

OnApp Cloud 6.3 Get Started Guide

1	Technical Details	5
2	Installation	6
3	Preparation	7
4	Use Cases	8
4.1	Pay-as-you-go Public Cloud	8
4.2	Virtual Private Cloud	8
4.3	VPS Cloud	9
4.4	Hybrid Cloud Hosting	9
4.5	Traditional VPS Model	10
4.6	OnApp Federation	10
5	Hardware Specifications	11
5.1	Processor	12
5.2	Memory	12
5.3	Disks	12
5.4	RAID	12
5.5	Network Interfaces	12
5.6	Processor	12
5.7	Memory	12
5.8	Disks	12
5.9	RAID	12
5.10	Network Interfaces	13
5.11	Processor	13
5.12	Memory	13
5.13	Disks	13
5.14	RAID	13
5.15	Network Interfaces	13
6	Software Specifications	15
6.1	CentOS Versions	15
6.2	Libvirt Versions	15
6.3	Requirements	15
7	Networking	17
7.1	Management Network	18
7.2	Appliance Network	19
7.3	Storage Network	20
7.4	Provisioning Network	20
7.5	External Network Connectivity	20
8	Storage	22
8.1	Centralized Storage (SAN)	22
8.1.1	Fabric Components - Network Fabric Between Compute Resources and SANs	22
8.1.2	Host Components - Compute Resource Connectivity to Storage Network	22
8.1.3	Storage Components - SAN Chassis, Controllers, and Disk Trays	23
8.2	Integrated Storage (OnApp Storage)	23
8.3	SolidFire Integration	24
8.4	StorPool Integration	25
9	Servers	26
9.1	Installation Requirements	26
9.2	Control Panel Server	27
9.3	Backup Server	27
9.4	Compute Resource Servers	27
9.5	CloudBoot Compute Resource Servers	28
10	ISOs	29
10.1	Mount ISO Locations	29
10.2	Enable Permissions in Control Panel UI	29

11	OVA s	31
11.1	Mount OVA Locations	31
11.2	Enable Permissions in Control Panel UI	31
11.3	Use OVA on CloudBoot.....	32
12	Support	33
12.1	24/7 Support	33
12.2	What Does OnApp Support in my Cloud?	33
12.3	Cloud Troubleshooting	33
12.4	Professional Services.....	33
12.5	OnApp Community.....	34
12.6	Knowledge Base	34
12.7	Documentation.....	34
13	What's Next	35

The guides in this section apply to installing the OnApp Cloud 6.3 version. For the release notes list, please refer to the [Release Notes](#) space.

1 Technical Details

- [Hardware Specifications](#)
- [Server Config Reminder](#)
- [Software Specifications](#)
- [Network Preparation](#)
- [Use Cases](#)

2 Installation

- [Installation Workflow](#)
- [Installation Guide](#)
- [ISOs](#)
- [OVAs](#)
- [Support](#)

3 Preparation

- [Networking](#)
- [Storage](#)
- [Servers](#)

4 Use Cases

OnApp provides different cloud models and deployment strategies for you to build an environment that is right for your needs. This document is an overview of the following services that you can build with OnApp:

- [Public Cloud](#)
- [Virtual Private Cloud](#)
- [VPS Cloud](#)
- [Hybrid Cloud Hosting](#)
- [Traditional VPS Model](#)
- [OnApp Federation](#)

See also:

[Hardware Specifications](#)

[Software Specifications](#)

[Networking](#)


[Storage](#)

[Servers](#)

4.1 Pay-as-you-go Public Cloud

You can use OnApp to set up a complete pay-as-you-go public cloud and provide your users with the following:

- Enable users to build virtual servers and other cloud resources.
- Set different prices for RAM, CPU, and storage.
- Create different compute zones with different pricing.




4.2 Virtual Private Cloud

Use OnApp to offer virtual private cloud services or run a private cloud alongside a public cloud.

- Group compute, network, and storage resources into a single private cloud resource for users.


- Provide your users with all the benefits of a private cloud enhanced by the resources of the public cloud.



4.3 VPS Cloud

Use OnApp to create a cloud hosting service with resources packaged as a pre-configured VPS:


- Group cloud resources into packages that you can offer on a monthly/plan billing basis.
- Provide your users with packages that are the building blocks for their VSs.
- Facilitate the transition of traditional VPS customers to the cloud.



4.4 Hybrid Cloud Hosting

This is where dedicated hosting meets the cloud. You can use OnApp to offer hybrid servers to users and compete with every dedicated server provider out there:


- Allocate compute resources on a one-to-one basis: each user gets a dedicated compute resource for their hosted service.
- Get a failover that is provided by the rest of the cloud (for example, one compute resource might act as failover for five *live* compute resources).



4.5 Traditional VPS Model

You can use OnApp to provide traditional VPS services based on local storage:

- OnApp doesn't demand that you have a SAN back-end.
- If you want to provide users with traditional VPSs using local storage, OnApp can handle this model.



4.6 OnApp Federation

OnApp Federation is a global network of clouds that you can use to add scale and reach to your own cloud service. It gives you instant access to global compute cloud and content delivery infrastructure. OnApp Federation enables you to:

- Expand your cloud to 170+ locations.
- Add global scale for compute and content delivery.
- Host customers close to their users.
- Host customers in specific locations (or outside specific locations) for compliance.
- Offer your cloud infrastructure within OnApp Federation. You can set a price and get paid when other members of Federation use your resources.

