onapp

OnApp

Cloud v.3.0 BETA

Installation Guide

Document version	1.0
Document release date	25th September 2012
	document revisions

ⓘ You must have a valid beta license before beginning your install!

ⓘ The installation processes described in this document require servers to have a base installation of CentOS with the standard repositories enabled (the versions of CentOS required are explained in the server config reminder chapter).

If you have additional repositories enabled, please disable them before continuing.

Contents

1. Server config reminder

OnApp Cloud runs on CentOS or (for the OnApp Controller Server) Red Hat Enterprise Linux Server. Please note that the RHEL/CentOS versions required can vary depending which virtualization method you choose, Xen or KVM:

- **XEN Hypervisors**
  RHEL/CentOS 5.x x64

- **KVM Hypervisors**
  RHEL/CentOS 5.x x64 or RHEL/CentOS 6.x x64

- **OnApp Controller Server**
  RHEL/CentOS 5.x x86/X64 or RHEL/CentOS 6.x x86/64

- **OnApp Backup Server**
  RHEL/CentOS 5.x x64

2. Control Panel installation

ⓘ *If mysql server is already installed, it must not have a password configured: this will be configured by our installer. Any password that is already configured will cause an installer error.*

ⓘ *Installer output is redirected to* ./onapp-cp-install.log

ⓘ *All installer critical errors are in* /var/log/messages

ⓘ *Once the installation of the control panel is complete, your default OnApp login will be* **admin / changeme.** *The password can be changed via the Control Panel's Users and Groups menu.*

ⓘ *If you're replacing an existing Control Panel with a new install, please dump your current mysql database. Once you've installed your new control panel, overwrite its database with the previous one. You can find details about the database by running* cat /onapp/interface/config/database.yml *and looking at the connection details located under 'production'.*

1. **Update your server using YUM:**

```
bash#> yum -y update
```

ⓘ *If anything was updated, reboot the server.*

2. **Download OnApp YUM repository file:**

```
bash#> rpm -Uvh http://rpm.repo.onapp.com/repo/centos/5/onapp-
repo.noarch.rpm
OnApp-CP#> rpm -ivh onapp-repo.noarch.rpm
```

Download the OnApp 3.0.0 repository config file:

```
bash# wget -N http://rpm.repo.onapp.com/repo/centos/5/OnApp-3.0.0.repo
-P /etc/yum.repos.d/

bash#> yum clean all
```

3. **Install OnApp Control Panel installer package:**

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

4. **Custom Control Panel configuration**

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

- OnApp to MySQL database connection data: connection timeout, pool, encoding, unix socket
- MySQL server configuration data (if MySQL is running on the same server as the CP): wait timeout, maximum number of connections
- The maximum number of requests queued to a listen socket (net.core.somaxconn value for sysctl.conf)
- The root of OnApp database backups directory (temporary directory on the CP box where MySQL backups are placed)

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

ⓘ Custom values must be set before the installer script runs.

5. **Run Control Panel installer:**

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

6. **Install CloudBoot Dependencies**

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

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).

ⓘ PLEASE NOTE: once you have entered a license it can take up to 15 minutes to activate.

3. VMware installation

To install and configure your cloud for VMware:

1. Install the VMWare vCenter server by following VMware documentation instructions.

    a. Create an administrator account on the vCenter server or use the default "administrator" account and specify login credentials.

    b. Create a vCenter virtual datacenter.

    c. On the datacenter, create a new cluster, turn on DRS and note the cluster name. Later, the cluster name will be used when you configure it as a hypervisor in the OnApp control panel.

    d. Open the following ports on the vCenter server:

    - TCP/UDP 902
    - TCP 443
    - TCP 80
    - TCP/UDP 598

    e. Install VMWare ESXi servers by following the VMware documentation instructions.

    f. Add all ESXi servers to the cluster.

    g. Attach all ESXi servers to the shared SAN storage. Remember the data store label. Later you'll use this data store name when configuring a data store in the OnApp control panel.

