

## 1 Customize Red Hat OpenStack Platform Director

The Director requires several disk images for provisioning Overcloud nodes. This includes:

- A discovery kernel and ramdisk – Used for bare metal system discovery over PXE boot.
- A deployment kernel and ramdisk – Used for system provisioning and deployment.
- An Overcloud kernel, ramdisk, and full image – A base Overcloud system that is written to the node's hard disk.

Other customization of the Director is detailed in this section.

### Download Overcloud Images

Obtain these images from the Red Hat OpenStack Platform downloads page on the Red Hat Customer Portal at [https://access.redhat.com/documentation/en/red-hat-openstack-platform/version-8/director-installation-and-usage/#sect-Obtaining\\_Images\\_for\\_Overcloud\\_Nodes](https://access.redhat.com/documentation/en/red-hat-openstack-platform/version-8/director-installation-and-usage/#sect-Obtaining_Images_for_Overcloud_Nodes).

**Note:** The OSP7 Discovery and Deployment images must be utilized for node discovery and deployment via OpenStack Ironic rather than the default Ironic Python Agent (IPA) in OSP8, as the IPA does not support discovering or deploying to servers with iSCSI backed root disks. For the latest information, see:

[https://bugzilla.redhat.com/show\\_bug.cgi?id=1283436](https://bugzilla.redhat.com/show_bug.cgi?id=1283436) and  
[https://bugzilla.redhat.com/show\\_bug.cgi?id=1317731](https://bugzilla.redhat.com/show_bug.cgi?id=1317731)

1. Browse to the following URL: [https://access.redhat.com/downloads/content/191/ver=7/rhel---7/7/x86\\_64/product-software](https://access.redhat.com/downloads/content/191/ver=7/rhel---7/7/x86_64/product-software)
2. Download the latest OSP7 based Discovery and Deployment Ramdisk to the images directory in the stack user's home directory on the Director host `/home/stack/images`.

**Note:** In our environment, the Discovery ramdisk is `discovery-ramdisk-7.3.1-59.tar`.

**Note:** In our environment, the Deployment ramdisk is `deploy-ramdisk-ironic-7.3.0-39.tar`.

3. Install the `rhosp-director-images` and `rhosp-director-images-ipa` packages.

```
[stack@osp-director images]$ sudo yum install rhosp-director-images rhosp-director-images-ipa
```

**Note:** In our environment, the Overcloud image is the one distributed as a part of the OSP8 GA. The IPA is needed as uploads to Glance in a later step will fail without it being present.

4. Copy the new image archives to the images directory on the stack user's home directory.

```
[stack@osp-director images]$ cp /usr/share/rhosp-director-images/overcloud-full-latest-8.0.tar ~/images/
[stack@osp-director images]$ cp /usr/share/rhosp-director-images/ironic-python-agent-latest-8.0.tar ~/images/
```

5. Verify that the images are in the `/home/stack/images` directory.

```
[stack@osp-director images]$ ls -lh
-rw-rw-r--. 1 stack stack 59M Mar  9 17:10 deploy-ramdisk-ironic-7.3.0-39.tar
-rw-rw-r--. 1 stack stack 152M Mar 17 11:15 discovery-ramdisk-7.3.1-59.tar
-rw-r--r--. 1 stack stack 334M Apr  7 21:59 ironic-python-agent-latest-8.0.tar
-rw-r--r--. 1 stack stack 1.1G Apr  7 21:58 overcloud-full-latest-8.0.tar
```

6. Extract the images from the archives.

```
[stack@osp-director images]$ for tarfile in *.tar; do tar -xf $tarfile; done
```

