

OnApp Cloud 6.1 Edge 1 Installation Guide

Table of Contents

1. Preparation	4
2. Installation	5
3. Post Install Configuration	6
4. Installation Workflow	7
4.1. Prepare Servers Configuration	7
4.2. Install Control Panel Server	7
4.3. Install Compute Resources	7
4.4. Install Data Stores	7
4.5. Install Backup Server	8
4.6. Configure Cloud	8
5. Install Control Panel Server	9
6. Install Compute Resources	19
6.1. Install CloudBoot Compute Resources	20
6.1.1. Configure CloudBoot Settings in BIOS	23
6.1.2. Configure InfiniBand	24
6.2. Install Static Compute Resources	24
7. Install Data Stores	31
7.1. Install LVM Data Store	31
7.2. Install Integrated Storage Data Store	32
7.3. Install SolidFire Data Store	33
8. Install Backup Server	35
8.1. Install Static Backup Server	35
8.2. Install CloudBoot Backup Server	38
9. Configure Cloud	43
9.1. Configure Control Panel Settings	43
9.2. Configure Compute Resources	43
9.3. Configure Data Stores	43
9.4. Configure Networks	44
9.5. Configure Backup Servers	44
9.6. Configure Relations Between Entities	44
9.7. Configure Templates	45
9.8. Configure ISOs	45
9.9. Configure OVAs	45
9.10. Configure Buckets	46
10. Configure vCloud Director Integration	47
11. Support	49
11.1. 24/7 Support	49
11.2. What Does OnApp Support in my Cloud?	49
11.3. Professional Services	49
11.4. OnApp Community	49
11.5. Knowledge Base	49
11.6. Documentation	50

This document describes how to install the 6.1 version of the OnApp Cloud. Please read each section carefully, as it is vital to the smooth operation of OnApp Cloud.

1. Preparation

1. Read the [Technical Details](#)
2. Read the [Preparation Guide](#)

2. Installation

1. [Install Control Panel Server](#)
2. [Install Compute Resources](#)
3. [Install Data Stores](#)
4. [Install Backup Server](#)

3. Post Install Configuration

1. [Configure Cloud](#)

4. Installation Workflow

The following workflow includes steps that you run to set up your OnApp installation. This is a standard workflow for OnApp Cloud. You can also integrate OnApp Cloud with VMware solutions, such as [vCloud Director](#) or [vCenter](#).

See also:

[Get Started Guide](#)

[Upgrade Guide](#)



4.1. Prepare Servers Configuration

Before you install OnApp, make sure your network, storage, and servers configuration meet the requirements.



4.2. Install Control Panel Server

The Control Panel server hosts the OnApp user interface and manages all the processes controlled by OnApp. You install the Control Panel server from a ready-made installer package provided by OnApp.



4.3. Install Compute Resources

Compute resources provide CPU, RAM, and storage resources for applications and virtual servers that you or your users run in the cloud. OnApp enables you to install and use Static and CloudBoot compute resources.



4.4. Install Data Stores

Data stores provide storage for templates, backups, ISOs, and virtual server disks. You can set up a separate server with NFS or SSH connection, use any block-based storage, or set up OnApp Integrated Storage.



4.5. Install Backup Server

Backup servers are used to store backups and templates of virtual servers running in the cloud to prevent data loss in the event of failure. You can install Static or CloudBoot backup servers.



4.6. Configure Cloud

After you finish setting up the servers, you can log in to your OnApp Control Panel user interface and configure relations between cloud entities.

5. Install Control Panel Server

Before You Begin

- Review [Get Started Guide](#) to ensure that your environment is ready for the installation.
- Use a corresponding [option](#) of the Control Panel installer in case MySQL is already installed and configured.
- The installer output is redirected to `./onapp-cp-install.log`
- All the installer critical errors are written to `/var/log/messages`
- If you want to reinstall an existing Control Panel, see [Control Panel Migration Guide](#) for instructions.
- If you plan to deploy [Accelerator](#), refer to the [RabbitMQ Configuration for Accelerator](#) document for more details.
- If you need to install other components, for example, OnApp Database Server, RabbitMQ Server, or Redis Server, see [OnApp Installation Components](#) for instructions. You can also [migrate an OnApp database](#) from MySQL to MariaDB, Percona Server, or Percona Cluster.

See also:

[Get Started Guide](#)

[Upgrade Guide](#)

To install the Control Panel server, run the following procedure:

1. Update your server:

```
bash# yum update
```

2. Download the OnApp YUM repository file:

```
bash# rpm -Uvh http://rpm.repo.onapp.com/repo/onapp-
repo-6.0.noarch.rpm
```

3. Install the OnApp Control Panel installer package:

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

4. **(Optional)** You can set custom configuration options for Control Panel. It is important to set custom values before the installer script runs.

The full list of custom configuration options for Control Panel.

Edit the `/onapp/onapp-cp.conf` file to set custom values for Control Panel:

#Template server URL

```
TEMPLATE_SERVER_URL='http://templates-manager.onapp.com'
```

IPs (separated with coma) list for the SNMP to trap. This is the list of Control Panel IP addresses on which the traps sent from the compute resources are processed.

```
SNMP_TRAP_IPS=""
```

OnApp Control Panel custom version

```
ONAPP_VERSION=""
```

OnApp MySQL/MariaDB connection data (database.yml)

```
ONAPP_CONN_WAIT_TIMEOUT=15
ONAPP_CONN_POOL=30
ONAPP_CONN_RECONNECT='true'
ONAPP_CONN_ENCODING='utf8'
```

MySQL/MariaDB server configuration data (in case of local server)

```
MYSQL_WAIT_TIMEOUT=604800
MYSQL_MAX_CONNECTIONS=500
MYSQL_LIMITNOFILE=8192
```

[Use MariaDB instead of MySQL as OnApp database server](#) (Deprecated parameter. If you set any values for this parameter, they will not take effect)

```
WITH_MARIADB=0
```

Configure the database server relative amount of available RAM

```
TUNE_DB_SERVER=1
```

The number of C data structures that can be allocated before triggering the garbage collector. It defaults to 8 million. Only change this value if you understand what it does.

```
RUBY_GC_MALLOC_LIMIT=16000000
```

sysctl.conf net.core.somaxconn value

```
NET_CORE_SOMAXCONN=2048
```

The root of OnApp database dump directory (on the Control Panel box)

```
ONAPP_DB_DUMP_ROOT=""
```

Remote server's (to store database dumps) IP, user, path, openssh connection options and number of dumps to keep

```

DB_DUMP_SERVER=""
DB_DUMP_USER="root"
DB_DUMP_SERVER_ROOT="/onapp/backups"
DB_DUMP_SERVER_SSH_OPT="-o StrictHostKeyChecking=no -o
UserKnownHostsFile=/dev/null -o PasswordAuthentication=no"
KEEP_DUMPS=168
DB_DUMP_CRON='40 * * * *'