    h. Create a Distributed Switch.

    i. Copy Firewall XML into each ESXi server to open the VNC ports. To do so, create a file called vnc.xml on the local ESXi host or on the central CP server host, and write the following into the file:

```
<ConfigRoot>
<service id='0000'>
<id>VNC</id>
<rule id = '5900'>
<direction>inbound</direction>
<protocol>tcp</protocol>
<porttype>dst</porttype>
<port>
<begin>5900</begin>
<end>5999</end>
</port>
```

```
</rule>
<rule id='0001'>
<direction>outbound</direction>
<protocol>tcp</protocol>
<porttype>dst</porttype>
<port>
<begin>0</begin>
<end>65535</end>
</port>
</rule>
<enabled>true</enabled>
<required>false</required>
</service>
</ConfigRoot>
```

Run the following commands:

```
echo "cp <LOCATION OF vnc.xml> /etc/vmware/firewall/vnc.xml" >> /etc/rc.local
echo "localcli network firewall refresh" >> /etc/rc.local
echo "esxcli network firewall refresh" >> /etc/rc.local
```

Restart the ESXi server for changes to come into effect.

j. Enable NTP on all ESXi server. For correct time synchronization, use the same NTP server for vCenter and CP.

k. Allow virtual machines to start and stop automatically with the system.

2. Configure firewall.

a. If running a virtual Vyatta on the vCenter you need to do some additional hypervisor configuration. On each hypervisor, create three virtual machine port groups:

- eth0 - for public access
- eth1 - for Vyatta communication with CP
- eth2 - fne for VLAN communication with all virtual machines. This port group should have VLAN 4095 on for trunking all VLAN traffic.

b. Create a new Vyatta instances on the vCenter with three network interfaces and attach on to each of the created port groups.

c. Install Vyatta v.6.4 or later from http://vyatta.org/ and install it by creating and booting the Vyatta LiveCD.

d. Login as user vyatta with password vyatta and run the "install image" command.

e. Remove the LiveCD.

f. Reboot system.

g. Log in using the vyatta user credentials.

h. Run the following commands:
(eth0 = public interface; eth1 = control panel communication interface; eth2 = VLAN communication interface. Adjust the commands below if your config is different)

- configure
- set service ssh
- set service ssh allow-root
- set interface ethernet eth0 address <OUTSIDEIPADDRESS/CIDR>
- set system gateway-address <OUTSIDEGATEWAYADDRESS>
- set interface ethernet eth1 address <COMMUNICATIONIPADDRESS/CIDR>
- set firewall state-policy established action accept
- set firewall state-policy related action accept
- set firewall state-policy invalid action drop
- set firewall state-policy invalid log enable
- set firewall name INSIDE_OUT
- set firewall name INSIDE_OUT default accept
- set vpn ipsec ipsec-interfaces interface eth0
- set system login user vyatta authentication plaintext-password <NEWPASSWORD>
- commit
- save

i. Configure the firewalls in OnApp/settings/firewalls.

3. Create a new Hypervisor in OnApp/settings/hypervisors/new.

   a. Select VMware hypervisor type.
   b. Specify vCenter login and password.

4. Create a new IP address pool in OnApp/settings/ip_address_pools/new

5. Create a new user. In OnApp/users/new

6. Create a new customer network in OnApp/customer_networks/new

7. Create a new data store zone in OnApp/data_store_zones/new

8. Create a new datastore in OnApp/settings/data_stores/new

9. Create a range of VLANs you want to use in OnApp/settings/customer_vlans/new

10. Assign customer network, network and datastore to the VMware hypervisor.

11. Specify vCenter cluster name in the System Configuration settings.

12. In the Default Settings configuration (OnApp/settings/edit#defaults) define the service account name that will be automatically created on all virtual machines to be able to communicate with them.

