


01. Openstack

Opensack

1.

1. OS : Rocky Linux 8.4 (RHEL / Centos 8)
2. Openstack : Wallaby

2. Component

1. Node list

h	p	p	C
o	u	r	o
s	b	i	m
t	l	v	p
n	i	a	o
a	c	t	n
m	i	e	e
e	p	i	n
		p	t
			r
			u
			l
			e

o	1	1	c
p	7	9	o
e	2	2	n
n	.	.	t
s	1	1	r
t	0	6	o
a	.	8	l
c	1	.	,
k	0	2	c
-	.	0	o
d	.	.	m
e	4	1	p
v	1	1	u
1			t
			e
o	1	1	c
p	7	9	o
e	2	2	m
n	.	.	p
s	1	1	u
t	0	6	t
a	.	8	e
c	1	.	
k	0	2	
-	.	0	
d	.	.	
e	4	1	
v	2	2	
2			

o	1	1	n
p	7	9	e
e	2	2	t
n	.	.	w
s	1	1	o
t	0	6	r
a	.	8	k
c	1	.	,
k	0	2	s
-	.	0	t
d	.	.	o
e	4	1	r
v	3	3	a
3			g
			e

2. Network Config

1. public : extnet - br-ex - ens3f0
2. private : physnet1 - br-vlan - ens3f1

3. Packstack

1.  

```
$ yum install glibc-langpack-en -y

$ vi /etc/environment
LANG=en_US.utf-8
LC_ALL=en_US.utf-8
```

2. NetworkManager / Firewall

```
$ dnf install network-scripts -y
$ systemctl disable firewalld
$ systemctl stop firewalld
$ systemctl disable NetworkManager
$ systemctl stop NetworkManager
$ systemctl enable network --now
```

```
$ systemctl start network
```

3. ☐☐☐☐☐ ☐☐ ☐ ☐☐☐ ☐☐

```
$ dnf config-manager --enable powertools
$ dnf install -y https://www.rdoproject.org/repos/rdo-release.el8.rpm
$ dnf update -y
$ dnf install -y openstack-packstack
```

4. packstack ☐☐☐ ☐☐ ☐☐☐☐ ☐☐

```
$ packstack --gen-answer-file=answer.txt
```

5. Sample ☐☐ - PW☐☐ ☐☐☐ {{ SET_PASS }} ☐☐ ☐☐☐☐

```
$ vi answer.txt
[general]

# Path to a public key to install on servers. If a usable key has not
# been installed on the remote servers, the user is prompted for a
# password and this key is installed so the password will not be
# required again.
CONFIG_SSH_KEY=/root/.ssh/id_rsa.pub

# Default password to be used everywhere (overridden by passwords set
# for individual services or users).
CONFIG_DEFAULT_PASSWORD=

# The amount of service workers/threads to use for each service.
# Useful to tweak when you have memory constraints. Defaults to the
# amount of cores on the system.
CONFIG_SERVICE_WORKERS=%{::processorcount}

# Specify 'y' to install MariaDB. ['y', 'n']
CONFIG_MARIADB_INSTALL=y

# Specify 'y' to install OpenStack Image Service (glance). ['y', 'n']
```

```
CONFIG_GLANCE_INSTALL=y
```

```
# Specify 'y' to install OpenStack Block Storage (cinder). ['y', 'n']
```

```
CONFIG_CINDER_INSTALL=y
```

```
# Specify 'y' to install OpenStack Shared File System (manila). ['y',
```

```
# 'n']
```

```
CONFIG_MANILA_INSTALL=n
```

```
# Specify 'y' to install OpenStack Compute (nova). ['y', 'n']
```

```
CONFIG_NOVA_INSTALL=y
```

```
# Specify 'y' to install OpenStack Networking (neutron) ['y']
```

```
CONFIG_NEUTRON_INSTALL=y
```

```
# Specify 'y' to install OpenStack Dashboard (horizon). ['y', 'n']
```

```
CONFIG_HORIZON_INSTALL=y
```

```
# Specify 'y' to install OpenStack Object Storage (swift). ['y', 'n']
```

```
CONFIG_SWIFT_INSTALL=y
```

```
# Specify 'y' to install OpenStack Metering (ceilometer). Note this
```

```
# will also automatically install gnocchi service and configures it as
```

```
# the metrics backend. ['y', 'n']
```

```
CONFIG_CEILOMETER_INSTALL=n
```

```
# Specify 'y' to install OpenStack Telemetry Alarming (Aodh). Note
```

```
# Aodh requires Ceilometer to be installed as well. ['y', 'n']
```

```
CONFIG_AODH_INSTALL=y
```

```
# Specify 'y' to install OpenStack Events Service (panko). ['y', 'n']
```

```
CONFIG_PANKO_INSTALL=n
```

```
# Specify 'y' to install OpenStack Data Processing (sahara). In case
```

```
# of sahara installation packstack also installs heat. ['y', 'n']
```

```
CONFIG_SAHARA_INSTALL=n
```

```
# Specify 'y' to install OpenStack Orchestration (heat). ['y', 'n']
```

```
CONFIG_HEAT_INSTALL=y
```

```
# Specify 'y' to install OpenStack Container Infrastructure
# Management Service (magnum). ['y', 'n']
CONFIG_MAGNUM_INSTALL=n

# Specify 'y' to install OpenStack Database (trove) ['y', 'n']
CONFIG_TROVE_INSTALL=n

# Specify 'y' to install OpenStack Bare Metal Provisioning (ironic).
# ['y', 'n']
CONFIG_IRONIC_INSTALL=n

# Specify 'y' to install the OpenStack Client packages (command-line
# tools). An admin "rc" file will also be installed. ['y', 'n']
CONFIG_CLIENT_INSTALL=y

# Comma-separated list of NTP servers. Leave plain if Packstack
# should not install ntpd on instances.
CONFIG_NTP_SERVERS= time.bora.net

# Comma-separated list of servers to be excluded from the
# installation. This is helpful if you are running Packstack a second
# time with the same answer file and do not want Packstack to
# overwrite these server's configurations. Leave empty if you do not
# need to exclude any servers.
EXCLUDE_SERVERS=

# Specify 'y' if you want to run OpenStack services in debug mode;
# otherwise, specify 'n'. ['y', 'n']
CONFIG_DEBUG_MODE=n

# Server on which to install OpenStack services specific to the
# controller role (for example, API servers or dashboard).
CONFIG_CONTROLLER_HOST=172.10.10..41

# List the servers on which to install the Compute service.
CONFIG_COMPUTE_HOSTS=172.10.10..41,172.10.10..42

# List of servers on which to install the network service such as
```

Compute networking (nova network) or OpenStack Networking (neutron).

CONFIG_NETWORK_HOSTS=172.10.10..43

Specify 'y' if you want to use VMware vCenter as hypervisor and

storage; otherwise, specify 'n'. ['y', 'n']

CONFIG_VMWARE_BACKEND=n

Specify 'y' if you want to use unsupported parameters. This should

be used only if you know what you are doing. Issues caused by using

unsupported options will not be fixed before the next major release.

['y', 'n']

CONFIG_UNSUPPORTED=n

Specify 'y' if you want to use subnet addresses (in CIDR format)

instead of interface names in following options:

CONFIG_NEUTRON_OVS_BRIDGE_IFACES,

CONFIG_NEUTRON_LB_INTERFACE_MAPPINGS, CONFIG_NEUTRON_OVS_TUNNEL_IF.

This is useful for cases when interface names are not same on all

installation hosts.

CONFIG_USE_SUBNETS=n

IP address of the VMware vCenter server.

CONFIG_VCENTER_HOST=

User name for VMware vCenter server authentication.

CONFIG_VCENTER_USER=

Password for VMware vCenter server authentication.

CONFIG_VCENTER_PASSWORD=

Comma separated list of names of the VMware vCenter clusters. Note:

if multiple clusters are specified each one is mapped to one

compute, otherwise all computes are mapped to same cluster.

