3.0 Get Started

3.0 Get Started

Author: Anonymous User
Version: 17
Date: 24-Apr-2015 11:09

Table of Contents

The guides in this section apply to installing and upgrading different editions of the OnApp Cloud 3.0 version (currently, from v3.0 to v3.0.8).

For the release notes list, please refer to the Release Notes space.

1 OnApp Cloud v.3.0.x Preparation Guide

This document describes how to prepare for the deployment different editions of the OnApp Cloud 3.0 version (currently, from v3.0 to v3.0.8). Please review the configuration details in each section carefully, as they are vital to the smooth operation of OnApp Cloud.

1.1 Overview

OnApp Cloud software enables service providers to turn their existing infrastructure (or any commodity hardware) into a single pool of resources - "a cloud" - which can then be sold to end users on a utility basis. It's a complete cloud deployment and management platform that's designed to make it easy for service providers to sell a wide range of cloud services.

This guide covers the three main areas you need to consider when preparing for deployment: network configuration, server configuration, and storage.

What's new in OnApp Cloud v3.0?
Preparing for deployment with OnApp Cloud v3.0 is basically the same as earlier versions, with three main exceptions:

- **SolidFire support**: OnApp Cloud v3.0 includes integration with the SolidFire storage management system. We've added a section that deals with the specifics of an OnApp /SolidFire cloud.
- **OnApp Storage integration**: OnApp Cloud v3.0 now includes OnApp Storage, our distributed block storage system. There are new sections that explain the storage and networking specifics of a cloud built with our integrated SAN.
- VMware support: OnApp Cloud v3.0 now supports VMware clouds, and these operate in a slightly different way to clouds built on Xen and KVM hypervisors. There are new sections on networking and hardware that deal with the changes here.

1.1.1 Getting support

24x7 support
OnApp customers with a full (paid) license can contact OnApp Support at any time:

- support@onapp.com
- http://onapp.com/support
- (+1) 888 876 8666

Forums

Visit http://forum.onapp.com to get support from the OnApp community. Members of OnApp's support and engineering teams also monitor the forums and contribute to discussions. To access the forums, log in with your OnApp Dashboard account details.

Documentation

For the latest OnApp documentation, see https://onappdev.atlassian.net/wiki/dashboard.action.

1.2 Network Configuration

The correct network configuration is important to ensure your cloud has optimal performance and stability. There are different recommended configs depending on your approach to backup, storage, and whether you're using VMware or not.
There are four core networks in a standard OnApp Cloud installation. The first part of this chapter explains each network requirements. The second part gives network diagrams for different deployment scenarios.

1.2.1 Network Roles in OnApp Cloud Using VMware

For proper VMware setup, your network confguration must meet the following requirements:

- An OnApp Cloud using VMware should have separate networks to isolate management, storage and VLAN traffic.
- The CP server needs to have access to the vCenter and all of the ESXi servers in the cluster used by CP on the management network.
- An optional vMotion network can be included, to give the cloud administrators the ability to hot-migrate VMs between hypervisors (using vCenter's management GUI).
- VMs network connectivity is performed through customer networks and VLANs.
- The VLAN networks are protected with a Vyatta firewall managed by the OnApp software. These can be Physical or Virtual firewalls running on an attached network.
- The OnApp Control Panel Server must be able to communicate with vCenter and Vyatta over the OnApp management network.
- The Vyatta installation will also connect to the OnApp appliance network.

1.2.2 Network Roles in OnApp Cloud Using Xen/KVM

There are four core networks: Storage, Management, Provisioning and Appliance. It is very important to separate these four core networks, either physically, using different switches, or with VLANs if your network supports it. The role of each network is explained below.

Appliance Network/VM Networking

The Appliance Network in OnApp is used for VM networking only: it provides network connectivity for virtual machines.
OnApp will bridge the public NIC and assign virtual interfaces to it when VMs are provisioned, and/or when additional network interfaces are added to VMs from the Web UI, or via the OnApp API. As the public interface is managed fully by OnApp, the public NIC requires a blank config - for example:

```
/etc/sysconfig/network-scripts/ifcfg-ethX
ONBOOT=no
BOOTPROTO=none
```

You should configure your network interface file accordingly. You will not need to add any configuration to this NIC, so no subnet, gateway or IP address details should be added.
The NIC could either be a standard physical interface (e.g. eth1) or a bonded interface (e.g. bond1). It *cannot* be a sub-interface (e.g. eth1:1) or a vlan sub-interface (e.g. eth1.101) so you should allow for this when you are designing your hypervisor, as you must make sure you have a physical NIC available.
This network should be a minimum of 1Gbit. You should also consider bonding on the Appliance Network to introduce redundancy at the network level.
You'll need to connect your Appliance Network to a switch trunk port if you want to use VLANs. VLANs allow a network administrator to segregate traffic for bandwidth or security purposes.
If you choose to VLAN your VM networking you'll need to associate your VLAN with the subnet when you add the VM networking range to OnApp.
Configuring a switch trunk port is the preferred method because it gives you additional flexibility and security. Alternatively you can configure a switch access port. If this is the case then you will not need to specify a VLAN when adding the range to OnApp.

Some hosting companies have limitations and the transfer of IP addresses between servers can sometimes require manual interventions - a change on their user portal, for example - so if you are leasing hosting server solutions, it is worth double-checking with your host that this will be possible.

If you are running VMware hypervisors, the method of VM networking will differ slightly as your Vyatta installation will manage some of the virtual routing.

Management Network

This network is responsible for a couple of different tasks. It provides incoming and outgoing connectivity to the servers, which means the Management Network should always be the default gateway.
If you are going to use Cloud Boot, this should be a local network behind a gateway device that is capable of bridging traffic to the Internet to allow the servers to perform tasks such as dns resolution, ntp updates and operating system updates. Also, you have to open the 5555 port for outgoing connections to the licensing server.

The control panel will need to have incoming traffic allowed to ports 80/443 & 30000->40000. This should again be configured at the gateway with incoming NAT. If your gateway device is not capable of supporting this then this network can also be an external network, but should always be firewalled at the gateway to block all incoming traffic with the exception of the ports listed above.
The Management Network also serves as a route for communication between the control panel server and the hypervisors for critical OnApp internal traffic. That means the stability of this network is critical: and you should always consider bonding to introduce network level redundancy, and the network should run at least 1Gbit.

Provisioning Network

The Provisioning Network is used to transfer backup and template data between the provisioning server and the primary storage volumes. The network will be used to transfer large amount of data, so we recommend that it runs at least 1Gbit. Ideally you should consider 10Gbit, FibreChannel, InfiniBand or aggregated 1Gbit links for maximum throughput.

Storage Network

The Storage Network provides the connection between storage devices (eg SANs) and the hypervisors.
The type of network will depend on what kind of connectivity your primary storage requires. For example, if you are using iSCSI or ATAoE, you will need to setup an ethernet network. If your SAN has fibre connectivity, then the Storage Network will be a fiber network.
The stability of the Storage Network is absolutely critical. You should always make redundancy your primary concern when designing this network. The Server configuration -> SAN -> fabric components section of this document discusses this in more detail.

- The Storage Network must be a local network.
- We recommend this network runs at 1Gbit, at least; however, you may want to consider 10Gbit, FibreChannel or InfiniBand to achieve maximum performance. (OnApp integrated Storage currently does not support InfiniBand)

- To achieve better performance and redundancy over 1Gbit you should consider NIC teaming/bonding and LACP or MPIO over multiple subnets if not using OnApp Integrated Storage.
- If your primary storage network is running over Ethernet, then it is important that the switch connecting the hypervisors to the SAN supports jumbo frames: the Storage Network on the hypervisors and the SAN(s) must have MTU set to 9000 to optimize performance.

We strongly recommend that you avoid NICs using Broadcom chipsets on the Storage Network due to known issues surrounding iSCSI and TCP offload in the Linux kernel modules.

OnApp Cloud v3.0 includes OnApp Storage functionality, which enables you to create your own distributed SAN. It also supports traditional centralized SANs. Deployment considerations for each approach to storage are explained in Chapter 3 .

Note: Emulex hardware is currently does not have support for 3.x Linux kernels, so is only compatible with CentOS 5.x

1.3 Recommended Network Configurations

This section lists the recommended network configurations for an OnApp Cloud installation.

1.3.1 For a Mixed Xen/KVM/VMware Cloud

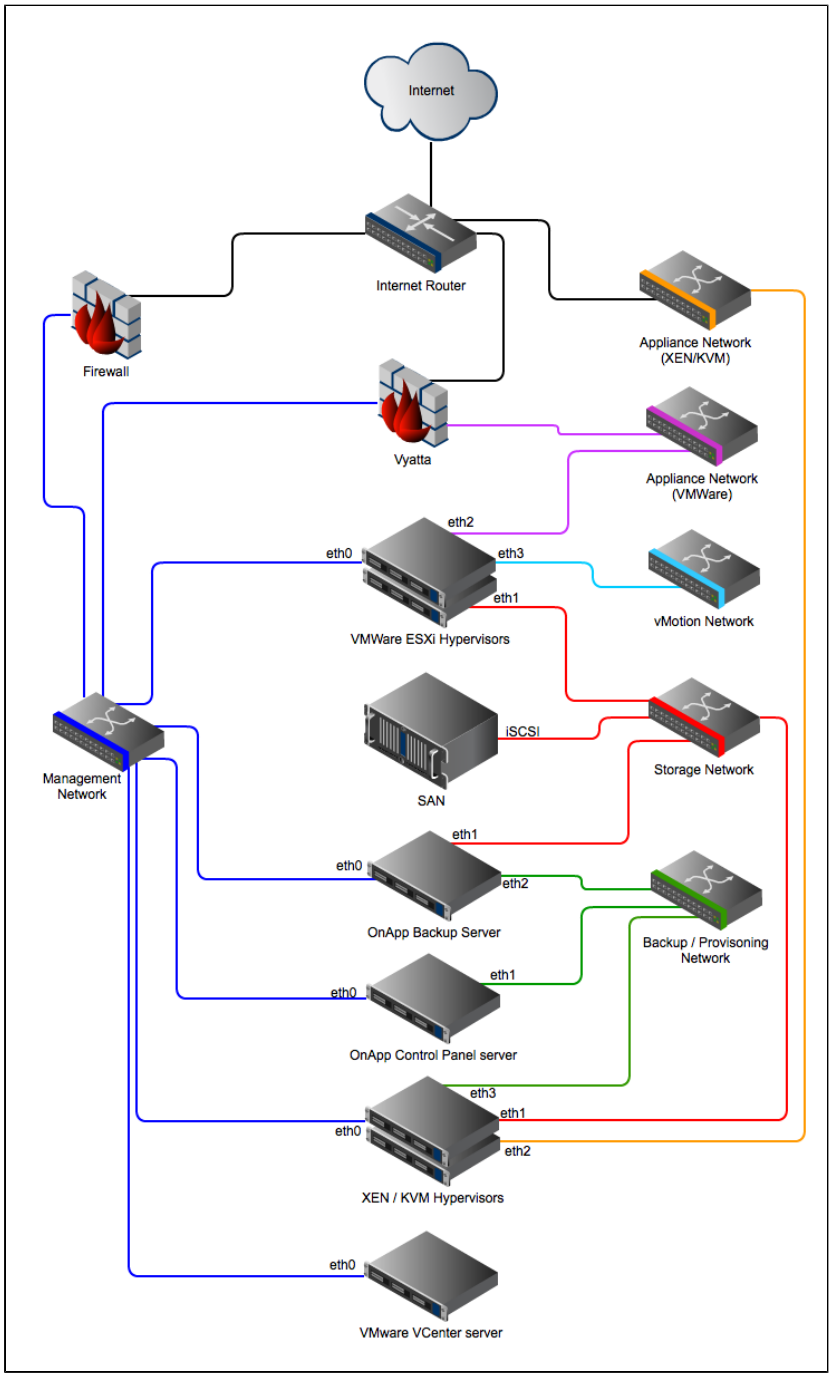

1.3.2 For a VMware loud

N.B. OnApp Storage, our integrated distributed SAN, does not currently support VMware.