13. Log in as the user you have created in step 5 and create a new VM in VMware.

4. Xen/KVM hypervisor installation

Once the control panel server has been installed successfully, you can follow one of 2 processes in order to set up Xen or KVM hypervisors: the Cloud Boot method, where hypervisors are installed over your network, or the standard, static install process to each hypervisor's local disk.

4.1 Cloud Boot hypervisor bootstrap method

Follow this method to enable cloudboot for your Hypervisors. This is a new feature that allows dynamic boot of Hypervisor servers without any persistent installation requirements.

ⓘ Servers must support and have PXE boot enabled on the Network Interface Card (setup in the BIOS if not already enabled by default).

1. **Enable Cloud Boot in the control panel**

   Settings -->Configuration -->CloudBoot

   Scroll down to the CloudBoot section and check the "enable" box.

2. **Enter IP addresses for static content target and control panel server cloud boot interface**

   Static content such as cloudboot images, kernels, VM templates can be hosted on a standalone NFS server if you wish. The default setting is to install everything on the control panel server.

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

3. **Add IP address range for Hypervisors**

   Settings -->Hypervisors -->CloudBootIPs -->New IP address

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

5. **Add servers to the control panel by selecting MAC addresses and assigning IP address**

   Settings -->Hypervisors -->Add a new CloudBoot Hypervisor

   ⓘ Note that if you want to expose drives in hypervisors to OnApp Storage, our integrated storage platform, then you must select them at this point.

   For more information on setting up and configuring Cloud Boot, see the [v3.0 beta Admin Guide].

6. **Generate SSH keys:**

ⓘ Iif you are going to use Cloud Boot with backup server configuration without integrated storage, set up SSH keys from the CP server to the backup server.

ⓘ 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.

- SSH into your Control Panel server.

- Download and run the script:

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

  ⓘ Do not specify passphrases - just leave them blank!

7. <u>Mount an NFS target:</u>

As you are using CloudBoot it is necessary to use the basic backup scheme rather then adding an advanced standalone backup server as covered in Section 6. A NFS target must be installed on each hypervisor under the /onapp/backup, /onapp/templates and /onapp/backups directory.

1. Setup an NFS target on a standalone CentOS installation, or create an NFS mount point on a NAS box.

2. Insert the following script into each hypervisor's execution script window in the control panel user interface, to allow it to mount the remote NFS target for backup:

   ```
   bash# mkdir -p /onapp/backup && mount -t nfs <IP ADDRESS>:/path/to/backup
   /onapp/backup
   ```

Do not configure sshfs for the cloudboot backups as the HVs do not have backup server keys pre-installed.

ⓘ Note that the control panel server can be used as a backup NFS target, and you can create a backup subdirectory under /onapp/templates on that server to avoid creation of new NFS target export entries.

In this case, run the following script to mount rules for HV:

```
bash# mkdir -p /onapp/backup && mount -t nfs <IP
ADDRESS>:/onapp/templates/backup/onapp/backup
```

8. <u>Download Templates:</u>

Log into the server providing the NFS mounts and run the following commands:

```
bash# cd /onapp/templates
bash# wget http://downloads.repo.onapp.com/install/bkinstall.sh
bash# sh ./bkinstall.sh
```

4.2 Static hypervisor installation method

ⓘ Base Centos 5.x (Xen) or Centos6 (KVM) templates must be installed on the local drive before hypervisor installation.

1. **<u>Add the hypervisor to your cloud using the OnApp Control Panel:</u>**

   Settings --> Hypervisors --> Add New Hypervisor

   Make sure the hypervisor is visible in the Control Panel, and at this point showing as inactive.

2. **<u>Make sure your OS is up to date:</u>**

```
bash#> yum -y update
```

3. **<u>Enable IPv6:</u>**

   ⓘ This step is required regardless of whether you'll be using IPv6 or not.

   Edit /etc/modprobe.conf and comment out the following strings:

```
alias ipv6 off
options ipv6 disable=1
```

   Next, edit /etc/sysconfig/network and replace

```
NETWORKING_IPV6=no
```

   with

```
NETWORKING_IPV6=yes
```

ⓘ These settings won't take effect until you reboot, but <u>do not</u> reboot now. We'll do that later.