```

[Enable monit - tool for managing and monitoring Unix systems](#)

```
ENABLE_MONIT=1
```

DEPRECATED: If enabled (the 1 value is set) - install (if local box) and configures RabbitMQ Server (messaging system) for the vCloud support. (Deprecated parameter. If you set any values for this parameter, they will not take effect)

```
ENABLE_RABBITMQ=1
```

Rotate transactions' log files created more than TRANS_LOGS_ROTATE_TIME day(s) ago

```
TRANS_LOGS_ROTATE_TIME=30
```

Maximum allowed for uploading file size in bytes, from 0 (meaning unlimited) to 2147483647 (2GB). Default is 0.

```
MAX_UPLOAD_SIZE=0
```

Timeout before ping Redis Server to check if it is started. Default is 10 sec.

```
REDIS_PING_TIMEOUT=10
```

OnApp Control Panel SSL certificates (please do not change if you aren't familiar with SSL certificates)

* The data below to generate self-signed PEM-encoded X.509 certificate

```

SSL_CERT_COUNTRY_NAME="UK"
SSL_CERT_ORGANIZATION_NAME='OnApp Limited'
SSL_CERT_ORGANIZATION_ALUNITNAME='OnApp Cloud'
SSL_CERT_COMMON_NAME="`hostname --fqdn 2>/dev/null`"

```

SSLCertificateFile, SSLCertificateKeyFile Apache directives' values
ssl_certificate, ssl_certificate_key Nginx directives' values

```
SSLCERTIFICATEFILE="/etc/pki/tls/certs/ca.crt"
SSLCERTIFICATECSRFILE="/etc/pki/tls/private/ca.csr"
SSLCERTIFICATEKEYFILE="/etc/pki/tls/private/ca.key"
```

```
# * PEM-encoded CA Certificate (if custom one exists)
# SSLCACertificateFile, SSLCertificateChainFile Apache directives' values
# ssl_client_certificate Nginx directives' values
```

```
SSLCACERTIFICATEFILE=""
SSLCERTIFICATECHAINFILE=""
```

```
# SSLCipherSuite, SSLProtocol Apache directives' values
# ssl_ciphers, ssl_protocols Nginx directives' values
```

```
SSLCIPHERSUITE=""
SSLPROTOCOL=""
```

```
bash# vi /onapp/onapp-cp.conf
```

For successful installation you need to accept the EULA.

Before installing the OnApp Control Panel package, apply the following changes to the `/onapp/onapp-cp.conf` file to avoid the RC4 Cipher Suite ([CVE-2015-2808](#)) vulnerability.

```
SSLCIPHERSUITE="EECDH+ECDSA+AESGCM:EECDH+aRSA+AESGCM:EECDH+ECDSA+SHA384
:EECDH+ECDSA+SHA256:EECDH+aRSA+SHA384:EECDH+aRSA+SHA256:EECDH+aRSA+RC4:
EECDH:EDH+aRSA:RC4:!aNULL:!eNULL:!LOW:!3DES:!MD5:!EXP:!PSK:!SRP:!DSS:!
RC4:!3DES"
SSLPROTOCOL="all -SSLv2 -SSLv3 -TLSv1"
```

5. Run the Control Panel installer:

```
bash#> /onapp/onapp-cp-install/onapp-cp-install.sh -i SNMP_TRAP_IPS
```

The full list of installer options for Control Panel.

Usage:

```
# /onapp/onapp-cp-install/onapp-cp-install.sh -h
Usage: /onapp/onapp-cp-install/onapp-cp-install.sh [-c CONFIG_FILE]
[--mariadb | --community | --percona | --percona-cluster] [-m
MYSQL_HOST] [--mysql-port=MYSQL_PORT] [--mysql-sock[=MYSQL SOCK] [-p
MYSQL_PASSWD] [-d MYSQL_DB] [-u MYSQL_USER] [-U ADMIN_LOGIN] [-P
ADMIN_PASSWD] [-F ADMIN_FIRSTNAME] [-L ADMIN_LASTNAME] [-E
ADMIN_EMAIL] [-v ONAPP_VERSION] [-i SNMP_TRAP_IPS] [--redis-
host=REDIS_HOST] [--redis-bind[=REDIS_BIND] [--redis-
passwd[=REDIS_PASSWD] [--redis-port=REDIS_PORT] [--redis-
sock[=REDIS SOCK] [--rbthost RBT_HOST] [--vcdlogin VCD_LOGIN] [--
vcdpasswd VCD_PASSWD] [--vcdvhost VCD_VHOST] [--rbtlogin RBT_LOGIN]
[--rbtpasswd RBT_PASSWD] [-a] [-y] [-D] [-t] [--noservices] [--ha-
install] [--rake=RAKE_TASKS] [--quick|--quick-update[=SERVICE] [--
accept-eula] [-w] [-h]
```

Where:	
Database server options:	Default database SQL server is MySQL Server. Please use one of the following option to install LOCALLY.
--mariadb	MariaDB Server
--community	MySQL Community Server
--percona	Percona Server
--percona-cluster	Percona Cluster
MYSQL_*	Options are useful if MySQL is already installed and configured.
-m MYSQL_HOST	MySQL host. Default is 'localhost'
--mysql-port=MYSQL_PORT	TCP port where MySQL Server serves connections. Default values is 3306 for the local installation
--mysql-sock[=MYSQL SOCK]	Unix socket on which MySQL Server serves connections. Default values is /var/lib/mysql/mysql.sock. Used if local server only. The socket is unset if the option's argument isn't specified.
-p MYSQL_PASSWD	MySQL password. Random is generated if is not set or specified.
-d MYSQL_DB	OnApp MySQL database name. Default is 'onapp'.
-u MYSQL_USER	MySQL user. Default is 'root'.
Redis Server options:	
REDIS_*	Options are useful if Redis Server is already installed and configured.

