

OnApp Cloud 5.8 Installation Guide

Table of Contents

1	Preparation	4
2	Installation	5
3	Post Install Configuration	6
4	OnApp Installation Walk-Through	7
4.11.	Prepare Servers Configuration	7
4.22.	Install Control Panel Server	7
4.33.	Install Compute Resources	7
4.44.	Install Data Stores	7
4.55.	Install Backup Server	7
4.66.	Configure Cloud	7
5	Install Control Panel Server	9
6	Install Compute Resources	17
6.1	Install CloudBoot Compute Resources	17
6.1.1	Configure CloudBoot Settings in BIOS	19
6.1.2	Configure InfiniBand	20
6.2	Install Static Compute Resources	21
7	Install Data Stores	26
7.1	Install LVM Data Store	26
7.2	Install Integrated Storage Data Store	28
7.3	Install SolidFire Data Store	29
8	Install Backup Server	31
8.1	Install Static Backup Server	31
8.2	Install CloudBoot Backup Server	34
9	Enable Recovery Mode for Baremetal Servers	38
10	Configure Cloud	39
10.11.	Configure Control Panel Settings	39
10.22.	Configure Compute Resources	39
10.33.	Configure Data Stores	39
10.44.	Configure Networks	40
10.55.	Configure Backup Servers	40
10.66.	Configure Relations Between Entities	40
10.77.	Configure Templates	40
10.88.	Configure ISOs	40
10.99.	Configure OVAs	41
10.1010.	Configure Buckets	41

OnApp version 5.8 is an edge release which is not designed to be installed on production environments.

This document describes how to install the 5.8 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](#)

Please do not change the default language settings during the installation process (en_US.UTF-8)!

See also:

[Get Started for Clouds with High Availability](#)

[Technical Details](#)

[OnApp Upgrade Guide](#)

[Get Started Guide](#)

4 OnApp Installation Walk-Through

Generally, the OnApp installation includes the following steps:

The workflow below describes a standard OnApp Installation. You can additionally configure [High Availability](#) for your cloud.

4.1 1. Prepare Servers Configuration

Before installing OnApp, it is required to make sure your network, storage, and servers configuration meets the requirements.

See also:

[Preparation Guide](#)

[OnApp Upgrade Guide](#)

[Get Started Guide](#)

4.2 2. Install Control Panel Server

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

4.3 3. Install Compute Resources

Compute resources provide system resources such as CPU, memory, and network, and control secure virtualization. After the Control Panel server installation, proceed to the compute resource installation. Depending on the desired cloud configuration, you can opt for Static compute resources or CloudBoot servers.

4.4 4. Install Data Stores

Make sure to install the appropriate 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 an OnApp Integrated storage.

4.5 5. Install Backup Server

Backup servers are servers responsible for storing backups and templates of virtual servers running in the cloud, in order to prevent data loss in the event of failure. You can install static or cloudboot backup server.

4.6 6. Configure Cloud

After you have set up the servers, log in to OnApp CP and configure the relations between the entities.

5 Install Control Panel Server

OnApp version 5.8 is an edge release which is not designed to be installed on production environment.

This section is the part of the OnApp installation procedure.

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

- Review the [Preparation Guide](#) to ensure that you have a suitable environment before starting.
- Use corresponding option of the Control Panel installer in case MySQL is already installed.
- Installer output is redirected to `./onapp-cp-install.log`
- All installer critical errors are in `/var/log/messages`
- If you consider deploying High Availability Clusters, refer to [Installation Guide for High Availability](#)
- If you're replacing an existing Control Panel with a new install, refer to [Control Panel Migration](#)
- If you plan to deploy [Accelerator](#), refer to the [RabbitMQ Configuration for Accelerator](#) document.
- If you need to install other components (OnApp Database Server, RabbitMQ Server, Redis Server), refer to their respective installation instructions. Also you can [migrate existing OnApp database from MySQL to MariaDB](#), [Percona](#)

See also:

[Technical Details](#)

[Preparation Guide](#)

[Get Started for Clouds with High Availability](#)

To install Control Panel server, perform the following procedure:

1. Update your server:

```
bash# yum update
```

2. Download OnApp YUM repository file:

```
# rpm -Uvh http://rpm.repo.onapp.com/repo/onapp-repo-5.8.noarch.rpm
```

3. Install OnApp Control Panel installer package:

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

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

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

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
```

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.

5. Run the Control Panel installer:

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

The full list of Control Panel installer options:

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] [-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_SOCKET]	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
REDIS_*	Options are useful if Redis Server is already installed and configured.
--redis-host=REDIS_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[=REDIS_BIND]	The IP address for Redis Server to serve connections (to listen). The option isn't mandatory.
--redis-port=REDIS_PORT	Redis Server listen port. Defaults are: 0 - if local server 6379 - if remote server
--redis-passwd[=REDIS_PASSWD]	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[=REDIS_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.
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_PASSWD	CP administrator password

-F ADMIN_FIRSTNAME	CP administrator first name
-L ADMIN_LASTNAME	CP administrator last name
-E ADMIN_EMAIL	CP administrator e-mail
--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_PASSWD	RabbitMQ/vCloud user password. The random password is generated if isn't specified.
--vcdvhost VCD_VHOST	RabbitMQ/vCloud vhost. Default value is '/'
RBT_*	Options are used to configure RabbitMQ manager account. If local RabbitMQ server.
--rbtlogin RBT_LOGIN	RabbitMQ manager login. The default value is 'rbtmgr'.
--rbtpasswd RBT_PASSWD	RabbitMQ manager password. The random password is generated if isn't specified.
--ha-install	Proceed with Control Panel and High Availability components installation.
--rake RAKE_TASKS	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_TRAP_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.
--noservices	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.
--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: <i>rpms</i> - for 'system packages' upgrade; <i>mysql</i> - for MySQL database upgrade and configuring; <i>redis</i> - for Redis Server upgrade and configuring; <i>rabbitmq</i> - for RabbitMQ Server upgrade and configuring; <i>monit</i> - 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

6. Install CloudBoot dependencies:

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

DISTRO - CentOS6 or 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 packages. Below you can find an example:

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

Any of the `ramdisk` packages will install the following packages automatically: `onapp-store-install onapp-ramdisk-tools`

7. Install OnApp license to activate the Control Panel. Enter a valid license key via the Web UI (you'll be prompted to do so). Your default OnApp login is **admin/changeme**. The password can be changed via the Control Panel's **Users** menu in the Control Panel.

Once you have entered a license it can take up to 15 minutes to activate.

8. Restart the OnApp service:


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

9. 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.

10. 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
```

Note that this script is applicable only if you are configuring the first compute resources in the cloud. If you already have some existing compute resources and are installing new ones, this script will set new SSH keys to all the compute resources, both new and existing.

11. 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
```

12. 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 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)
```

13. After you have installed the Control Panel server, configure your Cloud Settings. See [Configure Cloud](#) for details.

This section is the part of the OnApp installation procedure.

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

6 Install Compute Resources

OnApp version 5.8 is an edge release which is not designed to be installed on production environments.

This section is the part of the OnApp installation procedure.

[Install Control Panel Server](#) > **Install Compute Resources** > [Install Data Stores](#) > [Install Backup Server](#) > [Configure Compute Resources](#)

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.

On this page: [Install CloudBoot Compute Resources](#)
[Install Static Compute Resources](#)

See also:

[Install Control Panel Server](#)
[Install Data Stores](#)
[Install Backup Server](#)
[Technical Details](#)
[Preparation Guide](#)
[Get Started for Clouds with High Availability](#)

6.1 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:
 - Go to **Settings > Configuration > System > CloudBoot**
 - Scroll down to the CloudBoot section and check the **Enable** box.
2. Enable Storage in 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** 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 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

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/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 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

