchestration (heat) service. In addition to the standard composable service

parameters (see [Section 7.2.4, “Examining Role Parameters”](#)), the template contain a set of parameters specific to the container configuration.

puppet_config

Data to pass to Puppet when configuring the service. In the initial overcloud deployment steps, the director creates a set of containers used to configure the service before the actual containerized service runs. This parameter includes the following sub-parameters: +

- **config_volume** - The mounted docker volume that stores the configuration.
- **puppet_tags** - Tags to pass to Puppet during configuration. These tags are used in OpenStack Platform to restrict the Puppet run to a particular service’s configuration resource. For example, the OpenStack Identity (keystone) containerized service uses the **keystone_config** tag to ensure all required only the **keystone_config** Puppet resource run on the configuration container.
- **step_config** - The configuration data passed to Puppet. This is usually inherited from the referenced composable service.
- **config_image** - The container image used to configure the service.

kolla_config

A set of container-specific data that defines configuration file locations, directory permissions, and the command to run on the container to launch the service.

docker_config

Tasks to run on the service’s configuration container. All tasks are grouped into steps to help the director perform a staged deployment. The steps are: +

- **Step 1** - Load balancer configuration
- **Step 2** - Core services (Database, Redis)
- **Step 3** - Initial configuration of OpenStack Platform service
- **Step 4** - General OpenStack Platform services configuration
- **Step 5** - Service activation

host_prep_tasks

Preparation tasks for the bare metal node to accommodate the containerized service.

8.3. MODIFYING OPENSTACK PLATFORM CONTAINERS

Red Hat provides a set of pre-built container images through the Red Hat Container Catalog (registry.redhat.io). It is possible to modify these images and add additional layers to them. This is useful for adding RPMs for certified 3rd party drivers to the containers.



NOTE

To ensure continued support for modified OpenStack Platform container images, ensure that the resulting images comply with the ["Red Hat Container Support Policy"](#).

This example shows how to customize the latest **openstack-keystone** image. However, these instructions can also apply to other images:

1. Pull the image you aim to modify. For example, for the **openstack-keystone** image:

```
$ sudo docker pull registry.redhat.io/rhosp13/openstack-keystone:latest
```

2. Check the default user on the original image. For example, for the **openstack-keystone** image:

```
$ sudo docker run -it registry.redhat.io/rhosp13/openstack-keystone:latest whoami
root
```



NOTE

The **openstack-keystone** image uses **root** as the default user. Other images use specific users. For example, **openstack-glance-api** uses **glance** for the default user.

3. Create a **Dockerfile** to build an additional layer on an existing container image. The following is an example that pulls the latest OpenStack Identity (keystone) image from the Container Catalog and installs a custom RPM file to the image:

```
FROM registry.redhat.io/rhosp13/openstack-keystone
MAINTAINER Acme
LABEL name="rhosp13/openstack-keystone-acme" vendor="Acme" version="2.1"
release="1"

# switch to root and install a custom RPM, etc.
USER root
COPY custom.rpm /tmp
RUN rpm -ivh /tmp/custom.rpm

# switch the container back to the default user
USER root
```

4. Build and tag the new image. For example, to build with a local **Dockerfile** stored in the **/home/stack/keystone** directory and tag it to your undercloud's local registry:

```
$ docker build /home/stack/keystone -t "192.168.24.1:8787/rhosp13/openstack-keystone-acme:rev1"
```

5. Push the resulting image to the undercloud's local registry:

```
$ docker push 192.168.24.1:8787/rhosp13/openstack-keystone-acme:rev1
```

6. Edit your overcloud container images environment file (usually **overcloud_images.yaml**) and change the appropriate parameter to use the custom container image.

**WARNING**

The Container Catalog publishes container images with a complete software stack built into it. When the Container Catalog releases a container image with updates and security fixes, your existing custom container will **not** include these updates and will require rebuilding using the new image version from the Catalog.

8.4. DEPLOYING A VENDOR PLUGIN

To use third-party hardware as a Block Storage back end, you must deploy a vendor plugin. The following example demonstrates how to deploy a vendor plugin to use Dell EMC hardware as a Block Storage back end.

1. Log in to the **registry.connect.redhat.com** catalog:

```
$ docker login registry.connect.redhat.com
```

2. Download the plugin:

```
$ docker pull registry.connect.redhat.com/dellemc/openstack-cinder-volume-dellemc-rhosp13
```

3. Tag and push the image to the local undercloud registry using the undercloud IP address relevant to your OpenStack deployment:

```
$ docker tag registry.connect.redhat.com/dellemc/openstack-cinder-volume-dellemc-rhosp13  
192.168.24.1:8787/dellemc/openstack-cinder-volume-dellemc-rhosp13
```

```
$ docker push 192.168.24.1:8787/dellemc/openstack-cinder-volume-dellemc-rhosp13
```

4. Deploy the overcloud with an additional environment file that contains the following parameter:

```
parameter_defaults:  
  DockerCinderVolumeImage: 192.168.24.1:8787/dellemc/openstack-cinder-volume-dellemc-  
rhosp13
```

CHAPTER 9. BASIC NETWORK ISOLATION

This chapter show how to configure the overcloud with the standard network isolation configuration. This includes:

- The rendered environment file to enable network isolation (`/usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml`).
- A copied environment file to configure network defaults (`/usr/share/openstack-tripleo-heat-templates/environments/network-environment.yaml`).
- A `network_data` file to define network settings such as IP ranges, subnets, and virtual IPs. This example shows you how to create a copy of the default and edit it to suit your own network.
- Templates to define your NIC layout for each node. The overcloud core template collection contains a set of defaults for different use cases.
- An environment file to enable NICs. This example uses a default file located in the `environments` directory.
- Any additional environment files to customize your networking parameters.



NOTE

Run the `openstack overcloud netenv validate` command to validate the syntax of your `network-environment.yaml` file. This command also validates the individual nic-config files for compute, controller, storage, and composable roles network files. Use the `-f` or `--file` options to specify the file that you want to validate:

```
$ openstack overcloud netenv validate -f ~/templates/network-environment.yaml
```

The following content in this chapter shows how to define each of these aspects.

9.1. NETWORK ISOLATION

The overcloud assigns services to the provisioning network by default. However, director can divide overcloud network traffic into isolated networks. To use isolated networks, the overcloud contains an environment file that enables this feature. The `environments/network-isolation.j2.yaml` file in the core heat templates is a Jinja2 file that defines all ports and VIPs for each network in your composable network file. When rendered, it results in a `network-isolation.yaml` file in the same location with the full resource registry:

```
resource_registry:
  # networks as defined in network_data.yaml
  OS::TripleO::Network::Storage: ../network/storage.yaml
  OS::TripleO::Network::StorageMgmt: ../network/storage_mgmt.yaml
  OS::TripleO::Network::InternalApi: ../network/internal_api.yaml
  OS::TripleO::Network::Tenant: ../network/tenant.yaml
  OS::TripleO::Network::External: ../network/external.yaml

  # Port assignments for the VIPs
  OS::TripleO::Network::Ports::StorageVipPort: ../network/ports/storage.yaml
  OS::TripleO::Network::Ports::StorageMgmtVipPort: ../network/ports/storage_mgmt.yaml
  OS::TripleO::Network::Ports::InternalApiVipPort: ../network/ports/internal_api.yaml
```

```

OS::TripleO::Network::Ports::ExternalVipPort: ../network/ports/external.yaml
OS::TripleO::Network::Ports::RedisVipPort: ../network/ports/vip.yaml

# Port assignments by role, edit role definition to assign networks to roles.
# Port assignments for the Controller
OS::TripleO::Controller::Ports::StoragePort: ../network/ports/storage.yaml
OS::TripleO::Controller::Ports::StorageMgmtPort: ../network/ports/storage_mgmt.yaml
OS::TripleO::Controller::Ports::InternalApiPort: ../network/ports/internal_api.yaml
OS::TripleO::Controller::Ports::TenantPort: ../network/ports/tenant.yaml
OS::TripleO::Controller::Ports::ExternalPort: ../network/ports/external.yaml

# Port assignments for the Compute
OS::TripleO::Compute::Ports::StoragePort: ../network/ports/storage.yaml
OS::TripleO::Compute::Ports::InternalApiPort: ../network/ports/internal_api.yaml
OS::TripleO::Compute::Ports::TenantPort: ../network/ports/tenant.yaml

# Port assignments for the CephStorage
OS::TripleO::CephStorage::Ports::StoragePort: ../network/ports/storage.yaml
OS::TripleO::CephStorage::Ports::StorageMgmtPort: ../network/ports/storage_mgmt.yaml

```

The first section of this file has the resource registry declaration for the **OS::TripleO::Network::*** resources. By default these resources use the **OS::Heat::None** resource type, which does not create any networks. By redirecting these resources to the YAML files for each network, you enable the creation of these networks.

The next several sections create the IP addresses for the nodes in each role. The controller nodes have IPs on each network. The compute and storage nodes each have IPs on a subset of the networks.

Other functions of overcloud networking, such as [Chapter 10, Custom composable networks](#) and [Chapter 11, Custom network interface templates](#) rely on this network isolation environment file. As a result, you need to include the name of the rendered file with your deployment commands. For example:

```

$ openstack overcloud deploy --templates \
...
-e /usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml \
...

```

9.2. MODIFYING ISOLATED NETWORK CONFIGURATION

Copy the default **network_data.yaml** file and modify the copy to configure the default isolated networks.

Procedure

1. Copy the default **network_data** file:

```
$ cp /usr/share/openstack-tripleo-heat-templates/network_data.yaml /home/stack/.
```

2. Edit the local copy of the **network_data.yaml** file and modify the parameters to suit your networking requirements. For example, the Internal API network contains the following default network details:

```
- name: InternalApi
  name_lower: internal_api
```

```

vip: true
vlan: 201
ip_subnet: '172.16.2.0/24'
allocation_pools: [{'start': '172.16.2.4', 'end': '172.16.2.250'}]

```

Edit the following for each network:

- **vlan** defines the VLAN ID to use for this network.
- **ip_subnet** and **ip_allocation_pools** set the default subnet and IP range for the network..
- **gateway** sets the gateway for the network. Used mostly to define the default route for the External network, but can be used for other networks if necessary.

Include the custom **network_data** file with your deployment using the **-n** option. Without the **-n** option, the deployment command uses the default network details.

9.3. NETWORK INTERFACE TEMPLATES

The overcloud network configuration requires a set of the network interface templates. These templates are standard heat templates in YAML format. Each role requires a NIC template so that director can configure each node within that role correctly.

All NIC templates contain the same sections as standard Heat templates:

heat_template_version

The syntax version to use.

description

A string description of the template.

parameters

Network parameters to include in the template.

resources

Takes parameters defined in **parameters** and applies them to a network configuration script.

outputs

Renders the final script used for configuration.

The default NIC templates in **/usr/share/openstack-tripleo-heat-templates/network/config** take advantage of Jinja2 syntax to help render the template. For example, the following snippet from the **single-nic-vlans** configuration renders a set of VLANs for each network:

```

{%- for network in networks if network.enabled|default(true) and network.name in role.networks %}
- type: vlan
  vlan_id:
    get_param: {{network.name}}NetworkVlanID
  addresses:
  - ip_netmask:
    get_param: {{network.name}}IpSubnet
{%- if network.name in role.default_route_networks %}

```

For default Compute nodes, this only renders network information for the Storage, Internal API, and Tenant networks:

```

- type: vlan
  vlan_id:
    get_param: StorageNetworkVlanID
  device: bridge_name
  addresses:
  - ip_netmask:
      get_param: StorageIpSubnet
- type: vlan
  vlan_id:
    get_param: InternalApiNetworkVlanID
  device: bridge_name
  addresses:
  - ip_netmask:
      get_param: InternalApiIpSubnet
- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: bridge_name
  addresses:
  - ip_netmask:
      get_param: TenantIpSubnet

```

[Chapter 11, Custom network interface templates](#) explores how to render the default Jinja2-based templates to standard YAML versions, which you can use as a basis for customization.

9.4. DEFAULT NETWORK INTERFACE TEMPLATES

Director contains templates in `/usr/share/openstack-tripleo-heat-templates/network/config/` to suit most common network scenarios. The following table outlines each NIC template set and the respective environment file that you must use to enable the templates.




NOTE

Each environment file for enabling NIC templates uses the suffix `.j2.yaml`. This is the unrendered Jinja2 version. Make sure to include the rendered file name, which only uses the `.yaml` suffix, in your deployment.

NIC directory	Description	Environment file
single-nic-vlans	Single NIC (nic1) with control plane and VLANs attached to default Open vSwitch bridge.	environments/net-single-nic-with-vlans.j2.yaml
single-nic-linux-bridge-vlans	Single NIC (nic1) with control plane and VLANs attached to default Linux bridge.	environments/net-single-nic-linux-bridge-with-vlans
bond-with-vlans	Control plane attached to nic1 . Default Open vSwitch bridge with bonded NIC configuration (nic2 and nic3) and VLANs attached.	environments/net-bond-with-vlans.yaml

NIC directory	Description	Environment file
multiple-nics	Control plane attached to nic1 . Assigns each sequential NIC to each network defined in the network_data file. By default, this is Storage to nic2 , Storage Management to nic3 , Internal API to nic4 , Tenant to nic5 on the br-tenant bridge, and External to nic6 on the default Open vSwitch bridge.	environments/net-multiple-nics.yaml



NOTE

Environment files exist for using no external network (for example, **net-bond-with-vlans-no-external.yaml**) and using IPv6 (for example, **net-bond-with-vlans-v6.yaml**). These are provided for backwards compatibility and do not function with composable networks.

Each default NIC template set contains a **role.role.j2.yaml** template. This file uses Jinja2 to render additional files for each composable role. For example if your overcloud uses Compute, Controller, and Ceph Storage roles, the deployment renders new templates based on **role.role.j2.yaml**, such as

- **compute.yaml**
- **controller.yaml**
- **ceph-storage.yaml**.

9.5. ENABLING BASIC NETWORK ISOLATION

This procedure show how to enable basic network isolation using one of the default NIC templates. In this case, it is the single NIC with VLANs template (**single-nic-vlans**).

Procedure

1. When running the **openstack overcloud deploy** command, make sure to include the rendered environment file names for:
 - The custom **network_data** file.
 - The rendered file name of the default network isolation.
 - The rendered file name of the default network environment file.
 - The rendered file name of the default network interface configuration
 - Any additional environment files relevant to your configuration.

For example:

```
$ openstack overcloud deploy --templates \
```



```
...  
-n /home/stack/network_data.yaml \  
-e /usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml \  
-e /usr/share/openstack-tripleo-heat-templates/environments/network-environment.yaml \  
-e /usr/share/openstack-tripleo-heat-templates/environments/net-single-nic-with-vlans.yaml \  
...
```

CHAPTER 10. CUSTOM COMPOSABLE NETWORKS

This chapter follows on from the concepts and procedures outlined in [Chapter 9, Basic network isolation](#) and shows how to configure the overcloud with an additional composable network. This includes:

- The environment file to enable network isolation (**/usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml**).
- The environment file to configure network defaults (**/usr/share/openstack-tripleo-heat-templates/environments/network-environment.yaml**).
- A custom **network_data** file to create additional networks outside of the defaults.
- A custom **roles_data** file to assign custom networks to roles.
- Templates to define your NIC layout for each node. The overcloud core template collection contains a set of defaults for different use cases.
- An environment file to enable NICs. This example uses a default file located in the **environments** directory.
- Any additional environment files to customize your networking parameters. This example uses an environment file to customize OpenStack service mappings to composable networks.



NOTE

Run the **openstack overcloud netenv validate** command to validate the syntax of your **network-environment.yaml** file. This command also validates the individual nic-config files for compute, controller, storage, and composable roles network files. Use the **-f** or **--file** options to specify the file that you want to validate:

```
$ openstack overcloud netenv validate -f ~/templates/network-environment.yaml
```

The following content in this chapter shows how to define each of these aspects.

10.1. COMPOSABLE NETWORKS

The overcloud uses the following pre-defined set of network segments by default:

- Control Plane
- Internal API
- Storage
- Storage Management
- Tenant
- External
- Management (optional)

Composable networks allow you to add networks for various services. For example, if you have a network dedicated to NFS traffic, you can present it to multiple roles.

Director supports the creation of custom networks during the deployment and update phases. These additional networks can be used for ironic bare metal nodes, system management, or to create separate networks for different roles. You can also use them to create multiple sets of networks for split deployments where traffic is routed between networks.

A single data file (**network_data.yaml**) manages the list of networks to be deployed. You include this file with your deployment command using the **-n** option. Without this option, the deployment uses the default file (`/usr/share/openstack-tripleo-heat-templates/network_data.yaml`).

10.2. ADDING A COMPOSABLE NETWORK

Use composable networks to add networks for various services. For example, if you have a network that is dedicated to storage backup traffic, you can present the network to multiple roles.

Procedure

1. Copy the default **network_data** file:

```
$ cp /usr/share/openstack-tripleo-heat-templates/network_data.yaml /home/stack/.
```

2. Edit the local copy of the **network_data.yaml** file and add a section for your new network. For example:

```
- name: StorageBackup
  vip: true
  name_lower: storage_backup
  ip_subnet: '172.21.1.0/24'
  allocation_pools: [{'start': '171.21.1.4', 'end': '172.21.1.250'}]
  gateway_ip: '172.21.1.1'
```

- Sets the human readable name of the network. This parameter is the only mandatory parameter. If you want to normalize names for readability, use the **name_lower** parameter, for example, if you want to change **InternalApi** to **internal_api**. Do not modify the **name** parameter.
- **vip: true** creates a virtual IP address (VIP) on the new network. This IP is used as the target IP for services listed in the service-to-network mapping parameter (**ServiceNetMap**). Note that VIPs are only used by roles that use Pacemaker. The overcloud's load balancing service redirects traffic from these IPs to their respective service endpoint.
- **ip_subnet**, **allocation_pools**, and **gateway_ip** set the default IPv4 subnet, IP range, and gateway for the network.

Include the custom **network_data** file with your deployment using the **-n** option. Without the **-n** option, the deployment command uses the default set of networks.

10.3. INCLUDING A COMPOSABLE NETWORK IN A ROLE

You can assign composable networks to the overcloud roles defined in your environment. For example, you might include a custom **StorageBackup** network with your Ceph Storage nodes.

Procedure

1. If you do not already have a custom **roles_data** file, copy the default to your home directory:

```
$ cp /usr/share/openstack-tripleo-heat-templates/roles_data.yaml /home/stack/.
```

2. Edit the custom **roles_data** file.
3. Scroll to the role you want to add the composable network and add the network name to the list of **networks**. For example, to add the network to the Ceph Storage role, use the following snippet as a guide:

```
- name: CephStorage
  description: |
    Ceph OSD Storage node role
  networks:
    - Storage
    - StorageMgmt
    - StorageBackup
```

4. After adding custom networks to their respective roles, save the file.

When running the **openstack overcloud deploy** command, include the **roles_data** file using the **-r** option. Without the **-r** option, the deployment command uses the default set of roles with their respective assigned networks.

