



Deployment Guide

How to prepare your environment for an OnApp Cloud deployment.

Document version	1.12
Document release date	5 th September 2012
	document revisions

Contents

1. Overview	3
2. Network configuration	4
2.1 Storage Network	5
2.2 Management Network.....	5
2.3 Provisioning Network.....	6
2.4 Appliance Network/VM networking.....	6
3. Server configuration	8
3.1 Controller server	8
3.2 Backup server.....	8
3.3 Hypervisor servers	9
3.4 SAN.....	10
3.4.1 Host components – hypervisor connectivity to the Storage Network	10
3.4.2 Fabric components – the network fabric between hypervisors and SANs.....	10
3.4.3 Storage components – the SAN chassis, controllers and disk trays.....	10
3.5 Additional hardware considerations for CDN	11
4. Handing over	12
Appendix: document revisions	13

1. Overview

This document describes the OnApp Cloud deployment process.

There are four main activities involved in a deployment:

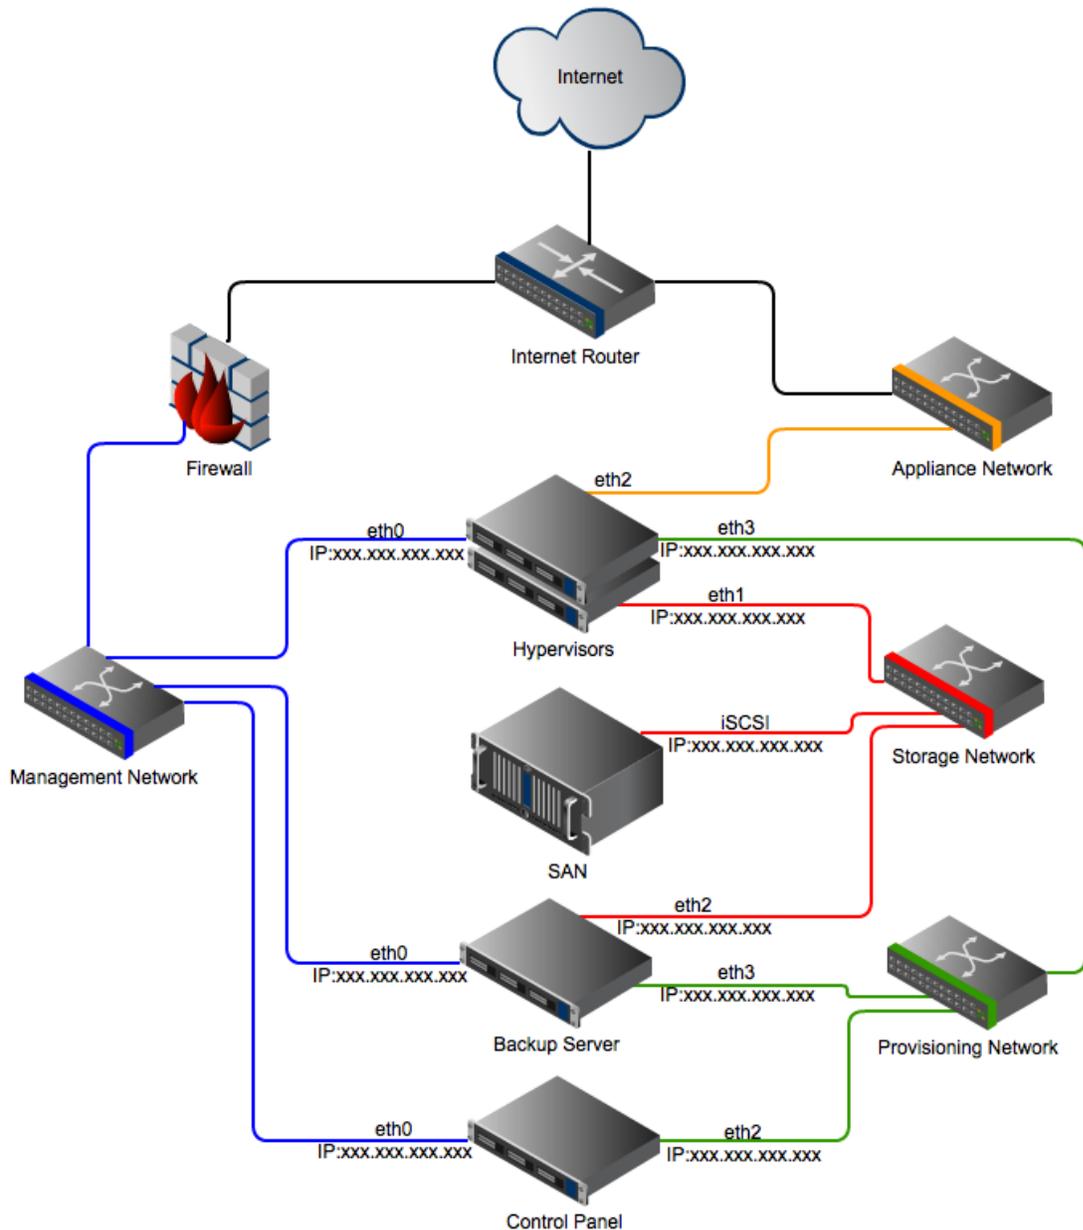
1. Network configuration (*you, with our help if required*)
2. Server configuration (*you, with our help if required*)
3. Handing over your configuration details (*you*)
4. Raising an installation ticket and performing the install (*OnApp*)

Please review the configuration details in each section carefully, as they are vital to the smooth operation of OnApp Cloud.

If you need help at any time, please get in touch with your OnApp integration specialist.