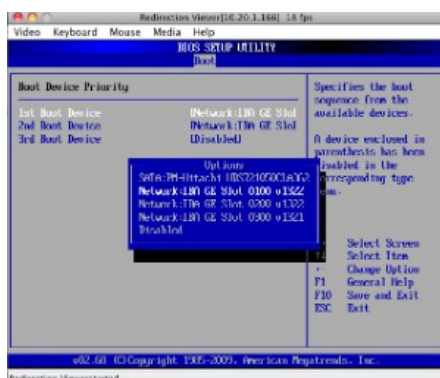
mkdir -p /onapp/templates && mount -t nfs
192.168.10.101:/onapp/templates /onapp/templates
```

6.1.1 Configure CloudBoot Settings in BIOS

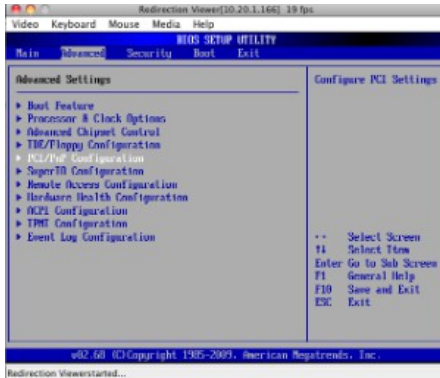
Your BIOS settings may vary from the example provided in this section.

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



- 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.

- Set local disk as a second boot device.

6.1.2 Configure InfiniBand

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

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

- Be aware that InfiniBand is not supported for CloudBoot CentOS7 image due to lack of integrated OFED stack.

To enable the InfiniBand mode:

- Go to **Control Panel > Settings > Configuration** in the OnApp CP
- Move the *Enable InfiniBand boot* slider in the CloudBoot section
- Click the **Save Configuration** button.

Current limitations:

- IB in Ethernet mode is only supported for Centos6/KVM nodes. It will not work with CentOS5 / Xen.

- InfiniBand is only supported for the SAN network, not PXE boot.

6.2 Install Static Compute Resources

Before you proceed

1. Install base CentOS packages on the local drive before compute resource installation, depending which virtualization method you choose:
 - Xen 4 compute resources: CentOS 6.x x64
 - KVM compute resources: CentOS 6.x x64 or CentOS 7.x x86/64
2. We recommend installing CentOS from the minimal CentOS ISO for static compute resources.
3. Disable CPU power-saving features in BIOS before you proceed to the compute resource installation.
4. 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_squas
5. Pay attention that smart and baremetal servers cannot be installed using the static compute resource installation method.

To install a compute resource:

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-5.8.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
```

The full list of OnApp compute resource custom values:

- o OnApp HV tools [custom version](#)

```
HV_VERSION=""
```

- o OnApp StorageAPI [custom version](#)

```
API_VERSION=""
```

- o Default server to sync time on the compute resource

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

- o 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
```

- o The number of loopback devices created

```
LOOPBACKS=128
```

- o 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=""
```

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

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

- o 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:

The full list of installer options:

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

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

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

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

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

8. 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>
```

The full list of 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] [-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 black; background-color: #ffffcc; padding: 5px; width: fit-content; margin: 5px auto;"> <p>If unsure, set the Control Panel server's management IP as CP_HOST_IP and FILE_TRANSFER_SERVER_IP.</p> </div>
-a	Install AoE
-s	Install sshfs
-?	Print this help info

9. Run the following commands:

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

10. Reboot the compute resource to complete the installation:

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

11. 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.

12. 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
```

Note that this script is applicable only if you are configuring the first compute resources in the cloud. If you already have some existing compute resources and are installing new ones, this script will set new SSH keys to all the compute resources, both new and existing.

13. 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:

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 or backup server 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:

- o `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.
- o `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.

This section is the part of the OnApp installation procedure.

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

7 Install Data Stores

OnApp version 5.8 is an edge release which is not designed to be installed on production environments.

This section is the part of the OnApp installation procedure.

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

Once the Control Panel server has been installed successfully, you can use one of the following processes to set up data stores:

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

7.1 Install LVM Data Store

PLEASE NOTE:

- To configure an [Integrated Storage data store](#), please consult the Admin guide.
- This process assumes you have already configured a compute resource to see the ISCSI/A will be shown when running a fdisk -l.
- All compute resources need access to the same data store. Ensure that you have the block
- Perform this procedure only once per data store!
- Take care when choosing the disk/partition you want to use for storing VM data!

