

# SUSE® OpenStack Cloud and Nutanix Deployment Guide

# Guide

[www.suse.com](http://www.suse.com)

**Deployment Guide**

SUSE OpenStack Cloud

**NUTANIX**



# Purpose of This Guide

**SUSE® OpenStack Cloud enables enterprises to easily deploy and manage heterogeneous software defined infrastructure. This guide provides SUSE OpenStack Cloud 7 operators with the information needed to integrate Nutanix into SUSE OpenStack Cloud.**

---

## Overview of the Solution

This deployment guide provides SUSE OpenStack Cloud 7 (SOC) operators with the information needed to integrate Nutanix Operating System (AOS) into a SUSE OpenStack Cloud environment. SUSE OpenStack Cloud supported hypervisors include KVM, Xen and VMware ESX additional support for hypervisors like Hyper-V and Acropolis are available from our partners. Support for mixed hypervisor environments offers more flexibility in designing cloud deployments and helps to optimize licensing costs and workload performance. Nutanix integration with SUSE OpenStack Cloud is supported through our established partnership across products.

Nutanix is the most popular hyperconverged infrastructure company, delivering a full infrastructure stack that integrates compute, virtualization, storage, networking and security to power any application, at any scale. SUSE is the world's first provider of an Enterprise Linux distribution that enables businesses to run their business-critical applications with the reliability and security that companies need. Together, Nutanix and SUSE software technology can deliver an enterprise solution with a broad range of functionality and flexibility to meet the needs of a constantly changing IT environment.

## How to Get Support

Product support for Nutanix AOS is available to licensed Nutanix customers from the [Nutanix Support Portal](#). For details, see their [Support Quick Reference Guide](#).

Product support for SUSE OpenStack Cloud is available to registered SUSE customers from [SUSE OpenStack Cloud Product Support](#).

SUSE and Nutanix will work together to resolve any interoperability issues that might occur in the deployment and operation of SOC 7 with the Nutanix drivers.

SUSE and Nutanix are both members of [TSAnet](#), a global, multi-vendor support community that enables members to collaborate and solve problems for joint customers. TSAnet facilitates support collaboration between vendors so that no matter who you contact initially, your issue will be addressed quickly by the responsible party.

## Prerequisites

- A running SOC 7 implementation
- A properly configured Nutanix installation
- The Nutanix-supplied OpenStack drivers
- Nutanix AHV licensing

## Software Configuration Tested

Nutanix Acropolis Operating System (AOS) 5.X and SUSE OpenStack Cloud 7 were tested.

Glance and Cinder were tested for image uploads and storage, respectively; and Neutron and Nova were tested for deployment of an instance.

---

## Getting Started

As mentioned in the prerequisites, in order for the information in this guide to be relevant, you must already have SUSE OpenStack Cloud 7 (SOC) deployed. If you do not have it deployed, follow the instructions in the Deployment Guide to get your infrastructure cloud-ready:

- [SUSE OpenStack Cloud 7 Download Page](#)
- [SUSE OpenStack Cloud 7 Deployment Guide](#)

In order for SOC 7 to be deployed quickly and successfully, the proper drivers must be installed and deployed. If you already have an OpenStack environment, the drivers will most likely be in place. This deployment requires Nova, Neutron, Glance and Cinder drivers installed in our OpenStack environment.

After SOC 7 is installed on the administration server and the compute and control nodes have been allocated, you should deploy the basic OpenStack services via SOC's crowbar barclamps. The following drivers are required: Cinder, Nova, Neutron and Glance services. You will need to deploy other barclamps before you can deploy these. This is detailed in the [SUSE Barclamp Deployment Guide](#).