10.4. ASSIGNING OPENSTACK SERVICES TO COMPOSABLE NETWORKS

Each OpenStack service is assigned to a default network type in the resource registry. These services are bound to IP addresses within the network type's assigned network. Although the OpenStack services are divided among these networks, the number of actual physical networks can differ as defined in the network environment file. You can reassign OpenStack services to different network types by defining a new network map in an environment file, for example, **/home/stack/templates/service-reassignments.yaml**. The **ServiceNetMap** parameter determines the network types that you want to use for each service.

For example, you can reassign the Storage Management network services to the Storage Backup Network by modifying the highlighted sections:

```
parameter_defaults:
  ServiceNetMap:
    SwiftMgmtNetwork: storage_backup
    CephClusterNetwork: storage_backup
```

Changing these parameters to **storage_backup** will place these services on the Storage Backup network instead of the Storage Management network. This means you only need to define a set of **parameter_defaults** for the Storage Backup network and not the Storage Management network.

The director merges your custom **ServiceNetMap** parameter definitions into a pre-defined list of defaults taken from **ServiceNetMapDefaults** and overrides the defaults. The director then returns the full list including customizations back to **ServiceNetMap**, which is used to configure network assignments for various services.

Service mappings only apply to networks that use **vip: true** in the **network_data** file for nodes that use Pacemaker. The overcloud's load balancer redirects traffic from the VIPs to the specific service endpoints.

**NOTE**

A full list of default services can be found in the **ServiceNetMapDefaults** parameter within `/usr/share/openstack-tripleo-heat-templates/network/service_net_map.j2.yaml`.

10.5. ENABLING CUSTOM COMPOSABLE NETWORKS

Enable custom composable networks using one of the default NIC templates. In this example, use the Single NIC with VLANs template (**net-single-nic-with-vlans**).


Procedure

1. When running the **openstack overcloud deploy** command, make sure to include:
 - The custom **network_data** file.
 - The custom **roles_data** file with network-to-role assignments.
 - The rendered file name of the default network isolation configuration.
 - The rendered file name of the default network environment file.
 - The rendered file name of the default network interface configuration.
 - Any additional environment files related to your network, such as the service reassignments.

For example:

```
$ openstack overcloud deploy --templates \
...
-n /home/stack/network_data.yaml \
-r /home/stack/roles_data.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/network-environment.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/net-single-nic-with-vlans.yaml \
-e /home/stack/templates/service-reassignments.yaml \
...
```

This deploys the composable networks, including your additional custom networks, across nodes in your overcloud.

**IMPORTANT**

Remember that you must render the templates again if you are introducing a new custom network, such as a management network. Simply adding the network name to the **roles_data.yaml** file is not sufficient.

10.6. RENAMING THE DEFAULT NETWORKS

You can use the **network_data.yaml** file to modify the user-visible names of the default networks:

- InternalApi
- External

- Storage
- StorageMgmt
- Tenant

To change these names, do not modify the **name** field. Instead, change the **name_lower** field to the new name for the network and update the ServiceNetMap with the new name.

Procedure

1. In your **network_data.yaml** file, enter new names in the **name_lower** parameter for each network that you want to rename:

```
- name: InternalApi  
  name_lower: MyCustomInternalApi
```

2. Include the default value of the **name_lower** parameter in the **service_net_map_replace** parameter:

```
- name: InternalApi  
  name_lower: MyCustomInternalApi  
  service_net_map_replace: internal_api
```

CHAPTER 11. CUSTOM NETWORK INTERFACE TEMPLATES

This chapter follows on from the concepts and procedures outlined in [Chapter 9, Basic network isolation](#). The purpose of this chapter is to demonstrate how to create a set of custom network interface template to suit nodes in your environment. This includes:

- The environment file to enable network isolation (**/usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml**).
- The environment file to configure network defaults (**/usr/share/openstack-tripleo-heat-templates/environments/network-environment.yaml**).
- Templates to define your NIC layout for each node. The overcloud core template collection contains a set of defaults for different use cases. In this situation, you will render a default a basis for your custom templates.
- A custom environment file to enable NICs. This example uses a custom environment file (**/home/stack/templates/custom-network-configuration.yaml**) that references your custom interface templates.
- Any additional environment files to customize your networking parameters.
- If using customizing your networks, a custom **network_data** file.
- If creating additional or custom composable networks, a custom **network_data** file and a custom **roles_data** file.



NOTE

Run the **openstack overcloud netenv validate** command to validate the syntax of your **network-environment.yaml** file. This command also validates the individual nic-config files for compute, controller, storage, and composable roles network files. Use the **-f** or **--file** options to specify the file that you want to validate:

```
$ openstack overcloud netenv validate -f ~/templates/network-environment.yaml
```

11.1. CUSTOM NETWORK ARCHITECTURE

The default NIC templates might not suit a specific network configuration. For example, you might want to create your own custom NIC template that suits a specific network layout. You might want to separate the control services and data services on to separate NICs. In this situation, you can map the service to NIC assignments in the following way:

- NIC1 (Provisioning):
 - Provisioning / Control Plane
- NIC2 (Control Group)
 - Internal API
 - Storage Management
 - External (Public API)

- NIC3 (Data Group)
 - Tenant Network (VXLAN tunneling)
 - Tenant VLANs / Provider VLANs
 - Storage
 - External VLANs (Floating IP/SNAT)
- NIC4 (Management)
 - Management

11.2. RENDERING DEFAULT NETWORK INTERFACE TEMPLATES FOR CUSTOMIZATION

To simplify the configuration of custom interface templates, render the Jinja2 syntax of a default NIC template and use the rendered templates as the basis for your custom configuration.

Procedure

1. Render a copy of the **openstack-tripleo-heat-templates** collection using the **process-templates.py** script:

```
$ cd /usr/share/openstack-tripleo-heat-templates
$ ./tools/process-templates.py -o ~/openstack-tripleo-heat-templates-rendered
```

This converts all Jinja2 templates to their rendered YAML versions and saves the results to **~/openstack-tripleo-heat-templates-rendered**.

If using a custom network file or custom roles file, you can include these files using the **-n** and **-r** options respectively. For example:

```
$ ./tools/process-templates.py -o ~/openstack-tripleo-heat-templates-rendered -n
/home/stack/network_data.yaml -r /home/stack/roles_data.yaml
```

2. Copy the multiple NIC example:

```
$ cp -r ~/openstack-tripleo-heat-templates-rendered/network/config/multiple-nics/
~/templates/custom-nics/
```

3. You can edit the template set in **custom-nics** to suit your own network configuration.

11.3. NETWORK INTERFACE ARCHITECTURE

The custom NIC templates that you render in [Section 11.2, “Rendering default network interface templates for customization”](#) contain the **parameters** and **resources** sections.

Parameters

The **parameters** section contains all network configuration parameters for network interfaces. This includes information such as subnet ranges and VLAN IDs. This section should remain unchanged as the Heat template inherits values from its parent template. However, you can modify the values for some

parameters using a network environment file.

Resources

The **resources** section is where the main network interface configuration occurs. In most cases, the **resources** section is the only one that requires editing. Each **resources** section begins with the following header:

```
resources:
  OsNetConfigImpl:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script
      config:
        str_replace:
          template:
            get_file: /usr/share/openstack-tripleo-heat-templates/network/scripts/run-os-net-config.sh
        params:
          $network_config:
            network_config:
```

This runs a script (**run-os-net-config.sh**) that creates a configuration file for **os-net-config** to use for configuring network properties on a node. The **network_config** section contains the custom network interface data sent to the **run-os-net-config.sh** script. You arrange this custom interface data in a sequence based on the type of device.



IMPORTANT

If creating custom NIC templates, you must set the **run-os-net-config.sh** script location to an absolute location for each NIC template. The script is located at **/usr/share/openstack-tripleo-heat-templates/network/scripts/run-os-net-config.sh** on the undercloud.

11.4. NETWORK INTERFACE REFERENCE

Network interface configuration contains the following parameters:

interface

Defines a single network interface. The configuration defines each interface using either the actual interface name ("eth0", "eth1", "enp0s25") or a set of numbered interfaces ("nic1", "nic2", "nic3").

For example:

```
- type: interface
  name: nic2
```

Table 11.1. interface options

Option	Default	Description
name		Name of the Interface
use_dhcp	False	Use DHCP to get an IP address

Option	Default	Description
use_dhcpv6	False	Use DHCP to get a v6 IP address
addresses		A list of IP addresses assigned to the interface
routes		A list of routes assigned to the interface. See routes .
mtu	1500	The maximum transmission unit (MTU) of the connection
primary	False	Defines the interface as the primary interface
defroute	True	Use a default route provided by the DHCP service. Only applies when use_dhcp or use_dhcpv6 is enabled.
persist_mapping	False	Write the device alias configuration instead of the system names
dhclient_args	None	Arguments to pass to the DHCP client
dns_servers	None	List of DNS servers to use for the interface
ethtool_opts		Set this option to " rx-flow-hash udp4 sdfn " to improve throughput when you use VXLAN on certain NICs.

vlan

Defines a VLAN. Use the VLAN ID and subnet passed from the **parameters** section.

For example:

```
- type: vlan
  vlan_id:{get_param: ExternalNetworkVlanID}
  addresses:
    - ip_netmask: {get_param: ExternalIpSubnet}
```

Table 11.2. vlan options

Option	Default	Description
vlan_id		The VLAN ID
device		The parent device to attach the VLAN. Use this parameter when the VLAN is not a member of an OVS bridge. For example, use this parameter to attach the VLAN to a bonded interface device.
use_dhcp	False	Use DHCP to get an IP address
use_dhcpv6	False	Use DHCP to get a v6 IP address
addresses		A list of IP addresses assigned to the VLAN
routes		A list of routes assigned to the VLAN. See routes .
mtu	1500	The maximum transmission unit (MTU) of the connection
primary	False	Defines the VLAN as the primary interface
defroute	True	Use a default route provided by the DHCP service. Only applies when use_dhcp or use_dhcpv6 is enabled.
persist_mapping	False	Write the device alias configuration instead of the system names
dhclient_args	None	Arguments to pass to the DHCP client
dns_servers	None	List of DNS servers to use for the VLAN

ovs_bond

Defines a bond in Open vSwitch to join two or more **interfaces** together. This helps with redundancy and increases bandwidth.

For example:

```
- type: ovs_bond
```

```

name: bond1
members:
- type: interface
  name: nic2
- type: interface
  name: nic3

```

Table 11.3. ovs_bond options

Option	Default	Description
name		Name of the bond
use_dhcp	False	Use DHCP to get an IP address
use_dhcpv6	False	Use DHCP to get a v6 IP address
addresses		A list of IP addresses assigned to the bond
routes		A list of routes assigned to the bond. See routes .
mtu	1500	The maximum transmission unit (MTU) of the connection
primary	False	Defines the interface as the primary interface
members		A sequence of interface objects to use in the bond
ovs_options		A set of options to pass to OVS when creating the bond
ovs_extra		A set of options to set as the OVS_EXTRA parameter in the bond's network configuration file
defroute	True	Use a default route provided by the DHCP service. Only applies when use_dhcp or use_dhcpv6 is enabled.
persist_mapping	False	Write the device alias configuration instead of the system names
dhclient_args	None	Arguments to pass to the DHCP client

Option	Default	Description
dns_servers	None	List of DNS servers to use for the bond

ovs_bridge

Defines a bridge in Open vSwitch, which connects multiple **interface**, **ovs_bond**, and **vlan** objects together.

The network interface type, **ovs_bridge**, takes a parameter **name**.



NOTE

If you have multiple bridges, you must use distinct bridge names other than accepting the default name of **bridge_name**. If you do not use distinct names, then during the converge phase, two network bonds are placed on the same bridge.

If you are defining an OVS bridge for the external tripleo network, then retain the values **bridge_name** and **interface_name** as your deployment framework automatically replaces these values with an external bridge name and an external interface name, respectively.

For example:

```
- type: ovs_bridge
  name: bridge_name
  addresses:
  - ip_netmask:
    list_join:
      - /
      - - {get_param: ControlPlaneIp}
        - {get_param: ControlPlaneSubnetCidr}
  members:
  - type: interface
    name: interface_name
- type: vlan
  device: bridge_name
  vlan_id:
    {get_param: ExternalNetworkVlanID}
  addresses:
  - ip_netmask:
    {get_param: ExternalIpSubnet}
```



NOTE

The OVS bridge connects to the Neutron server in order to get configuration data. If the OpenStack control traffic (typically the Control Plane and Internal API networks) is placed on an OVS bridge, then connectivity to the Neutron server gets lost whenever OVS is upgraded or the OVS bridge is restarted by the admin user or process. This will cause some downtime. If downtime is not acceptable under these circumstances, then the Control group networks should be placed on a separate interface or bond rather than on an OVS bridge:

- A minimal setting can be achieved, when you put the Internal API network on a VLAN on the provisioning interface and the OVS bridge on a second interface.
- If you want bonding, you need at least two bonds (four network interfaces). The control group should be placed on a Linux bond (Linux bridge). If the switch does not support LACP fallback to a single interface for PXE boot, then this solution requires at least five NICs.

Table 11.4. ovs_bridge options

Option	Default	Description
name		Name of the bridge
use_dhcp	False	Use DHCP to get an IP address
use_dhcpv6	False	Use DHCP to get a v6 IP address
addresses		A list of IP addresses assigned to the bridge
routes		A list of routes assigned to the bridge. See routes .
mtu	1500	The maximum transmission unit (MTU) of the connection
members		A sequence of interface, VLAN, and bond objects to use in the bridge
ovs_options		A set of options to pass to OVS when creating the bridge
ovs_extra		A set of options to to set as the OVS_EXTRA parameter in the bridge's network configuration file

Option	Default	Description
defroute	True	Use a default route provided by the DHCP service. Only applies when use_dhcp or use_dhcpv6 is enabled.
persist_mapping	False	Write the device alias configuration instead of the system names
dhclient_args	None	Arguments to pass to the DHCP client
dns_servers	None	List of DNS servers to use for the bridge

linux_bond

Defines a Linux bond that joins two or more **interfaces** together. This helps with redundancy and increases bandwidth. Make sure to include the kernel-based bonding options in the **bonding_options** parameter. For more information on Linux bonding options, see [7.7.1. Bonding Module Directives](#) in the Red Hat Enterprise Linux 7 Networking Guide.

For example:

```
- type: linux_bond
  name: bond1
  members:
  - type: interface
    name: nic2
    primary: true
  - type: interface
    name: nic3
  bonding_options: "mode=802.3ad"
```

Note that **nic2** uses **primary: true**. This ensures the bond uses the MAC address for **nic2**.

Table 11.5. linux_bond options

Option	Default	Description
name		Name of the bond
use_dhcp	False	Use DHCP to get an IP address
use_dhcpv6	False	Use DHCP to get a v6 IP address
addresses		A list of IP addresses assigned to the bond

Option	Default	Description
routes		A list of routes assigned to the bond. See routes .
mtu	1500	The maximum transmission unit (MTU) of the connection
primary	False	Defines the interface as the primary interface.
members		A sequence of interface objects to use in the bond
bonding_options		A set of options when creating the bond. For more information on Linux bonding options, see 7.7.1. Bonding Module Directives in the Red Hat Enterprise Linux 7 Networking Guide.
defroute	True	Use a default route provided by the DHCP service. Only applies when use_dhcp or use_dhcpv6 is enabled.
persist_mapping	False	Write the device alias configuration instead of the system names
dhclient_args	None	Arguments to pass to the DHCP client
dns_servers	None	List of DNS servers to use for the bond

linux_bridge

Defines a Linux bridge, which connects multiple **interface**, **linux_bond**, and **vlan** objects together. The external bridge also uses two special values for parameters:

- **bridge_name**, which is replaced with the external bridge name.
- **interface_name**, which is replaced with the external interface.

For example:

```
- type: linux_bridge
  name: bridge_name
  addresses:
    - ip_netmask:
      list_join:
```



```

- /
- - {get_param: ControlPlaneIp}
- {get_param: ControlPlaneSubnetCidr}
members:
- type: interface
  name: interface_name
- type: vlan
  device: bridge_name
  vlan_id:
  {get_param: ExternalNetworkVlanID}
addresses:
- ip_netmask:
  {get_param: ExternalIpSubnet}

```

Table 11.6. linux_bridge options

Option	Default	Description
name		Name of the bridge
use_dhcp	False	Use DHCP to get an IP address
use_dhcpv6	False	Use DHCP to get a v6 IP address
addresses		A list of IP addresses assigned to the bridge
routes		A list of routes assigned to the bridge. See routes .
mtu	1500	The maximum transmission unit (MTU) of the connection
members		A sequence of interface, VLAN, and bond objects to use in the bridge
defroute	True	Use a default route provided by the DHCP service. Only applies when use_dhcp or use_dhcpv6 is enabled.
persist_mapping	False	Write the device alias configuration instead of the system names
dhclient_args	None	Arguments to pass to the DHCP client
dns_servers	None	List of DNS servers to use for the bridge

routes

Defines a list of routes to apply to a network interface, VLAN, bridge, or bond.

For example:

```

- type: interface
  name: nic2
  ...
  routes:
  - ip_netmask: 10.1.2.0/24
    default: true
    next_hop:
      get_param: EC2MetadataIp

```

Option	Default	Description
ip_netmask	None	IP and netmask of the destination network.
default	False	Sets this this route to a default route. Equivalent to setting ip_netmask: 0.0.0.0/0 .
next_hop	None	The IP address of the router used to reach the destination network.