On this page:

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

See also:

[Install Control Panel Server](#)
[Install Backup Server](#)
[Technical Details](#)
[Preparation Guide](#)
[Get Started for Clouds with High Availability](#)

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** link at the bottom of the screen.
 - Follow the steps in the creation wizard:

Step 1 of 2

- 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 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)

- a. When you've finished configuring the 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's 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. SSH into 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 datastore page in the UI.

5. Test compute resource/volume group visibility:

Now 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 resources 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 that, you can proceed to the integrated storage data store creation.

To create a new integrated storage data store:

1. Go to your Control Panel's **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** link 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 username for cluster authorization

Password - specify 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.

This section is the part of the OnApp installation procedure.

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

8 Install Backup Server

OnApp version 5.8 is an edge release which is not designed to be installed on production environments.

This section is the part of the OnApp installation procedure.

[Install Control Panel Server](#) > [Install Compute Resources](#) > [Install Data Stores](#) > **Install Backup Server** > [Configure Backup Server](#)

Follow one of two processes to set up a backup server in your cloud:

- Install Static Backup Server
- Install CloudBoot Backup Server

8.1 Install Static Backup Server

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

On this page:

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

See also:

[Install Control Panel Server](#)
[Install Data Stores](#)
[Technical Details](#)
[Preparation Guide](#)
[Get Started for Clouds with High Availability](#)

1. Add a backup server via the Control Panel user interface:
 - Go to your Control Panel's **Settings** menu, then press **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-5.8.noarch.rpm
```

- Install the OnApp Backup Server installer package:

```
bash# yum install onapp-bk-install
```

- Check and set Backup Server default settings:
Edit Backup Server default settings by editing 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'
```

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

- 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] [-h]
```

Where:

-c CONFIG_FILE	Custom installer configuration file. Otherwise, preinstalled one is used.
-a	Do NOT be interactive. Processes with automatic installation.
-y	Update OS packages (except of OnApp provided) on the box with 'yum update'.
-t	Initiate Base, Load Balancer and CDN templates download. The download is initiated if '-a' option is used.
-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 CloudBooted KVM Compute resources that can be used as backup servers. Follow the step-by-step instructions provided in this chapter to configure CloudBoot backup servers in your cloud.

- 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 VS 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 the 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's **Settings** menu, then press **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 new CloudBoot compute resource: at the first step of the creation wizard choose 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's **Settings** menu, then press **Compute Resources** icon.
- Click the **Add New CloudBoot Compute Resource** button.
- In 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 intend 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
```

- Make folder for backups if it does not exist

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

- Mount the storage to /onapp/backups:

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

- Make folder for storing templates:

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

- Make folder 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
```

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

a. Find the database password:

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

b. Open the onapp database in MySQL:

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

c. Find the ID of the backup server:

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

d. 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
```

If your backups disappear after rebooting the CloudBoot backup server with LVM storage, a release.

To fix your custom config settings, use one of the following options provided in the example

1. 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
```

2. If you current array is detected as /dev/sda1 and currently everything is located in /onapp

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

This section is the part of the OnApp installation procedure.

[Install Control Panel Server](#) > [Install Compute Resources](#) > [Install Data Stores](#) > [Install Backup Server](#) > [Conf](#)

9 Enable Recovery Mode for Baremetal Servers

To enable recovery mode for baremetal servers, perform the following steps:

1. Download the following files:
<http://templates.repo.onapp.com/Linux/recovery-baremetal.kernel>
<http://templates.repo.onapp.com/Linux/recovery-baremetal.initrd>
2. Place the files into the `/tftpboot/images/ramdisk-recovery/` directory.
3. Create template file `/tftpboot/pxelinux.cfg/template-baremetal-recovery` with following contents:

```
default baremetal-recovery
label baremetal-recovery
kernel images/ramdisk-recovery/recovery-baremetal.kernel
append initrd=images/ramdisk-recovery/recovery-baremetal.initrd
root=live:/recovery-centos-3.2.iso rootfstype=auto ro liveimg
rd.luks=0 rd.md=0 rd.dm=0
```

4. Restart the OnApp services:

```
service onapp restart
service httpd restart
```

After that, recovery mode option will appear in the baremetal server's **Tools** menu:

The screenshot displays the zazaBM web interface for a Baremetal Server. At the top, there is a yellow notification box titled "Baremetal Server details" with an information icon and a close button. Below this, the "zazaBM" logo is visible on the left, and a green "Tools" button with a gear icon is on the right. The main content area is titled "Properties" and features a server icon labeled "CENTOS" and "CentOS 6.1 LAMP x64". To the right of the icon is an "Information" section with the following details:

- Hostname: zaza
- Hypervisor: Baremetal HV
- Login: root / password
- Owner: John Smith
- Price per Hour: \$0.00 (two entries)

On the right side of the "Information" section, there is a dark grey box showing "IP AD 109". A dropdown menu is open from the "Tools" button, listing the following options: "Delete Baremetal Server", "Recipes", "Custom Recipe Variables", and "Enable Recovery Mode". The "Enable Recovery Mode" option is highlighted with a blue background.

10 Configure Cloud

This section is the part of the OnApp installation procedure.

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

Once you've set up your hardware, the final step is to configure your cloud in your Control Panel. This section explains how to configure a basic cloud. If you complete these steps you should be in a position to create VSs.

10.11. Configure Control Panel Settings

Once you have installed OnApp, you need to make the necessary Control Panel configurations. Set the [Edit System Configuration](#), [Edit Backups/Templates Configuration](#), [Edit Interface Configuration](#) and [Edit Defaults Configuration](#) CP options.

10.22. Configure Compute Resources

To deploy virtual servers, you need to add [Create Compute Zone](#) and [Create Compute Resource](#) to your cloud. After that, [Add Compute Resource to Compute Zone](#) you've added:

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

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

10.33. Configure Data Stores

To provide your virtual servers with storage space, you need to configure [Create Data Store Zone](#) and data stores. Data stores can be [Create LVM Data Store](#) and [Create Integrated Storage Data Store](#). You should also [Attach Data Stores to Data Store Zone](#) you've added.

In case of Traditional storage you need to configure data store(s) on your compute resource. The commands below use `/dev/sda5` as an example. You can find the volume group identifier we're using in the second command, from the **Data Stores** screen in the Control Panel. Follow these steps for each local storage block on the compute resource:

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


10.44. Configure Networks

To provide IP address(es) to your future virtual servers, you need to perform the necessary network configurations. To do this, create [Create Network Zone](#) and [Create Network](#). When adding the network, select the network zone you've created. The network will be automatically attached to the network zone you chose during creation. You should also [Add IP Address to Network](#).

10.55. Configure Backup Servers

If you plan to use backup servers to store such items as, for example, templates, ISOs or backups, you need to add [Create Backup Server](#) and [Create Backup Server Zone](#) to your cloud. After that, [Add Backup Server to Backup Server Zone](#) you've added.

10.66. Configure Relations Between Entities

Once you've added all the necessary resources to your cloud, you need to associate them with the compute resource you've created in Step 2. For this, [assign data stores](#) and [networks](#) to the compute resource or compute zone (Step 2). You also need to assign backup server(s) (Step 5) to the [compute zone](#). 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:

- *VPC type*: to the compute zones and resources of this type you can assign external, org and vApp networks from the VPC type network zones and data stores from the VPC type data store zones. All vCloud Director related resources have the VPC type.
- *Smart type*: to the compute zones and resources of this type you can assign networks from the Smart type network zones and LVM, Integrated storage and SolidFire data stores from the Smart type data store zones.
- *Virtual type*: to the compute zones and resources of this type you can assign networks from the Virtual type network zones and LVM, Integrated storage, VMware and SolidFire data stores from the Virtual type data store zones.
- *Baremetal type*: to the compute zones and resources of this type you can assign networks from the Baremetal type network zones.

10.77. Configure Templates

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

10.88. Configure ISOs

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

10.99. Configure OVAs

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

10.10 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 do that using the Access Control. If you also want to set up pricing for resources, use the Rate Card where the prices and the number of free resources are set.