7. Verify that the extraction was successful.

```
[stack@osp-director images]$ ls -lh
```

```
-rw-rw-r--. 1 stack stack 59M Mar 9 17:10 deploy-ramdisk-ironic-7.3.0-39.tar
-rw-r--r--. 1 stack stack 54M Oct 7 2015 deploy-ramdisk-ironic.initramfs
-rwxr-xr-x. 1 stack stack 4.8M Oct 7 2015 deploy-ramdisk-ironic.kernel
-rw-rw-r--. 1 stack stack 152M Mar 17 11:15 discovery-ramdisk-7.3.1-59.tar
-rw-r--r--. 1 stack stack 148M Mar 11 15:53 discovery-ramdisk.initramfs
-rwxr-xr-x. 1 stack stack 5.0M Mar 11 15:53 discovery-ramdisk.kernel
-rw-r--r--. 1 stack stack 329M Apr 4 16:21 ironic-python-agent.initramfs
-rwxr-xr-x. 1 stack stack 5.0M Apr 4 16:21 ironic-python-agent.kernel
-rw-r--r--. 1 stack stack 334M Apr 7 21:59 ironic-python-agent-latest-8.0.tar
-rw-r--r--. 1 stack stack 39M Apr 4 16:46 overcloud-full.initrd
-rw-r--r--. 1 stack stack 1.1G Apr 7 21:58 overcloud-full-latest-8.0.tar
-rw-r--r--. 1 stack stack 982M Apr 4 16:49 overcloud-full.qcow2
-rwxr-xr-x. 1 stack stack 5.0M Apr 4 16:46 overcloud-full.vmlinuz
```

Move on to the next section to customize the Director and the overcloud-full.qcow2 file to support iSCSI boot LUNs hosted on the NetApp FAS.

## Customize Director and Overcloud Images

We must customize the following in order to support iSCSI backed root disks in the eventual OpenStack deployment:

- Modify the `/httpboot` directory for discovery image and the iPXE process so that OpenStack Ironic can discover the boot lun during introspection and discovery.
- Modify the Python `ipxe_config.template` source code so corrected information for recognizing iSCSI Boot LUNs is properly written to the `/tftpboot` directory during introspection and with the deployment image.
- Customize the overcloud-full.qcow2 image to support “iscsi” and “multipath” in the initial ramdisk so the system sees the remote LUN on bootup of RHEL 7.2.
- Set a password for the root user on the overcloud-full.qcow2 image for troubleshooting purposes.

1. Begin by installing packages on the Director server that we’ll use for the eventual customization.

```
[stack@osp-director images]$ sudo yum -y install libguestfs-tools-c libvirt
```

2. Restart the libvirtd service.

```
[stack@osp-director images]$ sudo systemctl restart libvirtd.service
```

3. Move the overcloud-full.qcow2 image to the `/var/lib/libvirt/images` directory.

```
[stack@osp-director images]$ sudo mv overcloud-full.qcow2 /var/lib/libvirt/images/
```

4. Use the `virt-edit` command to modify the `/etc/sysconfig/grub` file in the overcloud-full.qcow2 image.

```
[stack@osp-director images]$ sudo virt-edit -a /var/lib/libvirt/images/overcloud-full.qcow2 /etc/sysconfig/grub
```

5. Append the option `rd.iscsi.firmware=1` and `rd.iscsi.ibft=1` to the end of the `GRUB_CMDLINE_LINUX` line. Our environment looks like this:

```
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="console=tty0 console=ttyS0,115200n8 crashkernel=auto rd.iscsi.firmware=1
rd.iscsi.ibft=1"
GRUB_DISABLE_RECOVERY="true"
```

6. Save and close the file by using the `:wq` key combination.

7. Use the `virt-edit` command to modify the `/etc/fstab` file in the overcloud-full.qcow2 image.

```
[stack@osp-director images]$ sudo virt-edit -a /var/lib/libvirt/images/overcloud-full.qcow2 /etc/fstab
```

8. Append the option `_netdev` to the first uncommented line. Our environment looks like this:

```
#
# /etc/fstab
# Created by anaconda on Mon Apr  4 19:29:08 2016
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
LABEL=img-rootfs /                                xfs          defaults,_netdev    0 0
```