## Turning Off chef-client

SUSE tested this solution by disabling the chef-client process which overwrites manual configuration file changes. However, this method is not recommended nor supported for production ready environments. We ran the following on all SOC 7 nodes to disable the process:

```
systemctl stop chef-client
systemctl disable chef-client
```

As our changes will still work, the recommended method to make changes for custom configuration files should be done by creating files with a 3-digit number and dash. For example:

```
/etc/nova/nova.conf.d/200-nova.conf
```

[Please refer to section 13.2 in Deploying with Crowbar for more information.](#)

## Deploy an External Network

To communicate between AOS and SOC 7, you need to deploy an external network. This allows for API calls to Nutanix's Nova, Cinder, Neutron and Glance processes. To deploy a custom

network configuration, check our [OpenStack documentation](#). Section 7.5.8 provides information on configuring an external network for accessing AOS.

## Deploy a Nutanix Cluster

If you already have a Nutanix Cluster up and running, you can skip this section and continue to **“Creating a Nutanix Storage Container.”** Otherwise, your next step is to ensure that you have an active implementation of Nutanix software.

To access the Nutanix software, you must have a user account registered with a serial number or a service tag. This will verify your ability to use Nutanix software and documentation. After you have created an account in Nutanix Gatekeeper, you will have access to downloads, deployment guides, binary upgrades and relevant Nutanix material.

After you are logged in, you can download full versions of Nutanix AHV. All releases are provided in the portal. For this supported configuration you will need to use AHV 4.6, which you can select here: [Nutanix Gatekeeper Portal](#)

You can find the latest patches to upgrade your hypervisor binary here: [Hypervisor Binary Upgrade](#)

You also need to upgrade the Nutanix AOS software to ensure that you have the latest patches and features. You can find the software here: [Nutanix AOS Binary Upgrade](#)

Now that you have the required software downloads and upgrades, you can proceed to deploying a Nutanix Cluster.

To access the installation/deployment guides, you need to be logged in to the Nutanix portal. Nutanix installs AOS and the VM controller at the factory before shipping a node to customers. If for some reason you do not have AOS and the Nutanix Controller VM installed, follow the guide to its completion: [Nutanix AOS Installation Guide](#)

For any Nutanix information regarding download/documentation, check **portal.nutanix.com** to search for what you need.

The following are useful Nutanix documentation links:

- [AOS 5.6 documentation](#)
- [AOS 5.6 AHV documentation](#)

- [Nutanix Knowledge Base Articles](#) (can search for documents)

### Creating a Nutanix Storage Container

If you have a Nutanix cluster running, you should have access to the [Prism management web console](#), which allows you to create a storage pools and containers. In order for the next steps to be successful, you will need to create a storage container to house data. Follow this Nutanix Gatekeeper link to create a storage container: [Creating a Storage Container](#)

### Creating DNS Entries for Importing Images

You will be importing your images through FQDN image repos on the Nutanix cluster. It is mandatory to have DNS entries on the cluster before importing these images (in order to resolve any image downloads) or before contacting an FQDN server.

See the [Configuring Images](#) topic in the *Nutanix Prism Web Console Guide*.

### Getting Dependencies

For proper Nutanix driver deployment, you will need to ensure that you have *python-websocket* and NTP dependencies installed. Fortunately, SOC 7 comes with NTP natively installed. Ensure that you have python-websocket installed on all Nova compute nodes by running:

```
zypper in python-websocket
```

### Download the Nutanix OpenStack Drivers

Now that you have a storage container in the cluster and the external dependencies required for the OpenStack drivers, you can proceed to downloading and untarring Nutanix drivers from: [ftp://ftp.nutanix.com/nutanix\\_openstack/suse/](ftp://ftp.nutanix.com/nutanix_openstack/suse/)

**Username:** upgrade

**Password:** Nutanix/4u

Install the Nutanix drivers according to the following instructions on every SOC 7 node. After it is configured, SOC 7 will be able to communicate with AOS.

### Configure Nutanix Drivers

As you deploy the solution, you will want to configure the Nutanix side first and then configure SUSE OpenStack Cloud, to ensure consistency.