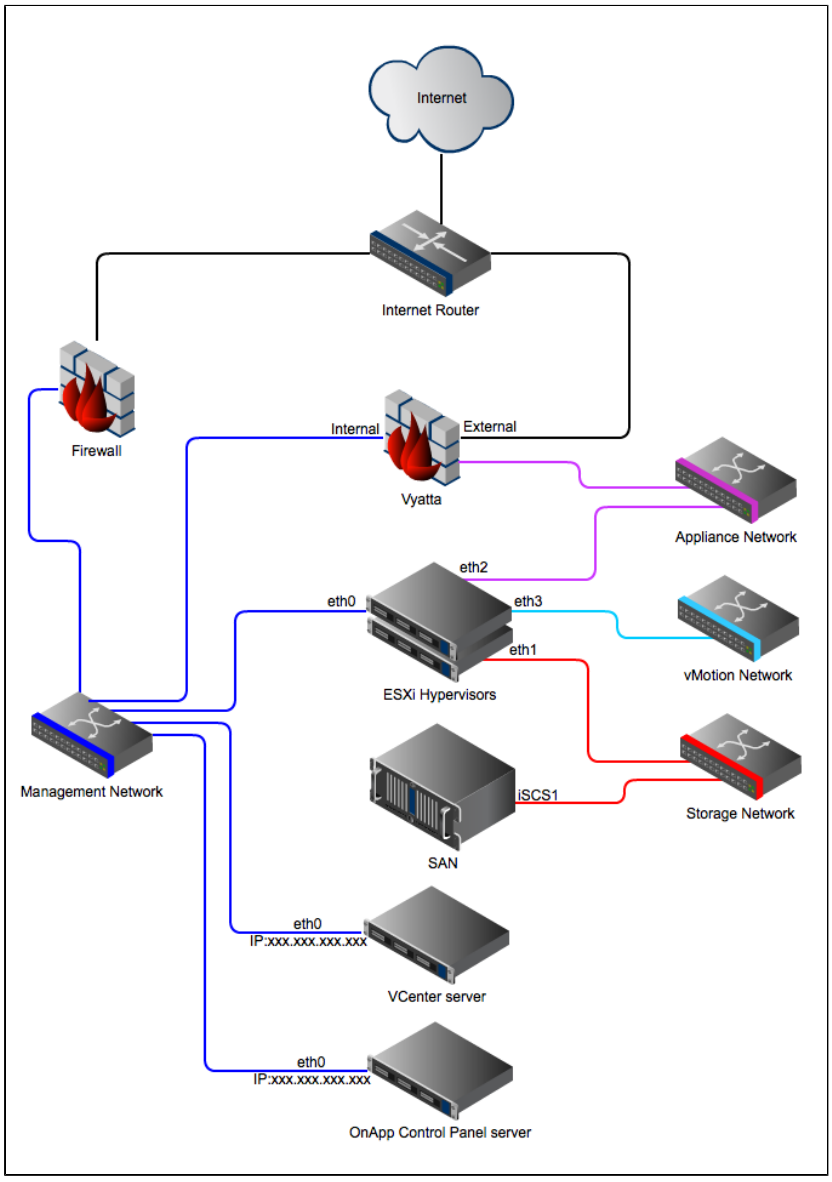
Internet
Internet Router
Firewall
Internal
External
Vyatta
Appliance Network
eth2
eth0
eth3
eth1
ESXi Hypervisors
vMotion Network
Management Network
iSCS1
Storage Network
SAN
eth0
IP:xxx.xxx.xxx.xxx
VCenter server
eth0
IP:xxx.xxx.xxx.xxx
OnApp Control Panel server

1.3.3 For a Xen/KVM Cloud

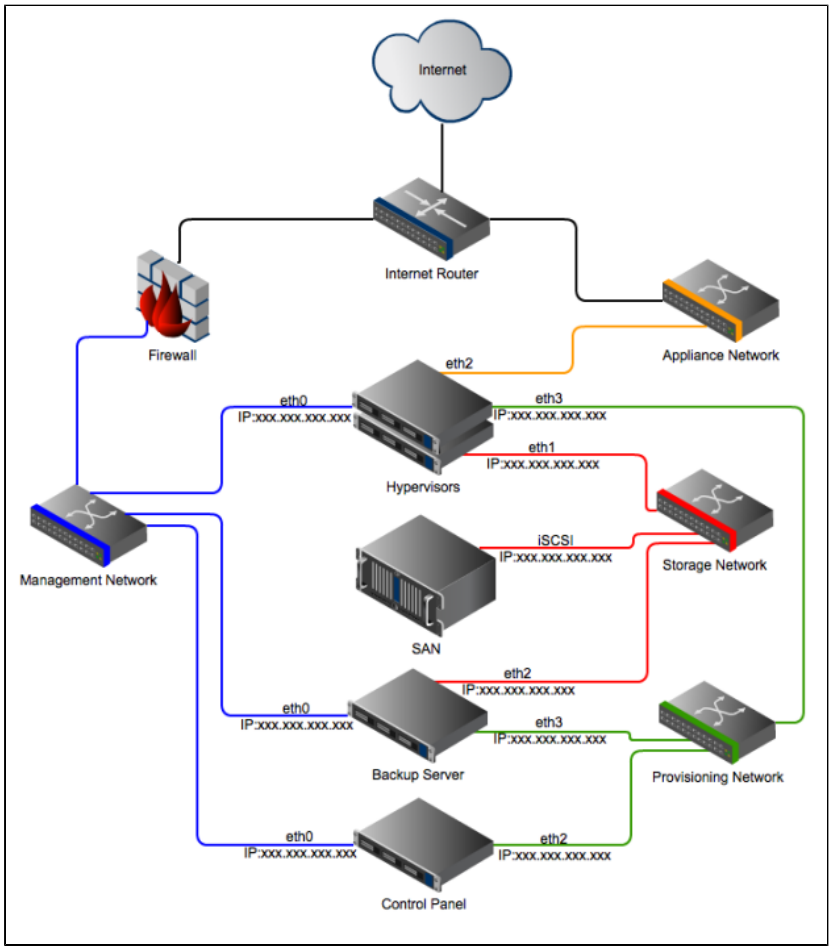

1.3.4 For a Xen/KVM cloud using OnApp Storage (integrated distributed SAN)

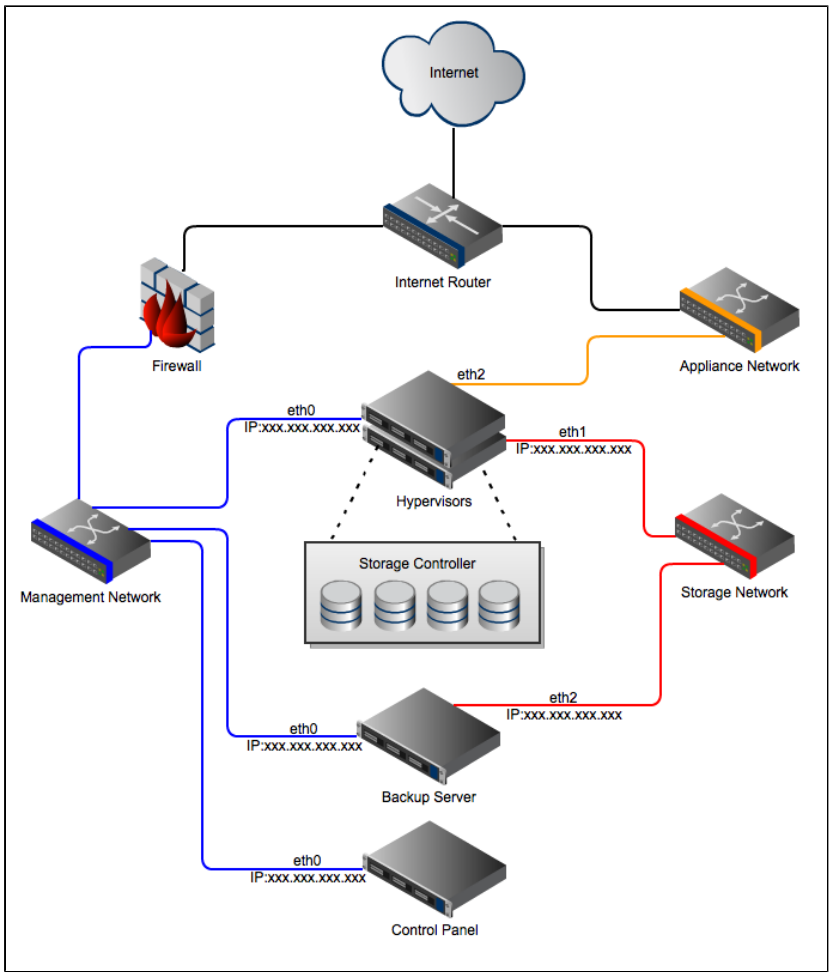

Provisioning network is not required for clouds using Integrated Storage with dedicated backup servers.

1.4 Server onfiguration

This section lists the server installation requirements needed for an OnApp Cloud installation. For minimum hardware specs, see http://www.onapp.com/onapp-requirements. OnApp will primarily run on CentOS or Red Hat but the version will depend on what virtualization method you will be running.

1.4.1 Supported server configuration

- **XEN Hypervisors** CentOS 5.x x64 or CentOS 6.x x64
- **KVM Hypervisors** CentOS 5.x x64 or CentOS 6.x x64

- **OnApp Control Panel Server** CentOS 5.x x86/X64 or CentOS 6.x x86/64
- **OnApp Backup Server** CentOS 5.x x64 or CentOS 6.x x64
- **Integrated Storage** CentOS 5.x x64 or CentOS 6.x x64

1.4.2 Recommended server configuration

We highly recommend using the following server configuration:

- **XEN 4.0 Hypervisors** CentOS 6.x x64,
- **KVM Hypervisors** CentOS 6.x x64
- **OnApp Control Panel Server** CentOS 6.x x86/64
- **OnApp Backup Server** CentOS 6.x x64

PLEASE NOTE: Cloud Boot is not compatible with CentOS 6 Xen hypervisors and CentOS 5 KVM hypervisors.

