

Install Compute Resources

This section is the part of the OnApp installation procedure.

[Install Control Panel Server](#) > [Install Compute Resources](#) > [Install Data Stores](#) > [Install Backup Server](#) > [Configure vCloud Director Integration](#) > [Configure Cloud](#)

Once the Control Panel server has been installed successfully, you can follow one of two processes to set up Xen or KVM compute resources:

- [Install CloudBoot Compute Resources](#) - the CloudBoot method where compute resources are installed over your network
- [Install Static Compute Resources](#) - standard static installation process to each compute resource's local disk



We strongly recommend that you avoid creating mixed compute zones:

- do not add CloudBoot and static boot compute resources to one compute zone
- do not add both XEN and KVM compute resources to one zone

The reason is that XEN VSs cannot migrate/failover to a KVM compute resource and KVM VSs cannot migrate/failover to a XEN compute resource.

See also:

[Install Control Panel Server](#)

[Install Data Stores](#)

[Install Backup Server](#)

[Technical Details](#)

[Preparation Guide](#)

[Configure vCloud Director Integration](#)

On this page:

- [Install CloudBoot Compute Resources](#)
 - [Configure CloudBoot Settings in BIOS](#)
 - [Configure InfiniBand](#)
- [Install Static Compute Resources](#)

Install CloudBoot Compute Resources

Follow this method to enable CloudBoot for your compute resources. CloudBoot compute resource installation enables dynamic boot of compute resource servers without any persistent installation requirements. The servers must support and have PXE boot enabled on the Network Interface Card (setup in the BIOS if not already enabled by default). See [Configure CloudBoot Settings in BIOS](#) for details. We strongly recommend you to deploy one or more backup servers for backups and VS provisioning when using CloudBoot functionality.

1. Enable CloudBoot in the Control Panel:

- a. Go to **Settings > Configuration > System > CloudBoot**
- b. Scroll down to the CloudBoot section and check the **Enable** box.

2. Enable Storage in the Control Panel:

- a. Go to **Settings > Configuration > System > OnApp Storage**
- b. Scroll down to the OnApp Storage section and check the **Enable OnApp Storage** box.
- c. Tick the **Use Local Read Path** check box to minimise the network throughput dependency for read heavy workloads.

3. Enter IP addresses for static content target and Control Panel server CloudBoot interface:

Static content, such as CloudBoot images, kernels, virtual server templates, can be hosted on a standalone NFS server if you wish. The default configuration is to install everything on the Control Panel server.

Enter the relevant IPs in **Settings > Configuration > System > CloudBoot**

4. Add IP address range for compute resources:

Settings > Compute resources > CloudBootIPs > New IP Address

5. Power on servers and allow them to boot the default image.

Add servers to the Control Panel by selecting MAC addresses and assigning IP address

Settings > Compute resources > Add a new CloudBoot Compute resource

If you want to expose drives in compute resources to OnApp Storage, our integrated storage platform, then you must select them at this point. For more information on setting up and configuring CloudBoot, see the [CloudBoot Compute resources](#) section of the Admin guide.

To increase dom0 memory for all new Xen compute resources, edit the dom0 value in the `/tftpboot/pxelinux.cfg/template-xen` file on the CP server.

To increase dom0 memory for a single Xen compute resource, edit the `/tftpboot/pxelinux.cfg/xx-xx-xx-xx-xx-xx` file, where you have to replace the x's with your compute resource's management NIC MAC address.

6. CloudBoot compute resources mount the following locations automatically at boot:

- `/tftpboot/export/centos5/xen` to `/.ro`
The path may vary depending on the compute resource template used.
- `/data` to `/onapp/tools/recovery`
- `/tftpboot/images/centos5/ramdisk-xen` to `/cloudboot/centos5/ramdisk-xen`
The path may vary depending on the compute resource template.

The NFS server from which these are mounted is defined by the **Static Config target** parameter (see [Edit System Configuration](#) section for details). You can set the default Control Panel server IP to any other server. This change will affect all CloudBoot compute resources.

The following paths must be available in the static config target to make it possible to use CloudBoot:

- `/tftpboot/export`
- `/data`
- `/tftpboot/images`

Compute resources will use local templates (mounted from Static Config target) during the server provisioning if the **Use SSH file transfer** configuration setting is disabled or the template has null `backup_server_id`.