4. **<u>Download the OnApp repository:</u>**

```
bash#> wget  http://rpm.repo.onapp.com/repo/centos/5/onapp-
repo.noarch.rpm
bash#> rpm  -ivh  onapp-repo.noarch.rpm
```

5. **<u>Download the OnApp 3.00 repository config file:</u>**

```
bash# wget -N http://rpm.repo.onapp.com/repo/centos/5/OnApp-3.0.0.repo
-P /etc/yum.repos.d/

bash#> yum clean all
```

6. Install the OnApp hypervisor installer package:

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

7. Edit custom hypervisor configuration

Edit the /onapp/onapp-hv.conf file to set hypervisor custom values, such as NTP time sync server, Xen Dom0 memory configuration data and number of loopback interfaces:

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

ⓘ Custom values must be set before the installer script runs.

8. Run the OnApp hypervisor installer script:

For Xen hypervisors:

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

For KVM hypervisors:

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

ⓘ Run the script with the -o option if you're installing hypervisors with Open vSwitch. The –o option is disabled by default. To get information about the installer and its properties, such as packages update, templates download and non-interactive mode, run the script with -h option.

```
bash# /onapp/onapp-hv-install/onapp-hv-xen-install.sh -h
Usage: /onapp/onapp-hv-install/onapp-hv-xen-install.sh [-c CONFIG_FILE]
[-a] [-y] [-o] [-t] [-h]
```

Options

`-c CONFIG_FILE`	Custom installer configuration file. Otherwise, the preinstalled one is used.
`-a`	Non-interactive mode. Automatic installation process.
`-y`	Update all packages on the box with 'yum update'. The update will be processed if the -a option is used.
`-o`	Xen + Open vSwitch installation
`-t`	Download recovery templates and ISO(s) used to provision FreeBSD guests.
`-h`	Print this info.

9. **Configure the hypervisor for your cloud:**

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

ⓘ Run the script with the -h option to configure FQDN or IP Address of the management server (CP box) which should receive all status.

ⓘ Run the script with the -p option to configure server (hypervisor) FQDN or IP Address which will serve all stats related and other requests send by the CP.

ⓘ FQDN or IP Address for Control Panel and Hypervisor servers are required for the new statistics receiver to work.

ⓘ Ignore any errors stating stats and vmon services aren't running. This is expected at this stage.

ⓘ <CP_HOST_IP> *is the IP addresses of the Control Panel server.*

ⓘ <HV_HOST_IP> *is the IP address of the Hypervisor.*

ⓘ <FILE_TRANSFER_SERVER_IP> *is the IP address of the server that will hold your backups and templates.*

10. **Update necessary sysctl variables**

Edit your /etc/sysctl.conf file. If the netfilter.ip_conntrack_max entry exists, update the value: if it doesn't exist, add it. You can increase the netfilter.ip_conntrack_max value if required.

```
bash#> net.ipv4.netfilter.ip_conntrack_max = 256000
```

11. **Reboot the hypervisor to complete the installation:**

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

12. **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.

- SSH into your Control Panel server.

- Download and run the script:

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

ⓘ Do not specify passphrases - just leave them blank!

5. Data store installation

ⓘ PLEASE NOTE:

- *To configure an integrated storage datastore, please consult the userguide*
- *This process assumes you have already configured a hypervisor to see the ISCSI/ATAoE block device it is connecting to, and that the SAN disk will be shown when running a fdisk -l.*
- *All hypervisors need access to the same datastore. Ensure that you have the block device visible on all hypervisors.*
- *VERY IMPORTANT: only perform this procedure once per data store!*
- *ALSO IMPORTANT: take care when choosing the disk/partition you wish to use for storing VM data!*

1. **Add the new data store to OnApp via the WebUI:**

To create a data store:

- Go to your Control Panel **Settings** menu.
- Click the **Data Stores** icon.
- Click the **Create Data Store** link at the bottom of the screen.
- On the screen that appears:
    - Enter a label and IP address for your data store.
    - 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.**
    - Set disk capacity in GB.
    - If required, you can also bind the data store with a local hypervisor. This is helpful if you wish that the data store and a hypervisor were located on the same physical server thus decreasing the time needed for a hypervisor-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)
    - Select the **lvm** data store type.
- 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 hypervisor or a hypervisor zone.