--redis-host=RE DIS_HOST	IP address/FQDN where Redis Server runs. It is used by Control Panel to connect to Redis Server. The Redis Server will be installed and configured on the current box if localhost/127.0.0.1 or box's public IP address (listed in SNMP_TRAP_IPS) is specified. Default value is 127.0.0.1. If local Redis, it will serve as well on the unix socket 'PORT' (if --redis-sock without argument isn't specified).
--redis-bind[=RE DIS_BIND]	The IP address for Redis Server to serve connections (to listen). The option isn't mandatory.
--redis-port=RE DIS_PORT	Redis Server listen port. Defaults are: 0 - if local server 6379 - if remote server
--redis-passwd[=REDIS_PASSWORD]	Redis Server password to authenticate. Random password is generated if the option's argument isn't specified. By default no password is used for local Redis.
--redis-sock[=RE DIS_SOCKET]	Path to the Redis Server's socket. Used if local server only. Default is /var/run/redis/redis.sock. The socket is unset if the option's argument isn't specified.
Options to manage the OnApp Control Panel administrator account:	
ADMIN_*	Options are used to configure OnApp Control Panel administrator data. Please note that these options are for NEW INSTALL only and not for upgrade
-P ADMIN_PASSWORD	CP administrator password
-F ADMIN_FIRSTNAME	CP administrator first name
-L ADMIN_LASTNAME	CP administrator last name
-E ADMIN_EMAIL	CP administrator e-mail
RabbitMQ Server and vCloud options:	
--rbthost RBT_HOST	IP address/FQDN where RabbitMQ Server runs. The RabbitMQ will be installed and configured on the current box if localhost/127.0.0.1 or box's public IP address (enlisted in SNMP_TRAP_IPS) Default value is 127.0.0.1.
VCD_*	Options are usefull if vCloud/RabbitMQ are already installed and configured.
--vcdlogin VCD_LOGIN	RabbitMQ/vCloud user. Default value is 'rbtvcd'.

-- vcdpasswd VCD_PAS SSWD	RabbitMQ/vCloud user password. The random password is generated if isn't specified.
-- vcdvhost VCD_VH OST	RabbitMQ/vCloud vhost. Default value is '/'
RBT_*	Options are used to configure RabbitMQ manager account. If local RabbitMQ server.
--rbtlogin RBT_LO GIN	RabbitMQ manager login. The default value is 'rbtmgr'.
-- rbtpassw d RBT_PAS SWD	RabbitMQ manager password. The random password is generated if isn't specified.
General options:	
--rake RAKE_T ASKS	List of OnApp Control Panel rake tasks (separated with space) to run at the very end of install or upgrade.
-v ONAPP_ VERSION	Install custom OnApp CP version
-i SNMP_T RAP_IPS	IP addresses separated with coma for snmp to trap
-y	Update OS packages (except of OnApp provided) on the box with 'yum update'.
-a	Is not interactive. Process with automatic installation. Please note, this will continue OnApp Control Panel install/upgrade even if there is transaction currently running.
-t	Add to the database and download Base Templates. For new installs only. If this option is not used, then only the following mandatory System Templates will be added by default during fresh install: OnApp CDN Appliance; Load Balancer Virtual Appliance; Application Server Appliance.
-- noservic es	Do not start OnApp services: monit, onapp and httpd Please note, crond and all OnApp's cron tasks remain running. They could be disabled by stopping crond service manually for your own risk.
-D	Do not make database dump, and make sure it is disabled in the cron and not running at the moment.
-w	Do not disable the iptables service. It is applicable on fresh installs only.

--quick --quick-update[=SERVICE]	Proceed with quick update procedure. This will skip update and configuration for services, such as system packages, MySQL database, Redis Server, RabbitMQ Server, and Monit service. Set the SERVICE parameter (space separated list of statements) to define services, which need to be updated. Possible reserved statements are: rpms - for 'system packages' upgrade; mysql - for MySQL database upgrade and configuring; redis - for ERedis Server upgrade and configuring; rabbitmq - for RabbitMQ Server upgrade and configuring; monit - for Monit upgrade and configuring.
--accept-eula	Automatically accept OnApp's End User License Agreement.
-c CONFIG_FILE	Custom installer configuration file. Otherwise, preinstalled one is used.
-h	Print this info

Perform the steps from six to nine (6-9) only if you are going to use CloudBoot and/or OnApp Storage.

6. Install CloudBoot dependencies:

Depending on a compute resource type, you should install `onapp-ramdisk-DISTRO-FLAVOR` package(s) where:

```
DISTRO: centos6, centos7
FLAVOR: xen, kvm
```

It is required to install `yum install onapp-ramdisk-centos7-default` together with `onappstore` packages.

It is recommended to install all the following packages:

```
bash#> # yum install onapp-ramdisk-centos6-kvm
bash#> # yum install onapp-ramdisk-centos6-xen
bash#> # yum install onapp-ramdisk-centos7-kvm
bash#> # yum install onapp-ramdisk-centos7-xen
bash#> # yum install onapp-ramdisk-centos7-default
bash#> /onapp/onapp-store-install/onapp-store-install.sh
```

Any of the `ramdisk` packages install the following packages automatically:

```
onapp-store-install
onapp-ramdisk-tools
```

7. Run the following utility to configure and customize ramdisks:

```
bash# /onapp/onapp-store-install/onapp-store-install.sh
```

8. Install an OnApp license to activate your Control Panel. Enter a valid license key via the OnApp UI. Your default OnApp credentials are **admin/changeme**. You can change a password via the Control Panel > **Users** menu.

After you enter a license key, it may take up to 15 minutes to activate the key.

9. Restart the OnApp service:

```
bash#> service onapp restart
```

Perform the following step (10) only if you plan to deploy Accelerator.

10. If you plan to configure an Accelerator, run the following command:

- o For all compute resources:

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

- o 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 IP addresses with a space.

The command above runs on compute resources that are online. If some compute resources are offline, you should run the command again when they are online. The *rabbitmq_host* parameter in the *on_app.yml* file should contain the real IP address of a server with RabbitMQ installed. The *rabbitmq_host* parameter should not be set to 'localhost' or '127.0.0.1'. The server with RabbitMQ installed should be available from the compute resources. For information on manual configuration for Accelerator, refer to [RabbitMQ Configuration for Accelerator](#).

11. Generate SSH keys that OnApp requires for you to access various elements of the cloud. The script provided generates and transfers keys as necessary. The script needs to be run on your Control Panel server. It overwrites any keys that already exist so if you have custom keys already installed, you need to add them again after running the script. You need to provide your login details to various servers during the script execution. Please follow the onscreen instructions.
12. If you install a new cloud instance, connect to your Control Panel server via SSH, download, and run the script:

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

Note that the script is applicable only if you configure your first compute resources in the cloud. If you already have some existing compute resources and install new ones, the script sets new SSH keys to all the compute resources, both new and existing.

13. If you add 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 repositories to compute resources. If your template and backup repositories are located on the Control Panel server, you can mount them as follows:

Add the repositories to `/etc/exports` on the Control Panel server and then restart the NFS service:

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

15. After you installed the Control Panel server, configure your [Cloud Settings](#).

6. Install Compute Resources

Please do not create mixed compute zones. Do not add CloudBoot and static compute resources to one compute zone, as well as Xen and KVM compute resources to one compute zone. The reason is that Xen virtual servers cannot be migrated to a KVM compute resource and vice versa.