11.5. EXAMPLE NETWORK INTERFACE LAYOUT

The following snippet for an example Controller node NIC template demonstrates how to configure the custom network scenario to keep the control group separate from the OVS bridge:

```

resources:
  OsNetConfigImpl:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script
      config:
        str_replace:
          template:
            get_file: /usr/share/openstack-tripleo-heat-templates/network/scripts/run-os-net-config.sh
        params:
          $network_config:
            network_config:

            # NIC 1 - Provisioning
            - type: interface
              name: nic1
              use_dhcp: false
              addresses:
                - ip_netmask:
                    list_join:
                      - /

```

```

    - - get_param: ControlPlaneIp
    - get_param: ControlPlaneSubnetCidr
  routes:
  - ip_netmask: 169.254.169.254/32
    next_hop:
      get_param: EC2MetadataIp

# NIC 2 - Control Group
- type: interface
  name: nic2
  use_dhcp: false
- type: vlan
  device: nic2
  vlan_id:
    get_param: InternalApiNetworkVlanID
  addresses:
  - ip_netmask:
      get_param: InternalApiIpSubnet
- type: vlan
  device: nic2
  vlan_id:
    get_param: StorageMgmtNetworkVlanID
  addresses:
  - ip_netmask:
      get_param: StorageMgmtIpSubnet
- type: vlan
  device: nic2
  vlan_id:
    get_param: ExternalNetworkVlanID
  addresses:
  - ip_netmask:
      get_param: ExternalIpSubnet
  routes:
  - default: true
    next_hop:
      get_param: ExternalInterfaceDefaultRoute

# NIC 3 - Data Group
- type: ovs_bridge
  name: bridge_name
  dns_servers:
    get_param: DnsServers
  members:
  - type: interface
    name: nic3
    primary: true
  - type: vlan
    vlan_id:
      get_param: StorageNetworkVlanID
    addresses:
    - ip_netmask:
        get_param: StorageIpSubnet
  - type: vlan
    vlan_id:
      get_param: TenantNetworkVlanID
    addresses:

```

```

- ip_netmask:
  get_param: TenantIpSubnet

# NIC 4 - Management
- type: interface
  name: nic4
  use_dhcp: false
  addresses:
  - ip_netmask: {get_param: ManagementIpSubnet}
  routes:
  - default: true
    next_hop: {get_param: ManagementInterfaceDefaultRoute}

```

This template uses four network interfaces and assigns a number of tagged VLAN devices to the numbered interfaces, **nic1** to **nic4**. On **nic3** it creates the OVS bridge that hosts the Storage and Tenant networks. As a result, it creates the following layout:

- NIC1 (Provisioning):
 - Provisioning / Control Plane
- NIC2 (Control Group)
 - Internal API
 - Storage Management
 - External (Public API)
- NIC3 (Data Group)
 - Tenant Network (VXLAN tunneling)
 - Tenant VLANs / Provider VLANs
 - Storage
 - External VLANs (Floating IP/SNAT)
- NIC4 (Management)
 - Management

11.6. NETWORK INTERFACE TEMPLATE CONSIDERATIONS FOR CUSTOM NETWORKS

When you use composable networks, the **process-templates.py** script renders the static templates to include networks and roles that you define in your **network_data.yaml** and **roles_data.yaml** files. Ensure that your rendered NIC templates contain the following items:

- Static file for each roles, including custom composable networks.
- Each static file for each role contains the correct network definitions.

Each static file requires all the parameter definitions for any custom networks even if the network is not used on the role. Check to make sure the rendered templates contain these parameters. For example, if a **StorageBackup** network is added to only the Ceph nodes, the **parameters** section in NIC

configuration templates for all roles must also include this definition:

```
parameters:
  ...
  StorageBackupIpSubnet:
    default: "
    description: IP address/subnet on the external network
    type: string
  ...
```

You can also include the **parameters** definitions for VLAN IDs and/or gateway IP, if needed:

```
parameters:
  ...
  StorageBackupNetworkVlanID:
    default: 60
    description: Vlan ID for the management network traffic.
    type: number
  StorageBackupDefaultRoute:
    description: The default route of the storage backup network.
    type: string
  ...
```

The **IpSubnet** parameter for the custom network appears in the parameter definitions for each role. However, since the Ceph role might be the only role that uses the **StorageBackup** network, only the NIC configuration template for the Ceph role would make use of the **StorageBackup** parameters in the **network_config** section of the template.

```
$network_config:
  network_config:
  - type: interface
    name: nic1
    use_dhcp: false
    addresses:
  - ip_netmask:
    get_param: StorageBackupIpSubnet
```

11.7. CUSTOM NETWORK ENVIRONMENT FILE

The custom network environment file (in this case, **/home/stack/templates/custom-network-configuration.yaml**) is a heat environment file that describes the overcloud network environment and points to the custom network interface configuration templates. You can define the subnets and VLANs for your network along with IP address ranges. You can then customize these values for the local environment.

The **resource_registry** section contains references to the custom network interface templates for each node role. Each resource registered uses the following format:

- **OS::TripleO::[ROLE]::Net::SoftwareConfig: [FILE]**

[ROLE] is the role name and **[FILE]** is the respective network interface template for that particular role. For example:

```
resource_registry:
```

```
OS::TripleO::Controller::Net::SoftwareConfig: /home/stack/templates/custom-nics/controller.yaml
```

The **parameter_defaults** section contains a list of parameters that define the network options for each network type.

11.8. NETWORK ENVIRONMENT PARAMETERS

The following table is a list of parameters that you can use in the **parameter_defaults** section of a network environment file to override the default parameter values in your NIC templates.

Parameter	Description	Type
ControlPlaneDefaultRoute	The IP address of the router on the Control Plane, which is used as a default route for roles other than the Controller nodes by default. Set to the undercloud IP if using IP masquerade instead of a router.	string
ControlPlaneSubnetCidr	The CIDR netmask of the IP network used on the Control Plane. If the Control Plane network uses 192.168.24.0/24, the CIDR is 24 .	string (though is always a number)
*NetCidr	The full network and CIDR netmask for a particular network. The default is automatically set to the network's ip_subnet setting in the network_data file. For example: InternalApiNetCidr: 172.16.0.0/24	string
*AllocationPools	"The IP allocation range for a particular network. The default is automatically set to the network's allocation_pools setting in the network_data file. For example: InternalApiAllocationPools: [{'start': '172.16.0.10', 'end': '172.16.0.200'}]	hash
*NetworkVlanID	The node's VLAN ID for on a particular network. The default is set automatically to the network's vlan setting in the network_data file. For example: InternalApiNetworkVlanID: 201	number

Parameter	Description	Type
*InterfaceDefaultRoute	The router address for a particular network, which you can use as a default route for roles or used for routes to other networks. The default is automatically set to the network's gateway_ip setting in the network_data file. For example: InternalApiInterfaceDefaultRoute: 172.16.0.1	string
DnsServers	A list of DNS servers added to resolv.conf. Usually allows a maximum of 2 servers.	comma delimited list
EC2MetadataIp	The IP address of the metadata server used to provision overcloud nodes. Set to the IP address of the undercloud on the Control Plane.	string
BondInterfaceOvsOptions	The options for bonding interfaces. For example: BondInterfaceOvsOptions: "bond_mode=balance-slb"	string
NeutronExternalNetworkBridge	Legacy value for the name of the external bridge to use for OpenStack Networking (neutron). This value is empty by default, which allows for multiple physical bridges to be defined in the NeutronBridgeMappings . This should not normally be overridden.	string
NeutronFlatNetworks	Defines the flat networks to configure in neutron plugins. Defaults to "datacentre" to permit external network creation. For example: NeutronFlatNetworks: "datacentre"	string

Parameter	Description	Type
NeutronBridgeMappings	The logical to physical bridge mappings to use. Defaults to mapping the external bridge on hosts (br-ex) to a physical name (datacentre). You would refer to the logical name when creating OpenStack Networking (neutron) provider networks or floating IP networks. For example NeutronBridgeMappings: "datacentre:br-ex,tenant:br-tenant"	string
NeutronPublicInterface	Defines the interface to bridge onto br-ex for network nodes when not using network isolation. Usually not used except in small deployments with only two networks. For example: NeutronPublicInterface: "eth0"	string
NeutronNetworkType	The tenant network type for OpenStack Networking (neutron). To specify multiple values, use a comma separated list. The first type specified is used until all available networks are exhausted, then the next type is used. For example: NeutronNetworkType: "vxlan"	string
NeutronTunnelTypes	The tunnel types for the neutron tenant network. To specify multiple values, use a comma separated string. For example: <i>NeutronTunnelTypes: 'gre,vxlan'</i>	string / comma separated list
NeutronTunnelIdRanges	Ranges of GRE tunnel IDs to make available for tenant network allocation. For example: NeutronTunnelIdRanges "1:1000"	string

Parameter	Description	Type
NeutronVniRanges	Ranges of VXLAN VNI IDs to make available for tenant network allocation. For example: NeutronVniRanges: "1:1000"	string
NeutronEnableTunnelling	Defines whether to enable or completely disable all tunnelled networks. Leave this enabled unless you are sure you will never want to create tunelled networks. Defaults to enabled.	Boolean
NeutronNetworkVLANRanges	The ML2 and Open vSwitch VLAN mapping range to support. Defaults to permitting any VLAN on the datacentre physical network. To specify multiple values, use a comma separated list. For example: NeutronNetworkVLANRanges: "datacentre:1:1000,tenant:100:299,tenant:310:399"	string
NeutronMechanismDrivers	The mechanism drivers for the neutron tenant network. Defaults to "openvswitch". To specify multiple values, use a comma-separated string. For example: NeutronMechanismDrivers: 'openvswitch,l2population'	string / comma separated list

11.9. EXAMPLE CUSTOM NETWORK ENVIRONMENT FILE

The following snippet is an example of an environment file that you can use to enable your NIC templates and set custom parameters.

```
resource_registry:
  OS::TripleO::BlockStorage::Net::SoftwareConfig:
    /home/stack/templates/nic-configs/cinder-storage.yaml
  OS::TripleO::Compute::Net::SoftwareConfig:
    /home/stack/templates/nic-configs/compute.yaml
  OS::TripleO::Controller::Net::SoftwareConfig:
    /home/stack/templates/nic-configs/controller.yaml
  OS::TripleO::ObjectStorage::Net::SoftwareConfig:
    /home/stack/templates/nic-configs/swift-storage.yaml
  OS::TripleO::CephStorage::Net::SoftwareConfig:
    /home/stack/templates/nic-configs/ceph-storage.yaml
```

```
parameter_defaults:
# Gateway router for the provisioning network (or Undercloud IP)
ControlPlaneDefaultRoute: 192.0.2.254
# The IP address of the EC2 metadata server. Generally the IP of the Undercloud
EC2MetadataIp: 192.0.2.1
# Define the DNS servers (maximum 2) for the overcloud nodes
DnsServers: ["8.8.8.8", "8.8.4.4"]
NeutronExternalNetworkBridge: ""
```

11.10. ENABLING NETWORK ISOLATION WITH CUSTOM NICs

To deploy the overcloud with network isolation and custom NIC templates, include all of the relevant networking environment files in the overcloud deployment command.

Procedure

1. When running the **openstack overcloud deploy** command, make sure to include:
 - The custom **network_data** file.
 - The rendered file name of the default network isolation.
 - The rendered file name of the default network environment file.
 - The custom environment network configuration that includes resource references to your custom NIC templates.
 - Any additional environment files relevant to your configuration.

For example:

```
$ openstack overcloud deploy --templates \
...
-n /home/stack/network_data.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/network-environment.yaml \
-e /home/stack/templates/custom-network-configuration.yaml \
...
```

- Include the **network-isolation.yaml** file first, then the **network-environment.yaml** file. The subsequent **custom-network-configuration.yaml** overrides the **OS::TripleO::[ROLE]::Net::SoftwareConfig** resources from the previous two files.
- If using composable networks, include the **network_data** and **roles_data** files with this command.

CHAPTER 12. ADDITIONAL NETWORK CONFIGURATION

This chapter follows on from the concepts and procedures outlined in [Chapter 11, Custom network interface templates](#) and provides some additional information to help configure parts of your overcloud network.

12.1. CONFIGURING CUSTOM INTERFACES

Individual interfaces might require modification. The following example shows the modifications that are necessary to use a second NIC to connect to an infrastructure network with DHCP addresses, and to use a third and fourth NIC for the bond:

```
network_config:
  # Add a DHCP infrastructure network to nic2
  - type: interface
    name: nic2
    use_dhcp: true
  - type: ovs_bridge
    name: br-bond
    members:
      - type: ovs_bond
        name: bond1
        ovs_options:
          get_param: BondInterfaceOvsOptions
        members:
          # Modify bond NICs to use nic3 and nic4
          - type: interface
            name: nic3
            primary: true
          - type: interface
            name: nic4
```

The network interface template uses either the actual interface name (**eth0**, **eth1**, **enp0s25**) or a set of numbered interfaces (**nic1**, **nic2**, **nic3**). The network interfaces of hosts within a role do not have to be exactly the same when using numbered interfaces (**nic1**, **nic2**, etc.) instead of named interfaces (**eth0**, **eno2**, etc.). For example, one host might have interfaces **em1** and **em2**, while another has **eno1** and **eno2**, but you can refer to both hosts' NICs as **nic1** and **nic2**.

The order of numbered interfaces corresponds to the order of named network interface types:

- **ethX** interfaces, such as **eth0**, **eth1**, etc. These are usually onboard interfaces.
- **enoX** interfaces, such as **eno0**, **eno1**, etc. These are usually onboard interfaces.
- **enX** interfaces, sorted alpha numerically, such as **enp3s0**, **enp3s1**, **ens3**, etc. These are usually add-on interfaces.


The numbered NIC scheme only takes into account the interfaces that are live, for example, if they have a cable attached to the switch. If you have some hosts with four interfaces and some with six interfaces, you should use **nic1** to **nic4** and only plug four cables on each host.

You can hardcode physical interfaces to specific aliases. This allows you to be pre-determine which physical NIC will be mapped as **nic1** or **nic2** and so on. You can also map a MAC address to a specified alias.

**NOTE**

Normally, **os-net-config** will only register interfaces that are already connected in an **UP** state. However, if you do hardcode interfaces using a custom mapping file, then the interface is registered even if it is in a **DOWN** state.

Interfaces are mapped to aliases with an environment file. In this example, each node has predefined entries for **nic1** and **nic2**.

**NOTE**

If you want to use the **NetConfigDataLookup** configuration, you must also include the **os-net-config-mappings.yaml** file in the **NodeUserData** resource registry.

```
resource_registry:
  OS::TripleO::NodeUserData: /usr/share/openstack/tripleo-heat-templates/firstboot/os-net-config-
  mappings.yaml
parameter_defaults:
  NetConfigDataLookup:
    node1:
      nic1: "em1"
      nic2: "em2"
    node2:
      nic1: "00:50:56:2F:9F:2E"
      nic2: "em2"
```

The resulting configuration is then applied by **os-net-config**. On each node, you can see the applied configuration in the **interface_mapping** section of the **/etc/os-net-config/mapping.yaml** file.

12.2. CONFIGURING ROUTES AND DEFAULT ROUTES

You can set the default route of a host in one of two ways. If the interface uses DHCP and the DHCP server offers a gateway address, the system uses a default route for that gateway. Otherwise, you can set a default route on an interface with a static IP.

Although the Linux kernel supports multiple default gateways, it only uses the one with the lowest metric. If there are multiple DHCP interfaces, this can result in an unpredictable default gateway. In this case, it is recommended to set **defroute: false** for interfaces other than the one using the default route.

For example, you might want a DHCP interface (**nic3**) to be the default route. Use the following YAML to disable the default route on another DHCP interface (**nic2**):

```
# No default route on this DHCP interface
- type: interface
  name: nic2
  use_dhcp: true
  defroute: false
# Instead use this DHCP interface as the default route
- type: interface
  name: nic3
  use_dhcp: true
```

**NOTE**

The **defroute** parameter only applies to routes obtained through DHCP.

To set a static route on an interface with a static IP, specify a route to the subnet. For example, you can set a route to the 10.1.2.0/24 subnet through the gateway at 172.17.0.1 on the Internal API network:

```
- type: vlan
  device: bond1
  vlan_id:
    get_param: InternalApiNetworkVlanID
  addresses:
    - ip_netmask:
        get_param: InternalApiIpSubnet
  routes:
    - ip_netmask: 10.1.2.0/24
      next_hop: 172.17.0.1
```

12.3. CONFIGURING JUMBO FRAMES

The Maximum Transmission Unit (MTU) setting determines the maximum amount of data transmitted with a single Ethernet frame. Using a larger value results in less overhead because each frame adds data in the form of a header. The default value is 1500 and using a higher value requires the configuration of the switch port to support jumbo frames. Most switches support an MTU of at least 9000, but many are configured for 1500 by default.

The MTU of a VLAN cannot exceed the MTU of the physical interface. Make sure to include the MTU value on the bond and/or interface.

The Storage, Storage Management, Internal API, and Tenant networks all benefit from jumbo frames. In testing, a project's networking throughput demonstrated substantial improvement when using jumbo frames in conjunction with VXLAN tunnels.

**NOTE**

It is recommended that the Provisioning interface, External interface, and any floating IP interfaces be left at the default MTU of 1500. Connectivity problems are likely to occur otherwise. This is because routers typically cannot forward jumbo frames across Layer 3 boundaries.

```
- type: ovs_bond
  name: bond1
  mtu: 9000
  ovs_options: {get_param: BondInterfaceOvsOptions}
  members:
    - type: interface
      name: nic3
      mtu: 9000
      primary: true
    - type: interface
      name: nic4
      mtu: 9000
```

The external interface should stay at default

```

- type: vlan
  device: bond1
  vlan_id:
    get_param: ExternalNetworkVlanID
  addresses:
    - ip_netmask:
        get_param: ExternalIpSubnet
  routes:
    - ip_netmask: 0.0.0.0/0
      next_hop:
        get_param: ExternalInterfaceDefaultRoute

# MTU 9000 for Internal API, Storage, and Storage Management
- type: vlan
  device: bond1
  mtu: 9000
  vlan_id:
    get_param: InternalApiNetworkVlanID
  addresses:
    - ip_netmask:
        get_param: InternalApiIpSubnet

```

12.4. CONFIGURING THE NATIVE VLAN ON A TRUNKED INTERFACE

If a trunked interface or bond has a network on the native VLAN, the IP addresses are assigned directly to the bridge and there is no VLAN interface.

For example, if the External network is on the native VLAN, a bonded configuration looks like this:

```

network_config:
- type: ovs_bridge
  name: bridge_name
  dns_servers:
    get_param: DnsServers
  addresses:
    - ip_netmask:
        get_param: ExternalIpSubnet
  routes:
    - ip_netmask: 0.0.0.0/0
      next_hop:
        get_param: ExternalInterfaceDefaultRoute
  members:
    - type: ovs_bond
      name: bond1
      ovs_options:
        get_param: BondInterfaceOvsOptions
      members:
        - type: interface
          name: nic3
          primary: true
        - type: interface
          name: nic4

```



NOTE

When moving the address (and possibly route) statements onto the bridge, remove the corresponding VLAN interface from the bridge. Make the changes to all applicable roles. The External network is only on the controllers, so only the controller template requires a change. The Storage network on the other hand is attached to all roles, so if the Storage network is on the default VLAN, all roles require modifications.

12.5. INCREASING THE MAXIMUM NUMBER OF CONNECTIONS THAT NETFILTER TRACKS

The Red Hat OpenStack Platform (RHOSP) Networking service (neutron) uses netfilter connection tracking to build stateful firewalls and to provide network address translation (NAT) on virtual networks. There are some situations that can cause the kernel space to reach the maximum connection limit and result in errors such as **nf_contrack: table full, dropping packet**. You can increase the limit for connection tracking (contrack) and avoid these types of errors. You can increase the contrack limit for one or more roles, or across all the nodes, in your RHOSP deployment.

Prerequisites

- A successful RHOSP undercloud installation.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the undercloud credentials file:

```
$ source ~/stackrc
```

3. Create a custom YAML environment file.

Example

```
$ vi /home/stack/templates/my-environment.yaml
```

4. Your environment file must contain the keywords **parameter_defaults** and **ExtraSysctlSettings**. Enter a new value for the maximum number of connections that netfilter can track in the variable, **net.nf_contrack_max**.