9. Save and close the file by using the `:wq` key combination.

10. Inject a NetApp specific `multipath.conf` file into the `overcloud-full.qcow2` image, which has been supplied as an exhibit from GitHub and downloaded earlier to the Director server. This should be located in the `/home/stack/flexpod-templates/netapp-extra` directory.

```
[stack@osp-director images]$ sudo virt-customize -a /var/lib/libvirt/images/overcloud-full.qcow2 --upload /home/stack/flexpod-templates/netapp-extra/multipath.conf:/etc/
[ 0.0] Examining the guest ...
[ 3.0] Setting a random seed
[ 3.0] Uploading: /home/stack/flexpod-templates/netapp-extra/multipath.conf to /etc/
[ 3.0] Finishing off
```

11. Boot the image in rescue mode so that we can re-create the initial ramdisk which is used to actually find the iSCSI LUN during bootup. We need to specifically add the “iscsi” and “multipath” modules in the `initramfs`, which is not there by default. If this is not done, the `systemd` daemon will freeze upon bootup.

```
[stack@osp-director images]$ sudo virt-rescue --rw /var/lib/libvirt/images/overcloud-full.qcow2
```

12. While in the rescue shell for the `overcloud-full.qcow2` image, mount `/dev/sda` to the `/sysroot` directory.

```
><rescue> mount /dev/sda /sysroot
[ 85.600177] SGI XFS with ACLs, security attributes, no debug enabled
[ 85.603879] XFS (sda): Mounting V4 Filesystem
[ 85.616669] XFS (sda): Ending clean mount
```

13. Change root to this newly mounted partition.

```
><rescue> chroot /sysroot
```

14. Regenerate the `initramfs` image by using `Dracut` to include necessary modules so that the system can find the remote LUN.

```
><rescue> dracut -f -v -N
```

15. Exit `virt-rescue`.

```
><rescue> exit
exit
><rescue> exit
exit

virt-rescue: Syncing the disk now before exiting ...

Rebooting.
[ 141.174395] Unregister pv shared memory for cpu 0
[ 141.174839] sd 2:0:1:0: [sdb] Synchronizing SCSI cache
[ 141.175347] sd 2:0:0:0: [sda] Synchronizing SCSI cache
[ 141.176426] Restarting system.
[ 141.176756] reboot: machine restart
```

16. Set a root password on the `overcloud-full.qcow2` image to aid in Troubleshooting.

```
[stack@osp-director images]$ sudo virt-customize -a /var/lib/libvirt/images/overcloud-full.qcow2 --root-password password:myNewRootPassword
```

17. Move the newly modified image back into the /home/stack/images directory on the Director server.

```
[stack@osp-director images]$ sudo mv /var/lib/libvirt/images/overcloud-full.qcow2 .
```

## Upload Images to Glance

The images must be uploaded to Glance using the `openstack overcloud image upload` command.

1. Import the images contained in the images directory into the Director.

**Note:** The `--old-deploy-image` parameter is critical to ensure that the OSP7 based deployment image is utilized whenever we deploy the Overcloud. Ignore the warning about using the bash-based ramdisk.

```
[stack@osp-director images]$ openstack overcloud image upload --image-path /home/stack/images/ --old-deploy-image
DEPRECATED: The old bash-based ramdisks are no longer supported. You should move to the agent-based ramdisk as soon as possible.
```

2. Verify the images are contained in the Image store.

```
[stack@osp-director images]$ openstack image list
+-----+-----+
| ID                                          | Name                               |
+-----+-----+
| 45529dd0-dd8c-411d-aeb6-fce0b113de65    | bm-deploy-ramdisk                |
| 35e03a6b-a468-4f34-bd33-b379de9dc09d    | bm-deploy-kernel                 |
| a87a2493-5d8a-4be6-9e13-b32095069914    | overcloud-full                   |
| 9ca4d072-0cf6-4bfd-95bc-547d61b4da00    | overcloud-full-initrd            |
| e5068c74-8e76-4bb8-a31d-c4d00a9d0991    | overcloud-full-vmlinuz           |
+-----+-----+
```

## Modify iPXE Discovery Configuration

We must configure iPXE to use the Discovery ramdisk and kernel from OSP7, as well as configure the PxE server options while performing introspection to account for the discovered process in OpenStack Ironi. The IPA will not work during the discovery/introspection phase of the Overcloud creation.