After you successfully installed the Control Panel server, you can follow the next procedures to set up Xen or KVM compute resources:

- [Install CloudBoot Compute Resources](#)
- [Install Static Compute Resources](#)

On this page:

[Install CloudBoot Compute Resources](#)

- [Configure CloudBoot Settings in BIOS](#)
- [Configure InfiniBand](#)

[Install Static Compute Resources](#)

See also:

[Install Control Panel Server](#)

[Install Data Stores](#)

[Install Backup Server](#)

6.1. Install CloudBoot Compute Resources

Follow this installation to enable CloudBoot for your compute resources. The 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 (set up in 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 virtual servers provisioning when using CloudBoot.

1. Enable CloudBoot on the Control Panel:
 - Go to **Settings > Configuration > System > CloudBoot**.
 - Scroll down to the CloudBoot section and check the **Enable** box.
2. Enable Storage on the Control Panel:
 - Go to **Settings > Configuration > System > OnApp Storage**.
 - Scroll down to the OnApp Storage section and check the **Enable OnApp Storage** box.
 - Tick the **Use Local Read Path** checkbox to minimize the network throughput dependency for read heavy workloads.
3. Enter IP addresses for static content target and Control Panel server CloudBoot interface. Enter the relevant IPs in **Settings > Configuration > System > CloudBoot**. Static content, such as CloudBoot images, kernels, virtual server templates, can be hosted on a standalone NFS server if you want. The default configuration is to install everything on the Control Panel server.
4. Add CloudBoot IP addresses for compute resources:
 - Create a CloudBoot IP Net: **Settings > Compute Resources > CloudBoot IPs tab > New IP Net**.

- Add an IP range to the new IP Net:
Settings > Compute Resources > CloudBoot IPs > Actions icon next to the *IP net* > **Add New IP Range**.
 - Add CloudBoot IP Addresses:
Settings > Compute Resources > CloudBoot IPs > Create 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 via **Settings > Compute Resources > Add 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-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/images/centos7/ramdisk-default/` to `/.ro`
The path may vary depending on the compute resource template used.
 - `/data` to `/onapp/tools/recovery`

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 repositories to compute resources. If your template and backup repositories are located on the Control Panel server, you can mount them as follows:

Add locations 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 locations to Custom Config on a 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
```

If you want to deploy Accelerator, run the following steps.

1. Run the following command:

```
su onapp
cd /onapp/interface
```

2. Run the following command on the Control Panel server:

- o For all compute resources:

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

- o 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 IP addresses with a space. 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{y31?s8l' #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
```

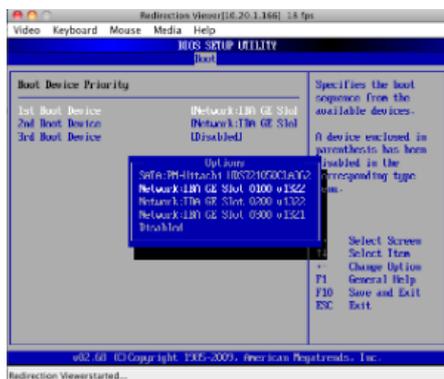
For information on manual configuration for Accelerator, refer to [RabbitMQ Configuration for Accelerator](#).

6.1.1. Configure CloudBoot Settings in BIOS

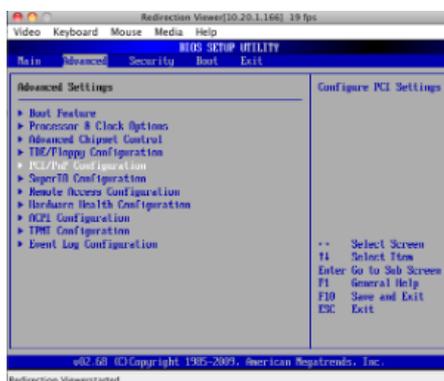
Your BIOS settings may vary from the example provided in this section.

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

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



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



4. Set a local disk as a second boot device.

6.1.2. Configure InfiniBand

Your hardware must meet the following requirements for the Ethernet mode utilization:

- VPI enabled switches (including a proper license key)
- VPI adapter cards (HCAs)

There are the following limitations for InfiniBand:

- Be aware that InfiniBand is not supported for CloudBoot CentOS 7 due to the lack of integrated OFED stack.
- InfiniBand in Ethernet mode is supported only for CentOS 6 KVM nodes.
- InfiniBand is supported only for the SAN network, not PXE boot.

To enable the InfiniBand mode:

1. Go to **Control Panel > Settings > Configuration** on the OnApp CP.
2. Move the *Enable InfiniBand Boot* slider in the CloudBoot section.
3. Click the **Save Configuration** button.

CentOS 6 KVM

6.2. Install Static Compute Resources

Before You Proceed

- Install base CentOS packages on the local drive before compute resource installation, depending on which virtualization method you choose:
 - Xen 4 compute resources: CentOS 6.x x86/64 or CentOS 7.x x86/64
 - KVM compute resources: CentOS 6.x x64 or CentOS 7.x x86/64
- We recommend installing CentOS from the minimal CentOS ISO for static compute resources.
- Disable CPU power-saving features in BIOS before you proceed to the compute resource installation.
- If you are not using a dedicated backup server in your cloud setup, configure NFS server with the following options to preserve files owner and group settings during template unpacking on NFS storage:
 - `no_root_squash`
 - `no_all_squash`
- Pay attention that smart and baremetal servers cannot be installed using the static compute resource installation method.
- If you plan to deploy baremetal servers, you can [enable recovery mode for baremetal servers](#).

To install a compute resource:

1. Add a compute resource to your cloud using the OnApp Control Panel: **Settings > Compute resources > Add New Compute Resource**. Make sure that the compute resource is visible in the Control Panel and at this point is shown 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-6.0.noarch.rpm
```

4. Install the OnApp compute resource installer package:

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

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

If you deploy Xen to a server running CentOS 6, it is important to specify a number for `XEN_DOM0_MAX_VCPUS`. We recommend to set the value 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
```

The full list of custom values for a compute resource.

- OnApp HV tools [custom version](#)

```
HV_VERSION=""
```

- OnApp StorageAPI [custom version](#)

```
API_VERSION=""
```

- Enable monit - tool for managing and monitoring Unix systems

```
ENABLE_MONIT=1
```

- 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 can be doubled (or even more, depending on the memory amount).

```
NET_IPV4_NETFILTER_IP_CONTRACK_MAX=""
```

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

```
hashsize = nf_contrack_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
```

6. Run the OnApp compute resource installer script, depending on the compute resource:
 - **For Xen**

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

The full list of Xen installer options.