Example

In this example, you can set the contrack limit across all hosts in your RHOSP deployment:

```
parameter_defaults:
  ExtraSysctlSettings:
    net.nf_contrack_max:
      value: 500000
```

Use the **<role>Parameter** parameter to set the contrack limit for a specific role:

```
parameter_defaults:
  <role>Parameters:
    ExtraSysctlSettings:
```

```
net.nf_conntrack_max:
  value: <simultaneous_connections>
```

- Replace **<role>** with the name of the role.
For example, use **ControllerParameters** to set the conntrack limit for the Controller role, or **ComputeParameters** to set the conntrack limit for the Compute role.
- Replace **<simultaneous_connections>** with the quantity of simultaneous connections that you want to allow.

Example

In this example, you can set the conntrack limit for only the Controller role in your RHOSP deployment:

```
parameter_defaults:
  ControllerParameters:
  ExtraSysctlSettings:
    net.nf_conntrack_max:
      value: 500000
```



NOTE

The default value for **net.nf_conntrack_max** is **500000** connections. The maximum value is: **4294967295**.

5. Run the deployment command and include the core heat templates, environment files, and this new custom environment file.



IMPORTANT

The order of the environment files is important as the parameters and resources defined in subsequent environment files take precedence.

Example

```
$ openstack overcloud deploy --templates \
-e /home/stack/templates/my-environment.yaml
```

Additional resources

- [Environment files](#)
- [Including Environment Files in Overcloud Creation](#)

CHAPTER 13. NETWORK INTERFACE BONDING

This chapter defines some of the bonding options you can use in your custom network configuration.

13.1. NETWORK INTERFACE BONDING AND LINK AGGREGATION CONTROL PROTOCOL (LACP)

You can bundle multiple physical NICs together to form a single logical channel known as a bond. You can configure bonds to provide redundancy for high availability systems or increased throughput.



IMPORTANT

Using single root I/O virtualization (SR-IOV) on bonded interfaces is not supported.

Red Hat OpenStack Platform supports Linux bonds, Open vSwitch (OVS) kernel bonds, and OVS-DPDK bonds.

The bonds can be used with the optional Link Aggregation Control Protocol (LACP). LACP is a negotiation protocol that creates a dynamic bond for load balancing and fault tolerance.

Red Hat recommends the use of Linux kernel bonds (bond type: `linux_bond`) over OVS kernel bonds (bond type: `ovs_bond`). User mode bonds (bond type: `ovs_dpdk_bond`) must be used with user mode bridges (type: `ovs_user_bridge`) as opposed to kernel mode bridges (type: `ovs_bridge`). However, do not combine `ovs_bridge` and `ovs_user_bridge` on the same node.

On control and storage networks, Red Hat recommends the use of Linux bonds with VLAN and LACP, because OVS bonds carry the potential for control plane disruption that can occur when OVS or the neutron agent is restarted for updates, hot fixes, and other events. The Linux bond/LACP/VLAN configuration provides NIC management without the OVS disruption potential.

Here is an example configuration of a Linux bond with one VLAN.

```
params:
  $network_config:
    network_config:

    - type: linux_bond
      name: bond_api
      bonding_options: "mode=active-backup"
      use_dhcp: false
      dns_servers:
        get_param: DnsServers
      members:
        - type: interface
          name: nic3
          primary: true
        - type: interface
          name: nic4

    - type: vlan
      vlan_id:
        get_param: InternalApiNetworkVlanID
      device: bond_api
```

```
addresses:
- ip_netmask:
  get_param: InternalApiIpSubnet
```

The following example shows a Linux bond plugged into the OVS bridge

```
params:
  $network_config:
  network_config:

- type: ovs_bridge
  name: br-tenant
  use_dhcp: false
  mtu: 9000
  members:
  - type: linux_bond
    name: bond_tenant
    bonding_options: "mode=802.3ad updelay=1000 miimon=100"
    use_dhcp: false
    dns_servers:
      get_param: DnsServers
    members:
    - type: interface
      name: p1p1
      primary: true
    - type: interface
      name: p1p2
  - type: vlan
    device: bond_tenant
    vlan_id: {get_param: TenantNetworkVlanID}
    addresses:
      -
      ip_netmask: {get_param: TenantIpSubnet}
```

The following example shows an OVS user space bridge:

```
params:
  $network_config:
  network_config:

- type: ovs_user_bridge
  name: br-ex
  use_dhcp: false
  members:
  - type: ovs_dpdk_bond
    name: dpdkbond0
    mtu: 2140
    ovs_options: {get_param: BondInterfaceOvsOptions}
    #ovs_extra:
    #- set interface dpdk0 mtu_request=$MTU
    #- set interface dpdk1 mtu_request=$MTU
    rx_queue:
      get_param: NumDpdkInterfaceRxQueues
    members:
    - type: ovs_dpdk_port
```

```

name: dpdk0
mtu: 2140
members:
- type: interface
  name: p1p1
- type: ovs_dpdk_port
  name: dpdk1
  mtu: 2140
  members:
  - type: interface
    name: p1p2

```

13.2. OPEN VSWITCH BONDING OPTIONS

The overcloud provides networking through Open vSwitch (OVS). Use the following table to understand support compatibility for OVS kernel and OVS-DPDK for bonded interfaces. The OVS/OVS-DPDK balance-tcp mode is available as a technology preview only.



NOTE

This support requires Open vSwitch 2.9 or later.

OVS Bond mode	Application	Notes	Compatible LACP options
active-backup	High availability (active-passive)		active, passive, or off
balance-slb	Increased throughput (active-active)	<ul style="list-style-type: none"> ● Performance is affected by extra parsing per packet. ● There is a potential for vhost-user lock contention. 	active, passive, or off

balance-tcp (tech preview only)	Not recommended (active-active)	<ul style="list-style-type: none"> ● Recirculation needed for L4 hashing has a performance impact. ● As with balance-slb, performance is affected by extra parsing per packet and there is a potential for vhost-user lock contention. ● LACP must be enabled. 	active or passive
----------------------------------	---------------------------------	---	-------------------

You can configure a bonded interface in the network environment file using the `BondInterfaceOvsOptions` parameter as shown in this example:

```
parameter_defaults:
  BondInterfaceOvsOptions: "bond_mode=balance-slb"
```

13.3. LINUX BONDING OPTIONS

You can use LACP with Linux bonding in your network interface templates:

```
- type: linux_bond
  name: bond1
  members:
  - type: interface
    name: nic2
  - type: interface
    name: nic3
  bonding_options: "mode=802.3ad lacp_rate=[fast|slow] updelay=1000 miimon=100"
```

- **mode** - enables LACP.
- **lacp_rate** - defines whether LACP packets are sent every 1 second, or every 30 seconds.
- **updelay** - defines the minimum amount of time that an interface must be active before it is used for traffic (this helps mitigate port flapping outages).
- **miimon** - the interval in milliseconds that is used for monitoring the port state using the driver's MIIMON functionality.

For more information on Linux bonding options, see [7.7.1. Bonding Module Directives](#) in the *Red Hat Enterprise Linux 7 Networking Guide*.

13.4. OVS BONDING OPTIONS

The following table provides some explanation of these options and some alternatives depending on your hardware.

Table 13.1. Bonding Options

bond_mode=balance-slb	Balances flows based on source MAC address and output VLAN, with periodic rebalancing as traffic patterns change. Bonding with balance-slb allows a limited form of load balancing without the remote switch's knowledge or cooperation. SLB assigns each source MAC and VLAN pair to a link and transmits all packets from that MAC and VLAN through that link. This mode uses a simple hashing algorithm based on source MAC address and VLAN number, with periodic rebalancing as traffic patterns change. This mode is similar to mode 2 bonds used by the Linux bonding driver. This mode can be used to provide load balancing even when the switch is not configured to use LACP.
bond_mode=active-backup	This mode offers active/standby failover where the standby NIC resumes network operations when the active connection fails. Only one MAC address is presented to the physical switch. This mode does not require any special switch support or configuration, and works when the links are connected to separate switches. This mode does not provide load balancing.
lACP=[active passive off]	Controls the Link Aggregation Control Protocol (LACP) behavior. Only certain switches support LACP. If your switch does not support LACP, use bond_mode=balance-slb or bond_mode=active-backup .
other-config:lACP-fallback-ab=true	Sets the LACP behavior to switch to <code>bond_mode=active-backup</code> as a fallback.
other_config:lACP-time=[fast slow]	Set the LACP heartbeat to 1 second (fast) or 30 seconds (slow). The default is slow.
other_config:bond-detect-mode=[miimon carrier]	Set the link detection to use miimon heartbeats (miimon) or monitor carrier (carrier). The default is carrier.
other_config:bond-miimon-interval=100	If using miimon, set the heartbeat interval in milliseconds.
bond_updelay=1000	Number of milliseconds a link must be up to be activated to prevent flapping.
other_config:bond-rebalance-interval=10000	Milliseconds between rebalancing flows between bond members. Set this value to zero to disable rebalancing flows between bond members.

CHAPTER 14. CONTROLLING NODE PLACEMENT

The default behavior for the director is to randomly select nodes for each role, usually based on their profile tag. However, the director provides the ability to define specific node placement. This is a useful method to:

- Assign specific node IDs e.g. **controller-0**, **controller-1**, etc
- Assign custom hostnames
- Assign specific IP addresses
- Assign specific Virtual IP addresses



NOTE

Manually setting predictable IP addresses, virtual IP addresses, and ports for a network alleviates the need for allocation pools. However, it is recommended to retain allocation pools for each network to ease with scaling new nodes. Make sure that any statically defined IP addresses fall outside the allocation pools. For more information on setting allocation pools, see [Section 11.7, “Custom network environment file”](#).

14.1. ASSIGNING SPECIFIC NODE IDS

This procedure assigns node ID to specific nodes. Examples of node IDs include **controller-0**, **controller-1**, **compute-0**, **compute-1**, and so forth.

The first step is to assign the ID as a per-node capability that the Compute scheduler matches on deployment. For example:

```
openstack baremetal node set --property capabilities='node:controller-0,boot_option:local' <id>
```

This assigns the capability **node:controller-0** to the node. Repeat this pattern using a unique continuous index, starting from 0, for all nodes. Make sure all nodes for a given role (Controller, Compute, or each of the storage roles) are tagged in the same way or else the Compute scheduler will not match the capabilities correctly.

The next step is to create a Heat environment file (for example, **scheduler_hints_env.yaml**) that uses scheduler hints to match the capabilities for each node. For example:


```
parameter_defaults:
  ControllerSchedulerHints:
    'capabilities:node': 'controller-%index%'
```

To use these scheduler hints, include the `scheduler_hints_env.yaml` environment file with the **overcloud deploy command** during Overcloud creation.

The same approach is possible for each role via these parameters:

- **ControllerSchedulerHints** for Controller nodes.
- **ComputeSchedulerHints** for Compute nodes.
- **BlockStorageSchedulerHints** for Block Storage nodes.

- **ObjectStorageSchedulerHints** for Object Storage nodes.
- **CephStorageSchedulerHints** for Ceph Storage nodes.
- **[ROLE]SchedulerHints** for custom roles. Replace **[ROLE]** with the role name.



NOTE

Node placement takes priority over profile matching. To avoid scheduling failures, use the default **baremetal** flavor for deployment and not the flavors designed for profile matching (**compute**, **control**, etc). For example:

```
$ openstack overcloud deploy ... --control-flavor baremetal --compute-flavor baremetal
...
```

14.2. ASSIGNING CUSTOM HOSTNAMES

In combination with the node ID configuration in [Section 14.1, "Assigning Specific Node IDs"](#), the director can also assign a specific custom hostname to each node. This is useful when you need to define where a system is located (e.g. **rack2-row12**), match an inventory identifier, or other situations where a custom hostname is desired.



IMPORTANT

Do not rename a node after it has been deployed. Renaming a node after deployment creates issues with instance management.

To customize node hostnames, use the **HostnameMap** parameter in an environment file, such as the `scheduler_hints_env.yaml`` file from [Section 14.1, "Assigning Specific Node IDs"](#). For example:

```
parameter_defaults:
  ControllerSchedulerHints:
    'capabilities:node': 'controller-%index%'
  ComputeSchedulerHints:
    'capabilities:node': 'compute-%index%'
  HostnameMap:
    overcloud-controller-0: overcloud-controller-prod-123-0
    overcloud-controller-1: overcloud-controller-prod-456-0
    overcloud-controller-2: overcloud-controller-prod-789-0
    overcloud-compute-0: overcloud-compute-prod-abc-0
```

Define the **HostnameMap** in the **parameter_defaults** section, and set each mapping as the original hostname that Heat defines using **HostnameFormat** parameters (e.g. **overcloud-controller-0**) and the second value is the desired custom hostname for that node (e.g. **overcloud-controller-prod-123-0**).

Using this method in combination with the node ID placement ensures each node has a custom hostname.

14.3. ASSIGNING PREDICTABLE IPS

For further control over the resulting environment, the director can assign Overcloud nodes with specific IPs on each network as well. Use the **environments/ips-from-pool-all.yaml** environment file in the core Heat template collection.

Copy this file to the **stack** user's **templates** directory.

```
$ cp /usr/share/openstack-tripleo-heat-templates/environments/ips-from-pool-all.yaml ~/templates/.
```

There are two major sections in the **ips-from-pool-all.yaml** file.

The first is a set of **resource_registry** references that override the defaults. These tell the director to use a specific IP for a given port on a node type. Modify each resource to use the absolute path of its respective template. For example:

```
OS::TripleO::Controller::Ports::ExternalPort: /usr/share/openstack-tripleo-heat-templates/network/ports/external_from_pool.yaml
OS::TripleO::Controller::Ports::InternalApiPort: /usr/share/openstack-tripleo-heat-templates/network/ports/internal_api_from_pool.yaml
OS::TripleO::Controller::Ports::StoragePort: /usr/share/openstack-tripleo-heat-templates/network/ports/storage_from_pool.yaml
OS::TripleO::Controller::Ports::StorageMgmtPort: /usr/share/openstack-tripleo-heat-templates/network/ports/storage_mgmt_from_pool.yaml
OS::TripleO::Controller::Ports::TenantPort: /usr/share/openstack-tripleo-heat-templates/network/ports/tenant_from_pool.yaml
```

The default configuration sets all networks on all node types to use pre-assigned IPs. To allow a particular network or node type to use default IP assignment instead, simply remove the **resource_registry** entries related to that node type or network from the environment file.

The second section is `parameter_defaults`, where the actual IP addresses are assigned. Each node type has an associated parameter:

- **ControllerIPs** for Controller nodes.
- **ComputeIPs** for Compute nodes.
- **CephStorageIPs** for Ceph Storage nodes.
- **BlockStorageIPs** for Block Storage nodes.
- **SwiftStorageIPs** for Object Storage nodes.
- **[ROLE]IPs** for custom roles. Replace **[ROLE]** with the role name.

Each parameter is a map of network names to a list of addresses. Each network type must have at least as many addresses as there will be nodes on that network. The director assigns addresses in order. The first node of each type receives the first address on each respective list, the second node receives the second address on each respective lists, and so forth.

For example, if an Overcloud will contain three Ceph Storage nodes, the `CephStorageIPs` parameter might look like:

```
CephStorageIPs:
  storage:
    - 172.16.1.100
    - 172.16.1.101
    - 172.16.1.102
  storage_mgmt:
```

- 172.16.3.100
- 172.16.3.101
- 172.16.3.102

The first Ceph Storage node receives two addresses: 172.16.1.100 and 172.16.3.100. The second receives 172.16.1.101 and 172.16.3.101, and the third receives 172.16.1.102 and 172.16.3.102. The same pattern applies to the other node types.

Make sure the chosen IP addresses fall outside the allocation pools for each network defined in your network environment file (see [Section 11.7, “Custom network environment file”](#)). For example, make sure the **internal_api** assignments fall outside of the **InternalApiAllocationPools** range. This avoids conflicts with any IPs chosen automatically. Likewise, make sure the IP assignments do not conflict with the VIP configuration, either for standard predictable VIP placement (see [Section 14.4, “Assigning Predictable Virtual IPs”](#)) or external load balancing (see [Section 23.2, “Configuring External Load Balancing”](#)).



IMPORTANT

If an overcloud node is deleted, do not remove its entries in the IP lists. The IP list is based on the underlying Heat indices, which do not change even if you delete nodes. To indicate a given entry in the list is no longer used, replace the IP value with a value such as **DELETED** or **UNUSED**. Entries should never be removed from the IP lists, only changed or added.

To apply this configuration during a deployment, include the **ips-from-pool-all.yaml** environment file with the **openstack overcloud deploy** command.



IMPORTANT

If using network isolation, include the **ips-from-pool-all.yaml** file after the **network-isolation.yaml** file.

For example:

```
$ openstack overcloud deploy --templates \
-e /usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml \
-e ~/templates/ips-from-pool-all.yaml \
[OTHER OPTIONS]
```

14.4. ASSIGNING PREDICTABLE VIRTUAL IPS

In addition to defining predictable IP addresses for each node, the director also provides a similar ability to define predictable Virtual IPs (VIPs) for clustered services. To accomplish this, edit the network environment file from [Section 11.7, “Custom network environment file”](#) and add the VIP parameters in the **parameter_defaults** section:

```
parameter_defaults:
...
# Predictable VIPs
ControlFixedIPs: [{'ip_address':'192.168.201.101'}]
InternalApiVirtualFixedIPs: [{'ip_address':'172.16.0.9'}]
PublicVirtualFixedIPs: [{'ip_address':'10.1.1.9'}]
```

```
StorageVirtualFixedIPs: [{'ip_address':'172.18.0.9'}]  
StorageMgmtVirtualFixedIPs: [{'ip_address':'172.19.0.9'}]  
RedisVirtualFixedIPs: [{'ip_address':'172.16.0.8'}]
```

Select these IPs from outside of their respective allocation pool ranges. For example, select an IP address for **InternalApiVirtualFixedIPs** that is not within the **InternalApiAllocationPools** range.

This step is only for overclouds using the default internal load balancing configuration. If assigning VIPs with an external load balancer, use the procedure in the dedicated [External Load Balancing for the Overcloud](#) guide.

CHAPTER 15. ENABLING SSL/TLS ON OVERCLOUD PUBLIC ENDPOINTS

By default, the overcloud uses unencrypted endpoints for its services. This means that the overcloud configuration requires an additional environment file to enable SSL/TLS for its Public API endpoints. The following chapter shows how to configure your SSL/TLS certificate and include it as a part of your overcloud creation.



NOTE

This process only enables SSL/TLS for Public API endpoints. The Internal and Admin APIs remain unencrypted.

This process requires network isolation to define the endpoints for the Public API.

15.1. INITIALIZING THE SIGNING HOST

The signing host is the host that generates new certificates and signs them with a certificate authority. If you have never created SSL certificates on the chosen signing host, you might need to initialize the host so that it can sign new certificates.