### Glance Configuration on Nutanix

To configure Glance on Nutanix controller node(s):

- Untar the package in a temporary directory:
 

```
tar xvf nutanix-openstack-2017.4.0.tgz
```
- Copy the folder `nutanix_openstack` to `pythonpath`:
 

```
cp nutanix_openstack /usr/lib/python2.7/site-packages/
```
- Copy the JSON file in the config folder to `/etc/`

```
cp config/nutanix_openstack_config.json /etc/
```
- Edit `/etc/nutanix_openstack_config.json` to add the cluster config:

```
{
  "username": "<prism_admin_username>",
  "num_vcpus_per_core": "4",
  "vnc": "65000",
  "name": "<cluster_name>",
  "ip": "<cluster_vip>",
  "container_name": "<container_name>",
  "services": [
    "compute",
    "volume",
    "network"
  ],
  "password": "prism_admin_password"
}
```

Arguments specific to the environment to be completed:

- **username:** *Prism admin User name*
- **password:** *Prism Admin User password*
- **ip:** *Prism Virtual IP Address*
- **container\_name:** *Name of the storage container on the cluster.*
- **vnc:** *default to 65000*

### Glance Configuration on Cloud 7

To configure Glance on OpenStack controller node(s):

- In `/etc/glance/glance-api.conf`, configure the following parameters with the indicated value:

```
[glance_store] stores=glance.store.acropolis.  
Store default_store=http
```

- b. In `/usr/lib/python2.7/site-packages/glance_store-<xyz>-py2.7.egg-info/entry_points.txt`, configure the following parameters with the indicated value:

```
[glance_store.drivers] glance.store.acropolis.  
Store=nutanix_openstack.glance:Store  
http=nutanix_openstack.glance:Store
```

- c. Restart Glance services:

```
systemctl restart openstack-glance-api  
systemctl restart openstack-glance-registry
```

## Neutron Configuration on Nutanix

To configure Neutron on Nutanix controller node(s):

- a. Untar the package in a temporary directory:  
`tar xvf nutanix-openstack-2017.4.0.tgz`
- b. Copy the folder `nutanix_openstack` to `pythonpath`:  
`cp nutanix_openstack /usr/lib/python2.7/site-packages/`
- c. Copy the JSON file in the config folder to `/etc/`:  
`cp config/nutanix_openstack_config.json /etc/`
- d. Edit `/etc/nutanix_openstack_config.json` to add the cluster config:

```
{  
  "username": "<prism_admin_username>",  
  "num_vcpus_per_core": "4",  
  "vnc": "65000",  
  "name": "<cluster_name>",  
  "ip": "<cluster_vip>",  
  "container_name": "<container_name>",  
  "services": [  
    "compute",  
    "volume",  
    "network"  
  ],  
  "password": "prism_admin_password"  
}
```

Arguments specific to the environment to be completed:

- **username:** *Prism Admin User name.*
- **password:** *Prism Admin User password.*
- **ip:** *Prism Virtual IP Address.*
- **container\_name:** *Name of the storage container on the cluster.*
- **vnc:** *This isn't needed for the neutron driver; nevertheless, default it to 65000.*

## Neutron Configuration on Cloud 7

To configure Neutron on OpenStack controller node(s):

- a. In `/etc/neutron/plugin.ini`, configure the following parameters with the indicated value:

```
type_drivers=vlan  
tenant_network_types = vlan  
mechanism_drivers = nutanix
```

- b. In `/usr/lib/python2.7/site-packages/neutron-<x.y.z>-py2.7.egg-info/entry_points.txt`, add the following parameters with the indicated value:

(**Note:** `<x.y.z>` is a placeholder. The package could be different based on your openstack distribution.)

```
[neutron.ml2.mechanism_drivers]  
nutanix = nutanix_openstack.neutron.  
driver:AcropolisNetworkDriver
```

- c. Restart the Neutron service:

```
systemctl restart neutron-server
```

## Nova Configuration on Nutanix

To configure Nova on Nutanix controller node(s):

- a. Untar the package in a temporary directory:  
`tar xvf nutanix-openstack-2017.4.0.tgz`
- b. Create a folder called `nutanix` in `nova /usr/lib/python2.7/site-packages/nova/virt/`:  
`mkdir -p /usr/lib/python2.7/site-packages/nova/virt/nutanix`
- c. Copy the contents of folder `nova` into `nova virt` path:  
`cp nutanix_openstack/nova/* /usr/lib/python2.7/site-packages/nova/virt/nutanix/`