Usage:

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

Where:

-c CONFIG_FILE	Custom installer configuration file. Otherwise, the pre-installed one is used.
-v HV_VERSION	Custom compute resource Tools version
-p API_VERSION	Custom StorageAPI version
-a	Do NOT be interactive. Process with automatic installation.
-y	Update OS packages (except for OnApp provided) with 'yum update'.
-t	Initiate Recovery templates and ISO(s), which are used to provision FreeBSD guests, download. The download is initiated if the '-a' option is used.
-s	Skip packages management: install, remove, upgrade. None of the packages are installed/updated/removed.
-x	Skip xen*, kernel*, and libvirt* (compute resource related RPM) packages management. Useful for update (not for fresh install).
-h	Print this info

○ **For KVM**

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

The full list of KVM installer options.

Usage:

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

Where:

-c CONFIG_FILE	Custom installer configuration file. Otherwise, the pre-installed 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
-t	Initiate Recovery templates and ISO(s), which are used to provision FreeBSD guests, download. The download is initiated if the '-a' option is used.
-y	Update OS packages (except those provided by OnApp for a compute resource) with 'yum update'. Useful for update (not for fresh install).
-s	Skip packages management: install, remove, upgrade. Useful for update (not for fresh install).
-x	Skip kvm, kernel, and libvirt (compute resource related RPM) packages management. Useful for update (not for fresh install).
-h	Print this info

7. 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> -b <HV_BSNET_IP>
```

The full list of compute resource (both Xen and KVM) configuration options.

Usage:

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

Where:

-h CP_HOST_IP	FQDN or IP Addresses (separated by a comma) of the management server, which should receive all status reports and are authoritative for this compute resource. Used by <i>snmpd</i> and <i>snmptrapd</i> .
-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 <i>snmpd</i> , <i>snmptrapd</i> , 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 If unsure, set the Control Panel server's management IP as CP_HOST_IP and FILE_TRANSFER_SERVER_IP.
-l LVM_HOST_ID	The lvmlockd sanlock host_id. The ID should be unique for each compute resource in the cloud. It's value for Static compute resources are between 257 and 2000.
-a	Install AoE
-s	Install sshfs
-?	Print this help info

8. Reboot the compute resource to complete the installation:

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

9. Generate SSH keys that OnApp requires for you to access various elements of the cloud. The script provided will generate and transfer keys as necessary. The script needs to be 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 your login details to various servers during the execution. Please follow the onscreen instructions.
10. If you install a new cloud instance, connect to your Control Panel server via SSH, download, and run the script:

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

11. If you are add 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
```

12. If you do not have a dedicated backup server, you must mount your template and backup repositories to compute resources. If your template and backup repositories are located on the Control Panel server, you can mount them as follows:

Add locations to `/etc/exports` on the Control Panel server and 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 locations 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
```

13. Mount ISO locations:

To build virtual servers from ISO images, it is required to mount and share the location where the ISOs are stored at the Control Panel server with all the compute resources. When virtual servers are booted from ISOs, the ISO image is taken from the compute resource server. The location is preconfigured in *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` user 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 a compute resource.

ISOs can be hosted on a dedicated server at any desired location with an arbitrary name if you want. 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 resources.

14. Reboot static compute resources.

To edit the configuration for Accelerator manually, perform the following steps:

1. Copy file:

```
cp /home/mq/onapp/messaging/credentials{_example,}.yml
```

2. Open `vi /home/mq/onapp/messaging/credentials.yml` and check the following details:

```
---
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{y31?s81' #RABBITMQ ACCESS PASSWORD
queue: 'hv-10.0.50.102' # hv-[IP Address of Compute Resource]
exchange:
  name: 'acceleration'
  type: 'direct'
  durable: True
```

3. Change an owner:

```
chown -R mq:mq /home/mq
```

4. Run the following commands:

```
service onapp-messaging start
monit monitor onapp-messaging
```

7. Install Data Stores

After you successfully installed the Control Panel server and compute resources, you can follow one of the next procedures to set up data stores:

- [Install LVM Data Store](#)
- [Install Integrated Storage Data Store](#)
- [Install SolidFire Data Store](#)

Before You Begin

- To configure an Integrated Storage data store, see the [Storage](#) guide.
- The following procedures assume that you have already configured a compute resource to see the iSCSI/ATAoE block device to connect to and that the SAN disk will be shown when running `fdisk -l`.
- All compute resources need access to the same data store. Ensure that you have the block device visible on all compute resources.
- OnApp doesn't support the 4 K block size for local LVM data stores.
- **VERY IMPORTANT:** Perform the procedures only once per data store! Take care when selecting the disk/partition you want to use for storing virtual servers data!

See also:

[Install Control Panel Server](#)

[Install Backup Server](#)

7.1. Install LVM Data Store

To install an LVM data store, follow the next procedure:

1. Add the new data store to OnApp via the Control Panel user interface:
 - Go to your Control Panel > **Settings** menu.
 - Click the **Data Stores** icon.
 - Click the **Create Data Store** button at the bottom of the screen.
 - Follow the steps in the creation wizard:

Step 1

- Enter a *Label* and *IP Address* for your data store.
- Select the data store type: **Lvm**.
- Move the slider to the right to enable a data store. When disabled, OnApp will not allow new disks to be created automatically on that data store. This

is useful to prevent an established data store from becoming too full. It also lets you prevent the automatic creation of root disks on 'special' data stores (high speed, etc).

- Click **Next**.

Step 2

- Set disk capacity in GB.
 - If required, you can also bind the data store with a local compute resource. This is helpful if you want that the data store and a compute resource were located on the same physical server, thus decreasing the time needed for a compute resource-data store connection.
 - If required, you can also assign the data store to a data store zone. The drop-down menu lists all data store zones set up in the cloud (to add or edit data store zones, see [Data Store Zones](#)).
- When you are finished configuring the data store, click the **Create Data Store** button.
To use the data store, you have to assign it either to a compute resource or a compute zone.
2. Find the data store unique identifier (this is needed to create your volume group in step #4): (Read the IDENTIFIER from the data stores screen: http://xxx.xxx.xxx.xxx/settings/data_stores)
 3. Connect via SSH to a compute resource that is able to connect to this data store. Create the physical volume:

```
bash#> pvcreate --metadatasize 50M /dev/xxx
```

Replace xxx with the real device.

4. Create the volume group:

```
bash#> vgcreate onapp-IDENTIFIER /dev/xxx
```

Replace xxx with the real device and *IDENTIFIER* with the info from the data store page on the UI.

5. Test compute resource/volume group visibility. Now when you have the new data store formatted, you should be able to see the volume group from all compute resources. To test this, run *pvscan* and *vgscan* on all compute resources. Make sure you can see all identifiers on all compute resources.

7.2. Install Integrated Storage Data Store

Before creating an Integrated Storage data store:

1. Create one or more Xen or KVM compute resource with Integrated Storage enabled to group their drives together into a virtual data store.
2. Create a compute zone.

3. Add your compute resources to the compute zone.

After running these steps, you can proceed to the Integrated Storage data store creation.

To create a new Integrated Storage data store:

1. Go to your Control Panel > **Integrated Storage** > **Data Stores** menu.
2. On the screen that appears, you'll see the list of all Integrated Storage data stores in the cloud.
3. To create a new data store, click the **Create New Integrated Storage Data Store** button, and complete the wizard that follows:

- *Name* - give your data store a name
- **Show advanced options** - select this check box to reveal the list of advanced settings:
 - *Replicas* - specify the number of data copies to increase the resilience to individual drive failure. You can specify 1, 2 or 4 replicas.
 - *Stripes* - specify the number of data splittings to increase the number of physical disks included to the virtual disk. You can specify 0, 2 or 4 stripes.
 - *Overcommit* - specify the over-provisioning percentage. You can set the following overcommit values: none (0%), 20%, 50% or unlimited (100%).

In order for your hard drives (nodes) to be detected and active, multicast traffic should be enabled on your switch, for the Onapp Integrated Storage Network/VLAN.

- **Storage Nodes**
 - *Filter by compute resource* - use this to filter the nodes (disks) available for inclusion in this data store, by specific compute resources.
 - *Filter by performance* - use this to filter the nodes available for inclusion in this data store by performance.

4. Click the **Save** button to create the data store. The data store must be assigned to a compute zone and data store zone before you can provision storage to a VS.

7.3. Install SolidFire Data Store

You can create one SolidFire data store per cloud that will represent the space available at the SolidFire side.

To create a SolidFire data store:

1. Go to your Control Panel > **Settings** menu.
2. Click the **Data Stores** icon.
3. Click the **Create Data Store** button at the bottom of the screen.
4. Follow the steps in the creation wizard:

Step 1 of 3

- Enter a data store label.
- Specify an IP address to be used for managing the data store via CP (Inasmuch SolidFire data stores have two interfaces, you'll have to specify the IP address for the cluster admin later.)
- Select a **SolidFire** data store type.
- Move the slider to the right to enable a data store. When disabled, OnApp will not allow new disks to be created automatically on that data store. This is useful to

prevent an established data store from becoming too full. It also lets you prevent the automatic creation of root disks on 'special' data stores (high speed, etc).

- Click **Next**.

Step 2 of 3

- Set disk capacity in GB.
- If required, you can also bind the data store with a local compute resource. This is helpful if you wish that the data store and a compute resource were located on the same physical server thus decreasing the time needed for a compute resource-data store connection.
- If required, you can also assign the data store to a data store zone. The drop-down menu lists all data store zones set up in the cloud (to add or edit data store zones, see the section on Data store zones in the Settings section of this guide).

Step 3

- Specify the cluster Admin settings:
 - *iSCSI IP* - iSCSI IP address
 - *Username* - specify a username for cluster authorization
 - *Password* - specify a password for cluster authorization
- Specify the SolidFire Account settings:
 - *Username* - specify SolidFire account username
 - *Initiator secret* - specify iSCSI initiator secret (optional)
 - *Target secret* - specify iSCSI initiator secret (optional)

Initiator secret and *target secret* are optional parameters. They are created automatically for a newly created account. For the new account they will be taken from the SolidFire database. If you specify target and initiator secrets for an existing user, they will be overwritten.

5. When you've finished configuring the store, click the **Create Data Store** button.

8. Install Backup Server

Follow one of the following procedures to set up a backup server in your cloud:

- [Install Static Backup Server](#)
- [Install CloudBoot Backup Server](#)

See also:

[Install Control Panel Server](#)

[Install Data Stores](#)

8.1. Install Static Backup Server

To install a static backup server, run the following procedure. Skip this section if you use CloudBoot. We recommend installing CentOS from the minimal CentOS ISO for static backup servers.

1. Add a backup server via the Control Panel user interface:
 - Go to your Control Panel > **Settings** menu and click the **Backup Servers** icon.
 - Click the **Create Backup Server** button.
 - Fill in the form that appears:
 - *Label* - give your backup server a label
 - *IP address* - enter the backup server IP address (IPv4)
 - *Backup IP address* - add a provisioning network IP address
 - *Capacity* - set the backup server capacity (in GB)
 - *Backup server zone* - select the backup server zone to which this backup server will be assigned.
- Move the **Enabled** slider to the right to enable the backup server.
- Click the **Add Backup Server** button.
- Update your server:


```
bash# yum update
```
- Download the OnApp repository:


```
bash# rpm -Uvh http://rpm.repo.onapp.com/repo/onapp-  
repo-6.0.noarch.rpm
```
- Install the OnApp Backup Server installer package:


```
bash# yum install onapp-bk-install
```
- Check and set the backup server default settings:

```
bash# vi /onapp/onapp-bk.conf
```

Edit the backup server default settings in the /onapp/onapp-bk.conf file.

OnApp BK tools custom version

```
BK_VERSION=""
```

OnApp StorageAPI custom version

```
API_VERSION=""
```

Default server to synch time on the HV

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

The number of retries for WGET to download the file

```
WGET_TRIES=5
```

OnApp templates directory

Please refer to the corresponding settings at OnApp Control Panel web interface

```
TEMPLATES_DIR='/onapp/templates'
```

OnApp backups directory

Please refer to the corresponding settings at OnApp Control Panel web interface

```
BACKUPSS_DIR='/onapp/backups'
```

- Run the installer. It is recommended to download Base, Load Balancer, and CDN templates while running the installer. You may rerun the installer later with the `-t` option.