2. Network configuration

OnApp Cloud requires a specific network configuration to ensure maximum stability. There are four core networks: Storage, Management, Provisioning and Appliance. The recommended config is shown below.



It is very important to separate the four core networks. In the configuration shown above, the networks are separated physically using different switches. It is also possible to isolate networks using VLANs if your network supports it.

2.1 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.
- To achieve better performance and redundancy over 1Gbit you should consider NIC teaming/bonding and LACP or MPIO over multiple subnets .
- 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.

2.2 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.

Ideally this network would be a local network and would be 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.

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.

2.3 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.

2.4 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
```

As part of the OnApp Cloud installation requirements we will need you to configure the network interface file as above, and include the name of the interface (eth2, for example) when you hand over to OnApp for installation (the [handover process](#) is explained later in this document). 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.

As part of the handover process you will need to tell OnApp how your Appliance Network connectivity is configured at the network level. You'll need to connect it 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 the VM networking then you will need to specify the VLAN in the handover template so we can associate it with the subnet when adding the VM networking range to OnApp.

Configuring a switch trunk port is the preferred method, as 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.

① As part of the handover you will also need to supply us with an Internet-facing IP address range for use on the public interfaces. Note that all hypervisors will have this public interface, so the IP address range must be portable between hypervisors.

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.

3. Server configuration

This section lists the server types needed for an OnApp Cloud installation. For minimum hardware specs, see <http://www.onapp.com/onapp-requirements>. All servers are installed with the latest version of CentOS 5 (x86_64) base installation only, fully updated with 'yum -y update'.

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

3.1 Controller server

The controller 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 controller 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 controller server spec in mind, you're very welcome to send it to your OnApp integrations specialist for review.

3.2 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.

3.3 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 SAN then the virtual machines' disks will live on the SAN: in this case, 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 choose not to run a SAN, you can instead run local storage in your hypervisors. If you plan to do this, talk to your integrations specialist, as you will need to allow for the additional disks when choosing the hypervisor's chassis.

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 (VDSs) 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 diagram](#) above 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!

3.4 SAN

Primary storage is absolutely 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.

3.4.1 Host components – hypervisor connectivity to the Storage Network

We currently develop against CentOS 5. 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).

3.4.2 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.

3.4.3 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.

3.5 Additional hardware considerations for CDN

The processes handled by CDN edge servers are fairly disk intensive. We recommend that all edge server virtual appliances are hosted on local SSD storage in hypervisors.

When you're spec'ing your hypervisors you should consider including an SSD disk in some or all of them. You can talk to your integrations specialist about how to configure local storage on SSDs.

SSDs are *strongly* recommended if you plan to sell your edge server resources on the OnApp CDN marketplace. We also require edge servers to have a minimum of 4GB of RAM and 100GB of disk space.

4. Handing over

It is essential that you've met all of the network and server requirements described in the previous sections. Next, you'll hand over details of your network and server config to us, so we can begin the installation process.

1. Download and complete the deployment template:
<http://onapp.com/files/docs/onapp-deployment-template.txt>
2. Email the completed template to your OnApp integration specialist.
3. Once we have double-checked your details we will open an installation ticket with our technical team. The installation will usually be completed within 48 hours.
4. If you have any questions, please let your OnApp integration specialist know.

Appendix: document revisions

v1.13, 5th September 2012

- Updated diagram in the [network configuration](#) section to reflect recommended network design for OnApp Cloud v2.3.3.

v1.12, 31st July 2012

- Updated [Storage Network](#) section to add FibreChannel/InfiniBand options and to clarify jumbo frames/MTU information
- Updated [Provisioning Network](#) section in line with OnApp Cloud v2.3.2 requirements and to add FibreChannel/InfiniBand options.
- Updated [Backup server](#) section in line with OnApp Cloud v2.3.2 requirements.

v1.11, 15th May 2012

- Added note clarifying Intel-VT/AMD-V BIOS settings in the [hypervisor server](#) configuration section.

v1.10, 14th May 2012

- Updated diagram in the [network configuration](#) section.

v1.09, 1st May 2012

- Edited [backup server sizing](#) information to reflect changes introduced in OnApp Cloud 2.3.2 (the ability to add multiple backup servers makes it easier to scale backup storage as your cloud scales)

v1.08, 10th April 2011

- Renamed networks in the [network configuration](#) section. From v2.3.2 of OnApp Cloud:
 - Public network -> is now called Appliance network
 - Backup network -> is now called Provisioning network
 - SAN network -> is now called Storage network
 - Management network -> is still called Management network

v1.07, 28th November 2011

- Updated all [network](#) and [server](#) configuration sections with more comprehensive information and advice.
- Added [CDN hardware considerations](#) section.

v1.06, 17th November 2011

- Updated server CentOS install version information in the [Server Configuration](#) section to state that all servers are installed with the latest version of CentOS 5, rather than a specific current version.

v1.05, 4th November 2011

- Corrected URL of deployment template in the [handing over](#) section because the host's directory structure has changed.

v1.04, 2nd November 2011

- Added a note about NICs with Broadcom chipsets to the [SAN Network](#) section.

v1.03, 31st August 2011

- Corrected port information typo in [Management Network](#) config section (port 5555 must be open for the Control Panel server, not 555 as stated in previous version).

v1.02, 23rd August 2011

- Updated port information in [Management Network](#) config section.

v1.01, 18th August 2011

- Cosmetic changes only.

v1, 15th August 2011

- Cosmetic changes only.

v0.3, 6th July 2011

- Clarified [SAN](#) and [Management Network](#) configuration information.
- Minor editing & formatting changes.

v0.2, 5th July 2011

- Minor editing & formatting changes.

V0.1, 1st July 2011

- First release.