CONFIG_VCENTER_CLUSTER_NAMES=

(Unsupported!) Server on which to install OpenStack services

specific to storage servers such as Image or Block Storage services.

CONFIG_STORAGE_HOST=172.10.10..41

```
# (Unsupported!) Server on which to install OpenStack services
# specific to OpenStack Data Processing (sahara).
CONFIG_SAHARA_HOST=172.10.10..41

# Comma-separated list of URLs for any additional yum repositories,
# to use for installation.
CONFIG_REPO=

# Specify 'y' to enable the RDO testing repository. ['y', 'n']
CONFIG_ENABLE_RDO_TESTING=n

# To subscribe each server with Red Hat Subscription Manager, include
# this with CONFIG_RH_PW.
CONFIG_RH_USER=

# To subscribe each server to receive updates from a Satellite
# server, provide the URL of the Satellite server. You must also
# provide a user name (CONFIG_SATELLITE_USERNAME) and password
# (CONFIG_SATELLITE_PASSWORD) or an access key (CONFIG_SATELLITE_AKEY)
# for authentication.
CONFIG_SATELLITE_URL=

# Specify a Satellite 6 Server to register to. If not specified,
# Packstack will register the system to the Red Hat server. When this
# option is specified, you also need to set the Satellite 6
# organization (CONFIG_RH_SAT6_ORG) and an activation key
# (CONFIG_RH_SAT6_KEY).
CONFIG_RH_SAT6_SERVER=

# To subscribe each server with Red Hat Subscription Manager, include
# this with CONFIG_RH_USER.
CONFIG_RH_PW=

# Specify 'y' to enable RHEL optional repositories. ['y', 'n']
CONFIG_RH_OPTIONAL=y

# HTTP proxy to use with Red Hat Subscription Manager.
CONFIG_RH_PROXY=
```


Specify a Satellite 6 Server organization to use when registering
the system.

CONFIG_RH_SAT6_ORG=

Specify a Satellite 6 Server activation key to use when registering
the system.

CONFIG_RH_SAT6_KEY=

Port to use for Red Hat Subscription Manager's HTTP proxy.

CONFIG_RH_PROXY_PORT=

User name to use for Red Hat Subscription Manager's HTTP proxy.

CONFIG_RH_PROXY_USER=

Password to use for Red Hat Subscription Manager's HTTP proxy.

CONFIG_RH_PROXY_PW=

User name to authenticate with the RHN Satellite server; if you
intend to use an access key for Satellite authentication, leave this
blank.

CONFIG_SATELLITE_USER=

Password to authenticate with the RHN Satellite server; if you
intend to use an access key for Satellite authentication, leave this
blank.

CONFIG_SATELLITE_PW=

Access key for the Satellite server; if you intend to use a user
name and password for Satellite authentication, leave this blank.

CONFIG_SATELLITE_AKEY=

Certificate path or URL of the certificate authority to verify that
the connection with the Satellite server is secure. If you are not
using Satellite in your deployment, leave this blank.

CONFIG_SATELLITE_CACERT=

Profile name that should be used as an identifier for the system in
RHN Satellite (if required).

CONFIG_SATELLITE_PROFILE=

```
# Comma-separated list of flags passed to the rhnreg_ks command.
# Valid flags are: novirtinfo, norhnsd, nopackages ['novirtinfo',
# 'norhnsd', 'nopackages']
CONFIG_SATELLITE_FLAGS=

# HTTP proxy to use when connecting to the RHN Satellite server (if
# required).
CONFIG_SATELLITE_PROXY=

# User name to authenticate with the Satellite-server HTTP proxy.
CONFIG_SATELLITE_PROXY_USER=

# User password to authenticate with the Satellite-server HTTP proxy.
CONFIG_SATELLITE_PROXY_PW=

# Specify filepath for CA cert file. If CONFIG_SSL_CACERT_SELFSIGN is
# set to 'n' it has to be preexisting file.
CONFIG_SSL_CACERT_FILE=/etc/pki/tls/certs/selfcert.crt

# Specify filepath for CA cert key file. If
# CONFIG_SSL_CACERT_SELFSIGN is set to 'n' it has to be preexisting
# file.
CONFIG_SSL_CACERT_KEY_FILE=/etc/pki/tls/private/selfkey.key

# Enter the path to use to store generated SSL certificates in.
CONFIG_SSL_CERT_DIR=~/.packstackca/

# Specify 'y' if you want Packstack to pregenerate the CA
# Certificate.
CONFIG_SSL_CACERT_SELFSIGN=y

# Enter the ssl certificates subject country.
CONFIG_SSL_CERT_SUBJECT_C=--

# Enter the ssl certificates subject state.
CONFIG_SSL_CERT_SUBJECT_ST=State

# Enter the ssl certificates subject location.
```

CONFIG_SSL_CERT_SUBJECT_L=City

Enter the ssl certificates subject organization.

CONFIG_SSL_CERT_SUBJECT_O=openstack

Enter the ssl certificates subject organizational unit.

CONFIG_SSL_CERT_SUBJECT_OU=packstack

Enter the ssl certificates subject common name.

CONFIG_SSL_CERT_SUBJECT_CN=openstack-dev1

CONFIG_SSL_CERT_SUBJECT_MAIL=admin@openstack-dev1

Service to be used as the AMQP broker. Allowed values are: rabbitmq

['rabbitmq']

CONFIG_AMQP_BACKEND=rabbitmq

IP address of the server on which to install the AMQP service.

CONFIG_AMQP_HOST=172.10.10..41

Specify 'y' to enable SSL for the AMQP service. ['y', 'n']

CONFIG_AMQP_ENABLE_SSL=n

Specify 'y' to enable authentication for the AMQP service. ['y',

'n']

CONFIG_AMQP_ENABLE_AUTH=n

Password for the NSS certificate database of the AMQP service.

CONFIG_AMQP_NSS_CERT

_PW=PW_PLACEHOLDER

User for AMQP authentication.

CONFIG_AMQP_AUTH_USER=amqp_user

Password for AMQP authentication.

CONFIG_AMQP_AUTH_PASSWORD=PW_PLACEHOLDER

IP address of the server on which to install MariaDB. If a MariaDB
installation was not specified in CONFIG_MARIADB_INSTALL, specify
the IP address of an existing database server (a MariaDB cluster can
also be specified).

CONFIG_MARIADB_HOST=172.10.10..41

User name for the MariaDB administrative user.

CONFIG_MARIADB_USER=root

Password for the MariaDB administrative user.

CONFIG_MARIADB_PW= {{ SET_PASS }}

Password to use for the Identity service (keystone) to access the
database.

CONFIG_KEYSTONE_DB_PW= {{ SET_PASS }}

Enter y if cron job to rotate Fernet tokens should be created.

CONFIG_KEYSTONE_FERNET_TOKEN_ROTATE_ENABLE=True

Default region name to use when creating tenants in the Identity
service.

CONFIG_KEYSTONE_REGION=RegionOne

Token to use for the Identity service API.

CONFIG_KEYSTONE_ADMIN_TOKEN={{ SET_PASS }}

Email address for the Identity service 'admin' user. Defaults to

CONFIG_KEYSTONE_ADMIN_EMAIL=root@localhost

User name for the Identity service 'admin' user. Defaults to

'admin'.

CONFIG_KEYSTONE_ADMIN_USERNAME=admin

Password to use for the Identity service 'admin' user.

CONFIG_KEYSTONE_ADMIN_PW={{ SET_PASS }}

Password to use for the Identity service 'demo' user.