- d. Copy the JSON file in config folder to /etc/:

```
cp config/nutanix_openstack_config.json /etc/
```

- e. Edit /etc/nutanix\_openstack\_config.json to add the cluster config:

```
{
  "username": "<prism_admin_username>",
  "num_vcpus_per_core": "4",
  "vnc": "<vnc_port_number>",
  "name": "<cluster_name>",
  "ip": "<cluster_vip>",
  "container_name": "<container_name>",
  "services": [
    "compute",
    "volume",
    "network"
  ],
  "password": "prism_admin_password"
}
```

Arguments specific to the environment to be completed:

- **username:** Prism Admin User name
- **password:** Prism Admin User password
- **ip:** Prism Virtual IP Address
- **container\_name:** Name of the storage container on the cluster
- **vnc:** default to 65000

### Nova Configuration on Cloud 7

To configure Nova on OpenStack controller node(s):

- a. In /etc/nova/nova.conf, dd the following parameters with the indicated value:

```
[DEFAULT]
compute_driver = nutanix.AcropolisComputeDriver
allow_resize_to_same_host=True
allow_migrate_to_same_host=True
config_drive_format=iso9660
security_group_api=neutron
metadata_host=<neutron_node_ip>
sync_power_state_interval=60
```

```
auth_strategy=keystone
rpc_backend=rabbit
vnc_enabled=true
ovncproxy_base_url=http://127.0.0.1:8888/
  console/vnc_auto.html
network_api_class=nova.network.neutronv2.api.API
security_group_api=neutron
my_ip=<nova_compute_ip>

[neutron]
url=http://<neutron_node_ip>:<neutron_port>
admin_auth_url=http://<neutron_node_ip>:<keystone_auth_port>/v2.0
admin_username=<neutron_username>
admin_password=<neutron_password>

[conductor]
use_local=false

[oslo_messaging_rabbit]
rabbit_host=<rabbit_host_ip>
rabbit_port=<rabbit_port>
rabbit_userid=<rabbit_userid>
rabbit_password=<rabbit_password>
```

Arguments specific to the environment to be completed:

- **neutron\_node\_ip:** IP of neutron node
- **neutron\_port:** port for neutron
- **neutron\_username:** username for neutron
- **neutron\_password:** password for neutron
- **nova\_compute\_ip:** IP of nova node
- **keystone\_auth\_port:** port for keystone
- **rabbit\_host\_ip:** IP of rabbit mq
- **rabbit\_port:** port of rabbit mq
- **rabbit\_userid:** username of rabbit mq
- **rabbit\_password:** password of rabbit mq

- b. Restart Nova compute:

```
systemctl restart openstack-nova-compute
```

## Configuring Nova VNC on Cloud 7

- a. Double check that python-websocket is installed in compute node by running:

```
zypper in python-websocket
```

- b. Copy the folder werkzeug to the python path:

```
cp werkzeug /usr/lib/python2.7/site-packages/
```

- c. Copy the binary prism\_vnc\_proxy to the bin directory:

```
cp prism_vnc_proxy /usr/bin/
```

- d. Copy the vnc folder in pythonpath:

```
cp vnc /usr/lib/python2.7/site-packages/
```

- e. Run the command below in the background:

```
/usr/bin/prism_vnc_proxy --bind_address=0.0.0.0  
--bind_port=<random_ephemeral_port>  
--prism_hostname=<cluster_VIP> --prism_  
username=<prism_admin_username> --prism_  
password=<prism_admin_password> --docroot=  
usr/share/nutanix_openstack/vnc/static &
```

Arguments specific to the environment to be completed:

- **bind\_port:** Any user-defined ephemeral port to proxy the vnc connection
- **prism\_hostname:** Cluster Virtual IP Address
- **prism\_username:** Prism Admin User name
- **prism\_password:** Prism Admin User password