7. If you do not have a Dedicated Backup Server in place, please use Custom Config to mount `/onapp/templates` and `/onapp/backup` from your Control Panel server or another NFS export.

8. After you have installed CloudBoot compute resource proceed to the [Configure CloudBoot Settings in BIOS](#) section.



If you do not have a dedicated backup server you must mount your Template and Backup repository to the Compute resource for VS provisioning and backups to work, for example from your Control Panel server:

Add to `/etc/exports` on the Control Panel server:

```
/onapp/templates 192.168.10.0/24(rw,no_root_squash)
```

```
/onapp/backups 192.168.10.0/24(rw,no_root_squash)
```

Add to Custom Config on the Compute resource and run them manually on the command line (In this example we are mounting from 192.168.10.101):

```
mkdir -p /onapp/backups && mount -t nfs 192.168.10.101:/onapp/backups /onapp/backups
```

```
mkdir -p /onapp/templates && mount -t nfs 192.168.10.101:/onapp/templates /onapp/templates
```

Perform the following steps if you plan to deploy Accelerator. Otherwise skip.

9. Run the following command on the CP server:

- For all compute resources:

```
rake hypervisor:messaging:configure
```

- For certain compute resources only:

```
rake hypervisor:messaging:configure['11.0.50.111 11.0.50.112']
```


To perform the configuration for a number of compute resources, separate their IPs with a space.

10. The command above should be run after every reboot. However, you can avoid the necessity to run the command repeatedly after every reboot by copying the following information (using your parameters) from `/home/mq/onapp/messaging/credentials.yml` to the custom config:

```

echo "---
host: 10.0.50.4# RABBITMQ SERVER IP/FQDN
port: 5672# RABBITMQ CONNECTION PORT(default: 5672)
vhost: '/'
user: accelerator-example # RABBITMQ USER NAME
password: 'e{y3!s8!}' #RABBITMQ ACCESS PASSWORD
queue: 'hv-10.0.50.102' # hv-[IP Address of Compute Resource]
exchange:
name: 'acceleration'
type: 'direct'
durable: True" > /home/mq/onapp/messaging/credentials.yml
chown -R mq:mq /home/mq
service onapp-messaging restart

```


 For information on manual configuration for Accelerator, refer to [RabbitMQ Configuration for Accelerator](#).

Configure CloudBoot Settings in BIOS

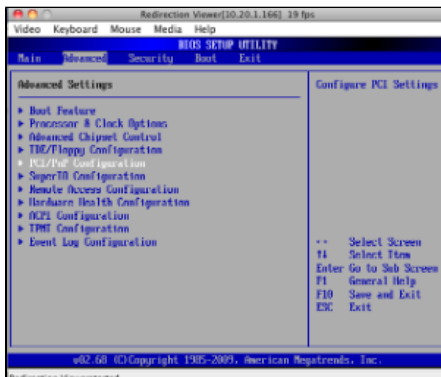
 Your BIOS settings may vary from the example provided in this section.

To use PXE boot, you have to make sure it is enabled in BIOS. To do so:

1. Select the required ethernet card supporting PXE as a boot device:



2. After that, go to the **Advanced** settings > **PCI/PnP configuration**.



3. In the **Advanced** settings, select the first/primary Onboard LAN/NIC Option ROM and press **Enter**.

Use up and down arrow keys to set Option ROM settings to enabled and press **Enter**.

Press **Escape** key to return to the **Advanced** menu.