CONFIG_KEYSTONE_DEMO_PW={{ SET_PASS }}

```
# Identity service API version string. ['v2.0', 'v3']
CONFIG_KEYSTONE_API_VERSION=v3

# Identity service token format (FERNET). Since Rocky, only FERNET is
# supported. ['FERNET']
CONFIG_KEYSTONE_TOKEN_FORMAT=FERNET

# Type of Identity service backend (sql or ldap). ['sql', 'ldap']
CONFIG_KEYSTONE_IDENTITY_BACKEND=sql

# URL for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_URL=ldap://172.10.10..41

# User DN for the Identity service LDAP backend. Used to bind to the
# LDAP server if the LDAP server does not allow anonymous
# authentication.
CONFIG_KEYSTONE_LDAP_USER_DN=

# User DN password for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_USER_PASSWORD=

# Base suffix for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_SUFFIX=

# Query scope for the Identity service LDAP backend. Use 'one' for
# onelevel/singleLevel or 'sub' for subtree/wholeSubtree ('base' is
# not actually used by the Identity service and is therefore
# deprecated). ['base', 'one', 'sub']
CONFIG_KEYSTONE_LDAP_QUERY_SCOPE=one

# Query page size for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_PAGE_SIZE=-1

# User subtree for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_USER_SUBTREE=

# User query filter for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_USER_FILTER=
```

```
# User object class for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_USER_OBJECTCLASS=

# User ID attribute for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_USER_ID_ATTRIBUTE=

# User name attribute for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_USER_NAME_ATTRIBUTE=

# User email address attribute for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_USER_MAIL_ATTRIBUTE=

# User-enabled attribute for the Identity service LDAP backend.
CONFIG_KEYSTONE_LDAP_USER_ENABLED_ATTRIBUTE=

# Bit mask integer applied to user-enabled attribute for the Identity
# service LDAP backend. Indicate the bit that the enabled value is
# stored in if the LDAP server represents "enabled" as a bit on an
# integer rather than a boolean. A value of "0" indicates the mask is
# not used (default). If this is not set to "0", the typical value is
# "2", typically used when
# "CONFIG_KEYSTONE_LDAP_USER_ENABLED_ATTRIBUTE = userAccountControl".
CONFIG_KEYSTONE_LDAP_USER_ENABLED_MASK=-1

# Value of enabled attribute which indicates user is enabled for the
# Identity service LDAP backend. This should match an appropriate
# integer value if the LDAP server uses non-boolean (bitmask) values
# to indicate whether a user is enabled or disabled. If this is not
# set as 'y', the typical value is "512". This is typically used when
# "CONFIG_KEYSTONE_LDAP_USER_ENABLED_ATTRIBUTE = userAccountControl".
CONFIG_KEYSTONE_LDAP_USER_ENABLED_DEFAULT=TRUE

# Specify 'y' if users are disabled (not enabled) in the Identity
# service LDAP backend (inverts boolean-enabled values). Some LDAP
# servers use a boolean lock attribute where "y" means an account is
# disabled. Setting this to 'y' allows these lock attributes to be
# used. This setting will have no effect if
# "CONFIG_KEYSTONE_LDAP_USER_ENABLED_MASK" is in use. ['n', 'y']
```

CONFIG_KEYSTONE_LDAP_USER_ENABLED_INVERT=n

Comma-separated list of attributes stripped from LDAP user entry
upon update.

CONFIG_KEYSTONE_LDAP_USER_ATTRIBUTE_IGNORE=

Identity service LDAP attribute mapped to default_project_id for
users.

CONFIG_KEYSTONE_LDAP_USER_DEFAULT_PROJECT_ID_ATTRIBUTE=

Specify 'y' if you want to be able to create Identity service users
through the Identity service interface; specify 'n' if you will
create directly in the LDAP backend. ['n', 'y']

CONFIG_KEYSTONE_LDAP_USER_ALLOW_CREATE=n

Specify 'y' if you want to be able to update Identity service users
through the Identity service interface; specify 'n' if you will
update directly in the LDAP backend. ['n', 'y']

CONFIG_KEYSTONE_LDAP_USER_ALLOW_UPDATE=n

Specify 'y' if you want to be able to delete Identity service users
through the Identity service interface; specify 'n' if you will
delete directly in the LDAP backend. ['n', 'y']

CONFIG_KEYSTONE_LDAP_USER_ALLOW_DELETE=n

Identity service LDAP attribute mapped to password.

CONFIG_KEYSTONE_LDAP_USER_PASS_ATTRIBUTE=

DN of the group entry to hold enabled LDAP users when using enabled
emulation.

CONFIG_KEYSTONE_LDAP_USER_ENABLED_EMULATION_DN=

List of additional LDAP attributes for mapping additional attribute
mappings for users. The attribute-mapping format is
<ldap_attr>:<user_attr>, where ldap_attr is the attribute in the
LDAP entry and user_attr is the Identity API attribute.

CONFIG_KEYSTONE_LDAP_USER_ADDITIONAL_ATTRIBUTE_MAPPING=

Group subtree for the Identity service LDAP backend.

CONFIG_KEYSTONE_LDAP_GROUP_SUBTREE=

Group query filter for the Identity service LDAP backend.

CONFIG_KEYSTONE_LDAP_GROUP_FILTER=

Group object class for the Identity service LDAP backend.

CONFIG_KEYSTONE_LDAP_GROUP_OBJECTCLASS=

Group ID attribute for the Identity service LDAP backend.

CONFIG_KEYSTONE_LDAP_GROUP_ID_ATTRIBUTE=

Group name attribute for the Identity service LDAP backend.

CONFIG_KEYSTONE_LDAP_GROUP_NAME_ATTRIBUTE=

Group member attribute for the Identity service LDAP backend.

CONFIG_KEYSTONE_LDAP_GROUP_MEMBER_ATTRIBUTE=

Group description attribute for the Identity service LDAP backend.

CONFIG_KEYSTONE_LDAP_GROUP_DESC_ATTRIBUTE=

Comma-separated list of attributes stripped from LDAP group entry

upon update.

CONFIG_KEYSTONE_LDAP_GROUP_ATTRIBUTE_IGNORE=

Specify 'y' if you want to be able to create Identity service

groups through the Identity service interface; specify 'n' if you

will create directly in the LDAP backend. ['n', 'y']

CONFIG_KEYSTONE_LDAP_GROUP_ALLOW_CREATE=n

Specify 'y' if you want to be able to update Identity service

groups through the Identity service interface; specify 'n' if you

will update directly in the LDAP backend. ['n', 'y']

CONFIG_KEYSTONE_LDAP_GROUP_ALLOW_UPDATE=n

Specify 'y' if you want to be able to delete Identity service

groups through the Identity service interface; specify 'n' if you

will delete directly in the LDAP backend. ['n', 'y']