2. **Find the data store's unique identifier (this is needed to create your volume group in step# 4):**

Rad the IDENTIFIER from the data stores screen:
http://xxx.xxx.xxx.xxx/settings/data_stores

3. **SSH into a hypervisor that is able to connect to this datastore. 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 hypervisor/volume group visibility:**

Now you have the new datastore formatted you should be able to see the volume group from all hypervisors. To test this, run *pvscan* and *vgscan* on all hypervisors. Make sure you can see all identifiers on all hypervisors.

6. Backup Server installation

ⓘ Skip this section if you are using a Cloud Boot method.

1. **Add a backup server to the webUI:**
   - Log into your Control Panel.
   - Go to the *Settings* menu and click the Backup Servers icon.
   - Click the Add New Backup Server button.
   - Fill in the form that appears:
     - Give your backup server a label.
     - Enter the backup server IP address (IPv4).
     - Set the backup server capacity (in GB).
   - Tick the Enabled box to enable the backup server.
   - Click the Add Backup Server button to finish

2. **Download the OnApp repository:**

```
bash# rpm –Uvh wget http://rpm.repo.onapp.com/repo/centos/5/onapp-
repo.noarch.rpm
bash# rpm -Uvh onapp-repo.noarch.rpm
```

Download the OnApp 3.0.0 repository config file:

```
bash# wget -N http://rpm.repo.onapp.com/repo/centos/5/OnApp-3.0.0.repo
-P /etc/yum.repos.d/

bash# yum clean all
```

3. **Install the OnApp Backup Server installer package:**

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

4. **Check and set Backup Server default settings**

   Edit Backup Server default settings (such as templates and backups directories, and ntp server) by editing the /onapp/onapp-bk.conf file:

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

5. **Run the installer:**

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

ⓘ To get the information about installer and its options, such as packages update, templates download and non-interactive mode, run the installer with '-h' option.

```
bash# /onapp/onapp-bk-install/onapp-bk-install.sh -h
Usage: /onapp/onapp-bk-install/onapp-bk-install.sh [-c CONFIG_FILE] [-
a] [-y] [-t] [-h]
```

Options

`-c CONFIG_FILE`	Custom installer configuration file. Otherwise, the preinstalled one is used.
`-a`	Non-interactive mode. Automatic installation process.
`-y`	Update all packages on the box with 'yum update'. The update will be processed if the -a option is used.
`-t`	Download of Base, Load Balancer and CDN templates. The download is initiated if '-a' option is used.
`-h`	Print this info.

ⓘ Use -y option carefully, as it updates all packages in the box with 'yum update'.

ⓘ 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.

ⓘ The -a option switches the installer into a non-interactive mode (nothing will be performed). This option also processes the packages update and templates download.

7. Control Panel cloud configuration

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

7.1 Create hypervisors and hypervisor zones

1. **Create a new hypervisor zone:**
    - Go to your Control Panel's *Settings* menu and click the *Hypervisor Zones* icon.
    - Click the *Add New Hypervisor Zone* button.
    - On the screen that follows, give your hypervisor zone a name (label).
    - Make sure that the *disable failover* option is selected.
    - Click the *Save* button to finish.