Full root access: please do not create the user 'onapp' since this is created as part of the RPM installation.

Note: Currently Emulex hardware does not support 3.x Linux kernels, so it is only compatible with CentOS 5.x.

VMware configuration requirements

To be able to run VMware ESXi servers through OnApp, you need to install vCenter and Vyatta externally. These can be installed on virtual machines if necessary but should not be hosted inside OnApp. See Additional hardware considerations for VMware section for details.

OnApp Cloud v3.0 supports VMWare vSphere 5, with all standard OnApp functionality available in line with the VMware VSPP 'Standard' licensing model. For use with OnApp, vCenter 5.x must be installed on Microsoft Windows.

1.4.3 Control Panel server

The control panel server is absolutely critical to the stability and performance of the cloud. There are a few things to consider when choosing hardware for this server. It is very simple to grow your cloud as you start to sell more resources, and as you add more hypervisors and SANs this puts more load on the control panel. Choosing the right hardware at the beginning is important and avoids having to take the server down for upgrades later down the line, causing interruption to customers.

The control panel server will become very mysql heavy as you add more hypervisors, so a fast disk array and lots of memory is recommended. A good example would be a 4xSAS RAID10 array with 24GB RAM and quad core Xeon CPU. SSD storage can also be considered.

If you have a control panel server spec in mind, you're very welcome to send it to your OnApp integrations specialist for review.

1.4.4 Backup server

The backup server stores virtual machine snapshots and virtual machine templates. It is also responsible for processing any disk transactions running in your cloud, such as provisioning virtual machines, taking backups or resizing disks.

The backup server must hold a backup storage volume. This can be a local disk array or can be mounted via NFS or iSCSI from a back end storage node. Note that the backup volume should not be presented from the same physical hardware that presents the primary storage volume to the hypervisors.

Unlike primary storage, performance is not so essential here – there is less need for RAID10 or a high volume of spindles. You can consider a RAID level that provides more space as opposed to redundancy and performance: RAID5 or 6 is usually ideal for the backup volume. Take care when configuring the SAN, however: a larger block size is recommended owing to the nature of the data being stored on this array.

Backup storage will be used to hold very large files, so we recommend that it's at least 1.5 - 2x larger than the primary storage volume(s) available in the cloud. Additional backup servers can be added to your cloud as needed.

NOTE: In the traditional/centralized SAN configuration, you have to bind all your data stores to the backup server.

In the OnApp cloud with cloud boot enabled, you have to use cloud booted backup servers instead of dedicated backup servers. To do so, you have to create a cloud boot hypervisor to be used as a backup server.

Starting from the 3.0 version, you are able to set up Cloud Boot backup servers and virtual dedicated backup servers to be used with the Integrated Storage functionality. The backup scheme remains unchanged.

1.4.5 Hypervisor servers

Hypervisors are where virtual machines live in your cloud. A small amount of hypervisor CPU, memory and disk resource is reserved for the OnApp engine: the remainder is available as virtual resources to allocate to virtual machines.
If you are using a centralized SAN, then the virtual machines' disks will live on that SAN, and, the hypervisor's own disk will simply be used to boot the hypervisor and run the OnApp engine. Performance here is not critical but we recommend introducing some redundancy: RAID1 SATA /SAS would be perfect.
If you are using OnApp Storage, our integrated SAN, you should obviously factor more disks into your hypervisor spec to enable the creation of a distributed SAN using those disks.
If you choose not to run a centralized SAN or OnApp Storage, it is possible to have storage running locally on hypervisors, though you lose the ability to failover from hypervisor to hypervisor: this is not recommended for an optimal cloud set-up.
When you are building your hardware it's important to take into consideration the specifications of the primary components that will be virtualized - the RAM and CPU.
Remember that while you can oversell CPU cores in OnApp, RAM is a dedicated resource, so the physical limitation to how many VMs you can fit on a single hypervisor is limited by the amount of RAM installed in that hypervisor.
Another limitation to consider is that the hypervisor's CPU is a shared resource: the physical cores are shared among the VMs running on a hypervisor. You don't want to overload the hypervisor with too many VMs, as this will stretch the available CPU time and degrade the performance of all VMs on that hypervisor.
It's also important to note that too many VMs could potentially saturate the SAN NICs on the hypervisor, which will also introduce instability and performance loss to VMs (see the server configuration -> SAN -> host components section for more detail).
In the network diagrams chapter you can see that OnApp requires at least 4 NICs on the hypervisors. Note that this does not take into consideration any bonding or multipath configurations, which we recommend for any production setup on most if not all of our networks. You should at least consider bonding on the Management Network and multipath on the Storage Network(s) to improve stability and performance.

You must have Intel-VT or AMD-V enabled in the BIOS of all hypervisors to enable you to provision Windows-based virtual machines on your OnApp cloud!

Cloud Boot hypervisors

Cloud Boot is a feature that enables fast provisioning of Xen and KVM Hypervisors without any pre-installation requirements. Using network/PXE boot methods, a new server can be plugged in and powered on, automatically discovered by the OnApp Control Panel Server, and installed over the network so it boots as a fully configured hypervisor, ready to host virtual machines.
The Control Panel Server manages IP address to hardware MAC assignment, and the booting of a Xen or KVM image on demand. Hypervisor images come pre-installed with all the ssh keys and any other settings specific to the node, to enable hypervisors to come online instantly.
Images are booted as a standalone ramdisk, so once bootstrapped they operate independently from other servers, but without any persistent installation dependency.
This enables booting of diskless blades, as well as booting hypervisors with the new integrated storage platform enabled (OnApp Storage) where all local storage drives are presented to the integrated SAN.

Dependencies:

- Network/PXE boot must be supported and enabled on the primary management NIC for the HV servers
- A secondary NIC is recommended for the Control Panel Server to provide a fully isolated network for the HV management subnet, including PXE boot and DHCP support for the HVs.

N.B. For resilience, a secondary static tftp server target can be configured to handle Controller Server failure and ensure hardware boot consistency in the event of such a failure.

The following Cloud Boot features are not yet available (they will be introduced in future releases):

- Bonded NICs for the management/boot interface

1.4.6 Additional hardware considerations for VMware

If you are looking to run VMware ESXi servers through OnApp then you will require an external installation of vCenter and Vyatta. These can be installed on virtual machines if necessary but should not be hosted inside OnApp.
The OnApp Control Panel Server will need to be able to communicate with vCenter and Vyatta

over the OnApp management network. The Vyatta installation will need connectivity to the OnApp appliance network.
In vCenter you should create a single datacenter cluster: OnApp will associate the name of the cluster with hypervisors when they're added to OnApp. The external Vyatta installation will be used to manage VLANs and route VM networking traffic in and out of OnApp.

1.4.7 Centralized Storage (SAN)

Primary storage is abolutely critical to your cloud, and your SAN will have a huge impact on the performance of the whole platform.
OnApp gives you a lot of flexibility in your primary storage technology. It supports anything that is capable of presenting a block device to hypervisors. This could be, for example, FiberChannel, SCSI or SAS HBA, iSCSI or ATAoE, or a InfiniBand HCA controller, since all of these present the block device directly. OnApp does not support services such as NFS for primary storage, because these present a filesystem and not the block device.
Beyond the type of block device there are three main things to consider in your SAN design: the host, fabric and storage components. You need to think about each very carefully and pay particular attention to performance, stability and throughput when planning your SAN.

Fabric components - the network fabric between hypervisors and SANs

You will need to think about redundancy, and whether you need to design a fault tolerant switching mesh to coincide with your multipath configurations at the host and SAN ends.
You should also think about future growth: as you add more hypervisors and SANs to the cloud you will need to be able to grow the physical connectivity without downtime on the Storage Network.

Host components - hypervisor connectivity to the Storage Network

You will need to make sure that your ethernet or HBA drivers are stable in this release. We recommend that you test this thoroughly before handing over to OnApp to deploy your cloud on your infrastructure.
You will also need to think about the throughput, and whether the connectivity on hypervisors will be suitable for the virtual machines they'll be running. A bottleneck here will cause major performance issues.
You should consider adding multiple HBAs or NICs if you plan to run a redundant switching mesh (see the fabric section below) as bonding or multipath will be required, unless the redundancy is built into the physical switch chassis (failover backplanes for example).

Storage components - the SAN chassis, controllers and disk trays

You need to take into consideration the size of storage required and the physical capacity you have to achieve this. This will give you a good idea on the size of disks you will be adding into the array and the RAID level you will choose.
As a general rule, more spindles in the array will give you better performance: you should avoid using a small number of large disks, or you will start to see I/O bottlenecks as you make increasing use of the storage in future.
You should also think about the physical storage hardware, and whether you'll be using SATA, SAS or SSD. Again, this will have a great impact on the I/O capabilities of the array.
It's also a good idea to consider RAID levels carefully and look into the advantages and disadvantages of each. We recommend RAID10. Although you will lose 50% of your capacity you will see good performance for both read and write, which is important for primary storage. RAID10 will also give you much better redundancy on the array.
Controller caching is another issue to consider. You should always aim to have both read and write caching. If you are looking at write caching you should also look at battery backups for the write cache. Some controllers also support SSD caching which can be a great advantage.
As with the host components, you should also take your HBA and ethernet connectivity into consideration, to ensure you have both the redundancy and throughput required for your cloud infrastructure.

SolidFire integration

Starting with the 3.0 version, OnApp is integrated with the SolidFire storage management system. With the Solid Fire integration it is possible to utilize the SF SAN directly within the OnApp cloud and manage the SolidFire cluster via the SolidFire API.

You can perform the following options with SolidFire:

- Utilize SolidFire SAN in the OnApp cloud.
- Allocate dedicated LUNs from the SF cluster per virtual machine disk, when creating a VM. (LUN is created per each VM disk, with a separate lun per swap disk.)
- Manage SolidFire LUNs automatically via API.
- Create virtual machines without the swap disk.
- Implement backups / snapshots using SF CloneVolume method

There is a disk dependency between OnApp and SolidFire - when a new disk is created on the OnApp side, a new LUN is created automatically on the SF side, using the CreateVolume API call. The LUNs on the SolidFire are managed automatically vis SolidFire API.

Inasmuch SolidFire data store has two interfaces: OnApp and SolidFire, you have to specify two IP addresses when creating a SolidFire data store.

To be able to use the SF volume, you have to enable export to this device (hypervisor or a data store). To do that, you need to send an account username and initiator password to the iscsi_ip address. You will be able to use this device after the authorization.

The following options are not available under SolidFire:

- It is not possible to migrate SolidFire disks, as SF virtualizes the storage layer.
- SolidFire does not support live disk resize. To reside disk, you need to shut down the virtual machine first and use the CloneVolume functionality to increase the disk size. After the disk resize operation is complete, the original volume will be replaced with the new one and deleted, after that the VM will be booted.

To be able to utilize SolidFire in the cloud, you need to install the SolidFire storage system first.