CONFIG_KEYSTONE_LDAP_GROUP_ALLOW_DELETE=n


```
# List of additional LDAP attributes used for mapping additional
# attribute mappings for groups. The attribute=mapping format is
# <ldap_attr>:<group_attr>, where ldap_attr is the attribute in the
# LDAP entry and group_attr is the Identity API attribute.
CONFIG_KEYSTONE_LDAP_GROUP_ADDITIONAL_ATTRIBUTE_MAPPING=

# Specify 'y' if the Identity service LDAP backend should use TLS.
# ['n', 'y']
CONFIG_KEYSTONE_LDAP_USE_TLS=n

# CA certificate directory for Identity service LDAP backend (if TLS
# is used).
CONFIG_KEYSTONE_LDAP_TLS_CACERTDIR=

# CA certificate file for Identity service LDAP backend (if TLS is
# used).
CONFIG_KEYSTONE_LDAP_TLS_CACERTFILE=

# Certificate-checking strictness level for Identity service LDAP
# backend; valid options are: never, allow, demand. ['never', 'allow',
# 'demand']
CONFIG_KEYSTONE_LDAP_TLS_REQ_CERT=demand

# Password to use for the Image service (glance) to access the
# database.
CONFIG_GLANCE_DB_PW={{ SET_PASS }}

# Password to use for the Image service to authenticate with the
# Identity service.
CONFIG_GLANCE_KS_PW={{ SET_PASS }}

# Storage backend for the Image service (controls how the Image
# service stores disk images). Valid options are: file or swift
# (Object Storage). The Object Storage service must be enabled to use
# it as a working backend; otherwise, Packstack falls back to 'file'.
# ['file', 'swift']
CONFIG_GLANCE_BACKEND=file

# Password to use for the Block Storage service (cinder) to access
```

```
# the database.
CONFIG_CINDER_DB_PW={{ SET_PASS }}

# Enter y if cron job for removing soft deleted DB rows should be
# created.
CONFIG_CINDER_DB_PURGE_ENABLE=True

# Password to use for the Block Storage service to authenticate with
# the Identity service.
CONFIG_CINDER_KS_PW={{ SET_PASS }}

# Storage backend to use for the Block Storage service; valid options
# are: lvm, gluster, nfs, vmdk, netapp, solidfire. ['lvm', 'gluster',
# 'nfs', 'vmdk', 'netapp', 'solidfire']
CONFIG_CINDER_BACKEND=nfs

# Specify 'y' to create the Block Storage volumes group. That is,
# Packstack creates a raw disk image in /var/lib/cinder, and mounts it
# using a loopback device. This should only be used for testing on a
# proof-of-concept installation of the Block Storage service (a file-
# backed volume group is not suitable for production usage). ['y',
# 'n']
CONFIG_CINDER_VOLUMES_CREATE=y

# Specify a custom name for the lvm cinder volume group
CONFIG_CINDER_VOLUME_NAME=cinder-volumes

# Size of Block Storage volumes group. Actual volume size will be
# extended with 3% more space for VG metadata. Remember that the size
# of the volume group will restrict the amount of disk space that you
# can expose to Compute instances, and that the specified amount must
# be available on the device used for /var/lib/cinder.
CONFIG_CINDER_VOLUMES_SIZE=20G

# A single or comma-separated list of Red Hat Storage (gluster)
# volume shares to mount. Example: 'ip-address:/vol-name', 'domain
# :/vol-name'
CONFIG_CINDER_GLUSTER_MOUNTS=
```

A single or comma-separated list of NFS exports to mount. Example:

'ip-address:/export-name'

CONFIG_CINDER_NFS_MOUNTS=192.168.20.13:/data

Administrative user account name used to access the NetApp storage

system or proxy server.

CONFIG_CINDER_NETAPP_LOGIN=

Password for the NetApp administrative user account specified in

the CONFIG_CINDER_NETAPP_LOGIN parameter.

CONFIG_CINDER_NETAPP_PASSWORD=

Hostname (or IP address) for the NetApp storage system or proxy

server.

CONFIG_CINDER_NETAPP_HOSTNAME=

The TCP port to use for communication with the storage system or

proxy. If not specified, Data ONTAP drivers will use 80 for HTTP and

443 for HTTPS; E-Series will use 8080 for HTTP and 8443 for HTTPS.

Defaults to 80.

CONFIG_CINDER_NETAPP_SERVER_PORT=80

Storage family type used on the NetApp storage system; valid

options are ontap_7mode for using Data ONTAP operating in 7-Mode,

ontap_cluster for using clustered Data ONTAP, or E-Series for NetApp

E-Series. Defaults to ontap_cluster. ['ontap_7mode',

'ontap_cluster', 'eseries']

CONFIG_CINDER_NETAPP_STORAGE_FAMILY=ontap_cluster

The transport protocol used when communicating with the NetApp

storage system or proxy server. Valid values are http or https.

Defaults to 'http'. ['http', 'https']

CONFIG_CINDER_NETAPP_TRANSPORT_TYPE=http

Storage protocol to be used on the data path with the NetApp

storage system; valid options are iscsi, fc, nfs. Defaults to nfs.

['iscsi', 'fc', 'nfs']

CONFIG_CINDER_NETAPP_STORAGE_PROTOCOL=nfs

Quantity to be multiplied by the requested volume size to ensure
enough space is available on the virtual storage server (Vserver) to
fulfill the volume creation request. Defaults to 1.0.

CONFIG_CINDER_NETAPP_SIZE_MULTIPLIER=1.0

Time period (in minutes) that is allowed to elapse after the image
is last accessed, before it is deleted from the NFS image cache.
When a cache-cleaning cycle begins, images in the cache that have
not been accessed in the last M minutes, where M is the value of
this parameter, are deleted from the cache to create free space on
the NFS share. Defaults to 720.

CONFIG_CINDER_NETAPP_EXPIRY_THRES_MINUTES=720

If the percentage of available space for an NFS share has dropped
below the value specified by this parameter, the NFS image cache is
cleaned. Defaults to 20.

CONFIG_CINDER_NETAPP_THRES_AVL_SIZE_PERC_START=20

When the percentage of available space on an NFS share has reached
the percentage specified by this parameter, the driver stops
clearing files from the NFS image cache that have not been accessed
in the last M minutes, where M is the value of the
CONFIG_CINDER_NETAPP_EXPIRY_THRES_MINUTES parameter. Defaults to 60.

CONFIG_CINDER_NETAPP_THRES_AVL_SIZE_PERC_STOP=60

Single or comma-separated list of NetApp NFS shares for Block
Storage to use. Format: ip-address:/export-name. Defaults to ''.

CONFIG_CINDER_NETAPP_NFS_SHARES=

File with the list of available NFS shares. Defaults to
'/etc/cinder/shares.conf'.

CONFIG_CINDER_NETAPP_NFS_SHARES_CONFIG=/etc/cinder/shares.conf

This parameter is only utilized when the storage protocol is
configured to use iSCSI or FC. This parameter is used to restrict
provisioning to the specified controller volumes. Specify the value
of this parameter to be a comma separated list of NetApp controller
volume names to be used for provisioning. Defaults to ''.

CONFIG_CINDER_NETAPP_VOLUME_LIST=

The vFiler unit on which provisioning of block storage volumes will
be done. This parameter is only used by the driver when connecting
to an instance with a storage family of Data ONTAP operating in
7-Mode Only use this parameter when utilizing the MultiStore feature
on the NetApp storage system. Defaults to ''.

CONFIG_CINDER_NETAPP_VFILER=

The name of the config.conf stanza for a Data ONTAP (7-mode) HA
partner. This option is only used by the driver when connecting to
an instance with a storage family of Data ONTAP operating in 7-Mode,
and it is required if the storage protocol selected is FC. Defaults
to ''.

CONFIG_CINDER_NETAPP_PARTNER_BACKEND_NAME=

This option specifies the virtual storage server (Vserver) name on
the storage cluster on which provisioning of block storage volumes
should occur. Defaults to ''.

CONFIG_CINDER_NETAPP_VSERVER=

Restricts provisioning to the specified controllers. Value must be
a comma-separated list of controller hostnames or IP addresses to be
used for provisioning. This option is only utilized when the storage
family is configured to use E-Series. Defaults to ''.

CONFIG_CINDER_NETAPP_CONTROLLER_IPS=

Password for the NetApp E-Series storage array. Defaults to ''.

CONFIG_CINDER_NETAPP_SA_PASSWORD=

This option is used to define how the controllers in the E-Series
storage array will work with the particular operating system on the
hosts that are connected to it. Defaults to 'linux_dm_mp'

CONFIG_CINDER_NETAPP_ESERIES_HOST_TYPE=linux_dm_mp

Path to the NetApp E-Series proxy application on a proxy server.