**NOTE:** The bind\_port port number should be the same as the one in /etc/nutanix\_openstack\_config.json cluster vnc port number

## Cinder Configurations on Nutanix

To configure Cinder on Nutanix controller node(s):

- a. Untar the package in a temporary directory:

```
tar xvf nutanix-openstack-2017.4.0.tgz
```

- b. Copy the folder nutanix\_openstack to pythonpath:

```
cp nutanix_openstack /usr/lib/python2.7/  
site-packages/
```

- c. Copy the json file in config folder to /etc/

```
cp config/nutanix_openstack_config.json /etc/
```

- d. Edit /etc/nutanix\_openstack\_config.json to add the cluster config:

```
{  
  "username": "<prism_admin_username>",  
  "num_vcpus_per_core": "4",  
  "vnc": "65000",  
  "name": "<cluster_name>",  
  "ip": "<cluster_vip>",  
  "container_name": "<container_name>",  
  "services": [  
    "compute",  
    "volume",  
    "network"  
  ],  
  "password": "prism_admin_password"  
}
```

Arguments specific to the environment to be completed:

- **username:** Prism Admin User name
- **password:** Prism Admin User password
- **ip:** Prism Virtual IP Address
- **container\_name:** Name of the storage container on the cluster

## Cinder Configurations on Cloud 7

To configure Cinder on OpenStack controller node(s):

- a. In /etc/cinder/cinder.conf, configure the following parameters with the indicated values:

```
[DEFAULT]  
enabled_backends=nutanix_openstack glance_host  
$service_ip  
glance_api_servers $service_ip:$glance_api_port  
[nutanix_openstack]  
volume_driver nutanix_openstack.cinder.driver.  
AcropolisVolumeDriver
```

- b. In /usr/lib/python2.7/site-packages/cinder-<x.y.z>-py2.7.egg-info/SOURCES.txt, add the following line:  
nutanix\_openstack/cinder/driver.py

c. Restart the services below:

```
systemctl restart openstack-cinder-api.service
systemctl restart openstack-cinder-scheduler.
service
systemctl restart openstack-cinder-volume.
service
systemctl restart openstack-cinder-backup.
service
```

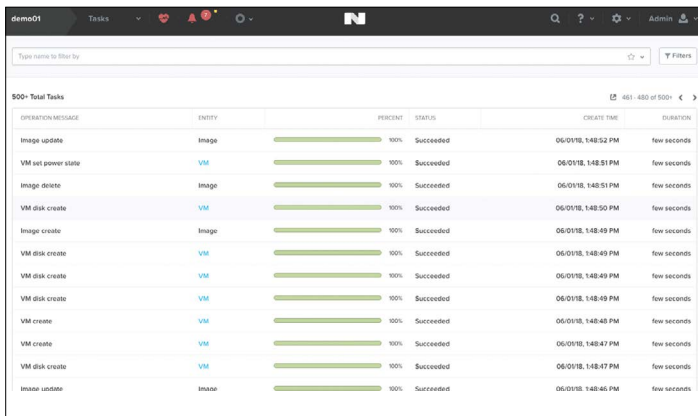
### Checking Your Work

#### Nutanix Checks

This concludes the instructions for deploying AOS drivers to operate with SO 7.

You will now want to verify that you have correctly deployed this solution by running the following checks through Nutanix.

- Log in to the Prism Element web console, then select “Tasks” in the main menu drop-down list.
- Check to see if the call for the operation has reached the cluster successfully (example below)



You can identify whether it was successful in the log file of the OpenStack service. If the call did not reach the Nutanix cluster, there might be an issue with the Nutanix driver. Check the driver log files of the OpenStack component or service, which can be found in the following locations:

- **Nova:** `/var/log/nova/nova-compute.log`  
(Note: this is on the nova-compute VM and not Nova-ControlPlane)
- **Neutron:** `/var/log/nova/server.log`
- **Cinder:** `/var/log/cinder/volume.log`
- **Glance:** `/var/log/glance/api.log`

Then check to see if a stack trace is available, set debug=True and capture the logs. You now have the payload or API and stack trace information.