The `/etc/pki/CA/index.txt` file stores records of all signed certificates. Check if this file exists. If it does not exist, create an empty file:

```
$ sudo touch /etc/pki/CA/index.txt
```

The `/etc/pki/CA/serial` file identifies the next serial number to use for the next certificate to sign. Check if this file exists. If it does not exist, create a new file with a new starting value:

```
$ echo '1000' | sudo tee /etc/pki/CA/serial
```

15.2. CREATING A CERTIFICATE AUTHORITY

Normally you sign your SSL/TLS certificates with an external certificate authority. In some situations, you might aim to use your own certificate authority. For example, you might aim to have an internal-only certificate authority.

For example, generate a key and certificate pair to act as the certificate authority:

```
$ sudo openssl genrsa -out ca.key.pem 4096  
$ sudo openssl req -key ca.key.pem -new -x509 -days 7300 -extensions v3_ca -out ca.crt.pem
```

The `openssl req` command asks for certain details about your authority. Enter these details.

This creates a certificate authority file called `ca.crt.pem`.

15.3. ADDING THE CERTIFICATE AUTHORITY TO CLIENTS

For any external clients aiming to communicate using SSL/TLS, copy the certificate authority file to each client that requires access your Red Hat OpenStack Platform environment. Once copied to the client, run the following command on the client to add it to the certificate authority trust bundle:

```
$ sudo cp ca.crt.pem /etc/pki/ca-trust/source/anchors/
$ sudo update-ca-trust extract
```

For example, the undercloud requires a copy of the certificate authority file so that it can communicate with the overcloud endpoints during creation.

15.4. CREATING AN SSL/TLS KEY

Run the following commands to generate the SSL/TLS key (**server.key.pem**), which we use at different points to generate our undercloud or overcloud certificates:

```
$ openssl genrsa -out server.key.pem 2048
```

15.5. CREATING AN SSL/TLS CERTIFICATE SIGNING REQUEST

This next procedure creates a certificate signing request for the overcloud. Copy the default OpenSSL configuration file for customization.

```
$ cp /etc/pki/tls/openssl.cnf .
```

Edit the custom **openssl.cnf** file and set SSL parameters to use for the overcloud. An example of the types of parameters to modify include:

```
[req]
distinguished_name = req_distinguished_name
req_extensions = v3_req

[req_distinguished_name]
countryName = Country Name (2 letter code)
countryName_default = AU
stateOrProvinceName = State or Province Name (full name)
stateOrProvinceName_default = Queensland
localityName = Locality Name (eg, city)
localityName_default = Brisbane
organizationalUnitName = Organizational Unit Name (eg, section)
organizationalUnitName_default = Red Hat
commonName = Common Name
commonName_default = 10.0.0.1
commonName_max = 64

[ v3_req ]
# Extensions to add to a certificate request
basicConstraints = CA:FALSE
keyUsage = nonRepudiation, digitalSignature, keyEncipherment
subjectAltName = @alt_names

[alt_names]
IP.1 = 10.0.0.1
DNS.1 = 10.0.0.1
DNS.2 = myovercloud.example.com
```

Set the **commonName_default** to one of the following:

- If using an IP to access over SSL/TLS, use the Virtual IP for the Public API. Set this VIP using the **PublicVirtualFixedIPs** parameter in an environment file. For more information, see [Section 14.4, “Assigning Predictable Virtual IPs”](#). If you are not using predictable VIPs, the director assigns the first IP address from the range defined in the **ExternalAllocationPools** parameter.
- If using a fully qualified domain name to access over SSL/TLS, use the domain name instead.

Include the same Public API IP address as an IP entry and a DNS entry in the **alt_names** section. If also using DNS, include the hostname for the server as DNS entries in the same section. For more information about **openssl.cnf**, run **man openssl.cnf**.

Run the following command to generate certificate signing request (**server.csr.pem**):

```
$ openssl req -config openssl.cnf -key server.key.pem -new -out server.csr.pem
```

Make sure to include the SSL/TLS key you created in [Section 15.4, “Creating an SSL/TLS Key”](#) for the **-key** option.

Use the **server.csr.pem** file to create the SSL/TLS certificate in the next section.

15.6. CREATING THE SSL/TLS CERTIFICATE

The following command creates a certificate for your undercloud or overcloud:

```
$ sudo openssl ca -config openssl.cnf -extensions v3_req -days 3650 -in server.csr.pem -out server.crt.pem -cert ca.crt.pem -keyfile ca.key.pem
```

This command uses:

- The configuration file specifying the v3 extensions. Include this as the **-config** option.
- The certificate signing request from [Section 15.5, “Creating an SSL/TLS Certificate Signing Request”](#) to generate the certificate and sign it through a certificate authority. Include this as the **-in** option.
- The certificate authority you created in [Section 15.2, “Creating a Certificate Authority”](#), which signs the certificate. Include this as the **-cert** option.
- The certificate authority private key you created in [Section 15.2, “Creating a Certificate Authority”](#). Include this as the **-keyfile** option.

This results in a certificate named **server.crt.pem**. Use this certificate in conjunction with the SSL/TLS key from [Section 15.4, “Creating an SSL/TLS Key”](#) to enable SSL/TLS.

15.7. ENABLING SSL/TLS

Copy the **enable-tls.yaml** environment file from the Heat template collection:

```
$ cp -r /usr/share/openstack-tripleo-heat-templates/environments/ssl/enable-tls.yaml ~/templates/.
```

Edit this file and make the following changes for these parameters:

SSLCertificate

Copy the contents of the certificate file (**server.crt.pem**) into the **SSLCertificate** parameter. For example:

```
parameter_defaults:
  SSLCertificate: |
    -----BEGIN CERTIFICATE-----
    MIIDgzCCAmugAwIBAgIJAKk46qw6ncJaMA0GCSqGSIb3DQEBCwUAMFgxGzAJBgNV
    ...
    sFW3S2roS4X0Af/kSSD8mIBBTFTCMBAj6rtLBKLaQblxEplzrgvp
    -----END CERTIFICATE-----
```

SSLIntermediateCertificate

If you have an intermediate certificate, copy the contents of the intermediate certificate into the **SSLIntermediateCertificate** parameter:

```
parameter_defaults:
  SSLIntermediateCertificate: |
    -----BEGIN CERTIFICATE-----
    sFW3S2roS4X0Af/kSSD8mIBBTFTCMBAj6rtLBKLaQblxEplzrgvpBCwUAMFgxGzAJB
    ...
    MIIDgzCCAmugAwIBAgIJAKk46qw6ncJaMA0GCSqGSIb3DQE
    -----END CERTIFICATE-----
```



IMPORTANT

The certificate contents require the same indentation level for all new lines.

SSLKey

Copy the contents of the private key (**server.key.pem**) into the **SSLKey** parameter. For example:

```
parameter_defaults:
  ...
  SSLKey: |
    -----BEGIN RSA PRIVATE KEY-----
    MIIEowIBAAKCAQEAqVw8lnQ9Rbel1EdLN5PJP0IVO9hkJZnGP6qb6wtYUoy1bVP7
    ...
    ctIKn3rAAadyumi4JDjESAXHIKfjJNOLrBmpQyES4XpZUC7yhqPaU
    -----END RSA PRIVATE KEY-----
```



IMPORTANT

The private key contents require the same indentation level for all new lines.

OS::TripleO::NodeTLSData

Change the resource path for **OS::TripleO::NodeTLSData**: to an absolute path:

```
resource_registry:
  OS::TripleO::NodeTLSData: /usr/share/openstack-tripleo-heat-
  templates/puppet/extraconfig/tls/tls-cert-inject.yaml
```

15.8. INJECTING A ROOT CERTIFICATE

If the certificate signer is not in the default trust store on the overcloud image, you must inject the certificate authority into the overcloud image. Copy the **inject-trust-anchor.yaml** environment file from the heat template collection:

```
$ cp -r /usr/share/openstack-tripleo-heat-templates/environments/ssl/inject-trust-anchor.yaml
~/templates/.
```

Edit this file and make the following changes for these parameters:

SSLRootCertificate

Copy the contents of the root certificate authority file (**ca.crt.pem**) into the **SSLRootCertificate** parameter. For example:

```
parameter_defaults:
  SSLRootCertificate: |
    -----BEGIN CERTIFICATE-----
    MIIDgzCCAmugAwIBAgIJAKk46qw6ncJaMA0GCSqGSIb3DQEBCwUAMFgx CzAJBgNV
    ...
    sFW3S2roS4X0Af/kSSD8mIBBTFTCMBAj6rtLBKLaQbIx Eplzrgvp
    -----END CERTIFICATE-----
```



IMPORTANT

The certificate authority contents require the same indentation level for all new lines.

OS::TripleO::NodeTLSCAData

Change the resource path for **OS::TripleO::NodeTLSCAData**: to an absolute path:

```
resource_registry:
  OS::TripleO::NodeTLSCAData: /usr/share/openstack-tripleo-heat-
  templates/puppet/extraconfig/tls/ca-inject.yaml
```

If you want to inject multiple CAs, you can use the **inject-trust-anchor-hiera.yaml** environment file. For example, you can inject the CA for both the undercloud and overcloud simultaneously:

```
parameter_defaults:
  CAMap:
    undercloud-ca:
      content: |
        -----BEGIN CERTIFICATE-----
        ... cert content ...
        -----END CERTIFICATE-----
    overcloud-ca:
      content: |
        -----BEGIN CERTIFICATE-----
        ... cert content ...
        -----END CERTIFICATE-----
```

15.9. CONFIGURING DNS ENDPOINTS

If using a DNS hostname to access the overcloud through SSL/TLS, you will need to copy the **custom-domain.yaml** file into **/home/stack/templates**. You can find this file in **/usr/share/tripleo-heat-templates/environments/predictable-placement/**.

1. Configure the host and domain names for all fields, adding parameters for custom networks if needed:



NOTE

It is not possible to redeploy with a TLS-everywhere architecture if this environment file is not included in the initial deployment.

```
# title: Custom Domain Name
# description: |
# This environment contains the parameters that need to be set in order to
# use a custom domain name and have all of the various FQDNs reflect it.
parameter_defaults:
  # The DNS domain used for the hosts. This must match the overcloud_domain_name
  # configured on the undercloud.
  # Type: string
  CloudDomain: localdomain

  # The DNS name of this cloud. E.g. ci-overcloud.tripleo.org
  # Type: string
  CloudName: overcloud.localdomain

  # The DNS name of this cloud's provisioning network endpoint. E.g. 'ci-
  # overcloud.ctlplane.tripleo.org'.
  # Type: string
  CloudNameCtlplane: overcloud.ctlplane.localdomain

  # The DNS name of this cloud's internal_api endpoint. E.g. 'ci-
  # overcloud.internalapi.tripleo.org'.
  # Type: string
  CloudNameInternal: overcloud.internalapi.localdomain

  # The DNS name of this cloud's storage endpoint. E.g. 'ci-overcloud.storage.tripleo.org'.
  # Type: string
  CloudNameStorage: overcloud.storage.localdomain

  # The DNS name of this cloud's storage_mgmt endpoint. E.g. 'ci-
  # overcloud.storagemgmt.tripleo.org'.
  # Type: string
  CloudNameStorageManagement: overcloud.storagemgmt.localdomain
```

2. Add a list of DNS servers to use under parameter defaults, in either a new or existing environment file:

```
parameter_defaults:
  DnsServers: ["10.0.0.254"]
  ....
```

15.10. ADDING ENVIRONMENT FILES DURING OVERCLOUD CREATION

The deployment command (**openstack overcloud deploy**) uses the **-e** option to add environment files. Add the environment files from this section in the following order:

- The environment file to enable SSL/TLS (**enable-tls.yaml**)
- The environment file to set the DNS hostname (**cloudname.yaml**)
- The environment file to inject the root certificate authority (**inject-trust-anchor.yaml**)
- The environment file to set the public endpoint mapping:
 - If using a DNS name for accessing the public endpoints, use **/usr/share/openstack-tripleo-heat-templates/environments/ssl/tls-endpoints-public-dns.yaml**
 - If using a IP address for accessing the public endpoints, use **/usr/share/openstack-tripleo-heat-templates/environments/ssl/tls-endpoints-public-ip.yaml**

For example:

```
$ openstack overcloud deploy --templates [...] -e /home/stack/templates/enable-tls.yaml -e  
~/templates/cloudname.yaml -e ~/templates/inject-trust-anchor.yaml -e /usr/share/openstack-tripleo-  
heat-templates/environments/ssl/tls-endpoints-public-dns.yaml
```

15.11. UPDATING SSL/TLS CERTIFICATES

If you need to update certificates in the future:

- Edit the **enable-tls.yaml** file and update the **SSLCertificate**, **SSLKey**, and **SSLIntermediateCertificate** parameters.
- If your certificate authority has changed, edit the **inject-trust-anchor.yaml** file and update the **SSLRootCertificate** parameter.

Once the new certificate content is in place, rerun your deployment command. For example:

```
$ openstack overcloud deploy --templates [...] -e /home/stack/templates/enable-tls.yaml -e  
~/templates/cloudname.yaml -e ~/templates/inject-trust-anchor.yaml -e /usr/share/openstack-tripleo-  
heat-templates/environments/ssl/tls-endpoints-public-dns.yaml
```

CHAPTER 16. ENABLING SSL/TLS ON INTERNAL AND PUBLIC ENDPOINTS WITH IDENTITY MANAGEMENT

You can enable SSL/TLS on certain overcloud endpoints. Due to the number of certificates required, the director integrates with a Red Hat Identity Management (IdM) server to act as a certificate authority and manage the overcloud certificates. This process involves using **novajoin** to enroll overcloud nodes to the IdM server.

To check the status of TLS support across the OpenStack components, refer to the [TLS Enablement status matrix](#).

16.1. ADD THE UNDERCLOUD TO THE CA

Before deploying the overcloud, you must add the undercloud to the Certificate Authority (CA):

1. On the undercloud node, install the **python-novajoin** package:

```
$ sudo yum install python-novajoin
```

2. On the undercloud node, run the **novajoin-ipa-setup** script, adjusting the values to suit your deployment:

```
$ sudo /usr/libexec/novajoin-ipa-setup \
  --principal admin \
  --password <IdM admin password> \
  --server <IdM server hostname> \
  --realm <overcloud cloud domain (in upper case)> \
  --domain <overcloud cloud domain> \
  --hostname <undercloud hostname> \
  --precreate
```

In the following section, you will use the resulting One-Time Password (OTP) to enroll the undercloud.

16.2. ADD THE UNDERCLOUD TO IDM

This procedure registers the undercloud with IdM and configures novajoin. Configure the following settings in **undercloud.conf** (within the **[DEFAULT]** section):

1. The novajoin service is disabled by default. To enable it:

```
[DEFAULT]
enable_novajoin = true
```

2. You need set a One-Time Password (OTP) to register the undercloud node with IdM:

```
ipa_otp = <otp>
```

3. Ensure the overcloud's domain name served by neutron's DHCP server matches the IdM domain (your kerberos realm in lowercase):

```
overcloud_domain_name = <domain>
```

4. Set the appropriate hostname for the undercloud:

```
undercloud_hostname = <undercloud FQDN>
```

5. Set IdM as the nameserver for the undercloud:

```
undercloud_nameservers = <IdM IP>
```

6. For larger environments, you will need to review the novajoin connection timeout values. In **undercloud.conf**, add a reference to a new file called **undercloud-timeout.yaml**:

```
hieradata_override = /home/stack/undercloud-timeout.yaml
```

Add the following options to **undercloud-timeout.yaml**. You can specify the timeout value in seconds, for example, **5**:

```
nova::api::vendordata_dynamic_connect_timeout: <timeout value>
nova::api::vendordata_dynamic_read_timeout: <timeout value>
```

7. Save the **undercloud.conf** file.
8. Run the undercloud deployment command to apply the changes to your existing undercloud:

```
$ openstack undercloud install
```

Verification

1. Check the **keytab** files for a key entry for the undercloud:

```
[root@undercloud-0 ~]# klist -kt
Keytab name: FILE:/etc/krb5.keytab
KVNO Timestamp      Principal
-----
1 04/28/2020 12:22:06 host/undercloud-0.redhat.local@REDHAT.LOCAL
1 04/28/2020 12:22:06 host/undercloud-0.redhat.local@REDHAT.LOCAL
```

```
[root@undercloud-0 ~]# klist -kt /etc/novajoin/krb5.keytab
Keytab name: FILE:/etc/novajoin/krb5.keytab
KVNO Timestamp      Principal
-----
1 04/28/2020 12:22:26 nova/undercloud-0.redhat.local@REDHAT.LOCAL
1 04/28/2020 12:22:26 nova/undercloud-0.redhat.local@REDHAT.LOCAL
```

2. Test the system **/etc/krb.keytab** file with the host principle:

```
[root@undercloud-0 ~]# kinit -k
[root@undercloud-0 ~]# klist
Ticket cache: KEYRING:persistent:0:0
Default principal: host/undercloud-0.redhat.local@REDHAT.LOCAL

Valid starting    Expires          Service principal
05/04/2020 10:34:30 05/05/2020 10:34:30  krbtgt/REDHAT.LOCAL@REDHAT.LOCAL
```

```
[root@undercloud-0 ~]# kdestroy
Other credential caches present, use -A to destroy all
```

3. Test the novajoin `/etc/novajoin/krb.keytab` file with the nova principle:

```
[root@undercloud-0 ~]# kinit -kt /etc/novajoin/krb5.keytab 'nova/undercloud-
0.redhat.local@REDHAT.LOCAL'
[root@undercloud-0 ~]# klist
Ticket cache: KEYRING:persistent:0:0
Default principal: nova/undercloud-0.redhat.local@REDHAT.LOCAL

Valid starting    Expires          Service principal
05/04/2020 10:39:14  05/05/2020 10:39:14  krbtgt/REDHAT.LOCAL@REDHAT.LOCAL
```

16.3. CONFIGURE OVERCLOUD DNS

For automatic detection of your IdM environment, and easier enrollment, consider using IdM as your DNS server:

1. Connect to your undercloud:

```
$ source ~/stackrc
```

2. Configure the control plane subnet to use IdM as the DNS name server:

```
$ openstack subnet set ctlplane-subnet --dns-nameserver <idm_server_address>
```

3. Set the **DnsServers** parameter in an environment file to use your IdM server:

```
parameter_defaults:
  DnsServers: ["<idm_server_address>"]
```

This parameter is usually defined in a custom **network-environment.yaml** file.

16.4. CONFIGURE OVERCLOUD TO USE NOVAJOIN

1. To enable IdM integration, create a copy of the `/usr/share/openstack-tripleo-heat-templates/environments/predictable-placement/custom-domain.yaml` environment file:

```
$ cp /usr/share/openstack-tripleo-heat-templates/environments/predictable-
placement/custom-domain.yaml \
/home/stack/templates/custom-domain.yaml
```

2. Edit the `/home/stack/templates/custom-domain.yaml` environment file and set the **CloudDomain** and **CloudName*** values to suit your deployment. For example:

```
parameter_defaults:
  CloudDomain: lab.local
  CloudName: overcloud.lab.local
  CloudNameInternal: overcloud.internalapi.lab.local
```

```
CloudNameStorage: overcloud.storage.lab.local
CloudNameStorageManagement: overcloud.storageemgmt.lab.local
CloudNameCtlplane: overcloud.ctlplane.lab.local
```

3. Include the following environment files in the overcloud deployment process:

- **/usr/share/openstack-tripleo-heat-templates/environments/ssl/enable-internal-tls.yaml**
- **/usr/share/openstack-tripleo-heat-templates/environments/ssl/tls-everywhere-endpoints-dns.yaml**
- **/home/stack/templates/custom-domain.yaml**

For example:

```
openstack overcloud deploy \
  --templates \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ssl/enable-internal-
  tls.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ssl/tls-everywhere-
  endpoints-dns.yaml \
  -e /home/stack/templates/custom-domain.yaml \
```

As a result, the deployed overcloud nodes will be automatically enrolled with IdM.