The value is combined with the value of the

CONFIG_CINDER_NETAPP_TRANSPORT_TYPE, CONFIG_CINDER_NETAPP_HOSTNAME,

and CONFIG_CINDER_NETAPP_HOSTNAME options to create the URL used by

the driver to connect to the proxy application. Defaults to

```
# '/devmgr/v2'.
CONFIG_CINDER_NETAPP_WEBSERVICE_PATH=/devmgr/v2

# Restricts provisioning to the specified storage pools. Only dynamic
# disk pools are currently supported. The value must be a comma-
# separated list of disk pool names to be used for provisioning.
# Defaults to ''.
CONFIG_CINDER_NETAPP_STORAGE_POOLS=

# Cluster admin account name used to access the SolidFire storage
# system.
CONFIG_CINDER_SOLIDFIRE_LOGIN=

# Password for the SolidFire cluster admin user account specified in
# the CONFIG_CINDER_SOLIDFIRE_LOGIN parameter.
CONFIG_CINDER_SOLIDFIRE_PASSWORD=

# Hostname (or IP address) for the SolidFire storage system's MVIP.
CONFIG_CINDER_SOLIDFIRE_HOSTNAME=

# Password to use for OpenStack Bare Metal Provisioning (ironic) to
# access the database.
CONFIG_IRONIC_DB_PW=PW_PLACEHOLDER

# Password to use for OpenStack Bare Metal Provisioning to
# authenticate with the Identity service.
CONFIG_IRONIC_KS_PW=PW_PLACEHOLDER

# Enter y if cron job for removing soft deleted DB rows should be
# created.
CONFIG_NOVA_DB_PURGE_ENABLE=True

# Password to use for the Compute service (nova) to access the
# database.
CONFIG_NOVA_DB_PW=c09650d78c1b45db

# Password to use for the Compute service to authenticate with the
# Identity service.
CONFIG_NOVA_KS_PW=96fa26e1400749c6
```

Whether or not Packstack should manage a default initial set of
Nova flavors. Defaults to 'y'.
CONFIG_NOVA_MANAGE_FLAVORS=y

Overcommitment ratio for virtual to physical CPUs. Specify 1.0 to
disable CPU overcommitment.
CONFIG_NOVA_SCHED_CPU_ALLOC_RATIO=16.0

Overcommitment ratio for virtual to physical RAM. Specify 1.0 to
disable RAM overcommitment.
CONFIG_NOVA_SCHED_RAM_ALLOC_RATIO=1.5

Protocol used for instance migration. Valid options are: ssh and
tcp. Note that the tcp protocol is not encrypted, so it is insecure.
['ssh', 'tcp']
CONFIG_NOVA_COMPUTE_MIGRATE_PROTOCOL=ssh

PEM encoded certificate to be used for ssl on the https server,
leave blank if one should be generated, this certificate should not
require a passphrase. If CONFIG_HORIZON_SSL is set to 'n' this
parameter is ignored.
CONFIG_VNC_SSL_CERT=

SSL keyfile corresponding to the certificate if one was entered. If
CONFIG_HORIZON_SSL is set to 'n' this parameter is ignored.
CONFIG_VNC_SSL_KEY=

Enter the PCI passthrough array of hash in JSON style for
controller eg. [{"vendor_id":"1234", "product_id":"5678",
"name":"default"}, {...}]
CONFIG_NOVA_PCI_ALIAS=

Enter the PCI passthrough whitelist array of hash in JSON style for
controller eg. [{"vendor_id":"1234", "product_id":"5678",
"name":"default"}, {...}]
CONFIG_NOVA_PCI_PASSTHROUGH_WHITELIST=

The hypervisor driver to use with Nova. Can be either 'qemu' or

```
# 'kvm'. Defaults to 'qemu' on virtual machines and 'kvm' on bare
# metal hardware. For nested KVM set it explicitly to 'kvm'.
CONFIG_NOVA_LIBVIRT_VIRT_TYPE=%{::default_hypervisor}

# Password to use for OpenStack Networking (neutron) to authenticate
# with the Identity service.
CONFIG_NEUTRON_KS_PW={{ SET_PASS }}

# The password to use for OpenStack Networking to access the
# database.
CONFIG_NEUTRON_DB_PW={{ SET_PASS }}

# The name of the Open vSwitch bridge (or empty for linuxbridge) for
# the OpenStack Networking L3 agent to use for external traffic.
# Specify 'provider' if you intend to use a provider network to handle
# external traffic.
CONFIG_NEUTRON_L3_EXT_BRIDGE=br-ex

# Password for the OpenStack Networking metadata agent.
CONFIG_NEUTRON_METADATA_PW={{ SET_PASS }}

# Specify 'y' to install OpenStack Networking's L3 Metering agent
# ['y', 'n']
CONFIG_NEUTRON_METERING_AGENT_INSTALL=y

# Specify 'y' to configure OpenStack Networking's Firewall-
# as-a-Service (FWaaS). ['y', 'n']
CONFIG_NEUTRON_FWAAS=n

# Specify 'y' to configure OpenStack Networking's VPN-as-a-Service
# (VPNaaS). ['y', 'n']
CONFIG_NEUTRON_VPNAAS=n

# Comma-separated list of network-type driver entry points to be
# loaded from the neutron.ml2.type_drivers namespace. ['local',
# 'flat', 'vlan', 'gre', 'vxlan', 'geneve']
CONFIG_NEUTRON_ML2_TYPE_DRIVERS=geneve,flat,vxlan,vlan

# Comma-separated, ordered list of network types to allocate as
```



```
# tenant networks. The 'local' value is only useful for single-box
# testing and provides no connectivity between hosts. ['local',
# 'vlan', 'gre', 'vxlan', 'geneve']
CONFIG_NEUTRON_ML2_TENANT_NETWORK_TYPES=geneve,vxlan
```

```
# Comma-separated ordered list of networking mechanism driver entry
# points to be loaded from the neutron.ml2.mechanism_drivers
# namespace. ['logger', 'test', 'linuxbridge', 'openvswitch',
# 'hyperv', 'ncs', 'arista', 'cisco_nexus', 'mlnx', 'l2population',
# 'sriovnicswitch', 'ovn']
CONFIG_NEUTRON_ML2_MECHANISM_DRIVERS=openvswitch
```

```
# Comma-separated list of physical_network names with which flat
# networks can be created. Use * to allow flat networks with arbitrary
# physical_network names.
CONFIG_NEUTRON_ML2_FLAT_NETWORKS=*
```

```
# Comma-separated list of <physical_network>:<vlan_min>:<vlan_max> or
# <physical_network> specifying physical_network names usable for VLAN
# provider and tenant networks, as well as ranges of VLAN tags on each
# available for allocation to tenant networks.
CONFIG_NEUTRON_ML2_VLAN_RANGES=
```

```
# Comma-separated list of <tun_min>:<tun_max> tuples enumerating
# ranges of GRE tunnel IDs that are available for tenant-network
# allocation. A tuple must be an array with tun_max +1 - tun_min >
# 1000000.
CONFIG_NEUTRON_ML2_TUNNEL_ID_RANGES=
```

```
# Comma-separated list of addresses for VXLAN multicast group. If
# left empty, disables VXLAN from sending allocate broadcast traffic
# (disables multicast VXLAN mode). Should be a Multicast IP (v4 or v6)
# address.
CONFIG_NEUTRON_ML2_VXLAN_GROUP=
```

```
# Comma-separated list of <vni_min>:<vni_max> tuples enumerating
# ranges of VXLAN VNI IDs that are available for tenant network
# allocation. Minimum value is 0 and maximum value is 16777215.
CONFIG_NEUTRON_ML2_VNI_RANGES=10:100
```

```

# Name of the L2 agent to be used with OpenStack Networking.
# ['linuxbridge', 'openvswitch', 'ovn']
CONFIG_NEUTRON_L2_AGENT=openvswitch

# Comma-separated list of interface mappings for the OpenStack
# Networking ML2 SRIOV agent. Each tuple in the list must be in the
# format <physical_network>:<net_interface>. Example:
# physnet1:eth1,physnet2:eth2,physnet3:eth3.
CONFIG_NEUTRON_ML2_SRIOV_INTERFACE_MAPPINGS=

# Comma-separated list of interface mappings for the OpenStack
# Networking linuxbridge plugin. Each tuple in the list must be in the
# format <physical_network>:<net_interface>. Example:
# physnet1:eth1,physnet2:eth2,physnet3:eth3.
CONFIG_NEUTRON_LB_INTERFACE_MAPPINGS=

# Comma-separated list of bridge mappings for the OpenStack
# Networking Open vSwitch plugin. Each tuple in the list must be in
# the format <physical_network>:<ovs_bridge>. Example: physnet1:br-
# eth1,physnet2:br-eth2,physnet3:br-eth3
CONFIG_NEUTRON_OVS_BRIDGE_MAPPINGS=extnet:br-ex,physnet1:br-vlan

# Comma-separated list of colon-separated Open vSwitch
# <bridge>:<interface> pairs. The interface will be added to the
# associated bridge. If you desire the bridge to be persistent a value
# must be added to this directive, also
# CONFIG_NEUTRON_OVS_BRIDGE_MAPPINGS must be set in order to create
# the proper port. This can be achieved from the command line by
# issuing the following command: packstack --allinone --os-neutron-
# ovs-bridge-mappings=ext-net:br-ex --os-neutron-ovs-bridge-interfaces
# =br-ex:eth0
CONFIG_NEUTRON_OVS_BRIDGE_IFACES=br-ex:ens3f0,br-vlan:ens3f1

# Comma-separated list of Open vSwitch bridges that must be created
# and connected to interfaces in compute nodes when flat or vlan type
# drivers are enabled. These bridges must exist in
# CONFIG_NEUTRON_OVS_BRIDGE_MAPPINGS and
# CONFIG_NEUTRON_OVS_BRIDGE_IFACES. Example: --os-neutron-ovs-bridges-