If a stack trace is not available, the call has not reached the Nutanix drivers and you must troubleshoot the OpenStack control plane.

#### OpenStack Checks

Next, you will go through several checks to ensure that the drivers are working properly. Since you have adjusted the Nova, Neutron, Glance, and Cinder drivers, you will want to test each component.

In the Horizon dashboard, perform a simple check to see if the newly-added Nutanix cluster has been added.

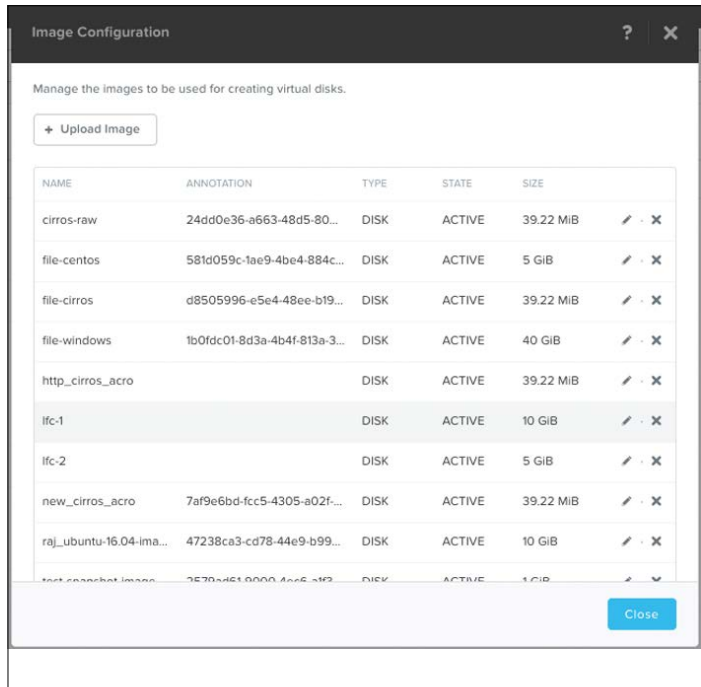
- In Glance, upload an image to the Nutanix cluster.  
For example:

```
openstack image create "cirros-raw" --disk-format
raw \ --container-format bare --public \
--file ~/images/cirros-0.3.5-x86_64-disk.img
```

- After creating an image, log in to the Prism web console to ensure that it was properly uploaded. If this was successful, you can delete the image:

```
openstack image delete <IMAGE>
```





- In Nova, first ensure that you can see the Nutanix cluster added to the hypervisor list and the communication between the two by running:

```
openstack hypervisor list
openstack service list
openstack network list
openstack network show <NETWORK_ID>
```

- If you see that the cluster was successfully added and proper communication is allowed, next allocate compute to ensure proper execution by running:

```
openstack server create --flavor <FLAVOR_ID>
--image <IMAGE_ID> \
--key-name <KEY_NAME> --user-data <USER_DATA_FILE>
--security-group \ <SEC_GROUP_NAME> --property
<KEY=VALUE> <INSTANCE_NAME>
```

- After you have created some compute, attempt to attach a cinder volume to the VM you created by running:

```
openstack volume create --size <SIZE_IN_GB> <NAME>
openstack server add volume <INSTANCE_ID>
<VOLUME_ID>
```

- If this was successful, you can now delete the attached volume to the VM and the newly- created VM by running:

```
openstack server remove volume <INSTANCE_ID>
<VOLUME_ID>
openstack volume delete <VOLUME_ID>
```

This completes to the instructions for integrating and deploying AOS with SUSE OpenStack Cloud 7. If you have any questions, you can look into the SUSE Support Forums for relevant information.

For more information on SUSE OpenStack Cloud, visit: [www.suse.com/products/suse-openstack-cloud/](http://www.suse.com/products/suse-openstack-cloud/).

Additional contact information and office locations:  
[www.suse.com](http://www.suse.com)

[www.suse.com](http://www.suse.com)