1.4.8 Integrated Storage (OnApp Storage)

OnApp Storage is a distributed block storage system that allows you to build a highly scalable and resilient SAN using local disks in hypervisors. With OnApp Storage you create a virtual data store that spans multiple physical drives in hypervisors, with RAID-like replication and striping across drives. The SAN is fully integrated into the hypervisor platform, and the platform is completely decentralized. There is no single point of failure: for example, if a hypervisor fails, the SAN reorganizes itself and automatically recovers the data.
The following requirements are recommended for integrated storage implementation:

- Any number of integrated storage drives can be grouped together across any HV
- SSD drives are recommended for best performance
- A least 1 dedicated NIC assigned per HV for the SAN
- Multiple NICs bonded or 10GBit/s ethernet (recommended)

The following Cloud Boot features are not yet available (they will be introduced in future releases):

- Bonded NICs for the management/boot interface

Currently the largest physical disk which can be used with OnApp Storage is 2TB.

PLEASE NOTE: To start using integrated storage, you must enable it in the system configuration first (**Settings > Configuration > OnApp Storage**).

1.4.9 Appendix: document revisions

v1.0, 20th March 2013

- First release

2 OnApp Cloud v.3.0.x Installation Guide

This document describes how to install different editions of the OnApp Cloud 3.0 version (currently, from v3.0 to v3.0.8). Please read each section carefully, as it is vital to the smooth operation of OnApp Cloud.

2.1 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.

2.1.1 Supported server configuration

- **XEN Hypervisors** CentOS 5.x x64 or CentOS 6.x x64
- **KVM Hypervisors** CentOS 5.x x64 or CentOS 6.x x64
- **OnApp Control Panel Server** CentOS 5.x x86/X64 or CentOS 6.x x86/64
- **OnApp Backup Server** CentOS 5.x x64 or CentOS 6.x x64
- **Integrated Storage** CentOS 5.x x64 or CentOS 6.x x64

2.1.2 Recommended server configuration

We highly recommend using the following server configuration:

- **XEN 4.0 Hypervisors** CentOS 6.x x64,
- **KVM Hypervisors** CentOS 6.x x64
- **OnApp Control Panel Server** CentOS 6.x x86/64
- **OnApp Backup Server** CentOS 6.x x64

PLEASE NOTE: Cloud Boot is not compatible with CentOS 6 Xen hypervisors and CentOS 5 KVM hypervisors.

2.2 Control Panel Installation

NOTE:

- 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/onapp-repo-3.0.noarch.rpm
```

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 Cloud Boot 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).

8. Restart licensing services:

   ```
   bash#> su onapp
   bash#> touch /onapp/interface/tmp/restart.txt
   bash#> exit
   ```

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

2.3 VMware Installation

Follow these guidelines to install and configure the VMware vCenter:

1. Install the VMWare vCenter server by following VMware documentation instructions.
2. Create an administrator account on the vCenter server or use the default “administrator” account and specify login credentials.
3. Create a vCenter virtual datacenter.
4. 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 on OnApp CP.
5. Open the following ports on the vCenter:
   - TCP/UDP 902
   - TCP 443
   - TCP 80
   - TCP/UDP 5988-5989

   For details, refer to the VMware documentation.
6. Install VMWare ESXi servers by following the VMware documentation instructions.
7. Add all ESXi servers to the cluster.
8. 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 CP.
9. Create a Distributed Switch.
10. Open the VNC ports on the ESXi host:
    - Enable SSH service on the ESXi host. To do so:
    - Enable SSH service on each ESXi host: **Configuration > Security Profile > SSH > Options > Start**
    - SSH into each HV in turn and run the following commands:

```
wget http://downloads.repo.onapp.com/onapp-vmware-
firewall-vnc-manage.sh
```

```
sh onapp-vmware-firewall-vnc-manage.sh
```

- Restart the ESXi server for changes to come into effect.

11. Enable NTP on all ESXi server. For correct time synchronization, use the same NTP server for vCenter and CP.
12. Allow virtual machines to start and stop automatically with the system.

2.3.1 Vyatta installation

PLEASE NOTE: You may experience compatibility issues when using the 6.6 version of Vyatta. We highly recommend using the 6.5 version.

To deploy the Vyatta as a virtual appliance running on the VMware cluster, you need to do the following::

1. On the distributed switch, create three virtual machine port groups:
   - Public - for communication between the Vyatta and external networks.
   - Management - for communication between the Vyatta and the OnApp Control Panel server
   - Appliance - for communication between the Vyatta and virtual machines.
     Set VLAN to 4095 when creating a normal port group, or to 1-4094 trunk range when creating a Distributed Switch port group.
2. Create a new Vyatta instances on the vCenter with three network interfaces and attach on to each of the created port groups.
3. Install Vyatta v.6.4 or later from http://vyatta.org/ and install it by creating and booting the Vyatta LiveCD.
4. Login as user vyatta with password vyatta and run the "install image" command.
5. Remove the LiveCD.
6. Reboot system.
7. Log in using the vyatta user credentials.
8. Run the following commands:

```
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
```

9. Configure the firewalls in Firewalls section.

2.3.2 CP configuration

Read the steps described in this section carefully to get a common notion of the VMware configuration within the OnApp cloud.

VMware implementation comprises several new features implemented in the OnApp cloud.

- **Customer VLANS** - VLANs are used to segment virtual network so that customer networks are isolated from one another as if they were on physically different segments. Each customer can have one VLAN with VMware VMs based on it. Configuring VLANs is essential, as it secures the network traffic and reduces the traffic overload.
- **Customer networks** - customer networks are used to isolate VMware virtual machines from other customers' VMs via VLAN. All the customer network traffic is handled by Vyatta to ensure high level of data protection. For detail, see Customer Networks section of the Admin guide.
- **IP Address Pools -** a range of IP addresses that you can associate with VLANs. You'll have to select an IP address pool during the customer network creation. See IP Address Pools section of the Admin guide for details.
- **Firewalls** (Vyatta firewall is used to manage VLANs and route VMware VM networking traffic in and out of OnApp. Because all customer VMs are running inside customer networks, firewalls are required as the VM gateways. See Firewalls section of the Admin guide for details. The default firewall rule should be set to INSIDE_OUT.

So, to configure the OnApp cloud with VMware, you have to configure Firewalls, create a VLAN to isolate your virtual machines, create an IP address pool and customer network. Aside from that, the rest of cloud configuration steps remains unchanged.

> You must enable the Contol Panel server network access to the vCenter and each cluster. The configuration instructions depend on your setup.

To configure VMware on CP:

1. Create new VMware hypervisor in the Hypervisors settings. See the Create VMware hypervisor section of the Admin guide for details how to do that.
2. Create new IP address pool.
3. Create a range of VLANs you want to use in the cloud.
4. Create new customer network selecting the IP address pool you have created at step 2.
5. Create new user.
6. Create new data store zone.
7. Create new VMware datastore and assign it to data store zone.
8. Assign customer network, network and datastore to the VMware hypervisor you have created at step 1.
9. Specify the vCenter cluster name in the System Configuration settings.
10. In the Default Settings configuration, define the service account name that will be automatically created on all virtual machines to be able to communicate with them. OnApp/settings/edit#defaults
11. Create a VMware VM templates by following the instructions in the Create Template for VMware Virtual Machine chapter.
12. Log in as this user you have created at step 4 and create a new VM in VMware.

2.4 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.

2.4.1 CloudBoot 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.

NOTE: Starting from the 3.0.7 version of the OnApp Cloud the default RAM value for Xen CloudBoot hypervisors (dom0 RAM) is set to 2 GB.

⚠ 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. Enable Storage in the control panel:

   Settings -->Configuration -->OnApp Storage

   Scroll down to the OnApp Storage section and check the "Enable OnApp Storage" box. Tick the "Use Local Read Path"checkbox to minimise the network throughput dependency for read heavy workloads.

3. 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**

4. Add IP address range for Hypervisors:

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

5. Power on servers and allow them to boot the default image
6. Add servers to the control panel by selecting MAC addresses and assigning IP address

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

   ⚠ 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 CloudBoot, see the CloudBoot Hypervisors section of the Admin guide.

If adding a KVM Hypervisor, add a FQDN for your HV to Custom Config to ensure Hot Migrate works correctly:

hostname HVx.mydomain.com

echo 'x.x.x.x HVx.mydomain.com HVx' >> /etc/hosts

To increase dom0 memory for all new Xen HVs - edit the dom0 value in /tftpboot/pxelinux.cfg/template-xen on the CP server

To increase dom0 memory for a single Xen HV - edit /tftpboot/pxelinux.cfg/xx-xx-xx-xx-xx-xx replacing the x's with your HV management NIC MAC address

7. Generate SSH keys:

If you are going to use CloudBoot 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.

a. SSH into your Control Panel server.

b. 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!

8. Configure backup server:

   As you are using CloudBoot, you can use either Cloud Boot backup servers or virtual backup servers instead of the advanced standalone backup server.

   > We strongly recommend you to deploy one or more backup servers for backups and VM provisioning when using a Cloud Boot functionality.

   To create a CloudBoot backup server:

   a. Update CloudBoot and CP server RPMs:

   ```
   yum update onapp-store-install
   yum update onapp-cp-install
   ```

   b. Configure CloudBoot settings:

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

   After that:

   a. Create new KVM hypervisor with an IP address from the dynamic range. See the CloudBoot Hypervisor Bootstrap Method chapter of the Admin guide for details. Wait till the hypervisor comes online.

   b. Go to your Control Panel's **Settings** menu, then press **Backup servers** icon.

   c. Click the **Create Backup Server** button.

   d. Fill in the form that appears:

      - *Label* - give your backup server a label.
      - *IP address* - enter the IP address of a hypervisor you have created at step 1.
      - *Backup IP address* - add a provisioning network IP address.
      - *Capacity* - set the backup server capacity (in GB).

   e. Tick the **Enabled** box to enable the backup server.

   f. 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 HV added in step 1 in the 'Backup IP address' field of each of your hypervisors.

> ⚠ PLEASE NOTE: You should configure some local or remote attached storage for persistent backups on the provisioning/backup server.

To create a virtual backup server:

Virtual backup servers serve for reducing IO load in Domain 0 of Xen hypervisors. Utilization of virtual servers helps to reduce load on hypervisor servers and improve their performance and may be used as an alternative to dedicated backup servers. The backup server can then be used to offload the backup activities from Dom0 and free up resources from the Hypervisor.

Virtual backup servers are included in the onappstore rpm and need to be configured via the CP terminal.

Create a virtual backup server via CLI:

a. Create a backup server

```
/onapp/backupServerAdmin
```

b. List Available Hvs + IP addresses:

```
backupServerAdmin list
```

c. Query Networks available to a given HV:

```
backupServerAdmin hvnetinfo
```

d. Create a config on a chosen HV:

```
backupServerAdmin create
```

e. Start the Backup server VM:

```
backupServerAdmin start
```

f. Go to the OnApp Control Panel and create new Xen CloudBoot hypervisor with the MAC IP address obtained at step 2.

g. Restart the backup server via CLI