```

```

# compute=br-vlan --os-neutron-ovs-bridge-mappings="extnet:br-
# ex,physnet1:br-vlan" --os-neutron-ovs-bridge-interfaces="br-ex:eth1
# ,br-vlan:eth2"
CONFIG_NEUTRON_OVS_BRIDGES_COMPUTE=br-vlan

# Name of physical network used for external network when enabling
# CONFIG_PROVISION_DEMO. Name must be one of the included in
# CONFIG_NEUTRON_OVS_BRIDGE_MAPPINGS. Example: --os-neutron-ovs-
# bridge-mappings="extnet:br-ex,physnet1:br-vlan" --os-neutron-ovs-
# bridge-interfaces="br-ex:eth1,br-vlan:eth2" --os-neutron-ovs-
# external-physnet="extnet"
CONFIG_NEUTRON_OVS_EXTERNAL_PHYSNET=extnet

# Interface for the Open vSwitch tunnel. Packstack overrides the IP
# address used for tunnels on this hypervisor to the IP found on the
# specified interface (for example, eth1).
CONFIG_NEUTRON_OVS_TUNNEL_IF=

# Comma-separated list of subnets (for example,
# 192.168.10.0/24,192.168.11.0/24) used for sending tunneling packets.
# This is used to configure IP filtering to accept tunneling packets
# from these subnets instead of specific IP addresses of peer nodes.
# This is useful when you add existing nodes to EXCLUDE_SERVERS
# because, in this case, packstack cannot modify the IP filtering of
# the existing nodes.
CONFIG_NEUTRON_OVS_TUNNEL_SUBNETS=

# VXLAN UDP port.
CONFIG_NEUTRON_OVS_VXLAN_UDP_PORT=4789

# Comma-separated list of bridge mappings for the OpenStack
# Networking Open Virtual Network plugin. Each tuple in the list must
# be in the format <physical_network>:<ovs_bridge>. Example: physnet1
# :br-eth1,physnet2:br-eth2,physnet3:br-eth3
CONFIG_NEUTRON_OVN_BRIDGE_MAPPINGS=extnet:br-ex

# Comma-separated list of colon-separated Open vSwitch
# <bridge>:<interface> pairs. The interface will be added to the
# associated bridge. If you desire the bridge to be persistent a value

```

```

# must be added to this directive, also
# CONFIG_NEUTRON_OVN_BRIDGE_MAPPINGS must be set in order to create
# the proper port. This can be achieved from the command line by
# issuing the following command: packstack --allinone --os-neutron-
# ovn-bridge-mappings=ext-net:br-ex --os-neutron-ovn-bridge-interfaces
# =br-ex:eth0
CONFIG_NEUTRON_OVN_BRIDGE_IFACES=

# Comma-separated list of Open vSwitch bridges that must be created
# and connected to interfaces in compute nodes when flat or vlan type
# drivers are enabled. These bridges must exist in
# CONFIG_NEUTRON_OVN_BRIDGE_MAPPINGS and
# CONFIG_NEUTRON_OVN_BRIDGE_IFACES. Example: --os-neutron-ovn-bridges-
# compute=br-vlan --os-neutron-ovn-bridge-mappings="extnet:br-
# ex,physnet1:br-vlan" --os-neutron-ovn-bridge-interfaces="br-ex:eth1
# ,br-vlan:eth2"
CONFIG_NEUTRON_OVN_BRIDGES_COMPUTE=

# Name of physical network used for external network when enabling
# CONFIG_PROVISION_DEMO. Name must be one of the included in
# CONFIG_NEUTRON_OVN_BRIDGE_MAPPINGS. Example: --os-neutron-ovn-
# bridge-mappings="extnet:br-ex,physnet1:br-vlan" --os-neutron-ovn-
# bridge-interfaces="br-ex:eth1,br-vlan:eth2" --os-neutron-ovn-
# external-physnet="extnet"
CONFIG_NEUTRON_OVN_EXTERNAL_PHYSNET=extnet

# Interface for the Open vSwitch tunnel. Packstack overrides the IP
# address used for tunnels on this hypervisor to the IP found on the
# specified interface (for example, eth1).
CONFIG_NEUTRON_OVN_TUNNEL_IF=

# Comma-separated list of subnets (for example,
# 192.168.10.0/24,192.168.11.0/24) used for sending tunneling packets.
# This is used to configure IP filtering to accept tunneling packets
# from these subnets instead of specific IP addresses of peer nodes.
# This is useful when you add existing nodes to EXCLUDE_SERVERS
# because, in this case, packstack cannot modify the IP filtering of
# the existing nodes.
CONFIG_NEUTRON_OVN_TUNNEL_SUBNETS=

```

Password to use for the OpenStack File Share service (manila) to
access the database.

CONFIG_MANILA_DB_PW=PW_PLACEHOLDER

Password to use for the OpenStack File Share service (manila) to
authenticate with the Identity service.

CONFIG_MANILA_KS_PW=PW_PLACEHOLDER

Backend for the OpenStack File Share service (manila); valid
options are: generic, netapp, glusternative, or glusternfs.

['generic', 'netapp', 'glusternative', 'glusternfs']

CONFIG_MANILA_BACKEND=generic

Denotes whether the driver should handle the responsibility of
managing share servers. This must be set to false if the driver is
to operate without managing share servers. Defaults to 'false'

['true', 'false']

CONFIG_MANILA_NETAPP_DRV_HANDLES_SHARE_SERVERS=false

The transport protocol used when communicating with the storage
system or proxy server. Valid values are 'http' and 'https'.

Defaults to 'https'. ['https', 'http']

CONFIG_MANILA_NETAPP_TRANSPORT_TYPE=https

Administrative user account name used to access the NetApp storage
system. Defaults to ''.

CONFIG_MANILA_NETAPP_LOGIN=admin

Password for the NetApp administrative user account specified in
the CONFIG_MANILA_NETAPP_LOGIN parameter. Defaults to ''.

CONFIG_MANILA_NETAPP_PASSWORD=

Hostname (or IP address) for the NetApp storage system or proxy
server. Defaults to ''.