4. This only sets TLS for the internal endpoints. For the external endpoints you can use the normal means of adding TLS with the **/usr/share/openstack-tripleo-heat-templates/environments/ssl/enable-tls.yaml** environment file (which must be modified to add your custom certificate and key). Consequently, your **openstack deploy** command would be similar to this:

```
openstack overcloud deploy \
  --templates \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ssl/enable-internal-tls.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ssl/tls-everywhere-endpoints-
  dns.yaml \
  -e /home/stack/templates/custom-domain.yaml \
  -e /home/stack/templates/enable-tls.yaml
```

5. Alternatively, you can also use IdM to issue your public certificates. In that case, you need to use the **/usr/share/openstack-tripleo-heat-templates/environments/services/haproxy-public-tls-certmonger.yaml** environment file. For example:

```
openstack overcloud deploy \
  --templates \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ssl/enable-internal-tls.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ssl/tls-everywhere-endpoints-
  dns.yaml \
  -e /home/stack/templates/custom-domain.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/services/haproxy-public-tls-
  certmonger.yaml
```

CHAPTER 17. CONVERTING YOUR EXISTING DEPLOYMENT TO USE TLS

You can configure your existing overcloud and undercloud endpoints to use TLS encryption. This approach uses **novajoin** to integrate your deployment with Red Hat Identity Management (IdM), allowing access to DNS, Kerberos, and certmonger. Each overcloud node uses a certmonger client to retrieve certificates for each service.

For more information on TLS, see the [Security and Hardening Guide](#).

17.1. REQUIREMENTS

- For Red Hat OpenStack Platform 13, you must be running version **z8** or higher.
- You must have an existing IdM deployment, and it must also supply DNS services to the OpenStack deployment.
- The existing deployment must use FQDNs for public endpoints. Default configurations might use IP address-based endpoints, and will consequently generate IP address-based certificates; these must be changed to FQDNs before proceeding with these steps.



IMPORTANT

The overcloud and undercloud services will be unavailable for the duration of this procedure.

17.2. REVIEWING YOUR ENDPOINTS

By default, your existing Red Hat OpenStack Platform 13 overcloud does not encrypt certain endpoints with TLS. For example, this output includes URLs that use **http** instead of **https**; these are not encrypted:

```
+-----+-----+-----+-----+-----+-----+-----+
| ID           | Region | Service Name | Service Type | Enabled | Interface | URL
|
+-----+-----+-----+-----+-----+-----+-----+
| 0ad11e943e1f4ff988650cfba57b4031 | regionOne | nova      | compute  | True  | internal | http://172.16.2.17:8774/v2.1
| 1413eb9ef38a45b8bee1bee1b0dfe744 | regionOne | swift     | object-store | True  | public   | https://overcloud.lab.local:13808/v1/AUTH_%(tenant_id)s
| 1a54f13f212044b0a20468861cd06f85 | regionOne | neutron   | network   | True  | public   | https://overcloud.lab.local:13696
| 3477a3a052d2445697bb6642a8c26a91 | regionOne | placement | placement | True  | internal | http://172.16.2.17:8778/placement
| 3f56445c0dd14721ac830d6afb2c2cd4 | regionOne | nova      | compute   | True  | admin    | http://172.16.2.17:8774/v2.1
| 425b1773a55c4245bcbe3d051772ebba | regionOne | glance    | image     | True  | internal | http://172.16.2.17:9292
| 57cf09fa33ed446f8736d4228bdfa881 | regionOne | placement | placement | True  | public   | https://overcloud.lab.local:13778/placement
| 58600f3751e54f7e9d0a50ba618e4c54 | regionOne | glance    | image     | True  | public   | https://overcloud.lab.local:13292
```

```

| 5c52f273c3284b068f2dc885c77174ca | regionOne | neutron | network | True | internal |
http://172.16.2.17:9696 |
| 8792a4dd8bbb456d9dea4643e57c43dc | regionOne | nova | compute | True | public |
https://overcloud.lab.local:13774/v2.1 |
| 94bbea97580a4c4b844478aad5a85e84 | regionOne | keystone | identity | True | public |
https://overcloud.lab.local:13000 |
| acbf11b5c76d44198af49e3b78ffedcd | regionOne | swift | object-store | True | internal |
http://172.16.1.9:8080/v1/AUTH_%(tenant_id)s |
| d4a1344f02a74f7ab0a50c5a7c13ca5c | regionOne | keystone | identity | True | internal |
http://172.16.2.17:5000 |
| d86c241dc97642419ddc12533447d73d | regionOne | placement | placement | True | admin |
http://172.16.2.17:8778/placement |
| de7d6c34533e4298a2752852427a7030 | regionOne | glance | image | True | admin |
http://172.16.2.17:9292 |
| e82086062ebd4d4b9e03c7f1544bdd3b | regionOne | swift | object-store | True | admin |
http://172.16.1.9:8080 |
| f8134cd9746247bca6a06389b563c743 | regionOne | keystone | identity | True | admin |
http://192.168.24.6:35357 |
| fe29177bd29545ca8fdc0c777a7cf03f | regionOne | neutron | network | True | admin |
http://172.16.2.17:9696 |
+-----+-----+-----+-----+-----+-----+-----+
-----+

```

The following sections explain how to encrypt these endpoints using TLS.

17.3. APPLY WORKAROUND FOR KNOWN ISSUE

There is currently a known issue for TLS Everywhere in-place upgrades, where overcloud nodes are consequently unable to enroll in IdM. As a workaround, remove **/etc/ipa/ca.crt/** from all overcloud nodes before running the overcloud deploy. For more information, see https://bugzilla.redhat.com/show_bug.cgi?id=1732564.

For example, the following script is one way of applying the workaround. You might need to amend this to suit your deployment.

```

[stack@undercloud-0 ~]$ vi rm-ca.crt-dir.sh
#!/bin/bash

source /home/stack/stackrc
NODES=$(openstack server list -f value -c Networks|sed s/ctlplane=//g)

for NODE in $NODES
do
    ssh heat-admin@$NODE sudo rm -rf /etc/ipa/ca.crt/
Done

[stack@undercloud-0 ~]$ bash rm-ca.crt-dir.sh

```

17.4. CONFIGURING ENDPOINTS TO USE TLS

This section explains how to enable TLS endpoint encryption for an existing deployment, and then how to check that the endpoints have been correctly configured.

When enabling TLS everywhere, there are different upgrade paths available, depending on how your domains are structured. These examples use sample domain names to describe the upgrade paths:

- Reuse the existing public endpoint certificates, and enable TLS everywhere on the **internal** and **admin** endpoints where the overcloud domain (**lab.local**) matches the IdM domain (**lab.local**).
- Allow IdM to issue new public endpoints certificates, and enable TLS everywhere on the **internal** and **admin** endpoints where the overcloud domain (**lab.local**) matches the IdM domain (**lab.local**).
- Reuse existing public endpoint certificates, and enable TLS everywhere on the **internal** and **admin** endpoints where the overcloud domain (**site1.lab.local**) is a subdomain of the IdM domain (**lab.local**).
- Allow IdM to issue new public endpoints certificates, and enable TLS everywhere on the **internal** and **admin** endpoints where the overcloud domain (**site1.lab.local**) is a subdomain of the IdM domain (**lab.local**).

The following procedures in this section explain how to configure this integration using the various combinations described above.

17.4.1. Configuring undercloud integration for deployments using the same domain as IdM

This procedure describes how to configure undercloud integration for deployments that use the same domain as IdM.

Red Hat OpenStack Platform uses **novajoin** to integrate with Red Hat Identity Management (IdM), which then issues and manages encryption certificates. In this procedure, you register the undercloud with IdM, generate a token, enable the token in the undercloud configuration, then re-run the undercloud and overcloud deployment scripts. For example:

1. Install **python-novajoin** for integration with IdM:

```
[stack@undercloud-0 ~]$ sudo yum install python-novajoin
```

2. Run the **novajoin** configuration script and supply the configuration details for your IdM deployment. For example:

```
[stack@undercloud-0 ~]$ sudo novajoin-ipa-setup --principal admin --password
ComplexRedactedPassword \
  --server ipa.lab.local --realm lab.local --domain lab.local \
  --hostname undercloud-0.lab.local --precreate
...
0Uvua6NyIWVkfCSTOmwbDAobsqGH2GONRJRw24MoQ4wg
```

This output includes a one time password (OTP) for IdM, which will be a different value for your deployment.

3. Configure the undercloud to use **novajoin**, add the one-time password (OTP), use the IdM IP address for DNS, and describe the overcloud domain. You will need to adjust this example for your deployment:

```
[stack@undercloud ~]$ vi undercloud.conf
...
```

```

enable_novajoin = true
ipa_otp = 0Uvua6NyIWVkfCSTOmwbdaObsqGH2GONRJRw24MoQ4wg
undercloud_hostname = undercloud-0.lab.local
undercloud_nameservers = X.X.X.X
overcloud_domain_name = lab.local
...

```

4. Install the **novajoin** services in the undercloud:

```
[stack@undercloud ~]$ openstack undercloud install
```

5. Add the overcloud IP address to DNS. You will need to amend this example to suit your deployment:

Note: Check the overcloud's **network-environment.yaml**, and choose a VIP within each network's range.

```

[root@ipa ~]$ ipa dnsrecord-add lab.local overcloud --a-rec=10.0.0.101
[root@ipa ~]# ipa dnszone-add ctlplane.lab.local
[root@ipa ~]# ipa dnsrecord-add ctlplane.lab.local overcloud --a-rec 192.168.24.101
[root@ipa ~]# ipa dnszone-add internalapi.lab.local
[root@ipa ~]# ipa dnsrecord-add internalapi.lab.local overcloud --a-rec 172.17.1.101
[root@ipa ~]# ipa dnszone-add storage.lab.local
[root@ipa ~]# ipa dnsrecord-add storage.lab.local overcloud --a-rec 172.17.3.101
[root@ipa ~]# ipa dnszone-add storagegmt.lab.local
[root@ipa ~]# ipa dnsrecord-add storagegmt.lab.local overcloud --a-rec 172.17.4.101

```

6. Create a **public_vip.yaml** mapping for all the endpoints:

```

Parameter_defaults:
  PublicVirtualFixedIPs: [{'ip_address':'10.0.0.101'}]
  ControlFixedIPs: [{'ip_address':'192.168.24.101'}]
  InternalApiVirtualFixedIPs: [{'ip_address':'172.17.1.101'}]
  StorageVirtualFixedIPs: [{'ip_address':'172.17.3.101'}]
  StorageMgmtVirtualFixedIPs: [{'ip_address':'172.17.4.101'}]
  RedisVirtualFixedIPs: [{'ip_address':'172.17.1.102'}]

```

17.4.2. Configuring overcloud integration for deployments that use the same domain as IdM, and retain the existing public endpoint certificates

1. Make sure the following parameters exist in your **openstack overcloud deploy** command (with valid settings) and then re-run the deployment command:
 - `--ntp-server` - If not already set, specify the NTP server to suit your environment. The IdM server should be running ntp.
 - **cloud-names.yaml** - Contains the FQDNs (not IPs) from the initial deployment command.
 - **enable-tls.yaml** - Contains the new overcloud certificate. For an example, see <https://github.com/openstack/tripleo-heat-templates/blob/master/environments/ssl/enable-tls.yaml>.
 - **public_vip.yaml** - The maps the endpoints to a specific ip so dns can match.
 - **enable-internal-tls.yaml** - Enables TLS for internal endpoints.

- **tls-everywhere-endpoints-dns.yaml** - Configures TLS endpoints using DNS names. You can review the contents of this file to check the configuration scope.
- **haproxy-internal-tls-certmonger.yaml** - certmonger will manage the internal certs in haproxy.
- **inject-trust-anchor.yaml** - Adds the root certificate authority. This is only needed when the certificates rely on a CA chain that is not already part of the common set used by default; for example, when using self-signed.

For example:

```
[ stack@undercloud ~]$ openstack overcloud deploy \
...
--ntp-server 10.13.57.78 \
-e /home/stack/cloud-names.yaml \
-e /home/stack/enable-tls.yaml \
-e /home/stack/public_vip.yaml \
-e <tripleo-heat-templates>/environments/ssl/enable-internal-tls.yaml \
-e <tripleo-heat-templates>/environments/ssl/tls-everywhere-endpoints-dns.yaml \
-e <tripleo-heat-templates>/environments/services/haproxy-internal-tls-certmonger.yaml \
-e /home/stack/inject-trust-anchor.yaml
...
```



NOTE

Examples of these environment files can be found here:

<https://github.com/openstack/tripleo-heat-templates/tree/master/environments/ssl>.

17.4.3. Configuring overcloud integration for deployments that use the same domain as IdM, and replace the existing public endpoint certificates with an IdM generated certificate

1. Make sure the following parameters exist in your **openstack overcloud deploy** command (with valid settings) and then re-run the deployment command:
 - `--ntp-server`` - If not already set, specify the NTP server to suit your environment. The IdM server should be running ntp.
 - **cloud-names.yaml** - Contains the FQDNs (not IPs) from the initial deployment command.
 - **enable-tls.yaml** - Contains the new overcloud certificate. For an example, see <https://github.com/openstack/tripleo-heat-templates/blob/master/environments/ssl/enable-tls.yaml>.
 - **public_vip.yaml** - The maps the endpoints to a specific ip so dns can match.
 - **enable-internal-tls.yaml** - Enables TLS for internal endpoints.
 - **tls-everywhere-endpoints-dns.yaml** - Configures TLS endpoints using DNS names. You can review the contents of this file to check the configuration scope.
 - **haproxy-public-tls-certmonger.yaml** - certmonger will manage the internal and public certs in haproxy.

- **inject-trust-anchor.yaml** - Adds the root certificate authority. This is only needed when the certificates rely on a CA chain that is not already part of the common set used by default; for example, when using self-signed.

For example:

```
[stack@undercloud ~]$ openstack overcloud deploy \
...
--ntp-server 10.13.57.78 \
-e /home/stack/cloud-names.yaml \
-e /home/stack/enable-tls.yaml \
-e /home/stack/public_vip.yaml \
-e <tripleo-heat-templates>/environments/ssl/enable-internal-tls.yaml \
-e <tripleo-heat-templates>/environments/ssl/tls-everywhere-endpoints-dns.yaml \
-e <tripleo-heat-templates>/environments/services/haproxy-public-tls-certmonger.yaml \
-e /home/stack/inject-trust-anchor.yaml
...
```



NOTE

Examples of these environment files can be found at <https://github.com/openstack/tripleo-heat-templates/tree/master/environments/ssl>.



NOTE

The template **enable-internal-tls.j2.yaml** is referenced as **enable-internal-tls.yaml** in the overcloud deploy command.

In addition, the old public endpoint certificates in **enable-tls.yaml** will be replaced by certmonger with **haproxy-public-tls-certmonger.yaml**, however, this file must still be referenced in the upgrade process.

17.4.4. Configuring undercloud integration for deployments that use an IdM subdomain

This procedure explains how to configure undercloud integration for deployments that use an IdM subdomain.

Red Hat OpenStack Platform uses **novajoin** to integrate with Red Hat Identity Management (IdM), which then issues and manages encryption certificates. In this procedure, you register the undercloud with IdM, generate a token, enable the token in the undercloud configuration, then re-run the undercloud and overcloud deployment scripts. For example:

1. Install **python-novajoin** for integration with IdM:

```
[stack@undercloud-0 ~]$
```

2. Run the **novajoin** configuration script and supply the configuration details for your IdM deployment. For example:

```
[stack@undercloud-0 ~]$ sudo novajoin-ipa-setup --principal admin --password
ComplexRedactedPassword \
--server ipa.lab.local --realm lab.local --domain lab.local \
```

```
--hostname undercloud-0.site1.lab.local --precreate
...
0Uvua6NyIWVkfCSTOmwbdaObsqGH2GONRJRw24MoQ4wg
```

This output includes a one time password (OTP) for IdM, which will be a different value for your deployment.

3. Configure the undercloud to use **novajoin**, and add the OTP, IdM IP for DNS and NTP, and overcloud domain:

```
[stack@undercloud ~]$ vi undercloud.conf
...
[DEFAULT]
undercloud_ntp_servers=X.X.X.X
hieradata_override = /home/stack/hiera_override.yaml
enable_novajoin = true
ipa_otp = 0Uvua6NyIWVkfCSTOmwbdaObsqGH2GONRJRw24MoQ4wg
undercloud_hostname = undercloud-0.site1.lab.local
undercloud_nameservers = X.X.X.X
overcloud_domain_name = site1.lab.local
...
```

4. Configure the undercloud to use **novajoin**, and add the OTP, IdM IP for DNS, and overcloud domain:

```
[stack@undercloud-0 ~]$ vi hiera_override.yaml
nova::metadata::novajoin::api::ipa_domain: site1.lab.local
...
```

5. Install the **novajoin** services in the undercloud:

```
[stack@undercloud ~]$ openstack undercloud install
```

6. Add the overcloud IP address to DNS. You will need to amend this example to suit your deployment:

Note: Check the overcloud's **network-environment.yaml**, and choose a VIP within each network's range.

```
[root@ipa ~]$ ipa dnsrecord-add site1.lab.local overcloud --a-rec=10.0.0.101
[root@ipa ~]# ipa dnszone-add site1.ctlplane.lab.local
[root@ipa ~]# ipa dnsrecord-add site1.ctlplane.lab.local overcloud --a-rec 192.168.24.101
[root@ipa ~]# ipa dnszone-add site1.internalapi.lab.local
[root@ipa ~]# ipa dnsrecord-add site1.internalapi.lab.local overcloud --a-rec 172.17.1.101
[root@ipa ~]# ipa dnszone-add site1.storage.lab.local
[root@ipa ~]# ipa dnsrecord-add site1.storage.lab.local overcloud --a-rec 172.17.3.101
[root@ipa ~]# ipa dnszone-add site1.storagegmt.lab.local
[root@ipa ~]# ipa dnsrecord-add site1.storagegmt.lab.local overcloud --a-rec 172.17.4.101
```

7. Create a **public_vip.yaml** mapping for each of the endpoints. For example:

```
Parameter_defaults:
  PublicVirtualFixedIPs: [{'ip_address':'10.0.0.101'}]
  ControlFixedIPs: [{'ip_address':'192.168.24.101'}]
  InternalApiVirtualFixedIPs: [{'ip_address':'172.17.1.101'}]
```

```
StorageVirtualFixedIPs: [{'ip_address':'172.17.3.101'}]
StorageMgmtVirtualFixedIPs: [{'ip_address':'172.17.4.101'}]
RedisVirtualFixedIPs: [{'ip_address':'172.17.1.102'}]
```

8. Create the **extras.yaml** mapping for each of the endpoints. For example:

```
parameter_defaults:
  MakeHomeDir: True
  IdMNoNtpSetup: false
  IdMDomain: redhat.local
  DnsSearchDomains: ["site1.redhat.local", "redhat.local"]
```

17.4.5. Configuring undercloud integration for deployments that use an IdM subdomain, and retain the existing public endpoint certificates

This procedure explains how to configure undercloud integration for deployments that use an IdM subdomain, and still retain the existing public endpoint certificates.