```
backupServerAdmin stop
backupServerAdmin start
```

The backup server IP address will appear in the **Selecting a backup server** drop-down menu.

After that:

a. Go to your Control Panel's **Settings** menu, then press **Backup servers** icon.

b. Click the **Create Backup Server** button.

c. Fill in the form that appears:

- *Label* - give your backup server a label
- *IP address* - enter the IP address of a virtual backup server
- Skip the *Backup IP address* field, as it is not required for the virtual backup server
- *Capacity* - set the backup server capacity (in GB)

d. Tick the **Enabled** box to enable the backup server.

9. Download Templates:

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

```
bash# wget http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash# rpm -Uvh onapp-repo.noarch.rpm
bash# yum install onapp-bk-install
bash# sh /onapp/onapp-bk-install/onapp-bk-install.sh -t
```

CloudBoot hypervisors mount the following locations automatically at boot:

- */tftpboot/export/centos5/xen* to */.ro*
  The path may vary depending on the hypervisor template used.

- */onapp/templates* to */.templates*
  This path is symlinked to */onapp/templates.*
- */data* to */onapp/tools/recovery*
- */tftpboot/images/centos5/ramdisk-xen* to */cloudboot/centos5/ramdisk-xen*
  The path may vary depending on the hypervisor template.

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

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

- /tftpboot/export
- /onapp/templates
- /data
- /tftpboot/images

Hypervisors 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.

PLEASE NOTE: To utilize Cloud Boot, make sure it is enabled in your BIOS settings. See Configuring Cloud Boot settings in BIOS section for details.

2.4.2 Static Hypervisor Installation Method

PLEASE NOTE: Base CentOS templates must be installed on the local drive before hypervisor installation, depending which virtualization method you choose:

- Xen 3 hypervisors: CentOS 5.x x64
- Xen 4 hypervisors: CentOS 6.x x64
- KVM hypervisors: CentOS 5.x x64 or CentOS 6.x x64

1. Add the hypervisor to your cloud using the OnApp Control Panel:

   Settings --> Hypervisors --> Add New Hypervisor

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

2. Make sure your OS is up to date:

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

3. Enable IPv6:

> ⚠ Skip this step if you are using the CentOS 6 template.

> ⚠ 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 do not reboot now. We'll do that later.

4. Download the OnApp repository:

```
bash#> wget  http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash#> rpm  -ivh  onapp-repo.noarch.rpm
bash#> yum clean all
```

5. Install the OnApp hypervisor installer package:

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

6. 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.

7. 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
```

> ⚠ 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.

8. Configure the hypervisor 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>
```

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

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

9. 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
```

10. Reboot the hypervisor 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.

- 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!

- Update autorized_keys file by running the following script on the hypervisor:

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

2.5 Data Store Installation

⚠ PLEASE NOTE:

- To configure an integrated storage datastore, please consult the Admin guide.
- 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.

2.6 Backup Server Installation

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

1. Add a backup server to the web UI:
   - 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 -i http://rpm.repo.onapp.com/repo/onapp-repo-3.0.noarch.rpm
```

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.

6. onfigure 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] -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 the Backup Server 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 Backup 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.

<BK_HOST_IP> is the IP address of the Backup Server.

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

2.6.1 Cloud Boot Backup Server Installation

We strongly recommend you to deploy one or more backup servers for backups and VM provisioning when using a Cloud Boot functionality.

Follow the step-by-step instructions provided in this chapter to configure cloud boot backup servers in your cloud.

1. Create a new KVM hypervisor via Cloudboot with an IP address from the dynamic range. See the Create Hypervisor chapter of this guide for details.
2. Ensure to choose the 'Backup' option and don't format disks.

   PLEASE NOTE: You should configure some local or remote attached storage for persistent backups on the provisioning/backup server using Custom Config when adding the KVM Hypervisor if you plan to hold templates or backups on the server.

3. Go to your Control Panel's **Settings** menu, then press **Backup servers** icon.
4. Click the **Create Backup Server** button.
5. Fill in the form that appears:
   - *Label* - give your backup server a label
   - *IP address* - enter the IP address ofa hypervisor you have created at step 1.
   - *Backup IP address* - add a provisioning network IP address
   - *Capacity* - set the backup server capacity (in GB)
6. Tick the **Enabled** box to enable the backup server.
7. 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 HV added in step 1 in the 'Backup IP address' field of each of your hypervisors.

2.6.2 Virtual Backup Server Installation

Virtual backup servers are designed to reduce IO load in Domain 0 of Xen hypervisors. Utilization of virtual servers helps to reduce load on hypervisor servers and improve their performance and may be used as an alternative to dedicated backup servers. The backup servers can then be used to offload the backup activities from Dom0 and free up resources from the Hypervisor. Once configured via the CloudBoot interface, virtual backup servers are managed exactly the same as dedicated physical backup servers.

For clouds using the backup scheme without dedicated backup servers, virtual backups appliance should be used rather than the standard procedure.

PLEASE NOTE: You need to configure a backup target for storing backups before using a virtual backup server.

If you reboot a hypervisor that functions as a virtual backup server without a target specified, all backups will be lost!

How can you tell if a backup server is virtual or dedicated?

In the UI, there will be an additional HV with the MAC address beginning with "de:be", that is available when selecting "Add a new CloudBoot appliance".

From the CP server side, running "backupServerAdmin list" command will show the presence of a backup server as per step 4 below.

Virtual backup servers are included in the onappstore rpm and need to be configured manually via the CP terminal.

You can execute the following commands:

```
Command backupServerAdmin
```

Usage:

```
backupServerAdmin list
backupServerAdmin create <HV MAC Addr> <RAM> <vCPUs>
<Bridge1,Bridge2,...BridgeN>
backupServerAdmin delete <HV MAC Addr> <VMname>
backupServerAdmin start <HV MAC Addr> <VMname>
backupServerAdmin stop <HV MAC Addr> <VMname>
backupServerAdmin move <Src HV MAC Addr> <Dst HV MAC Addr>
<VMname>
backupServerAdmin hvnetinfo <HV MAC Addr>
```

Where:

- HV MAC Addr - MAC address of a Xen hypervisor that is used for the virtual backup server
- RAM - virtual backup server RAM
- vCPUs - virtual backup server CPUS
- Bridge1,Bridge2,...BridgeN - bridge identifiers configured on the hypervisors
- Src HV MAC Addr - the MAC address of the HV we will move the virtual backup server from during the migration
- Dst HV MAC Addr - target MAC IP address during the virtual backup server migration
- VMname - virtual backup server name that is generated automatically during the creation process

backupServerAdmin list report example:

```
Node <MAC_ADDRESS> (<IP_ADDR>)
Backup Server '<BS_NAME>':
Memory: 'BS_RAM'
vCPUs: 'BS_CPU'
Networks: ['bridge=BRIDGENAME,vifname=VIFNAME,mac=MAC_ADDR'*]
Running: TRUE/FALSE
```

To add the virtual backup server via CLI:

1. List available hypervisors and IP addresses:

   ```
   backupServerAdmin list
   ```

2. Query Networks available to a given hypervisor:

```
backupServerAdmin hvnetinfo
```

3. Create a config on a chosen hypervisor:

```
backupServerAdmin create
```

4. Find the ID of the backup server:

```
backupServerAdmin list
```

5. Start the Backup server VM:

```
backupServerAdmin start
```

6. Wait for two minutes. After that, go to the OnApp Control Panel and create new Xen Cloud Boot hypervisor with the MAC of a virtual backup server address obtained with the 'backupServerAdmin list' command.
7. Restart the backup server from the CLI:

```
backupServerAdmin stop
backupServerAdmin start
```

To create new virtual backup server in the OnApp Control Panel:

1. Go to your Control Panel's **Settings** menu, then press **Backup servers** icon.
2. Click the **Create Backup Server** button.
3. Fill in the form that appears:
   - *Label* - give your backup server a label
   - *IP address* - enter the IP address of a virtual backup server
   - Skip the *Backup IP address* field, as it is not required for the virtual backup server
   - *Capacity* - set the backup server capacity (in GB)
4. Tick the **Enabled** box to enable the backup server.

2.7 Download and Configure Templates on the Template Server

This configuration should be applied in case you are not using the new backup server scheme.

1. Go to your Control Panel's **Settings** menu, click the **Configuration** icon, then choose **Backup/Templates**. On the screen that follows:
    a. Enable the **Use SSH File Transfe**r option.
    b. The Server IP should be the management IP address of your templates/backups server.
    c. Set the user to root and leave the other options default.
2. Login to the templates/backups server as root, and run:

```
bash# wget http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash# rpm -Uvh onapp-repo.noarch.rpm
bash# yum install onapp-bk-install
bash# sh /onapp/onapp-bk-install/onapp-bk-install.sh -t
```


PLEASE NOTE: Before creating a virtual machine, you must create at least one template group in the template store with the required templates. See Template Store section of the Admin guide for details.

2.8 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.

2.8.1 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**:**
    a. Go to your Control Panel's **Settings** menu and click the **Data store zones** icon.
    b. Click the **Add New Data store zone** button.
    c. 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:

    a. Go to your Control Panel's **Integrated Storage** menu.
    b. 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.)
    c. 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.
3. 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.

2.8.2 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:
    a. Go to your Control Panel's **Settings** menu and click the **Data store zones** icon.
    b. Click the **Add New Data store zone** button.
    c. 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

    To create an LVM data store:

    a. Go to your Control Panel **Settings** menu.
    b. Click the **Data Stores** icon.
    c. Click the **Create Data Store** link at the bottom of the screen.
    d. Follow the steps in the creation wizard:

    Step 1 of 2

    - 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.**

    Step 2

    - 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 */vm* data store type.

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

To create a SolidFire data store:

a. Go to your Control Panel **Settings** menu.

b. Click the **Data Stores** icon.

c. Click the **Create Data Store** link at the bottom of the screen.

d. 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 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).

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 Solid Fire Account settings:

  Username - specify SolidFire account username

  Initiator secret - specify iSCSI initiator secret (optional)

  Target secret - specify iSCSI initiator secret (optional)

  Initator 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.

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

To create a VMware data store:

a. Go to your Control Panel **Settings** menu.

b. Click the **Data Stores** icon.

c. Click the **Create Data Store** link at the bottom of the screen.

d. Follow the steps in the creation wizard:

Step 1 of 2

- Enter a label of the vCenter data store. The label of VMware data store must match the vCenter data store label!
- Leave the IP address field empty.
- 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 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 *vmware* data store type.

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

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

- 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
```