CONFIG_MANILA_NETAPP_SERVER_HOSTNAME=

The storage family type used on the storage system; valid values
are ontap_cluster for clustered Data ONTAP. Defaults to

```
# 'ontap_cluster'. ['ontap_cluster']
CONFIG_MANILA_NETAPP_STORAGE_FAMILY=ontap_cluster

# The TCP port to use for communication with the storage system or
# proxy server. If not specified, Data ONTAP drivers will use 80 for
# HTTP and 443 for HTTPS. Defaults to '443'.
CONFIG_MANILA_NETAPP_SERVER_PORT=443

# Pattern for searching available aggregates for NetApp provisioning.
# Defaults to '(.*)'.
CONFIG_MANILA_NETAPP_AGGREGATE_NAME_SEARCH_PATTERN=(.*)

# Name of aggregate on which to create the NetApp root volume. This
# option only applies when the option
# CONFIG_MANILA_NETAPP_DRV_HANDLES_SHARE_SERVERS is set to True.
CONFIG_MANILA_NETAPP_ROOT_VOLUME_AGGREGATE=

# NetApp root volume name. Defaults to 'root'.
CONFIG_MANILA_NETAPP_ROOT_VOLUME_NAME=root

# This option specifies the storage virtual machine (previously
# called a Vserver) name on the storage cluster on which provisioning
# of shared file systems should occur. This option only applies when
# the option driver_handles_share_servers is set to False. Defaults to
# ''.
CONFIG_MANILA_NETAPP_VSERVER=

# Denotes whether the driver should handle the responsibility of
# managing share servers. This must be set to false if the driver is
# to operate without managing share servers. Defaults to 'true'.
# ['true', 'false']
CONFIG_MANILA_GENERIC_DRV_HANDLES_SHARE_SERVERS=true

# Volume name template for Manila service. Defaults to 'manila-
# share-%s'.
CONFIG_MANILA_GENERIC_VOLUME_NAME_TEMPLATE=manila-share-%s

# Share mount path for Manila service. Defaults to '/shares'.
CONFIG_MANILA_GENERIC_SHARE_MOUNT_PATH=/shares
```

```
# Location of disk image for Manila service instance. Defaults to '
CONFIG_MANILA_SERVICE_IMAGE_LOCATION=https://www.dropbox.com/s/vi5oeh10q1qkckh/ub
untu_1204_nfs_cifs.qcow2

# User in Manila service instance.
CONFIG_MANILA_SERVICE_INSTANCE_USER=ubuntu

# Password to service instance user.
CONFIG_MANILA_SERVICE_INSTANCE_PASSWORD=ubuntu

# Type of networking that the backend will use. A more detailed
# description of each option is available in the Manila docs. Defaults
# to 'neutron'. ['neutron', 'nova-network', 'standalone']
CONFIG_MANILA_NETWORK_TYPE=neutron

# Gateway IPv4 address that should be used. Required. Defaults to ''.
CONFIG_MANILA_NETWORK_STANDALONE_GATEWAY=

# Network mask that will be used. Can be either decimal like '24' or
# binary like '255.255.255.0'. Required. Defaults to ''.
CONFIG_MANILA_NETWORK_STANDALONE_NETMASK=

# Set it if network has segmentation (VLAN, VXLAN, etc). It will be
# assigned to share-network and share drivers will be able to use this
# for network interfaces within provisioned share servers. Optional.
# Example: 1001. Defaults to ''.
CONFIG_MANILA_NETWORK_STANDALONE_SEG_ID=

# Can be IP address, range of IP addresses or list of addresses or
# ranges. Contains addresses from IP network that are allowed to be
# used. If empty, then will be assumed that all host addresses from
# network can be used. Optional. Examples: 10.0.0.10 or
# 10.0.0.10-10.0.0.20 or
# 10.0.0.10-10.0.0.20,10.0.0.30-10.0.0.40,10.0.0.50. Defaults to ''.
CONFIG_MANILA_NETWORK_STANDALONE_IP_RANGE=

# IP version of network. Optional. Defaults to '4'. ['4', '6']
CONFIG_MANILA_NETWORK_STANDALONE_IP_VERSION=4
```

```
# List of GlusterFS servers that can be used to create shares. Each
# GlusterFS server should be of the form [remoteuser@]<volserver>, and
# they are assumed to belong to distinct Gluster clusters.
CONFIG_MANILA_GLUSTERFS_SERVERS=

# Path of Manila host's private SSH key file.
CONFIG_MANILA_GLUSTERFS_NATIVE_PATH_TO_PRIVATE_KEY=

# Regular expression template used to filter GlusterFS volumes for
# share creation. The regex template can optionally (ie. with support
# of the GlusterFS backend) contain the #{size} parameter which
# matches an integer (sequence of digits) in which case the value
# shall be interpreted as size of the volume in GB. Examples: "manila-
# share-volume-d+$", "manila-share-volume-#{size}G-d+$"; with matching
# volume names, respectively: "manila-share-volume-12", "manila-share-
# volume-3G-13". In latter example, the number that matches "#{size}",
# that is, 3, is an indication that the size of volume is 3G.
CONFIG_MANILA_GLUSTERFS_VOLUME_PATTERN=

# Specifies the GlusterFS volume to be mounted on the Manila host.
# For e.g: [remoteuser@]<volserver>:/<volid>
CONFIG_MANILA_GLUSTERFS_TARGET=

# Base directory containing mount points for Gluster volumes.
CONFIG_MANILA_GLUSTERFS_MOUNT_POINT_BASE=

# Type of NFS server that mediate access to the Gluster volumes
# (Gluster or Ganesha).
CONFIG_MANILA_GLUSTERFS_NFS_SERVER_TYPE=gluster

# Path of Manila host's private SSH key file.
CONFIG_MANILA_GLUSTERFS_PATH_TO_PRIVATE_KEY=

# Remote Ganesha server node's IP address.
CONFIG_MANILA_GLUSTERFS_GANESHA_SERVER_IP=

# Specify 'y' to set up Horizon communication over https. ['y', 'n']
CONFIG_HORIZON_SSL=n
```



```
# Secret key to use for Horizon Secret Encryption Key.
CONFIG_HORIZON_SECRET_KEY=0275c625ea7d401fa655bbc49c4ef0fb

# PEM-encoded certificate to be used for SSL connections on the https
# server. To generate a certificate, leave blank.
CONFIG_HORIZON_SSL_CERT=

# SSL keyfile corresponding to the certificate if one was specified.
# The certificate should not require a passphrase.
CONFIG_HORIZON_SSL_KEY=

CONFIG_HORIZON_SSL_CACERT=

# Password to use for the Object Storage service to authenticate with
# the Identity service.
CONFIG_SWIFT_KS_PW={{ SET_PASS }}

# Comma-separated list of devices to use as storage device for Object
# Storage. Each entry must take the format /path/to/dev (for example,
# specifying /dev/vdb installs /dev/vdb as the Object Storage storage
# device; Packstack does not create the filesystem, you must do this
# first). If left empty, Packstack creates a loopback device for test
# setup.
CONFIG_SWIFT_STORAGES=

# Number of Object Storage storage zones; this number MUST be no
# larger than the number of configured storage devices.
CONFIG_SWIFT_STORAGE_ZONES=1

# Number of Object Storage storage replicas; this number MUST be no
# larger than the number of configured storage zones.
CONFIG_SWIFT_STORAGE_REPLICAS=1

# File system type for storage nodes. ['xfs', 'ext4']
CONFIG_SWIFT_STORAGE_FSTYPE=ext4

# Custom seed number to use for swift_hash_path_suffix in
# /etc/swift/swift.conf. If you do not provide a value, a seed number
```

is automatically generated.

CONFIG_SWIFT_HASH=b0bd07663dff4aa6

Size of the Object Storage loopback file storage device.

CONFIG_SWIFT_STORAGE_SIZE=2G

Password used by Orchestration service user to authenticate against
the database.

CONFIG_HEAT_DB_PW={{ SET_PASS }}

Encryption key to use for authentication in the Orchestration
database (16, 24, or 32 chars).

CONFIG_HEAT_AUTH_ENC_KEY={{ SET_PASS }}

Password to use for the Orchestration service to authenticate with
the Identity service.

CONFIG_HEAT_KS_PW={{ SET_PASS }}

Specify 'y' to install the Orchestration CloudFormation API. ['y',
'n']

CONFIG_HEAT_CFN_INSTALL=y

Name of the Identity domain for Orchestration.

CONFIG_HEAT_DOMAIN=heat

Name of the Identity domain administrative user for Orchestration.

CONFIG_HEAT_DOMAIN_ADMIN=heat_admin

Password for the Identity domain administrative user for
Orchestration.