Perform the following steps.

1. Move the IPA-based agent files to a different location so they are not used.

```
[stack@osp-director httpboot]$ cd /httpboot/
[stack@osp-director httpboot]$ sudo mv agent.kernel agent.kernel.old
[stack@osp-director httpboot]$ sudo mv agent.ramdisk agent.ramdisk.old
[stack@osp-director httpboot]$ sudo mv inspector.ipxe inspector.ipxe.old
```

2. Copy the OSP7 based discovery kernel and ramdisk to the original IPA-based locations.

```
[stack@osp-director httpboot]$ cd /home/stack/images/
[stack@osp-director images]$ sudo cp discovery-ramdisk.kernel /httpboot/agent.kernel
[stack@osp-director images]$ sudo cp discovery-ramdisk.initramfs /httpboot/agent.ramdisk
```

3. Copy the `inspector.pxe.old` file to `inspector.ipxe` and modify it to suit your environment. Our lab `inspector.ipxe` is listed here for reference. You may substitute the IP address of your respective Director server and place the following into `/httpboot/inspector.ipxe`:

```
#!ipxe
dhcp
```

```
kernel http://172.21.19.18:8088/agent.kernel
discoverd_callback_url=http://172.21.19.18:5050/v1/continue ip=${ip}:${next-
server}:${gateway}:${netmask} BOOTIF=${mac} rd.iscsi.firmware=1
initrd http://172.21.19.18:8088/agent.ramdisk
boot
```

## Modify iPXE Deployment Configuration

We must also modify the source code for iPXE to pass options intended to help find the iSCSI LUN used for the root disk of the server itself to the bash-based deployment disk used for the Openstack creation.

The IPA will not work during the deployment phase of the Openstack creation, so care must be made as shown in the section titled “Upload Images to Glance” to ensure that the `--old-deploy-image` parameter is passed whenever uploading the deployment kernel and ramdisk to OpenStack Glance.

1. From the Director server, open the following file with a text editor.

```
[stack@osp-director ~]$ sudo vim /usr/lib/python2.7/site-
packages/ironic/drivers/modules/ipxe_config.template
```

2. Insert the `rd.iscsi.firmware=1` to the end of the stanza with the `:deploy` and `:kernel` lines. Our file looks like the following after modification:

```
#!/ipxe

dhcp

goto deploy

:deploy
kernel {{ pxe_options.deployment_aki_path }} selinux=0 disk={{ pxe_options.disk }}
iscsi_target_iqn={{ pxe_options.iscsi_target_iqn }} deployment_id={{ pxe_options.deployment_id }}
deployment_key={{ pxe_options.deployment_key }} ironic_api_url={{ pxe_options.ironic_api_url }}
troubleshoot=0 text {{ pxe_options.pxe_append_params|default("", true) }} boot_option={{
pxe_options.boot_option }} ip=${ip}:${next-server}:${gateway}:${netmask} BOOTIF=${mac} {% if
pxe_options.root_device %}root_device={{ pxe_options.root_device }}{% endif %} ipa-api-url={{
pxe_options['ipa-api-url'] }} ipa-driver-name={{ pxe_options['ipa-driver-name'] }} boot_mode={{
pxe_options['boot_mode'] }} initrd=deploy_ramdisk coreos.configdrive=0 rd.iscsi.firmware=1

initrd {{ pxe_options.deployment_ari_path }}
boot

:boot_partition
kernel {{ pxe_options.aki_path }} root={{ ROOT }} ro text {{
pxe_options.pxe_append_params|default("", true) }} rd.iscsi.firmware=1
initrd {{ pxe_options.ari_path }}
boot

:boot_whole_disk
sanboot --no-describe
```

3. Restart OpenStack Ironic services.

```
[stack@osp-director ~]$ sudo systemctl restart openstack-ironic-{api,conductor,inspector-
dnsmasq,inspector}
```