2. **Add your new hypervisor to the control panel:**
    - Go to your Control Panel's *Settings* menu and click the *Hypervisors* icon.
    - Click the *Add New Hypervisor* button and fill in the form on the screen that appears:
        - The hypervisor's IP address should be its IP on the management network
          The memory overhead is determined by the output of the hypervisor installation script.
          ⓘ *The value you noted when installing the hypervisor.*
        - Make sure that "disable failover" is selected
        - Make sure that you select the "Enable" option
    - Click the *Add Hypervisor* button to finish. You can view the hypervisor under the main *Hypervisors* menu.

3. **Add that hypervisor to your new hypervisor zone:**
    - Go to your Control Panel's *Settings* menu and click the *Hypervisor Zones* icon.
    - Click the label of the zone you want to add a hypervisor to.
    - The screen that appears will show you all hypervisors in the cloud, organized into two lists – those assigned to the zone already, and those that are unassigned.
    - In the unassigned list, find the hypervisor you want to add to the zone, and click the *Add* icon next to it.

7.2 Create networks and network zones

1. **Create a new network zone**
   - Go to your Control Panel's *Settings* menu and click the *Network zones* icon.
   - Click the *Add New Network zone* button.
   - On the screen that follows, give your network zone a name (label) and then click the *Save* button.

2. **Create a new network**
   - Go to your Control Panel's *Settings* menu and click the *Networks* icon.
   - Click the *Add New Network* button at the end of the list.
   - On the screen that follows, give the new network a name (label), a VLAN number, and assign it to a network zone if required.
   - Click the *Add Network* button to finish.

ⓘ The network label is simply your choice of a human-readable name – "public", "external", "1Gb", "10Gb" etc.

ⓘ The VLAN field only needs to be given a value if you are tagging the IP addresses you will add to this network with a VLAN ID (IEEE 802.1Q). If you plan to tag IP addresses in this way, you need to make sure the link to the public interface on the hypervisors is a trunked network port. If you are not VLAN tagging addresses, this field can be left blank and the public port on the hypervisor can be an access port.

3. **Add that network to your new network zone**
   - Go to your Control Panel's *Settings* menu and click the *Network Zones* icon.
   - Click the label of the zone you want to add a network to.
   - The screen that appears will show you all networks in the cloud, organized into two lists – those assigned to the zone already, and those that are unassigned.
   - In the unassigned list, find the network you want to add to the zone, and click the *Add* icon next to it.

4. **Add a range of IP addresses to the new network**
   - Go to your Control Panel's *Settings* menu.
   - Click the *Networks* icon: the screen that appears shows every network available in your cloud.
   - Click the name (label) of the network you want to add addresses to. On the screen that follows you'll see a list of all IP addresses currently assigned to this network.
   - Click the *Add New IP Address* button at the bottom of the screen, and complete the form that appears:
     - *IP Address* – add a range of addresses. For example:
       - '192.168.0.2-254' or '192.168.0.2-192.168.0.254' (IPv4)
         '2001:db8:8:800:200C:417A-427A' (IPv6).

- ○ *Netmask* – for example: '255.255.255.0' (IPv4) or '24' (IPv6).
- ○ *Gateway* – enter a single IP to specify a gateway. If you leave this blank the address will be added without a gateway.
- ○ *Don't use as primary during VM build* – If you tick this box, the IP addresses you add will never be assigned as primary IPs. Primary IPs are only allocated to VMs when the VM is built, so with this box ticked, the address range will never be assigned to a newly built VM.

- Click the *Add New IP Address* button to finish.

ⓘ You can add up to 1,000 IP addresses at once. To add more than 1,000 addresses, repeat the procedure again.

7.3 Create data stores & data store zones (traditional/centralized SAN)

Use this information to set up data stores based on traditional/centralized storage.

1. **Create a new data store zone**
   - Go to your Control Panel's *Settings* menu and click the *Data store zones* icon.
   - Click the *Add New Data store zone* button.
   - On the screen that follows, give your data store zone a name (label) and then click the *Save* button.