CONFIG_HEAT_DOMAIN_PASSWORD=3da8995d89de4321

Specify 'y' to provision for demo usage and testing. ['y', 'n']

CONFIG_PROVISION_DEMO=n

Specify 'y' to configure the OpenStack Integration Test Suite
(tempest) for testing. The test suite requires OpenStack Networking
to be installed. ['y', 'n']

CONFIG_PROVISION_TEMPEST=n

```
# CIDR network address for the floating IP subnet.
CONFIG_PROVISION_DEMO_FLOATRANGE=172.24.4.0/24

# Allocation pools in the floating IP subnet.
CONFIG_PROVISION_DEMO_ALLOCATION_POOLS=[]

# The name to be assigned to the demo image in Glance (default
# "cirros").
CONFIG_PROVISION_IMAGE_NAME=cirros

# A URL or local file location for an image to download and provision
# in Glance (defaults to a URL for a recent "cirros" image).
CONFIG_PROVISION_IMAGE_URL=https://download.cirros-cloud.net/0.3.5/cirros-0.3.5-x86_64-
disk.img

# Format for the demo image (default "qcow2").
CONFIG_PROVISION_IMAGE_FORMAT=qcow2

# Properties of the demo image (none by default).
CONFIG_PROVISION_IMAGE_PROPERTIES=

# User to use when connecting to instances booted from the demo
# image.
CONFIG_PROVISION_IMAGE_SSH_USER=cirros

# Name of the uec image created in Glance used in tempest tests
# (default "cirros-uec").
CONFIG_PROVISION_UEC_IMAGE_NAME=cirros-uec

# URL of the kernel image copied to Glance image for uec image
# (defaults to a URL for a recent "cirros" uec image).
CONFIG_PROVISION_UEC_IMAGE_KERNEL_URL=https://download.cirros-cloud.net/0.3.5/cirros-
0.3.5-x86_64-kernel

# URL of the ramdisk image copied to Glance image for uec image
# (defaults to a URL for a recent "cirros" uec image).
CONFIG_PROVISION_UEC_IMAGE_RAMDISK_URL=https://download.cirros-cloud.net/0.3.5/cirros-
0.3.5-x86_64-initramfs
```

```
# URL of the disk image copied to Glance image for uec image
# (defaults to a URL for a recent "cirros" uec image).
CONFIG_PROVISION_UEC_IMAGE_DISK_URL=https://download.cirros-cloud.net/0.3.5/cirros-0.3.5-
x86_64-disk.img

CONFIG_TEMPEST_HOST=

# Name of the Integration Test Suite provisioning user. If you do not
# provide a user name, Tempest is configured in a standalone mode.
CONFIG_PROVISION_TEMPEST_USER=

# Password to use for the Integration Test Suite provisioning user.
CONFIG_PROVISION_TEMPEST_USER_PW=PW_PLACEHOLDER

# CIDR network address for the floating IP subnet.
CONFIG_PROVISION_TEMPEST_FLOATRANGE=172.24.4.0/24

# Primary flavor name to use in Tempest.
CONFIG_PROVISION_TEMPEST_FLAVOR_NAME=m1.nano

# Primary flavor's disk quota in Gb.
CONFIG_PROVISION_TEMPEST_FLAVOR_DISK=1

# Primary flavor's ram in Mb.
CONFIG_PROVISION_TEMPEST_FLAVOR_RAM=128

# Primary flavor's vcpus number.
CONFIG_PROVISION_TEMPEST_FLAVOR_VCPUS=1

# Alternative flavor name to use in Tempest.
CONFIG_PROVISION_TEMPEST_FLAVOR_ALT_NAME=m1.micro

# Alternative flavor's disk quota in Gb.
CONFIG_PROVISION_TEMPEST_FLAVOR_ALT_DISK=1

# Alternative flavor's ram in Mb.
CONFIG_PROVISION_TEMPEST_FLAVOR_ALT_RAM=128
```

```
# Alternative flavor's vcpus number.
CONFIG_PROVISION_TEMPEST_FLAVOR_ALT_VCPUS=1

# Specify 'y' to run Tempest smoke test as last step of installation.
CONFIG_RUN_TEMPEST=n

# Test suites to run, example: "smoke dashboard TelemetryAlarming".
# Optional, defaults to "smoke".
CONFIG_RUN_TEMPEST_TESTS=smoke

# Tests to skip, example: "test_basic_scenario test_volume".
# Optional, defaults to "".
CONFIG_SKIP_TEMPEST_TESTS=

# Specify 'y' to configure the Open vSwitch external bridge for an
# all-in-one deployment (the L3 external bridge acts as the gateway
# for virtual machines). ['y', 'n']
CONFIG_PROVISION_OVS_BRIDGE=y

# Password to use for Gnocchi to access the database.
CONFIG_GNOCCHI_DB_PW=PW_PLACEHOLDER

# Password to use for Gnocchi to authenticate with the Identity
# service.
CONFIG_GNOCCHI_KS_PW=PW_PLACEHOLDER

# Secret key for signing Telemetry service (ceilometer) messages.
CONFIG_CEILOMETER_SECRET={{ SET_PASS }}

# Password to use for Telemetry to authenticate with the Identity
# service.
CONFIG_CEILOMETER_KS_PW=PW_PLACEHOLDER

# Ceilometer service name. ['httpd', 'ceilometer']
CONFIG_CEILOMETER_SERVICE_NAME=httpd

# Backend driver for Telemetry's group membership coordination.
# ['redis', 'none']
CONFIG_CEILOMETER_COORDINATION_BACKEND=redis
```

Whether to enable ceilometer middleware in swift proxy. By default
this should be false to avoid unnecessary load.
CONFIG_ENABLE_CEILOMETER_MIDDLEWARE=n

IP address of the server on which to install the Redis server.
CONFIG_REDIS_HOST=172.10.10..41

Port on which the Redis server listens.
CONFIG_REDIS_PORT=6379

Password to use for Telemetry Alarming to authenticate with the
Identity service.
CONFIG_AODH_KS_PW=PW_PLACEHOLDER

Password to use for Telemetry Alarming (AODH) to access the
database.
CONFIG_AODH_DB_PW=PW_PLACEHOLDER

Password to use for Panko to access the database.
CONFIG_PANKO_DB_PW=PW_PLACEHOLDER

Password to use for Panko to authenticate with the Identity
service.
CONFIG_PANKO_KS_PW=PW_PLACEHOLDER

Password to use for OpenStack Database-as-a-Service (trove) to
access the database.
CONFIG_TROVE_DB_PW=PW_PLACEHOLDER

Password to use for OpenStack Database-as-a-Service to authenticate
with the Identity service.
CONFIG_TROVE_KS_PW=PW_PLACEHOLDER

User name to use when OpenStack Database-as-a-Service connects to
the Compute service.
CONFIG_TROVE_NOVA_USER=trove

Tenant to use when OpenStack Database-as-a-Service connects to the

```
# Compute service.
CONFIG_TROVE_NOVA_TENANT=services

# Password to use when OpenStack Database-as-a-Service connects to
# the Compute service.
CONFIG_TROVE_NOVA_PW=PW_PLACEHOLDER

# Password to use for OpenStack Data Processing (sahara) to access
# the database.
CONFIG_SAHARA_DB_PW=PW_PLACEHOLDER

# Password to use for OpenStack Data Processing to authenticate with
# the Identity service.
CONFIG_SAHARA_KS_PW=PW_PLACEHOLDER

# Password to use for the Magnum to access the database.
CONFIG_MAGNUM_DB_PW=PW_PLACEHOLDER

# Password to use for the Magnum to authenticate with the Identity
# service.
CONFIG_MAGNUM_KS_PW=PW_PLACEHOLDER
```

6. Packstack

```
$ packstack --answer-file=answer.txt
```

reference

- <https://www.rdoproject.org/install/packstack/>

Revision #5

Created 17 July 2022 17:58:53 by artop0420

Updated 24 December 2023 02:49:01 by artop0420