1. Make sure the following parameters exist in your **openstack overcloud deploy** command (with valid settings) and then re-run the deployment command:
 - `--ntp-server`` - If not already set, specify the NTP server to suit your environment. The IdM server should be running ntp.
 - **cloud-names.yaml** - Contains the FQDNs (not IPs) from the initial deployment command.
 - **enable-tls.yaml** - Contains the new overcloud certificate. For an example, see <https://github.com/openstack/tripleo-heat-templates/blob/master/environments/ssl/enable-tls.yaml>.
 - **public_vip.yaml** - Contains endpoint maps to a specific ip so dns can match.
 - `extras.yaml`` - Contains settings for pam to make home directories on login, no ntp setup, the base IdM domain, and the dns search for resolv.conf.
 - **enable-internal-tls.yaml** - Enables TLS for internal endpoints.
 - **tls-everywhere-endpoints-dns.yaml** - Configures TLS endpoints using DNS names. You can review the contents of this file to check the configuration scope.
 - **haproxy-internal-tls-certmonger.yaml** - certmonger will manage the internal certs in haproxy.
 - **inject-trust-anchor.yaml** - Adds the root certificate authority. This is only needed when the certificates rely on a CA chain that is not already part of the common set used by default; for example, when using self-signed.

For example:

```
[ stack@undercloud ~]$ openstack overcloud deploy \
...
--ntp-server 10.13.57.78 \
-e /home/stack/cloud-names.yaml \
-e /home/stack/enable-tls.yaml \
-e /home/stack/public_vip.yaml \
-e /home/stack/extras.yaml \
```

```

-e <tripleo-heat-templates>/environments/ssl/enable-internal-tls.yaml \
-e <tripleo-heat-templates>/environments/ssl/tls-everywhere-endpoints-dns.yaml \
-e <tripleo-heat-templates>/environments/services/haproxy-internal-tls-certmonger.yaml \
\
-e /home/stack/inject-trust-anchor.yaml
...

```



NOTE

Examples of these environment files can be found here:

<https://github.com/openstack/tripleo-heat-templates/tree/master/environments/ssl>.

17.4.6. Configuring undercloud integration for deployments that use an IdM subdomain, and replace the existing public endpoint certificates with an IdM generated certificate

This procedure explains how to configure undercloud integration for deployments that use an IdM subdomain, and how to replace the existing public endpoint certificates with an IdM generated certificate.

1. Make sure the following parameters exist in your **openstack overcloud deploy** command (with valid settings) and then re-run the deployment command:
 - `--ntp-server`` - If not already set, specify the NTP server to suit your environment. The IdM server should be running ntp.
 - **cloud-names.yaml** - Contains the FQDNs (not IPs) from the initial deployment command.
 - **enable-tls.yaml** - Contains the new overcloud certificate. For an example, see <https://github.com/openstack/tripleo-heat-templates/blob/master/environments/ssl/enable-tls.yaml>.
 - **public_vip.yaml** - The maps the endpoints to a specific ip so dns can match.
 - `extras.yaml`` - Contains settings for pam to make home directories on login, no ntp setup, the base IdM domain, and the dns search for resolv.conf.
 - **enable-internal-tls.yaml** - Enables TLS for internal endpoints.
 - **tls-everywhere-endpoints-dns.yaml** - Configures TLS endpoints using DNS names. You can review the contents of this file to check the configuration scope.
 - **haproxy-public-tls-certmonger.yaml** - certmonger will manage the internal and public certs in haproxy.
 - **inject-trust-anchor.yaml** - Adds the root certificate authority. This is only needed when the certificates rely on a CA chain that is not already part of the common set used by default; for example, when using self-signed.

For example:

```

[ stack@undercloud ~]$ openstack overcloud deploy \
...
--ntp-server 10.13.57.78 \
-e /home/stack/cloud-names.yaml \

```

```
-e /home/stack/enable-tls.yaml \
-e /home/stack/public_vip.yaml \
-e /home/stack/extras.yaml \
-e <tripleo-heat-templates>/environments/ssl/enable-internal-tls.yaml \
-e <tripleo-heat-templates>/environments/ssl/tls-everywhere-endpoints-dns.yaml \
-e <tripleo-heat-templates>/environments/services/haproxy-public-tls-certmonger.yaml \
-e /home/stack/inject-trust-anchor.yaml
...

```



NOTE

Examples of these environment files can be found here:
<https://github.com/openstack/tripleo-heat-templates/tree/master/environments/ssl>.



NOTE

In this example, the template **enable-internal-tls.j2.yaml** is referenced as **enable-internal-tls.yaml** in the **overcloud deploy** command. In addition, the old public endpoint certificates in **enable-tls.yaml** will be replaced by certmonger using **haproxy-public-tls-certmonger.yaml**, however, this file must still be referenced in the upgrade process.

17.5. CHECKING TLS ENCRYPTION

Once the overcloud re-deployment has completed, check that all endpoints are now encrypted with TLS. In this example, all endpoints are configured to use **https**, indicating that they are using TLS encryption:

```
+-----+-----+-----+-----+-----+-----+
| ID           | Region  | Service Name | Service Type | Enabled | Interface | URL
|
+-----+-----+-----+-----+-----+-----+
| 0fee4efdc4ae4310b6a139a25d9c0d9c | regionOne | neutron      | network      | True   | public   | https://overcloud.lab.local:13696
| 220558ab1d2445139952425961a0c89a | regionOne | glance       | image        | True   | public   | https://overcloud.lab.local:13292
| 24d966109ffa419da850da946f19c4ca | regionOne | placement    | placement    | True   | admin    | https://overcloud.internalapi.lab.local:8778/placement
| 27ac9e0d22804ee5bd3cd8c0323db49c | regionOne | nova         | compute      | True   | internal | https://overcloud.internalapi.lab.local:8774/v2.1
| 31d376853bd241c2ba1a27912fc896c6 | regionOne | swift        | object-store | True   | admin    | https://overcloud.storage.lab.local:8080
| 350806234c784332bfb8615e721057e3 | regionOne | nova         | compute      | True   | admin    | https://overcloud.internalapi.lab.local:8774/v2.1
| 49c312f4db6748429d27c60164779302 | regionOne | keystone     | identity     | True   | public   | https://overcloud.lab.local:13000
| 4e535265c35e486e97bb5a8bc77708b6 | regionOne | nova         | compute      | True   | public   | https://overcloud.lab.local:13774/v2.1
| 5e93dd46b45f40fe8d91d3a5d6e847d3 | regionOne | keystone     | identity     | True   | admin    | https://overcloud.ctlplane.lab.local:35357
| 6561984a90c742a988bf3d0acf80d1b6 | regionOne | swift        | object-store | True   | public   | https://overcloud.lab.local:13808/v1/AUTH_%(tenant_id)s

```



```

| 76b8aad0bdda4313a02e4342e6a19fd6 | regionOne | placement | placement | True | public
| https://overcloud.lab.local:13778/placement |
| 96b004d5217c4d87a38cb780607bf9fb | regionOne | placement | placement | True | internal |
| https://overcloud.internalapi.lab.local:8778/placement |
| 98489b4b107f4da596262b712c3fe883 | regionOne | glance | image | True | internal |
| https://overcloud.internalapi.lab.local:9292 |
| bb7ab36f30b14b549178ef06ec74ff84 | regionOne | glance | image | True | admin |
| https://overcloud.internalapi.lab.local:9292 |
| c1547f7bf9a14e9e85eaaaeaa26413b7 | regionOne | neutron | network | True | admin |
| https://overcloud.internalapi.lab.local:9696 |
| ca66f499ec544f42838eb78a515d9f1e | regionOne | keystone | identity | True | internal |
| https://overcloud.internalapi.lab.local:5000 |
| df0181358c07431390bc66822176281d | regionOne | swift | object-store | True | internal |
| https://overcloud.storage.lab.local:8080/v1/AUTH_%(tenant_id)s |
| e420350ef856460991c3edbfbae917c1 | regionOne | neutron | network | True | internal |
| https://overcloud.internalapi.lab.local:9696 |
+-----+-----+-----+-----+-----+-----+-----+
-----+

```

CHAPTER 18. DEBUG MODES

You can enable and disable the **DEBUG** level logging mode for certain services in the overcloud. To configure debug mode for a service, set the respective debug parameter.

For example, OpenStack Identity (keystone) uses the **KeystoneDebug** parameter. Create a **debug.yaml** environment file to store debug parameters and set the **KeystoneDebug** parameter in the **parameter_defaults** section:

```
parameter_defaults:  
  KeystoneDebug: True
```

After you have set the **KeystoneDebug** parameter to **True**, the **/var/log/containers/keystone/keystone.log** standard keystone log file is updated with **DEBUG** level logs.

For a full list of debug parameters, see "[Debug Parameters](#)" in the *Overcloud Parameters* guide.

CHAPTER 19. STORAGE CONFIGURATION

This chapter outlines several methods of configuring storage options for your Overcloud.




IMPORTANT

By default, the overcloud uses local ephemeral storage provided by OpenStack Compute (nova) and LVM block storage provided by OpenStack Storage (cinder). However, these options are not supported for enterprise-level overclouds. Instead, use one of the storage options in this chapter.

19.1. CONFIGURING NFS STORAGE

This section describes how to configure the overcloud to use an NFS share. The installation and configuration process is based on the modification of an existing environment file in the core heat template collection.



IMPORTANT

Red Hat recommends that you use a certified storage back end and driver. Red Hat does not recommend that you use NFS that comes from the generic NFS back end, because its capabilities are limited when compared to a certified storage back end and driver. For example, the generic NFS back end does not support features such as volume encryption and volume multi-attach. For information about supported drivers, see the [Red Hat Ecosystem Catalog](#).



NOTE

There are several director heat parameters that control whether an NFS back end or a NetApp NFS Block Storage back end supports a NetApp feature called NAS secure:

- CinderNetappNasSecureFileOperations
- CinderNetappNasSecureFilePermissions
- CinderNasSecureFileOperations
- CinderNasSecureFilePermissions

Red Hat does not recommend that you enable the feature, because it interferes with normal volume operations. Director disables the feature by default, and Red Hat OpenStack Platform does not support it.



NOTE

For Block Storage and Compute services, you must use NFS version 4.0 or later.

The core heat template collection contains a set of environment files in `/usr/share/openstack-tripleo-heat-templates/environments/`. With these environment files you can create customized configuration of some of the supported features in a director-created overcloud. This includes an environment file designed to configure storage. This file is located at `/usr/share/openstack-tripleo-heat-templates/environments/storage-environment.yaml`.

1. Copy the file to the **stack** user's template directory:

```
$ cp /usr/share/openstack-tripleo-heat-templates/environments/storage-environment.yaml
~/templates/.
```

2. Modify the following parameters:

CinderEnableIscsiBackend

Enables the iSCSI backend. Set to **false**.

CinderEnableRbdBackend

Enables the Ceph Storage backend. Set to **false**.

CinderEnableNfsBackend

Enables the NFS backend. Set to **true**.

NovaEnableRbdBackend

Enables Ceph Storage for Nova ephemeral storage. Set to **false**.

GlanceBackend

Define the back end to use for glance. Set to **file** to use file-based storage for images. The overcloud saves these files in a mounted NFS share for glance.

CinderNfsMountOptions

The NFS mount options for the volume storage.

CinderNfsServers

The NFS share to mount for volume storage. For example, 192.168.122.1:/export/cinder.

GlanceNfsEnabled

When **GlanceBackend** is set to **file**, **GlanceNfsEnabled** enables images to be stored through NFS in a shared location so that all Controller nodes have access to the images. If disabled, the overcloud stores images in the file system of the Controller node. Set to **true**.

GlanceNfsShare

The NFS share to mount for image storage. For example, 192.168.122.1:/export/glance.

GlanceNfsOptions

The NFS mount options for the image storage.

The environment file contains parameters that configure different storage options for the Red Hat OpenStack Platform Block Storage (cinder) and Image (glance) services. This example demonstrates how to configure the overcloud to use an NFS share.

The options in the environment file should look similar to the following:

```
parameter_defaults:
  CinderEnableIscsiBackend: false
  CinderEnableRbdBackend: false
  CinderEnableNfsBackend: true
  NovaEnableRbdBackend: false
  GlanceBackend: file

  CinderNfsMountOptions: rw,sync,context=system_u:object_r:cinder_var_lib_t:s0
  CinderNfsServers: 192.0.2.230:/cinder

  GlanceNfsEnabled: true
  GlanceNfsShare: 192.0.2.230:/glance
  GlanceNfsOptions: rw,sync,context=system_u:object_r:glance_var_lib_t:s0
```

These parameters are integrated as part of the heat template collection. Setting them as shown in the example code creates two NFS mount points for the Block Storage and Image services to use.



IMPORTANT

Include the **context=system_u:object_r:glance_var_lib_t:s0** option in the **GlanceNfsOptions** parameter to allow the Image service to access the **/var/lib** directory. Without this SELinux content, the Image service cannot to write to the mount point.

3. Include the file when you deploy the overcloud.

19.2. CONFIGURING CEPH STORAGE

The director provides two main methods for integrating Red Hat Ceph Storage into an Overcloud.

Creating an Overcloud with its own Ceph Storage Cluster

The director has the ability to create a Ceph Storage Cluster during the creation on the Overcloud. The director creates a set of Ceph Storage nodes that use the Ceph OSD to store the data. In addition, the director install the Ceph Monitor service on the Overcloud's Controller nodes. This means if an organization creates an Overcloud with three highly available controller nodes, the Ceph Monitor also becomes a highly available service. For more information, see the [Deploying an Overcloud with Containerized Red Hat Ceph](#) guide.

Integrating a Existing Ceph Storage into an Overcloud

If you already have an existing Ceph Storage Cluster, you can integrate this during an Overcloud deployment. This means you manage and scale the cluster outside of the Overcloud configuration. For more information, see the [Integrating an Overcloud with an Existing Red Hat Ceph Cluster](#) guide.

19.3. USING AN EXTERNAL OBJECT STORAGE CLUSTER

You can reuse an external Object Storage (swift) cluster by disabling the default Object Storage service deployment on the controller nodes. Doing so disables both the proxy and storage services for Object Storage and configures haproxy and keystone to use the given external Swift endpoint.



NOTE

User accounts on the external Object Storage (swift) cluster have to be managed by hand.

You need the endpoint IP address of the external Object Storage cluster as well as the **authtoken** password from the external Object Storage **proxy-server.conf** file. You can find this information by using the **openstack endpoint list** command.

To deploy director with an external Swift cluster:

1. Create a new file named **swift-external-params.yaml** with the following content:
 - Replace **EXTERNAL.IP:PORT** with the IP address and port of the external proxy and
 - Replace **AUTHTOKEN** with the **authtoken** password for the external proxy on the **SwiftPassword** line.

```
parameter_defaults:
  ExternalPublicUrl: 'https://EXTERNAL.IP:PORT/v1/AUTH_%(tenant_id)s'
  ExternalInternalUrl: 'http://192.168.24.9:8080/v1/AUTH_%(tenant_id)s'
  ExternalAdminUrl: 'http://192.168.24.9:8080'
  ExternalSwiftUserTenant: 'service'
  SwiftPassword: AUTHTOKEN
```

2. Save this file as **swift-external-params.yaml**.
3. Deploy the overcloud using these additional environment files.

```
openstack overcloud deploy --templates \
-e [your environment files]
-e /usr/share/openstack-tripleo-heat-templates/environments/swift-external.yaml
-e swift-external-params.yaml
```

19.4. CONFIGURING THE IMAGE IMPORT METHOD AND SHARED STAGING AREA

The default settings for the OpenStack Image service (glance) are determined by the Heat templates used when OpenStack is installed. The Image service Heat template is **tht/puppet/services/glance-api.yaml**.

The interoperable image import allows two methods for image import:

- web-download and
- glance-direct.

The **web-download** method lets you import an image from a URL; the **glance-direct** method lets you import an image from a local volume.

19.4.1. Creating and Deploying the glance-settings.yaml File

You use an environment file to configure the import parameters. These parameters override the default values established in the Heat template. The example environment content provides parameters for the interoperable image import.

```
parameter_defaults:
  # Configure NFS backend
  GlanceBackend: file
  GlanceNfsEnabled: true
  GlanceNfsShare: 192.168.122.1:/export/glance

  # Enable glance-direct import method
  GlanceEnabledImportMethods: glance-direct,web-download

  # Configure NFS staging area (required for glance-direct import method)
  GlanceStagingNfsShare: 192.168.122.1:/export/glance-staging
```

The **GlanceBackend**, **GlanceNfsEnabled**, and **GlanceNfsShare** parameters are defined in the [Storage Configuration](#) section in the *Advanced Overcloud Customization Guide*.

Two new parameters for interoperable image import define the import method and a shared NFS staging area.

GlanceEnabledImportMethods

Defines the available import methods, web-download (default) and glance-direct. This line is only necessary if you wish to enable additional methods besides web-download.

GlanceStagingNfsShare

Configures the NFS staging area used by the glance-direct import method. This space can be shared amongst nodes in a high-availability cluster setup. Requires GlanceNfsEnabled be set to true.

To configure the settings:


1. Create a new file called, for example, glance-settings.yaml. The contents of this file should be similar to the example above.
2. Add the file to your OpenStack environment using the **openstack overcloud deploy** command:

```
$ openstack overcloud deploy --templates -e glance-settings.yaml
```

For additional information about using environment files, see the [Including Environment Files in Overcloud Creation](#) section in the *Advanced Overcloud Customization Guide*.

19.5. CONFIGURING CINDER BACK END FOR THE IMAGE SERVICE

The **GlanceBackend** parameter sets the back end that the Image service uses to store images.



IMPORTANT

The default maximum number of volumes you can create for a project is 10.