```
bash# sh /onapp/onapp-bk-install/onapp-bk-install.sh
```

The full list of installer options.

Usage:

```
/onapp/onapp-bk-install/onapp-bk-install.sh [-c CONFIG_FILE] [-a] [-y] [-t] [-b] [-v BK_VERSION] [-p API_VERSION] [-h]
```

Where:

-c CONFIG_FILE	Custom installer configuration file. Otherwise, the pre-installed one is used.
-a	Do NOT be interactive. Process with automatic installation.
-v BK_VERSION	Custom backup tools version
-p API_VERSION	Custom StorageAPI version
-y	Update OS packages (except for the ones provided by OnApp) with 'yum update'.
-t	Initiate the Load Balancer Virtual Appliance, OnApp CDN Appliance, and Application Server Appliance templates download. The download is initiated if the '-a' option is used.
-b	Initiate the Base templates download
-h	Print this info

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

```
bash#> /onapp/onapp-bk-install/onapp-bk-config.sh -h <CP_HOST_IP> -p <BK_HOST_IP>
```

The full list of configuration options.

Usage:

```
/onapp/onapp-bk-install/onapp-bk-config.sh [-h CP_HOST_IP] [ -p BK_HOST_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 backup server.
-p BK_HOST_IP	FQDN or IP Address of Backup Server which will serve all stats related and other requests send by the CP_HOST_IP. Used by snmpd and StorageAPI.
-a	Install AoE
-i [USER:PASSWD]	Install iSCSI utils and configure with USER and PASSWD (if specified)
-s	Install sshfs
-?	Print this help info

- Run the following commands:

```
# yum install gdisk fuse fuse-libs fuse-devel qemu-img libvmdk libvmdk-tools
```

- 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
```

- Mount the locations for OVA. If you enable building VSs from [OVAs](#) in your cloud, you should mount the source location to which OVA is uploaded on CP with a destination location where OVA is stored on a backup server. The location is preconfigured in the *onapp.yml* config file:
 - **Source Path** (*ova_path*) - specifies the location where OVAs are downloaded/uploaded on the Control Panel server. By default the location is */data*.
 - **Destination Path** (*ova_path*) - specifies the location where OVAs are stored on backup servers. By default the location is */data*.

To mount the locations:

- Add locations (e.g. */data*) to */etc/exports* on the Control Panel server and then restart the nfs service:

```
/data 192.168.10.0/24(rw,no_root_squash)
```

- Add locations to */etc/rc.local* on the backup server and run them manually on the command line (in this example we are mounting from 192.168.10.101):

```
mkdir -p /data && mount -t nfs 192.168.10.101:/data /data
```

8.2. Install CloudBoot Backup Server

CloudBoot backup servers are CloudBoot KVM compute resources that can be used as backup servers. Follow the step-by-step instructions provided in this section to configure CloudBoot backup servers in your cloud.

Before You Begin

- You should configure some local or remote attached storage for persistent backups on the provisioning/backup server. We strongly recommend you to deploy one or more backup servers for backups and virtual servers provisioning when using a CloudBoot functionality.
- We strongly recommend you to deploy one or more backup servers on your cloud. Incremental backups are only supported with a dedicated backup server.
- CentOS now defaults to NFSv4. This is known to cause compatibility issues so we strongly recommend that you use NFSv3 for all mounts. This can be done by passing `-t nfs -o vers=3` in any mount commands.
- To use a backup server, you have to add it to a [backup server zone](#) and assign it either to a [compute resource](#) or a [compute zone](#).

To create a CloudBoot backup server:

1. Add a backup server via the Control Panel user interface:
 - Go to your Control Panel > **Settings** menu and click the **Backup Servers** icon.
 - Click the **Create Backup Server** button.
 - Fill in the form that appears:
 - *Label* - give your backup server a label
 - *IP address* - enter the backup server IP address (IPv4)
 - *Backup IP address* - add a provisioning network IP address
 - *Capacity* - set the backup server capacity (in GB)
 - *Backup server zone* - select the backup server zone to which this backup server will be assigned.
- Move the **Enabled** slider to the right to enable the backup server.
- Click the **Add Backup Server** button.
- Update CloudBoot and CP server RPMs:


```
yum update onapp-store-install yum update onapp-cp-install
```
- Configure CloudBoot settings:


```
/onapp/onapp-store-install/onapp-store-install.sh
```
- Create a new CloudBoot compute resource: at the first step of the creation wizard, select the type - Backup (CloudBoot Provisioning and Backup Resource). Refer to the [Create CloudBoot Compute Resource](#) section of this guide for details.
- Go to your Control Panel > **Settings** menu and then click the **Compute Resources** icon.
- Click the **Add New CloudBoot Compute Resource** button.

- In the wizard menu, click Backup (CloudBoot Provisioning and Backup Resource).
- Select right MAC address of it.
- Fill in the form that appears:
 - *Label* - give the compute resource a name
 - *Pxe IP address* - select an IP address for this compute resource from the address pool available
 - *Enabled* - move the slider to the right to allow VSs to be installed/booted on this compute resource
 - *Compute Zone* - select the compute zone, to which this compute resource will be assigned, from the drop-down list
 - *Custom Config* - specify any custom commands you want to run when compute resource is booted
- At this step, specify the Backup Server properties:
 - *Label* - give your backup server a label
 - *IP address* - enter the IP address of a Compute resource you have created at step 1
 - *Backup IP address* - add a provisioning network IP address
 - *Capacity* - set the backup server capacity (in GB)
- Then, you shouldn't assign any drive, but it needs right NIC to be applied to the storage network.
- After that, assign your backup server to the backup server zone. If you want to attach LVM-based storage and create backups, you should also add the IP address of the KVM Compute resource added in step 1 in the 'Backup IP address' field of each of your compute resources.

Further steps:

1. Format and mount the local storage:

- SSH to the backup server
- Format the storage with your preferred filesystem type, e.g.:

```
bash#> mkfs.ext4 /dev/sda
```

- Create a directory for backups if it does not exist

```
bash#> mkdir /backupstorage
```

- Mount the storage to /onapp/backups:

```
bash#> mount /dev/sda /backupstorage
```

- Create a directory for storing templates:

```
bash#> mkdir /backupstorage/templates
```

- Create a directory for storing backups:

```
bash#> mkdir /backupstorage/backups
```

- Create symbolic links in /onapp:

```
bash#> ln -s /backupstorage/backups /onapp/backups
bash#> unlink /onapp/templates
bash#> ln -s /backupstorage/templates /onapp/templates
bash#> ln -s /onapp/tools/recovery /data
```

- Add the following to custom config file:

```
mkdir /backupstorage
mount /dev/sda /backupstorage
ln -s /backupstorage/backups /onapp/backups
unlink /onapp/templates
ln -s /backupstorage/templates /onapp/templates
ln -s /onapp/tools/recovery /data
```

- Another example of the custom config files:

- If you have a separate partition for backups and templates (/dev/sda1 and /dev/sda2):

```
mkdir /backupstorage
mkdir /backupstorage/templates
mkdir /backupstorage/backups
mount /dev/sda1 /backupstorage/backups
mount /dev/sda2 /backupstorage/templates
ln -s /backupstorage/backups /onapp/backups
unlink /onapp/templates
ln -s /backupstorage/templates /onapp/
templates
```

- If you current array is detected as /dev/sda1 and currently everything is located in /onapp within templates and backup directories within:

```
mkdir /backupstorage
mount /dev/sda1 /backupstorage/
ln -s /backupstorage/backups /onapp/backups

unlink /onapp/templates
ln -s /backupstorage/templates /onapp/
templates
```

2. Update the database so that the location of the templates is known:

- Find the database password:

```
cat /onapp/interface/config/database.yml |grep password
```

- Open the onapp database in MySQL:

```
bash#> mysql -p bash#> use onapp;
```

- Find the ID of the backup server:

```
bash#> select * from backup_servers;
```

- For all of the templates, set the required backup_server_id:

```
bash#> update templates set
backup_server_id='[your_id]';
```

3. To download the base templates during the installation to your Control Panel, download and run the following script:

```
bash#> wget http://downloads.repo.onapp.com/get_template.sh  
bash#> /bin/sh get_template.sh
```

9. Configure Cloud

After you set up your infrastructure, the final step is to configure your cloud on your Control Panel user interface. This section provides information on how to configure a basic cloud. After you complete the following steps, you can create virtual servers in your cloud.



9.1. Configure Control Panel Settings

After you installed OnApp, you can apply the necessary Control Panel configurations. For more information, see the [Configuration Settings](#) section.



9.2. Configure Compute Resources

To deploy virtual servers, you need to create the [compute zone](#) and [compute resources](#). Next, you [add compute resources to compute zones](#) as follows:

- You can add Xen compute resources to compute zones of the Virtual and Baremetal types
- You can add KVM compute resources to compute zones of the Virtual and Smart types
- You can add vCloud Director compute resources to compute zones of the VPC type
- You can add vCenter compute resources to compute zones of the Virtual type

If you want to use Integrated Storage, select a compute zone as a storage API endpoint and make sure to enable Integrated Storage in the **Settings > Configuration** to group compute resource drives together into a virtual data store.



9.3. Configure Data Stores

To provide your virtual servers with storage space, you need to create [data store zones](#) and [data stores](#). You can [create LVM](#) and [Integrated Storage](#) data stores. You should also [attach data stores to data store zone](#). If you deploy traditional storage, you need to configure data stores on your compute resource. The commands below use `/dev/sda5` as an example. You can find a *volume group identifier* we're using in the second command on the **Data Stores** page in the Control Panel. Follow these steps for each local storage block on a compute resource:

```
bash#> pvcreate --metadatasize=50M /dev/sda5
bash#> vgcreate onapp-ar0akk2wyer3tf /dev/sda5
```



9.4. Configure Networks

To provide IP addresses to your future virtual servers, you need to perform the necessary network configurations. Create the [network zone](#) and [network](#). When adding the network, select the network zone you created. The network will be automatically attached to the network zone.



9.5. Configure Backup Servers

If you plan to use backup servers to store templates, ISO images, or backups, you need to add the [backup server](#) and [backup server zone](#) to your cloud. Next you add the [backup server to backup server zone](#).



9.6. Configure Relations Between Entities

After you added all the necessary resources to your cloud, you need to associate them with the compute resource you created. For this, assign the [data store](#) and [network](#) to the compute resource or compute zone. You also need to assign a backup server to [compute resources](#) or [compute zones](#). All data stores, networks, and backup servers you assign to a compute zone or resource should be from zones with the same type as the compute zone or resource. Zones can have one of the following types: *Virtual*, *Baremetal*, *Smart*, or *VPC*.

- To *Virtual* compute zones and resources, you can assign networks from the Virtual network zones and LVM, Integrated Storage, VMware and SolidFire data stores from the Virtual data store zones.
- To *Baremetal* compute zones and resources, you can assign networks from the Baremetal network zones.

- To *Smart* compute zones and resources, you can assign networks from the Smart network zones and LVM, Integrated Storage, and SolidFire data stores from the Smart data store zones.
- To *VPC* compute zones and resources, you can assign external, org, and vApp networks from the VPC network zones, data stores, and data store zones. All vCloud Director related resources have the *VPC* type.



9.7. Configure Templates

To build Linux virtual servers, you need to download templates using the UI downloader. For this, you [install templates](#) and [template store](#). You should also add the installed templates to the Template Store.



9.8. Configure ISOs

To be able to build and boot virtual servers from ISO images, additional steps are required. For more information, refer to the [Additional Considerations for ISOs](#) section.



9.9. Configure OVAs

To be able to build and boot virtual servers from OVA images, additional steps are required. For more information, refer to the [Additional Considerations for OVAs](#) section.



9.10. Configure Buckets

To control user access to resources that are available in your cloud, create and configure [Buckets](#). Buckets enable you to set up resources allocation and pricing separately. If you want to configure resources to which a user has access, you can edit Access Control. If you want to set up pricing for resources, edit Rate Card where the prices and the number of free resources are set.

10. Configure vCloud Director Integration

You can integrate OnApp with VMware solutions such as vCloud Director and vCenter. For detailed instructions, see [OnApp and vCloud Director Configuration Guide](#) and [vCenter Implementation Guide](#).

11.Support

11.1.24/7 Support

OnApp provides 24/7 support that includes remote installation, free upgrades, and 24/7 global support by telephone and email. If you need our help, you can reach to us as follows:

- [Submit a request by email](#)
- **Call (+1) 888 876 8666**

11.2.What Does OnApp Support in my Cloud?

OnApp provides support for everything directly related to our core products such as [OnApp Cloud](#), [OnApp CDN](#), [OnApp for VMware](#), and [other products](#). We take responsibility for the software, bug fixes, patches, and general maintenance of our products.

Although there are some areas that are not covered under the standard OnApp support, we attempt to offer suggestions or put you in touch with our Professional Services team. Typically the following areas are not covered by the standard OnApp support:

- Switch, router, and firewall configuration.
- SAN configuration and optimization.
- Attaching, removing, and resizing LUNs.
- The compute resource and Control Panel server hardware support.
- Operating System installation and support.
- Maintenance of your passwords or [Whitelists](#).
- Configuration and troubleshooting inside virtual servers.
- VMware vSphere installation and configuration.
- Known bugs and limitations within virtualization platforms.
- Third party integrations.
- Alpha and Beta releases.
- Coding for [Recipes](#).
- Coding for custom configs on CloudBoot compute resources.

11.3.Professional Services

Get in touch with our [Professional Services](#) to get expert help from launch to production and beyond.

11.4.OnApp Community

Visit [OnApp Community](#) that is a public forum where you can share your feedback and product ideas. Only OnApp customers can give suggestions but anyone can explore them.

11.5.Knowledge Base

Visit [Knowledge Base](#) where you can find a lot of how-to articles to resolve questions that you may face while working with OnApp.

11.6.Documentation

You are now located in one of the OnApp documentation spaces that is [Installation Guide](#). Go to [Documentation Home](#) to browse other available spaces.