2. **Create a new data store**

- Go to your Control Panel's *Settings* menu and click the *Data Stores* icon.
- Click the *Add New Data* Store link at the bottom of the screen that appears.
- Fill in the form on the next screen:
  - Label – give the data store an appropriate name (eg HV1-SSD)
  - Size – you can get this by doing fdisk -l /dev/sdaX, for example
  - IP address
  - Select "local hypervisor" (since this will be local storage)
  - Leave the hypervisor zone field blank
  - Don't disable the data store
  - Assign the data store to the data store zone you just created using the drop-down menu.
- When you've finished configuring the store, click the *Add Data Store* button.

ⓘ Follow these steps for each local storage block on the hypervisor.

3. **Configure the data store on your hypervisor**

ⓘ 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 DataStores screen in the Control Panel.

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

4. **Update necessary sysctl variables and reload**

Edit your /etc/sysctl.conf file. If the netfilter.ip_conntrack_max entry exists, update the value: if it doesn't exist, add it. You can increase the netfilter.ip_conntrack_max value if required.

```
net.ipv4.netfilter.ip_conntrack_max = 256000
bash#> sysctl -p
```

7.4 Create data stores & data store zones (OnApp Storage/integrated SAN)

Use this information to set up data stores based on OnApp Storage, our integrated distributed SAN.

1. Create a new data store zone

- Go to your Control Panel's *Settings* menu and click the *Data store zones* icon.
- Click the *Add New Data store zone* button.
- On the screen that follows, give your data store zone a label and then click the *Save* button.

2. Create a new data store

Once some hypervisors have been added (Xen or KVM) with integrated storage enabled, you can group their drives together into a virtual data store.

To create new integrated data store:

1. Go to your Control Panel's **Integrated Storage** menu.
2. On the screen that appears, you'll see the list of all distributed storage data stores in the cloud, Click the the **Integrated Storage** menu item, and a graphical list of all storage nodes available in your distributed SAN (i.e. all drives on hypervisors.)
3. To create a new data store, click the **Create New Integrated Storage Datastore** button, and complete the wizard that follows:

   Name- give your datastore a name

   Advanced settings - check this to reveal the Advanced settings below:

   - *Redundancy* - increasing the number of copies increases resilience to invidivual drive failure.
   - *Stripes* - increasing the number of stripes increases the number of physical disks involved in any single virtual disk.

   Nodes

   - *HV filter* - use this to filter the nodes (disks) available for inclusion in this data store, by specific hypervisors.
   - *Performance* - use this to filter the nodes available for inclusion in this data store by performance.
   - *[disk volume/ performance]* - individual disks are displayed according to the filters above: select which disks you want to include in this data store. Disks that do not meet the filter settings are greyed out.

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

7.5 Join networks and datastores to hypervisors

1. **Join datastores to hypervisors**
   - Go to your Control Panel's *Settings* menu and click the *Hypervisors* icon.
   - Click the label of the hypervisor you want to manage data stores for.
   - On the screen that appears, click the *Manage Data Stores* link in the *Actions* section.
   - On the screen that follows, you'll see a list of all data stores currently associated with this hypervisor:
     - To add a data store join, choose a data store from the drop-down menu and click the *Add Data Store* button.
     - To remove a data store join, click the *Delete* icon next to it. You'll be asked for confirmation before the store is removed.

2. **Join networks to hypervisors**
   - Go to your Control Panel's *Settings* menu and click the *Hypervisors* icon.
   - Click the label of the hypervisor you want to manage networks for.
   - On the screen that appears, click the *Manage Networks* link in the *Actions* section.
   - On the screen that follows, you'll see a list of all networks currently associated with this hypervisor:
     - To add a new network join, choose a network from the drop-down menu, enter its interface name (eth0, eth1) and click the *Add Network* button.
     - To remove a network join, click the *Delete* icon next to it. You'll be asked for confirmation before the network is removed.

ⓘ Note that when you join the network to a hypervisor you must specify the relevant NIC: this should be a dedicated NIC with a blank config that is patched to route the network in question.

Appendix: document revisions

v1.0, 25th September 2012

- First release