1. Add the compute resource to your cloud using the OnApp Control Panel: **Settings > Compute resources > Add New Compute resource**. Make sure the compute resource is visible in the Control Panel, and at this point showing as inactive.
2. Update your server:

```
bash# yum update
```

3. Download the OnApp repository:

```
bash#> rpm -Uvh http://rpm.repo.onapp.com/repo/onapp-repo-5.1.noarch.rpm
```

4. Install the OnApp compute resource installer package:

```
bash#> yum install onapp-hv-install
```

5. Update OS components using the following command:

For XEN

```
bash# /onapp/onapp-hv-install/onapp-hv-xen-install.sh -y
```

or

For KVM

```
bash# /onapp/onapp-hv-install/onapp-hv-kvm-install.sh -y
```

6. Edit custom compute resource configuration. Custom values must be set before the installer script runs.



If deploying XEN onto a server running CentOS 6, it is important to specify a number for `XEN_DOM0_MAX_VCPUS`. We recommend that this is set to 2 if the compute resource has 12 cores or less. Or 4 if the compute resource has more than 12 cores.

```
#vi /onapp/onapp-hv.conf
```

- OnApp HV tools [custom version](#)

```
HV_VERSION=" "
```

- OnApp StorageAPI [custom version](#)

```
API_VERSION=" "
```

- Default server to sync time on the compute resource

```
NTP_TIME_SERVER='pool.ntp.org'
```

- Xen HV (Domain-0) related configuration

```
XEN_DOM0_MEM_MIN=409600
XEN_DOM0_MEM_DEVISOR=48
XEN_DOM0_MAX_VCPUS=" "
XEN_DOM0_VCPUS_PIN_ENABLE=0
XEN_DOM0_SCHEDULER_WEIGHT=65535
XEN_DOM0_SCHEDULER_CAP=200
```

4.2.x and higher versions only:

```
XEN_DOM0_SCHEDULER_RATELIMIT_US=100
XEN_DOM0_SCHEDULER_TIMESLICE_MS=5
```

- The number of loopback devices created

```
LOOPBACKS=128
```

- The maximum size of the connection tracking table.



The value can't be greater than 65536 if the total memory of Xen Domain-0 or KVM is less than 1Gb.
The value could be doubled (or even more, depends on memory amount).

```
NET_IPV4_NETFILTER_IP_CONTRACK_MAX=" "
```

- The divisor to calculate the hash table. The recommended value is 8.

```
hashsize = nf_conntrack_max / 8  
CONTRACK_TO_HASHSIZE=8
```

- Outdated Xen compute resource's (Domain-0) configuration parameters

```
XEN_DOM0_MEM_OVERHEAD_MIN=262144  
P_TO_VCPUS=4
```

7. Run the OnApp compute resource installer script:

Usage:

```
/onapp/onapp-hv-install/onapp-hv-xen-install.sh [-c CONFIG_FILE] [-a] [-y] [-t] [-s] [-v HV_VERSION]  
[-p API_VERSION] [-h]
```

Where:

-c CONFIG_FILE	custom installer configuration file. Otherwise, preinstalled one is used.
-a	do NOT be interactive. Process with automatic installation.
-v HV_VERSION	custom Compute resource Tools version
-p API_VERSION	custom StorageAPI version
-y	update OS packages (except for OnApp provided) on the box with 'yum update'.
-t	initiate Recovery templates and ISO(s), which are used to provision FreeBSD guests, download The download is initiated if '-a' option is used
-s	skip packages management: install, remove, upgrade
-h	print this info

- Run the OnApp compute resource installer script for Xen compute resources:

```
bash#> /onapp/onapp-hv-install/onapp-hv-xen-install.sh
```

- Run the OnApp compute resource installer script for KVM compute resources:

```
bash#> /onapp/onapp-hv-install/onapp-hv-kvm-install.sh
```


- Configure the compute resource for your cloud. This step is also required for the SNMP statistics receiver configuration:

```
bash#> /onapp/onapp-hv-install/onapp-hv-config.sh -h <CP_HOST_IP> -p [HV_HOST_IP] -f  
<FILE_TRANSFER_SERVER_IP> -b <HV_BSNET_IP>
```

Usage:

```
/onapp/onapp-hv-install/onapp-hv-config.sh[-h CP_HOST_IP] [-p HV_HOST_IP] [-b HV_BSNET_IP] [-f  
FTS_IP] [-a|-i [USER:PASSWD]] [-s] -?
```

Where:

-h CP_HOST_IP	FQDN or IP Address of the management server which should receive all status reports and is authoritative for this compute resource
-p HV_HOST_IP	FQDN or IP Address of Server (the Compute resource) which will serve all stats related and other requests send by the CP_HOST_IP. Used by snmpd, snmptrapd and StorageAPI.
-b HV_BSNET_IP	Compute resource's IP Address from Backup Servers' network Used to bind the SCSI target daemon.
-f FTS_IP	File Transfer Server FQDN or IP address, used for daily cron update recovery ISO by recovery.sh <div style="border: 1px solid yellow; padding: 5px;"> If unsure, set the Control Panel server's management IP as CP_HOST_IP and FILE_TRANSFER_SERVER_IP</div>
-a	Install AoE
-s	Install sshfs
-?	Print this help info

- Run the following commands:

```
# yum install gdisk lsblk-wrapper
```

- Reboot the compute resource to complete the installation:

```
bash#> shutdown -r now
```

- Generate SSH keys:

OnApp requires SSH keys to access various elements of the cloud. The script provided will generate and transfer keys as necessary. The script needs to run on your Control Panel server. It will overwrite any keys that already exist, so if you have custom keys already installed you will need to add them again after running the script. The script will ask you for login details to various servers during the execution. Please follow the onscreen instructions.

- If you are installing a new cloud, SSH into your Control Panel server then download and run the script:

```
bash#> wget http://downloads.repo.onapp.com/install-all-keys.sh  
bash#> /bin/sh install-all-keys.sh
```

- If you are adding additional compute resources to an existing cloud, update the authorized_keys file by running the following script on the Control Panel server:

```
bash#> ssh-copy-id -i /home/onapp/.ssh/id_rsa.pub root@HV_HOST_IP
```

14. Mount the locations for templates and backups:

If you do not have a dedicated backup server you must mount your Template and Backup repository to the compute resource for VS provisioning and backups to work, for example from your Control Panel server:

Add to /etc/exports on the Control Panel server then reboot:

```
/onapp/templates 192.168.10.0/24(rw,no_root_squash)
```

```
/onapp/backups 192.168.10.0/24(rw,no_root_squash)
```

Add to /etc/rc.local on the Compute resource and run them manually on the command line (In this example we are mounting from 192.168.10.101):

```
mkdir -p /onapp/backups && mount -t nfs 192.168.10.101:/onapp/backups /onapp/backups
```

```
mkdir -p /onapp/templates && mount -t nfs 192.168.10.101:/onapp/templates /onapp/templates
```

15. Mount ISO locations:

To rebuild a VS from ISO, it is required to mount and share the location where the ISOs are stored at CP with all the compute resources. When the virtual servers are booted from the ISOs, the ISO is taken from the compute resource server. The location is preconfigured at onapp.yml config file:

- `iso_path_on_cp` - specifies the location where ISOs are stored on the Control Panel server. By default the location is /data. You can change it to any other suitable location. Make sure that this location is shared with the specified `iso_path_on_hv` location.
- `iso_path_on_hv` - specifies the location where ISOs are located on the compute resource servers. By default the location is /data. You can change it to any other suitable location with the `onappowner` and `read/write` access. Make sure that this location is mounted to the specified `iso_path_on_cp` location.

CloudBoot compute resources mount the /data location automatically at boot to the /onapp/tools/recovery on compute resource.

ISOs can be hosted on a dedicated server at any desired location with an arbitrary name if you wish. In this case it is necessary to mount the ISOs' location on this server to the Control Panel `iso_path_on_cp` directory and all the compute resources' `iso_path_on_hv` locations. This can be a backup server to avoid the excess usage of the Control Panel's space.

16. Reboot static compute resources.



If you do not have the /home/mq/onapp/messaging/credentials.yml file on your compute resources and plan to deploy an Accelerator, run the following command on the CP server:

- For all compute resources:

```
rake hypervisor:messaging:configure
```

- For certain compute resources only:

```
rake hypervisor:messaging:configure[ '11.0.50.111 11.0.50.112' ]
```

To perform the configuration for a number of compute resources, separate their IPs with a space.

For information on manual configuration for Accelerator, refer to [RabbitMQ Configuration for Accelerator](#).

This section is the part of the OnApp installation procedure.

[Install Control Panel Server](#) > [Install Compute Resources](#) > [Install Data Stores](#) > [Install Backup Server](#) > [Configure vCloud Director Integration](#) > [Configure Cloud](#)