Procedure

1. To configure **cinder** as the Image service back end, add the following to the environment file:

```
parameter_defaults:
  GlanceBackend: cinder
```

2. If the **cinder** back end is enabled, the following parameters and values are set by default:

```
cinder_store_auth_address = http://172.17.1.19:5000/v3
cinder_store_project_name = service
cinder_store_user_name = glance
cinder_store_password = ****secret****
```

3. To use a custom user name, or any custom value for the **cinder_store** parameters, add the **ExtraConfig** settings to **parameter_defaults** and pass the custom values:

```
ExtraConfig:
  glance::config::api_config:
    glance_store/cinder_store_auth_address:
      value: "%{hiera('glance::api::authtoken::auth_url')}/v3"
    glance_store/cinder_store_user_name:
      value: <user-name>
```

```
glance_store/cinder_store_password:
  value: "%{hiera('glance::api::authtoken::password')}}"
glance_store/cinder_store_project_name:
  value: "%{hiera('glance::api::authtoken::project_name')}}"
```

19.6. CONFIGURING THE MAXIMUM NUMBER OF STORAGE DEVICES TO ATTACH TO ONE INSTANCE

By default, you can attach an unlimited number of storage devices to a single instance. To limit the maximum number of devices, add the **max_disk_devices_to_attach** parameter to your Compute environment file. The following example shows how to change the value of **max_disk_devices_to_attach** to "30":

```
parameter_defaults:
  ComputeExtraConfig:
    nova::config::nova_config:
      compute/max_disk_devices_to_attach:
        value: '30'
```

Guidelines and considerations

- The number of storage disks supported by an instance depends on the bus that the disk uses. For example, the IDE disk bus is limited to 4 attached devices.
- Changing the **max_disk_devices_to_attach** on a Compute node with active instances can cause rebuilds to fail if the maximum number is lower than the number of devices already attached to instances. For example, if instance **A** has 26 devices attached and you change **max_disk_devices_to_attach** to 20, a request to rebuild instance **A** will fail.
- During cold migration, the configured maximum number of storage devices is only enforced on the source for the instance that you want to migrate. The destination is not checked before the move. This means that if Compute node A has 26 attached disk devices, and Compute node B has a configured maximum of 20 attached disk devices, a cold migration of an instance with 26 attached devices from Compute node A to Compute node B succeeds. However, a subsequent request to rebuild the instance in Compute node B fails because 26 devices are already attached which exceeds the configured maximum of 20.
- The configured maximum is not enforced on shelved offloaded instances, as they have no Compute node.
- Attaching a large number of disk devices to instances can degrade performance on the instance. You should tune the maximum number based on the boundaries of what your environment can support.
- Instances with machine type Q35 can attach a maximum of 500 disk devices.

19.7. CONFIGURING THIRD PARTY STORAGE

The director include a couple of environment files to help configure third-party storage providers. This includes:

Dell EMC Storage Center

Deploys a single Dell EMC Storage Center back end for the Block Storage (cinder) service.

The environment file is located at **`/usr/share/openstack-tripleo-heat-templates/environments/cinder-dellsc-config.yaml`**.

See the [Dell Storage Center Back End Guide](#) for full configuration information.

Dell EMC PS Series

Deploys a single Dell EMC PS Series back end for the Block Storage (cinder) service.

The environment file is located at **`/usr/share/openstack-tripleo-heat-templates/environments/cinder-dellps-config.yaml`**.

See the [Dell EMC PS Series Back End Guide](#) for full configuration information.

NetApp Block Storage

Deploys a NetApp storage appliance as a back end for the Block Storage (cinder) service.

The environment file is located at **`/usr/share/openstack-tripleo-heat-templates/environments/storage/cinder-netapp-config.yaml`**.

See the [NetApp Block Storage Back End Guide](#) for full configuration information.

CHAPTER 20. SECURITY ENHANCEMENTS

The following sections provide some suggestions to harden the security of your overcloud.

20.1. MANAGING THE OVERCLOUD FIREWALL

Each of the core OpenStack Platform services contains firewall rules in their respective composable service templates. This automatically creates a default set of firewall rules for each overcloud node.

The overcloud Heat templates contain a set of parameters to help with additional firewall management:

ManageFirewall

Defines whether to automatically manage the firewall rules. Set to **true** to allow Puppet to automatically configure the firewall on each node. Set to **false** if you want to manually manage the firewall. The default is **true**.

PurgeFirewallRules

Defines whether to purge the default Linux firewall rules before configuring new ones. The default is **false**.

If **ManageFirewall** is set to **true**, you can create additional firewall rules on deployment. Set the **tripleo::firewall::firewall_rules** hieradata using a configuration hook (see [Section 4.5, “Puppet: Customizing Hieradata for Roles”](#)) in an environment file for your overcloud. This hieradata is a hash containing the firewall rule names and their respective parameters as keys, all of which are optional:

port

The port associated to the rule.

dport

The destination port associated to the rule.

sport

The source port associated to the rule.

proto

The protocol associated to the rule. Defaults to **tcp**.

action

The action policy associated to the rule. Defaults to **accept**.

jump

The chain to jump to. If present, it overrides **action**.

state

An Array of states associated to the rule. Defaults to **['NEW']**.

source

The source IP address associated to the rule.

iface

The network interface associated to the rule.

chain

The chain associated to the rule. Defaults to **INPUT**.

destination

The destination CIDR associated to the rule.

The following example demonstrates the syntax of the firewall rule format:

```
ExtraConfig:
  tripleo::firewall::firewall_rules:
    '300 allow custom application 1':
      port: 999
      proto: udp
      action: accept
    '301 allow custom application 2':
      port: 8081
      proto: tcp
      action: accept
```

This applies two additional firewall rules to all nodes through **ExtraConfig**.



NOTE

Each rule name becomes the comment for the respective **iptables** rule. Note also each rule name starts with a three-digit prefix to help Puppet order all defined rules in the final **iptables** file. The default OpenStack Platform rules use prefixes in the 000 to 200 range.

20.2. CHANGING THE SIMPLE NETWORK MANAGEMENT PROTOCOL (SNMP) STRINGS

The director provides a default read-only SNMP configuration for your overcloud. It is advisable to change the SNMP strings to mitigate the risk of unauthorized users learning about your network devices.

Set the following hieradata using the **ExtraConfig** hook in an environment file for your overcloud:

SNMP traditional access control settings

snmp::ro_community

IPv4 read-only SNMP community string. The default value is **public**.

snmp::ro_community6

IPv6 read-only SNMP community string. The default value is **public**.

snmp::ro_network

Network that is allowed to **RO query** the daemon. This value can be a string or an array. Default value is **127.0.0.1**.

snmp::ro_network6

Network that is allowed to **RO query** the daemon with IPv6. This value can be a string or an array. The default value is **::1/128**.

tripleo::profile::base::snmp::snmpd_config

Array of lines to add to the *snmpd.conf* file as a safety valve. The default value is `[]`. See the [SNMP Configuration File](#) web page for all available options.

For example:

```
parameter_defaults:
  ExtraConfig:
    snmp::ro_community: mysecurestring
    snmp::ro_community6: myv6securestring
```

This changes the read-only SNMP community string on all nodes.

SNMP view-based access control settings (VACM)

`snmp::com2sec`

IPv4 security name.

`snmp::com2sec6`

IPv6 security name.

For example:

```
parameter_defaults:
  ExtraConfig:
    snmp::com2sec: mysecurestring
    snmp::com2sec6: myv6securestring
```

This changes the read-only SNMP community string on all nodes.

For more information, see the `snmpd.conf` man page.

20.3. CHANGING THE SSL/TLS CIPHER AND RULES FOR HAPROXY

If you enabled SSL/TLS in the overcloud (see [Chapter 15, Enabling SSL/TLS on Overcloud Public Endpoints](#)), you might want to harden the SSL/TLS ciphers and rules used with the HAProxy configuration. This helps avoid SSL/TLS vulnerabilities, such as the [POODLE vulnerability](#).

Set the following hieradata using the **ExtraConfig** hook in an environment file for your overcloud:

`tripleo::haproxy::ssl_cipher_suite`

The cipher suite to use in HAProxy.

`tripleo::haproxy::ssl_options`

The SSL/TLS rules to use in HAProxy.

For example, you might aim to use the following cipher and rules:

- Cipher: **ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-POLY1305:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:DHE-RSA-AES128-GCM-SHA256:DHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES128-SHA256:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-AES256-SHA384:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES256-SHA384:ECDHE-ECDSA-AES256-SHA:ECDHE-RSA-AES256-SHA:DHE-RSA-AES128-SHA256:DHE-RSA-AES128-SHA:DHE-RSA-AES256-SHA256:DHE-RSA-AES256-SHA:ECDHE-ECDSA-DES-CBC3-SHA:ECDHE-RSA-DES-CBC3-SHA:EDH-RSA-DES-CBC3-SHA:AES128-GCM-SHA256:AES256-GCM-SHA384:AES128-SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:DES-CBC3-SHA:!DSS**
- Rules: **no-sslv3 no-tls-tickets**

Create an environment file with the following content:

```
parameter_defaults:
  ExtraConfig:
```

```
tripleo::haproxy::ssl_cipher_suite: ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-
CHACHA20-POLY1305:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-
SHA256:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:DHE-RSA-
AES128-GCM-SHA256:DHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES128-
SHA256:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-AES256-
SHA384:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES256-SHA384:ECDHE-ECDSA-AES256-
SHA:ECDHE-RSA-AES256-SHA:DHE-RSA-AES128-SHA256:DHE-RSA-AES128-SHA:DHE-RSA-
AES256-SHA256:DHE-RSA-AES256-SHA:ECDHE-ECDSA-DES-CBC3-SHA:ECDHE-RSA-DES-
CBC3-SHA:EDH-RSA-DES-CBC3-SHA:AES128-GCM-SHA256:AES256-GCM-SHA384:AES128-
SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:DES-CBC3-SHA:!DSS
tripleo::haproxy::ssl_options: no-sslv3 no-tls-tickets
```



NOTE

The cipher collection is one continuous line.


Include this environment file with your overcloud creation.

20.4. USING THE OPEN VSWITCH FIREWALL

You can configure security groups to use the Open vSwitch (OVS) firewall driver in Red Hat OpenStack Platform director. The **NeutronOVSFirewallDriver** parameter allows you to specify which firewall driver to use:

- **iptables_hybrid** - Configures neutron to use the iptables/hybrid based implementation.
- **opendvswitch** - Configures neutron to use the OVS firewall flow-based driver.

The **opendvswitch** firewall driver includes higher performance and reduces the number of interfaces and bridges used to connect guests to the project network.



IMPORTANT

Multicast traffic is handled differently by the Open vSwitch (OVS) firewall driver than by the iptables firewall driver. With iptables, by default, VRRP traffic is denied, and you must enable VRRP in the security group rules for any VRRP traffic to reach an endpoint. With OVS, all ports share the same OpenFlow context, and multicast traffic cannot be processed individually per port. Because security groups do not apply to all ports (for example, the ports on a router), OVS uses the **NORMAL** action and forwards multicast traffic to all ports as specified by RFC 4541.



NOTE

The **iptables_hybrid** option is not compatible with OVS-DPDK.

Configure the **NeutronOVSFirewallDriver** parameter in the **network-environment.yaml** file:

```
NeutronOVSFirewallDriver: opendvswitch
```

- **NeutronOVSFirewallDriver** : Configures the name of the firewall driver to use when implementing security groups. Possible values depend on your system configuration. Examples include: **noop**, **opendvswitch**, **iptables_hybrid**. The default value, an empty string, equates to **iptables_hybrid**.

20.5. USING SECURE ROOT USER ACCESS

The overcloud image automatically contains hardened security for the **root** user. For example, each deployed overcloud node automatically disables direct SSH access to the **root** user. You can still access the **root** user on overcloud nodes through the following method:

1. Log into the undercloud node's **stack** user.
2. Each overcloud node has a **heat-admin** user account. This user account contains the undercloud's public SSH key, which provides SSH access without a password from the undercloud to the overcloud node. On the undercloud node, log into the chosen overcloud node through SSH using the **heat-admin** user.
3. Switch to the **root** user with **sudo -i**.

Reducing Root User Security

Some situations might require direct SSH access to the **root** user. In this case, you can reduce the SSH restrictions on the **root** user for each overcloud node.



WARNING

This method is intended for debugging purposes only. It is not recommended for use in a production environment.

The method uses the first boot configuration hook (see [Section 4.1, "First Boot: Customizing First Boot Configuration"](#)). Place the following content in an environment file:

```
resource_registry:
  OS::TripleO::NodeUserData: /usr/share/openstack-tripleo-heat-
  templates/firstboot/userdata_root_password.yaml

parameter_defaults:
  NodeRootPassword: "p@55w0rd!"
```

Note the following:

- The **OS::TripleO::NodeUserData** resource refers to the a template that configures the **root** user during the first boot **cloud-init** stage.
- The **NodeRootPassword** parameter sets the password for the **root** user. Change the value of this parameter to your desired password. Note the environment file contains the password as a plain text string, which is considered a security risk.

Include this environment file with the **openstack overcloud deploy** command when creating your overcloud.

CHAPTER 21. CONFIGURING NETWORK PLUGINS

The director includes environment files to help configure third-party network plugins:

21.1. FUJITSU CONVERGED FABRIC (C-FABRIC)

You can enable the Fujitsu Converged Fabric (C-Fabric) plugin using the environment file located at **/usr/share/openstack-tripleo-heat-templates/environments/neutron-ml2-fujitsu-cfab.yaml**.

1. Copy the environment file to your **templates** subdirectory:

```
$ cp /usr/share/openstack-tripleo-heat-templates/environments/neutron-ml2-fujitsu-cfab.yaml
/home/stack/templates/
```

2. Edit the **resource_registry** to use an absolute path:

```
resource_registry:
  OS::TripleO::Services::NeutronML2FujitsuCfab: /usr/share/openstack-tripleo-heat-
  templates/puppet/services/neutron-plugin-ml2-fujitsu-cfab.yaml
```

3. Review the **parameter_defaults** in **/home/stack/templates/neutron-ml2-fujitsu-cfab.yaml**:

- **NeutronFujitsuCfabAddress** - The telnet IP address of the C-Fabric. (string)
- **NeutronFujitsuCfabUserName** - The C-Fabric username to use. (string)
- **NeutronFujitsuCfabPassword** - The password of the C-Fabric user account. (string)
- **NeutronFujitsuCfabPhysicalNetworks** - List of **<physical_network>:<vfab_id>** tuples that specify **physical_network** names and their corresponding vfab IDs. (comma_delimited_list)
- **NeutronFujitsuCfabSharePprofile** - Determines whether to share a C-Fabric pprofile among neutron ports that use the same VLAN ID. (boolean)
- **NeutronFujitsuCfabPprofilePrefix** - The prefix string for pprofile name. (string)
- **NeutronFujitsuCfabSaveConfig** - Determines whether to save the configuration. (boolean)

4. To apply the template to your deployment, include the environment file in the **openstack overcloud deploy** command. For example:

```
$ openstack overcloud deploy --templates -e /home/stack/templates/neutron-ml2-fujitsu-
cfab.yaml [OTHER OPTIONS] ...
```

21.2. FUJITSU FOS SWITCH

You can enable the Fujitsu FOS Switch plugin using the environment file located at **/usr/share/openstack-tripleo-heat-templates/environments/neutron-ml2-fujitsu-fossw.yaml**.

1. Copy the environment file to your **templates** subdirectory:

```
$ cp /usr/share/openstack-tripleo-heat-templates/environments/neutron-ml2-fujitsu-fossw.yaml /home/stack/templates/
```

2. Edit the **resource_registry** to use an absolute path:

```
resource_registry:  
  OS::TripleO::Services::NeutronML2FujitsuFossw: /usr/share/openstack-tripleo-heat-templates/puppet/services/neutron-plugin-ml2-fujitsu-fossw.yaml
```

3. Review the **parameter_defaults** in **/home/stack/templates/neutron-ml2-fujitsu-fossw.yaml**:
 - **NeutronFujitsuFosswIps** - The IP addresses of all FOS switches. (comma_delimited_list)
 - **NeutronFujitsuFosswUserName** - The FOS username to use. (string)
 - **NeutronFujitsuFosswPassword** - The password of the FOS user account. (string)
 - **NeutronFujitsuFosswPort** - The port number to use for the SSH connection. (number)
 - **NeutronFujitsuFosswTimeout** - The timeout period of the SSH connection. (number)
 - **NeutronFujitsuFosswUdpDestPort** - The port number of the VXLAN UDP destination on the FOS switches. (number)
 - **NeutronFujitsuFosswOvsdbVlanidRangeMin** - The minimum VLAN ID in the range that is used for binding VNI and physical port. (number)
 - **NeutronFujitsuFosswOvsdbPort** - The port number for the OVSDDB server on the FOS switches. (number)
4. To apply the template to your deployment, include the environment file in the **openstack overcloud deploy** command. For example:

```
$ openstack overcloud deploy --templates -e /home/stack/templates/neutron-ml2-fujitsu-fossw.yaml [OTHER OPTIONS] ...
```


CHAPTER 22. CONFIGURING IDENTITY

The director includes parameters to help configure Identity Service (keystone) settings:

22.1. REGION NAME

By default, your overcloud's region will be named **regionOne**. You can change this by adding a **KeystoneRegion** entry your environment file. This setting cannot be changed post-deployment:

```
parameter_defaults:  
  KeystoneRegion: 'SampleRegion'
```

CHAPTER 23. OTHER CONFIGURATIONS

23.1. CONFIGURING THE KERNEL ON OVERCLOUD NODES

OpenStack Platform director includes parameters that configure the kernel on overcloud nodes.

ExtraKernelModules

Kernel modules to load. The modules names are listed as a hash key with an empty value:

```
ExtraKernelModules:
  <MODULE_NAME>: {}
```

ExtraKernelPackages

Kernel-related packages to install prior to loading the kernel modules from **ExtraKernelModules**. The package names are listed as a hash key with an empty value.

```
ExtraKernelPackages:
  <PACKAGE_NAME>: {}
```

ExtraSysctlSettings

Hash of sysctl settings to apply. Set the value of each parameter using the **value** key.

```
ExtraSysctlSettings:
  <KERNEL_PARAMETER>:
    value: <VALUE>
```

This example shows the syntax of these parameters in an environment file:

```
parameter_defaults:
  ExtraKernelModules:
    iscsi_target_mod: {}
  ExtraKernelPackages:
    iscsi-initiator-utils: {}
  ExtraSysctlSettings:
    dev.scsi.logging_level:
      value: 1
```

23.2. CONFIGURING EXTERNAL LOAD BALANCING

An Overcloud uses multiple Controllers together as a high availability cluster, which ensures maximum operational performance for your OpenStack services. In addition, the cluster provides load balancing for access to the OpenStack services, which evenly distributes traffic to the Controller nodes and reduces server overload for each node. It is also possible to use an external load balancer to perform this distribution. For example, an organization might use their own hardware-based load balancer to handle traffic distribution to the Controller nodes.

For more information about configuring external load balancing, see the dedicated [External Load Balancing for the Overcloud](#) guide for full instructions.

23.3. CONFIGURING IPV6 NETWORKING

As a default, the Overcloud uses Internet Protocol version 4 (IPv4) to configure the service endpoints. However, the Overcloud also supports Internet Protocol version 6 (IPv6) endpoints, which is useful for organizations that support IPv6 infrastructure. The director includes a set of environment files to help with creating IPv6-based Overclouds.

For more information about configuring IPv6 in the Overcloud, see the dedicated [IPv6 Networking for the Overcloud](#) guide for full instructions.