- 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
```

2.8.3 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.
    - Press "+" button or click the **Add a New Hypervisor** button underneath the list of Hypervisors on the screen.
    - On the screen that appears:
        - Enter a hypervisor label.
        - Add an IP address.
        - Add a backup IP address.
        - Choose a hypervisor type (Xen, KVM or VMware).
        - Move the slide to the right to enable a hypervisor. Hypervisors that are not enabled cannot be used to host VMs.
        - Move the slider to the right to collect statistics for this hypervisor.
        - Move the slider to the right to disable failover on this hypervisor (failover is automatic VM migration to another hypervisor if this one goes down).
    - Click the **Save** button to finish. The hypervisor will be added to the system. 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.

2.8.4 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.

2.8.5 Join Networks and Datastores to Hypervisors

1. Join datastores to hypervisors:
    a. Go to your Control Panel's **Settings** menu and click the **Hypervisors** icon.
    b. Click the label of the hypervisor you want to manage data stores for.
    c. On the screen that appears, click the **Manage Data Stores** link in the **Actions** section.
    d. 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:
    a. Go to your Control Panel's **Settings** menu and click the **Hypervisors** icon.
    b. Click the label of the hypervisor you want to manage networks for.
    c. On the screen that appears, click the **Manage Networks** link in the **Actions** section.
    d. 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.

2.8.6 Download and Configure Templates

1. Go to your Control Panel's **Settings** menu, click the **Configuration** icon, then choose **Backups/Templates**. On the screen that follows:
   a. Enable the **Use SSH File Transfer** option.
   b. The Server IP should be the management IP address of your Control Panel server.
   c. Set the user to root and leave the other options default.
2. Login to the OnApp Control Panel server as root, and run:

```
bash# wget http://rpm.repo.onapp.com/repo/centos/5/onapp-repo.
noarch.rpm
bash# rpm -Uvh onapp-repo.noarch.rpm
bash# yum install onapp-bk-install
bash# sh /onapp/onapp-bk-install/onapp-bk-install.sh -t
```

⚠ PLEASE NOTE: Before creating a virtual machine, you must create at least one template group in the template store with the required templates. See Template Store section of the Admin guide for details.

2.9 Support

OnApp customers with a full (paid) license can contact support at any time:

support@onapp.com

+1 (888) 876 8666

http://onapp.com/support

Response times

OnApp offers a 15 minute SLA for critical issues.

2.10 Appendix:document revisions

v1.0, 15th May 2013

- First release

3 OnApp Cloud v3.0.8 - 3.0.10 to v3.0.11 Upgrade Guide

This guide explains how to upgrade to OnApp Cloud v3.0.11 from the 3.0.8 - 3.0.10 versions.

3.0.11 version of the OnApp Cloud comprises only Integrated Storage improvements and fixes. Do not upgrade to the 3.0.11 version if you are not running the Integrated Storage.

With this version, only CloudBoot packages are updated. The Control Panel version stays 3.0.8.

PLEASE READ THE INTRODUCTION AND IMPORTANT NOTES CHAPTER BEFORE YOU BEGIN!

3.1 ntroduction and important notes

This guide explains the 3.0.8-3.0.10 versions of the OnApp Cloud to the new version, OnApp Cloud v3.0.11.

Before starting an upgrade process, stop the OnApp daemon on the Control Panel!

```
service onapp stop
```

After running the service onapp stop command, make sure that there are no OnApp daemons running:

```
ps ax |grep onapp_daemon |grep -v grep
```

In a case there are running transactions in your cloud, wait until all transactions are complete.

> ⚠ Please close all Rails console connections during the upgrade! Make sure no control panel files are open for editing under the root user account.

To upgrade your cloud, please follow the upgrade process for your CloudBoot hypervisors.

> ⚠ Important - if you are using a 3rd party billing platform please ensure that this is compatible with OnApp 3.0.11 before proceeding with the upgrade! The latest WHMCS modules can be found here.

3.1.1 Upgrade to the v3.0.11 from older versions

To upgrade from older versions of OnApp, you have to upgrade to v3.0.8 version first, and then perform and upgrade to the v.3.0.11 by following the instructions provided in this guide.

3.1.2 Getting support for your upgrade

You can use the information in this document to perform your own upgrade to OnApp Cloud v3.0.11. However, if you have a full (paid) OnApp Cloud license, you are entitled to free upgrade support from the OnApp Support team.

If you would prefer to have the Support team perform the upgrade for you, just raise a ticket in the normal way. Please be aware, however, that there may be a queue!

For help with your upgrade, visit the OnApp community forum: http://forum.onapp.com.

3.2 Upgrade CloudBoot hypervsors

Use one of the following CloudBoot hypervisor upgrade paths when upgrading to the 3.0.11 version of the OnApp Cloud.

- Live upgrade CloudBoot hypervisors
- Upgrade CloudBoot hypervisors by rebooting them

3.2.1 Lve upgrade CloudBoot hypervsors.

> ⊘ Do not make any changes to the cloud during the upgrade!

CloudBoot hypervisors must meet the following requirements for the live migration:

- Hypervisor must be installed and running
- HVs must be running a 2.6.18-308 (Xen) or 2.6.32-279 (kvm) kernel
- Minimum 512 NBD connections must be probed and loaded on the hypervisor

NOTE: Windows virtual machines must be powered off before the upgrade.

1. Make sure no disks are out of sync. To do so, log in into a hypervisor and run the following command:

```
bash#> cd /usr/pythoncontroller/
bash#>  ./getdegradeddisks
```

Repair all the degraded disks before proceeding to the upgrade process.

NOTE: VDisks can be degraded because of the incorrect blkback location. Check that correct version of blkback is running by looking for blkback:Init messages on the hypervisor in dmesg.

2. Stop the OnApp service:

```
service onapp stop
```

3. Download and install the latest OnApp YUM repository file:

```
bash#> rpm -Uvh http://rpm.repo.onapp.com/repo/centos/5/onapp-repo.noarch.rpm
bash#> yum clean all
```

This command may return an error message during the upgrade if the package is already installed (this occurs on upgrades between 3.0.8 and 3.0.9).

4. Install Cloud Boot dependencies:

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

5. Run the following command from the Control Panel server terminal to display the list of hypervisors with their IP addresses:

```
liveUpdate listHVs
```

This command will also show whether hypervisors are eligible for live upgrade.

> If the command liveUpdate is not available then it may be located in the sbin directory instead (cd /usr/local/sbin).

6. Add the following line to /etc/lvm/lvm.conf on each of the Hypervisor Servers if it is not already present:

```
filter = [ "r|/dev/nbd|","r|/dev/mapper|","r|/dev/dm-|" ]
```

7. Run lvmdiskscan from each of the Hypervisors to enable those changes.
8. Run the following commands from the Control Panel server terminal for each hypervisor:

```
/usr/local/sbin
liveUpdate updateToolstack <HV IP Addr>
```

The synchronization will take approximately three minutes for each HV.

9. Once each hypervisor's Toolstack is upgraded, run:

```
cd /usr/pythoncontroller/
wget http://downloads.repo.onapp.com/diskutil.pyc
wget http://downloads.repo.onapp.com/controllerRestart.pyc
```

Make sure that checksum is correct for diskutil.pyc and controllerRestart.pyc:

```
md5sum diskutil.pyc
96cfaebae1ade8df33a547f6d93ec01e  diskutil.pyc
md5sum controllerRestart.pyc
da845fe62577b5ec99238410e59b4e1f
```

If modifying the memory for the storage controllers please see the section at the end of the document for further info.

10. Run the following command for every hypervisor in turn

```
/usr/local/sbin
liveUpdate restartControllers <HV IP Addr>
```

⚠ Before restarting the controller, check that groupmon / redis-server etc are working by running an onappstore list command.

At this stage, an error message about degraded disks may be displayed. VDisks should still be unpaused, but may be degraded. Check the number of degraded disks after restarting the controller.

11. Make sure that the package versions are upgraded by running the following command on each HV:

```
cat /onappstore/package-version.txt | grep Source
```

12. Check that the storage controllers have been started cleanly by running the following command on each HV:

```
ifconfig onappstoresan
log into storagenodes
uptime
```

13. After that, upgrade the hypervisor drivers. You can perform this step after the upgrade process is completed.

    (it is also required if the kernel is updated between onapp versions):

```
liveUpdate updateDrivers <HV IP Addr>
```

14. Restart the OnApp service.

⚠

Please wait at least 2 minutes after the last HV has come online such that the storage system can stabilise before restarting the service.

```
service onapp restart
```

PLEASE NOTE:

- the default RAM value for Xen CloudBoot hypervisors (dom0 RAM) is set to 2 GB.
- the number of NBD connections is increased to 512 by default. Please remove the following lines from your custom config file:

```
modprobe -r nbd
modprobe nbd nbds_max=256
```

Increasing memory during the upgrade

This should be performed before the controller restart such that the changed memory values are detected.

To increase memory during the upgrade, change the /onappstore/onappstore.conf ramperctl on hypervisor to the required value. After that, change the memory for all storagenodes in /onappstore/VMConfigs/NODEX-STORAGENODEY to the same value.

The CP server can sometimes show that an HV is offline during the upgrade process.

Please contact support if hypervisors are displayed as offline or report I/O errors after 5 minutes and are not accessible via ping or ssh.

3.2.2 Upgrade CloudBoot hypervisors by rebootng them

Before upgrading the CloudBoot hypervisors, you need to download the OnApp YUM repository and install the CloudBoot dependencies. After that, you need to simply reboot the hypervisors to upgrade them. You do not need to perform any hypervisor upgrade operations using console.

1. Download and install the latest OnApp YUM repository file:

```
bash#> rpm -Uvh http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash#> yum clean all
```

2. Install Cloud Boot dependencies:

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

Once you have upgraded the CloudBoot dependencies, you have to reboot your Cloud Boot hypervisors to update the Cloud Boot RPM.

To do so:

1. Migrate all the virtual machines from the CloudBoot hypervisor to another hypervisor. Follow the instructions described in the Migrate VM section of the Admin guide to migrate virtual machines.
2. After that, Go to your Control Panel **Settings** menu.
3. Click the **Appliances** icon.
4. Click the label of the CloudBoot hypervisor you have migrated all VMs from.
5. On the hypervisor details screen, click the **Actions** button, then click **Reboot Hypervisor.**

> PLEASE NOTE: Rebooting a hypervisor assigned to a data store with a single replica (single-replica HV) or degraded virtual disks may result in data loss.

6. A new screen will open asking for confirmation (via two checkboxes) before reboot:
   - **Stop all virtual machines that cannot be migrated to another hypervisor?** Check this box if you want VMs that cannot be migrated to be powered off. When a hypervisor is scheduled for a reboot, OnApp will first attempt to hot migrate all

VMs it hosts. If hot migration is not possible for a VM, OnApp will attempt to cold migrate that VM. With this box checked, if cold migration fails, the VM will be stopped so the reboot may proceed. If you don't check this box, OnApp will attempt to hot and then cold migrate all VMs hosted by the hypervisor being rebooted – but will stop the migration process if any VM cannot be migrated.

- **Are you sure you want to reboot this hypervisor?** A simple confirmation to confirm that you want the hypervisor to reboot.

7. When you're certain you want to proceed with the reboot, click the **Reboot** button.
8. On the hypervisor is booted, repair the disk that were degraded during the reboot.
9. Repeat these steps for all CloudBoot hypervisors in your cloud.

Starting from the 3.0.7 version of the OnApp Cloud the default RAM value for Xen CloudBoot hypervisors (dom0 RAM) is set to 2 GB.

PLEASE NOTE: In the 3.0.7 version of the OnApp Cloud the number of NBD connections is increased to 512 by default. Please remove the following lines from your custom config file:

```
modprobe -r nbd
modprobe nbd nbds_max=256
```

4 Hypervisor Kernel Upgrade for the 3.0.8 Version

Follow the instructions below to upgrade your Xen 4 static hypervisors to the latest kernel version, kernel-3.4.53-8.el6.centos.alt.x86_64. You only need to perform an upgrade if you experience stability issues with CentOS6/Xen.

This upgrade only applies to Xen 4 hypervisors. New 3.0.8 OnApp Cloud installation package already comes with the updated kernel.

1. Make sure your hypervisor is visible and online in the Control Panel.
2. Download the OnApp repository:

```
bash#> wget http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash#> rpm -Uvh onapp-repo.noarch.rpm
bash#> yum clean all
```

3. Update YUM repository configuration file for CentOS 6.x with Xen 4.x packages: Skip the step if the hypervisor isn't CentOS 6.x with Xen4:

```
# yum --disablerepo=Xen4CentOS update centos-xen-repo
```

4. Install the OnApp hypervisor installer package:

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

5. 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.

6. Run the OnApp hypervisor installer script:

   For Xen hypervisors:

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

> 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 (Not supported in 3.0)
-t	Download recovery templates and ISO(s) used to provision FreeBSD guests.
-h	Print this info.

> Please continue to complete the remaining steps as they are important in activating your cloud.

7. Configure the hypervisor for your cloud. This step is also required to enable 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>
```


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

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 recovery templates.

8. Reboot the hypervisor.

5 OnApp Cloud v3.0.x to v3.0.8 (Onapp-Store 3.0.11) Upgrade Guide

This guide explains how to upgrade to OnApp Cloud v3.0.8 (Integrated Storage v3.0.11) from earlier v3.0.x versions.

PLEASE READ THE INTRODUCTION AND IMPORTANT NOTES CHAPTER BEFORE YOU BEGIN!

VMware clusters do not require any update procedures during the 3.0.x to 3.0.8 upgrade.

5.1 Introduction and Important Notes

This guide explains how to upgrade different editions of the OnApp Cloud 3.0 version (currently, versions 3.0 - 3.0.7) to the new version, OnApp Cloud v3.0.8 (Integrated Storage v3.0.11).

Before starting an upgrade process, stop the OnApp daemon on the Control Panel!

```
service onapp stop
```

After running the service onapp stop command, make sure that there are no OnApp daemons running:

```
ps ax |grep onapp_daemon |grep -v grep
```

In a case there are running transactions in your cloud, wait until all transactions are complete.

⚠ Please close all Rails console connections during the upgrade! Make sure no control panel files are open for editing under the root user account.

To upgrade your v3.0 cloud, please follow the upgrade process for your static hypervisors, then backup servers, then CloudBoot hypervisors and the control panel server. You must follow the upgrade instructions in the correct order!

⚠ Important - if you are using a 3rd party billing platform please ensure that this is compatible with OnApp 3.0.8 before proceeding with the upgrade! The latest WHMCS modules can be found here.

5.1.1 Upgrade to the v3.0.8 from older versions

To upgrade from older versions of OnApp, you have to upgrade to v2.3.3 version first, and then perform an upgrade to the v.3.0.8 by following the instructions provided in this guide.

For upgrade instructions for earlier versions, refer to the upgrade documentation corresponding to the OnApp version you are upgrading from:

- If you are running a 2.3.1 version of the OnApp cloud, refer to the OnApp Cloud v2.3.1 to v2.3.3 Upgrade Guide.
- If you are running a 2.3.2 version of the OnApp cloud, refer to OnApp Cloud v2.3.2 to v2.3.3 Upgrade Guide.

5.1.2 Getting support for your upgrade

You can use the information in this document to perform your own upgrade to OnApp Cloud v3.0.8. However, if you have a full (paid) OnApp Cloud license, you are entitled to free upgrade support from the OnApp Support team.

If you would prefer to have the Support team perform the upgrade for you, just raise a ticket in the normal way. Please be aware, however, that there may be a queue!

For help with your upgrade, visit the OnApp community forum: http://forum.onapp.com.

5.2 Upgrade Static Hypervisors

Follow the instructions below to upgrade your static hypervisors.

1. Make sure your hypervisor is visible and online in the Control Panel.
2. Download the OnApp repository:

```
bash#> rpm -Uvh http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash#> yum clean all
```

3. Update YUM repository configuration file for CentOS 6.x with Xen 4.x packages: Skip the step if the hypervisor isn't CentOS 6.x with Xen4:

```
# yum --disablerepo=Xen4CentOS update centos-xen-repo
```

4. Install the OnApp hypervisor installer package:

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

5. 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.

6. 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
```

   ⚠

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 (Not supported in 3.0)
-t	Download recovery templates and ISO(s) used to provision FreeBSD guests.
-h	Print this info.

Please continue to complete the remaining steps as they are important in activating your cloud

7. Configure the hypervisor for your cloud. This step is also required to enable 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>
```

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

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 recovery templates.

OnApp 3.0 introduces CentOS 6 with Xen4 support and improved KVM integration. Unfortunately, there is no inline upgrade between CentOS 5 and CentOS 6, so if you wish to run CentOS 6 hypervisors in your cloud, you will have to move the virtual machines from the hypervisor to other HVs before reinstalling it, and then reinstalling each hypervisor in turn.

We suggest you upgrade the whole cloud, including current hypervisors to OnApp 3.0 before considering reinstalling hypervisors. There is no need for this to be completed within a single maintenance window.

Depending on your setup, you can continue using your current hypervisors running CentOS 5 for now and add new HVs to your cloud with CentOS 6. You may also wish to consider using the new Cloudboot functionality to avoid the hypervisors being manually reinstalled.

5.3 Upgrade Static Backup Servers

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

1. Download the OnApp repository:

```
bash# wget http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash# rpm -Uvh onapp-repo.noarch.rpm
bash# yum clean all
```

2. Install the OnApp Backup Server installer package:

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

3. 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
```

4. 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.

⚠ FQDN or IP Address for Control Panel and Backup Servers are required for the new statistics receiver to work.

ⓘ You can configure Cloud Boot backup servers and virtual dedicated backup servers to be used with the Integrated Storage functionality. The backup scheme remains unchanged.

5.4 Upgrade CloudBoot Hypervisors.

Before upgrading the CloudBoot hypervisors, you need to download the OnApp YUM repository and install the CloudBoot dependencies. After that, you need to simply reboot the hypervisors to upgrade them. You do not need to perform any hypervisor upgrade operations using console.

1. Download and install the latest OnApp YUM repository file:

```
bash#> rpm -Uvh http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash#> yum clean all
```

2. Install Cloud Boot dependencies:

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

Once you have upgraded the CloudBoot dependencies to the 3.0.8 version, you have to reboot your Cloud Boot hypervisors to update the Cloud Boot RPM.

To do so:

1. Migrate all the virtual machines from the CloudBoot hypervisor to another hypervisor. Follow the instructions described in the Migrate VM section of the Admin guide to migrate virtual machines.
2. After that, Go to your Control Panel **Settings** menu.
3. Click the **Appliances** icon.
4. Click the label of the CloudBoot hypervisor you have migrated all VMs from.
5. On the hypervisor details screen, click the **Actions** button, then click **Reboot Hypervisor.**

> PLEASE NOTE: Rebooting a hypervisor assigned to a data store with a single replica (single-replica HV) or degraded virtual disks may result in data loss.

6. A new screen will open asking for confirmation (via two checkboxes) before reboot:
   - **Stop all virtual machines that cannot be migrated to another hypervisor?** Check this box if you want VMs that cannot be migrated to be powered off. When a hypervisor is scheduled for a reboot, OnApp will first attempt to hot migrate all VMs it hosts. If hot migration is not possible for a VM, OnApp will attempt to cold migrate that VM. With this box checked, if cold migration fails, the VM will be stopped so the reboot may proceed. If you don't check this box, OnApp will attempt to hot and then cold migrate all VMs hosted by the hypervisor being rebooted – but will stop the migration process if any VM cannot be migrated.
   - **Are you sure you want to reboot this hypervisor?** A simple confirmation to confirm that you want the hypervisor to reboot.
7. When you're certain you want to proceed with the reboot, click the **Reboot** button.
8. On the hypervisor is booted, repair the disk that were degraded during the reboot.

9. Repeat these steps for all CloudBoot hypervisors in your cloud.
10. Once all CloudBoot HVs are rebooted, proceed to the Control Panel server upgrade.

> NOTE: In the 3.0.7 version of the OnApp Cloud the default RAM value for Xen CloudBoot hypervisors (dom0 RAM) is set to 2 GB.

5.5 Upgrade Control Panel Server(s) .

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

> All installer critical errors are in /var/log/messages

1. Download and install the latest OnApp YUM repository file:

```
bash#> rpm -Uvh http://rpm.repo.onapp.com/repo/onapp-repo-3.0.noarch.rpm
bash#> yum clean all
```

2. Install OnApp Control Panel installer package:

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

3. 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.

4. Run Control Panel installer:

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

5. In the OnApp UI navigate to **Settings > Configuration** and click **Save** to complete the process.

6 OnApp Cloud v2.3.3 to v3.0.8 (OnApp-Store 3.0.11) Upgrade Guide

This guide explains how to upgrade OnApp Cloud v2.3.3. to the new version, OnApp Cloud v3.0.8 (3.0.11 version of the Integrated Storage). Please read the introduction and important notes chapter before you begin!

CDN billing statistics are not displayed after the upgrade to the 3.0 version

Currently the CDN billing statistics are not displayed in UI after the upgrade to the 3.0 version. This happens because the CDN billing statistics gathering has changed: in the previous versions data were collected per edge group, and starting from 3.0 version they are collected per CDN resource. The CDN billing statistics data are present in the database.

We strongly recommend that you test all your custom scripts before upgrading your production environment to the 3.0 version of OnApp Cloud.

The version of the Ruby on Rails used in OnApp Cloud has changed from v3.0.7 to v3.0.20 for security reasons. As a consequence, the upgrade may affect the operation of scripts that are currently in use. Unfortunately, OnApp will not be able to troubleshoot or fix these issues.

PLEASE NOTE: Make sure that See list of all template groups (image_template_groups.list) and See details of any template group (image_template_groups.read) permissions are enabled for the default User role after the upgrade, otherwise users will not be able to create virtual machines. To assign these permissions to any role which has permission to create a virtual machine, you can run:

```
cd /onapp/interface
RAILS_ENV=production rake permissions:enable_image_template_groups
```

6.1 Introduction and Important Notes.

This guide explains how to upgrade OnApp Cloud v2.3.3 version of the OnApp cloud to the new version, OnApp Cloud v3.0.8 (OnApp Storage 3.0.11).

Before starting an upgrade process, stop the OnApp daemon on the Control Panel!

```
service onapp stop
```

After running the service onapp stop command, make sure that there are no running OnApp daemons:

```
ps ax |grep onapp_daemon |grep -v grep
```

In a case there are running transactions in your cloud, wait until all transactions are complete.

Please close all Rails console connections during the upgrade! Make sure no control panel files are open for editing under the root user account.

To upgrade your v2.3.3 cloud, please follow the upgrade process for your hypervisors, then backup servers, then your control panel server. You must follow the upgrade instructions in the correct order!

Make sure that the 161 and 162 ports are open on the Control Panel server, hypervisors and backup servers before starting an upgrade, as this is obligatory for SNMP utilization!

Important - if you are using a 3rd party billing platform please ensure that this is compatible with OnApp 3.0.9 before proceeding with the upgrade! The latest WHMCS modules can be found here

6.1.1 Upgrade to the v3.0.8 from older versions

To upgrade from older versions of OnApp, you have to upgrade to v2.3.3 version first, and then perform and upgrade to the v.3.0.8 by following the instructions provided in this guide.

For upgrade instructions for earlier versions, refer to the upgrade documentation corrsponding to the OnApp version you are upgrading from:

- If you are running a 2.3.1 version of the OnApp cloud, refer to the OnApp Cloud v2.3.1 to v2.3.3 Upgrade Guide.
- If you are running a 2.3.2 version of the OnApp cloud, refer to OnApp Cloud v2.3.2 to v2.3.3 Upgrade Guide.

6.1.2 Getting support for your upgrade

You can use the information in this document to perform your own upgrade to OnApp Cloud v3.0.8. However, if you have a full (paid) OnApp Cloud license, you are entitled to free upgrade support from the OnApp Support team.

If you would prefer to have the Support team perform the upgrade for you, just raise a ticket in the normal way. Please be aware, however, that there may be a queue!

For help with your upgrade, visit the OnApp community forum: http://forum.onapp.com.

6.2 Upgrade Static Hypervisors.

Follow the instructions below to upgrade your static hypervisors.

1. Make sure your hypervisor is visible and online in the Control Panel.
2. Download the OnApp repository:

```
bash#> wget http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash#> rpm -Uvh onapp-repo.noarch.rpm
bash#> yum clean all
```

3. Update YUM repository configuration file for CentOS 6.x with Xen 4.x packages: Skip the step if the hypervisor isn't CentOS 6.x with Xen4:

```
# yum --disablerepo=Xen4CentOS update centos-xen-repo
```

4. Install the OnApp hypervisor installer package:

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

5. 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.

6. 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
   ```

   ⚠ 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 (Not supported in 3.0)
-t	Download recovery templates and ISO(s) used to provision FreeBSD guests.
-h	Print this info.

Please continue to complete the remaining steps as they are important in activating your cloud.

7. Configure the hypervisor for your cloud. This step is also required to enable 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>
```

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

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 recovery templates.

OnApp v3.0 provides CentOS 6 with Xen4 support and improved KVM integration. Unfortunately, there is no inline upgrade between CentOS 5 and CentOS 6, so if you wish to run CentOS 6 hypervisors in your cloud, you will have to move the virtual machines from the hypervisor to other HVs before reinstalling it, and then reinstalling each hypervisor in turn.

We suggest you upgrade the whole cloud, including current hypervisors to OnApp 3.0 before considering reinstalling hypervisors. There is no need for this to be completed within a single maintenance window.

Depending on your setup, you can continue using your current hypervisors running CentOS 5 for now and add new HVs to your cloud with CentOS 6. You may also wish to consider using the new Cloudboot functionality to avoid the hypervisors being manually reinstalled.

6.3 Upgrade Backup Servers.

1. Download the OnApp repository:

```
bash# wget http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash# rpm -Uvh onapp-repo.noarch.rpm
bash# yum clean all
```

2. Install the OnApp Backup Server installer package:

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

3. 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
```

4. 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.

5. onfigure 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] -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 the Backup Server 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 Backup 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.

⚠ <BK_HOST_IP> is the IP address of the Backup Server.

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

You can configure Cloud Boot backup servers and virtual dedicated backup servers to be used with the Integrated Storage functionality. The backup scheme remains unchanged.

6.4 Upgrade Control Panel Server(s).

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

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

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. Make sure your OS is up to date:

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

2. Download OnApp YUM repository file:

```
bash#> wget http://rpm.repo.onapp.com/repo/onapp-repo-3.0.
noarch.rpm
bash#> rpm -Uvh onapp-repo.noarch.rpm
bash#> yum clean all
```

3. Install OnApp Control Panel installer package:

```
bash#> yum update 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. Restart licensing services:

   ```
   bash#> su onapp
   bash#> touch /onapp/interface/tmp/restart.txt
   bash#> exit
   ```

7. Install CloudBoot dependencies:

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

8. In the OnApp UI navigate to **Settings > Configuration** and click **Save** to complete the process.

6.5 Configure the Template Store.

> ⚠ PLEASE NOTE: Before creating a virtual machine, you must to create at least one template group in the template store with the required templates. See Template Store section of the Admin guide for details.

To add a template group:

1. Go to your Control Panel's **Template Store** menu.
2. On the page that follows, click the **"+"** button next to the required template group's label, then select **Add Child.**
3. Give a name to your group.
4. Specify the Windows Licensing type: MAK, KMS, or User license.
5. For KMS licensing, set the following parameters:
    - *Server label* – the name of the KMS server
    - *KMS server host* – the hostname of the licensing server
    - *KMS server port* – the port used to connect to the licensing server
6. Click **Save.**
7. On the page that appears, you'll be prompted to assign a template to a group.

To assign a template to a template group:

1. Go to your Control Panel's **Template Store** menu.
2. Click the **"+"** button next to the required child group's label, then select **Add Template.**
3. Choose the template from the drop-down box at the **Add a template** section.
4. Click **Add a template to a group** button to confirm.

6.6 Upgrade CDN Edge Groups

Since the logic has changed form 2.3.3 to 3.0, an additional step is required at the end of the upgrade process aimed at synchronizing CDN edge groups. Follow the instructions provided in this chapter to synchronize your CDN groups.

After your license becomes valid, run the following script:

```
bash#> cd /onapp/interface
bash#> RAILS_ENV=production rake cdn:upgrade
```

Look through the output to see what CDN groups were not upgraded because of duplicated locations, and notify their owners. See the output example given below for more details.

The output provides the following details:

- CDN Hostname
- CDN Resouce id
- CDN Publisher's name
- CDN Publisher's email address

You can also find these details in /onapp/interface/log/production_aflexi.log and production_log

Output example

```
[root@vl interface]# RAILS_ENV=production rake cdn:upgrade
[2012-12-24 19:18:57 +0100] ==  EdgeGroupRefactoring: migrating
==========================================
[2012-12-24 19:18:57 +0100] -- add_column(:edge_groups, :
aflexi_id, :integer)
[2012-12-24 19:18:57 +0100]    -> 0.0066s
[2012-12-24 19:18:57 +0100] -- add_index(:edge_groups, :
aflexi_id, {:unique=>true})
[2012-12-24 19:18:57 +0100]    -> 0.0059s
[2012-12-24 19:18:59 +0100] Created Aflexi::EdgeGroup 862784544,
bind to EdgeGroup 1
[2012-12-24 19:19:00 +0100] Created Aflexi::EdgeGroup 364166690,
bind to EdgeGroup 8
[2012-12-24 19:19:00 +0100] Created Aflexi::EdgeGroup 601047047,
bind to EdgeGroup 9
[2012-12-24 19:19:01 +0100] Created Aflexi::EdgeGroup 5247118,
bind to EdgeGroup 10
[2012-12-24 19:19:02 +0100] Created Aflexi::EdgeGroup 457930977,
bind to EdgeGroup 11
[2012-12-24 19:19:03 +0100] Created Aflexi::EdgeGroup 108576807,
bind to EdgeGroup 12
[2012-12-24 19:19:05 +0100] Created Aflexi::EdgeGroup 369131778,
bind to EdgeGroup 13
[2012-12-24 19:19:06 +0100] Created Aflexi::EdgeGroup 923762372,
bind to EdgeGroup 14
```

```
[2012-12-24 19:19:08 +0100] Created Aflexi::EdgeGroup 710939515,
bind to EdgeGroup 15
[2012-12-24 19:19:09 +0100] Created Aflexi::EdgeGroup 836280135,
bind to EdgeGroup 16
[2012-12-24 19:19:11 +0100] Created Aflexi::EdgeGroup 357564912,
bind to EdgeGroup 17
[2012-12-24 19:19:12 +0100] Created Aflexi::EdgeGroup 812074754,
bind to EdgeGroup 18
[2012-12-24 19:19:13 +0100] Linked edge groups [862784544, 3641666
90] to cdn resource (id = 5, aflexi_id = 838680885)
[2012-12-24 19:19:14 +0100] Linked edge groups [364166690, 6010470
47] to cdn resource (id = 6, aflexi_id = 694056113)
[2012-12-24 19:19:16 +0100] Error. CDN Hostname: oh-dup1.tst, CDN
Resouce id: 7, CDN Publisher: OHdupl OHdupl, ohdupl@cv.co.
Resource must not have overlapping locations
[2012-12-24 19:19:18 +0100] Error. CDN Hostname: oh-dup2.tst, CDN
Resouce id: 8, CDN Publisher: OHdupl OHdupl, ohdupl@cv.co.
Resource must not have overlapping locations
[2012-12-24 19:19:19 +0100] Linked edge groups [369131778] to cdn
resource (id = 9, aflexi_id = 312009919)
[2012-12-24 19:19:21 +0100] Linked edge groups [357564912] to cdn
resource (id = 10, aflexi_id = 122435022)
[2012-12-24 19:19:23 +0100] Linked edge groups [812074754] to cdn
resource (id = 11, aflexi_id = 983084551)
[2012-12-24 19:19:24 +0100] Error. CDN Hostname: ohdupl.new, CDN
Resouce id: 13, CDN Publisher: OHdupl OHdupl, ohdupl@cv.co.
Resource must not have overlapping locations
[2012-12-24 19:19:24 +0100] ==  EdgeGroupRefactoring: migrated (26
.9292s) ================================
[2012-12-24 19:19:24 +0100] Finish upgrade
```

For your convenience, you can find the required strings in the output of this command in a log, for example:

cat /onapp/interface/production.log | grep '[CDN]'

⚠ Note: If the CDN edge group with overlapped locations has been initially tied to the CDN resource (this scenario is possible in 2.3.3), it will remain in the old scheme after the upgrade (tied directly to locations, but not to the edge group), and will function normally. To switch these CDN resources to the new scheme, you have to tie them to edge groups, edit and save the changes. Required validation during the editing process will make you select non-overlapped edge groups for this CDN resource.

6.7 Modify CDN Edge Server Creation Permissions

To avoid billing ambiguities, we highly recommend you to restrict the ability to create CDN edge servers and CDN storage servers to cloud administrator only.

In the OnApp 3.0.8 you can modify roles' permissions to limit users ability to create CDN servers with a rake task.

Use the following rake task to disable edge/storage server creation permissions for all roles, except for that specified in the square brackets:

```
RAILS_ENV=production rake permissions:
restrict_cdn_servers_creation_except_admins['1 2 3']
```

Where you have to specify IDs of roles that should be able to create CDN servers.

In case you do not specify roles' IDs, the CDN server creation permissions will be disabled for all roles, except for the default Administrator role.

6.8 Appendix: document revsions

v1.0, 15th May 2013

- First release