5 Hardware Specifications

To run an OnApp installation, you need to set up one Control Panel server, at least two compute resource servers, and one backup server. These server instances are required to handle your environment as follows:

- **Control Panel Server**

The Control Panel server hosts the OnApp user interface and database, as well as handles all cloud management processes. You do not need to set up a separate database server for your OnApp installation. OnApp provides full compatibility of Control Panel UI with Google Chrome and Firefox web browsers where JavaScript is enabled.

- **Compute Resource Servers**

The compute resource servers provide CPU, RAM, and storage resources for applications and virtual servers that you or your users run in the cloud. A certain amount of CPU and RAM on each server is reserved for a compute resource (hypervisor) and the storage controller system. The remaining resources are available for allocation to virtual servers. You can set up at least two compute resource servers and scale out based on your needs.

- **Backup Server**

The backup server stores virtual server backups and templates from which you can create virtual servers. It also handles disk related transactions, such as provisioning virtual servers, taking backups, and resizing disks. The backup server may be optional for a staging environment, however, it is critical for an environment running production workloads. The backup server should include a backup storage volume that can be a local disk array or a mounted SAN/NAS storage.

This document provides hardware suggestions that you can use to set up each of the servers in your environment. Here you can also find information about [OnApp Integrated Storage](#) and requirements for local storage.

If you have any questions regarding hardware specifications, please contact our [support team](#) to get assistance.

See also:

[Use Cases](#)

[Software Specifications](#)

[Networking](#)

[Storage](#)

[Servers](#)

- Control Panel Server

To set up a Control Panel server, you can use hardware with the following specifications:

5.1 Processor

8 Cores CPUs
Intel Xeon e5-2640 v3 or similar

5.2 Memory

32+ GB DDR4 RAM

5.3 Disks

2 x 500 GB SSD

5.4 RAID

RAID 1

5.5 Network Interfaces

2 x 1 Gbps+

- Compute Resources Server

To set up a compute resource server, you can use hardware with the following specifications:

5.6 Processor

8 Core CPUs
Intel Xeon e5-2640 v3 or similar

5.7 Memory

From 196 to 999 GB DDR4 RAM

5.8 Disks

2-4 x 500 GB - 1 TB SATA/SAS/SSD (for VS storage)
1 x 500+ GB NVMe (for caching)

5.9 RAID

PCIe Gen3
PERC H730 1 GB Cache or similar
Pass-through / JBOD Mode

5.10 Network Interfaces

2 x 1 Gbps+ and 2 x 10 Gbps+

- Backup Server

To set up a backup server, you can use hardware with the following specifications:

5.11 Processor

8 Core CPUs
Intel Xeon e5-2620 v3 or similar

5.12 Memory

32+ GB DDR4 RAM

5.13 Disks

12 x 2 TB SAS

5.14 RAID

RAID 5/6/10

5.15 Network Interfaces

1 x 1 Gbps+ and 2 x 10 Gbps+

- Network Hardware

To set up a network, you can use hardware with the following specifications:

2 x High performance switches with 48 x 10 GbE and 4 x 40 GbE ports.

- Integrated Storage

[Integrated Storage](#) enables you to build a scalable and resilient distributed SAN by polling disks that are attached to compute resources. As a result, you can create one or more virtual data stores that span multiple physical drives on compute resources with RAID-like replication and striping across drives.

The following requirements are recommended to implement Integrated Storage:

- Integrated Storage can group together any number of drives across any compute resource. We recommend at least two drives per compute resource to enable redundant data store configurations.

- At least one dedicated NIC assigned per compute resource for the storage network (SAN).
- IGMP snooping must be disabled on a storage switch for a storage network.
- Enterprise-grade SSD drives for the best performance and reliability.
- Every compute resource in a zone should have identical disks (same model and capacity) and storage network bandwidth.
- Local Storage

If you want to use local storage, the following requirements are recommended to implement:

- At least one dedicated partition on each compute resource.
- A separate disk from the primary OS drive is recommended.

6 Software Specifications

OnApp Cloud runs on CentOS. The CentOS versions can vary, depending on Xen or KVM virtualization.

6.1 CentOS Versions

The following table lists CentOS versions for compute resource, Control Panel, and backup servers. The version in bold is a recommended version.

See also:

[Hardware Specifications](#)

[Networking](#)

[Storage](#)

Server	Supported Versions
Xen Compute Resource Server	CentOS 6.x x64 or CentOS 7.x x86/64
KVM Compute Resource Server	CentOS 6.x x64 or CentOS 7.x x86/64
Control Panel Server	CentOS 6.x x86/64 or CentOS 7.x x86/64
Backup Server	CentOS 6.x x64 or CentOS 7.x x86/64

6.2 Libvirt Versions

The following table provides the libvirt versions compatible with your system, depending on the virtualization type and version.

Xen		KVM	
CentOS 6.x	libvirt 3.2.1	CentOS 6.x	libvirt 0.10.2
CentOS 7.x	libvirt 4.1.0	CentOS 7.x	libvirt 3.9.0

6.3 Requirements

There are the following requirements to Control Panel, compute resource, and backup servers based on CentOS:

- Install CentOS from the minimal CentOS ISO for Control Panel servers, static backup servers and static compute resources.
- The minimum running services are listed on the box:

```
network    0:off  1:off  2:on   3:on   4:on   5:on   6:off
sshd       0:off  1:off  2:on   3:on   4:on   5:on   6:off
```

- The *network* should be configured with an ability to access rpm.repo.onapp.com and templates.repo.onapp.com.

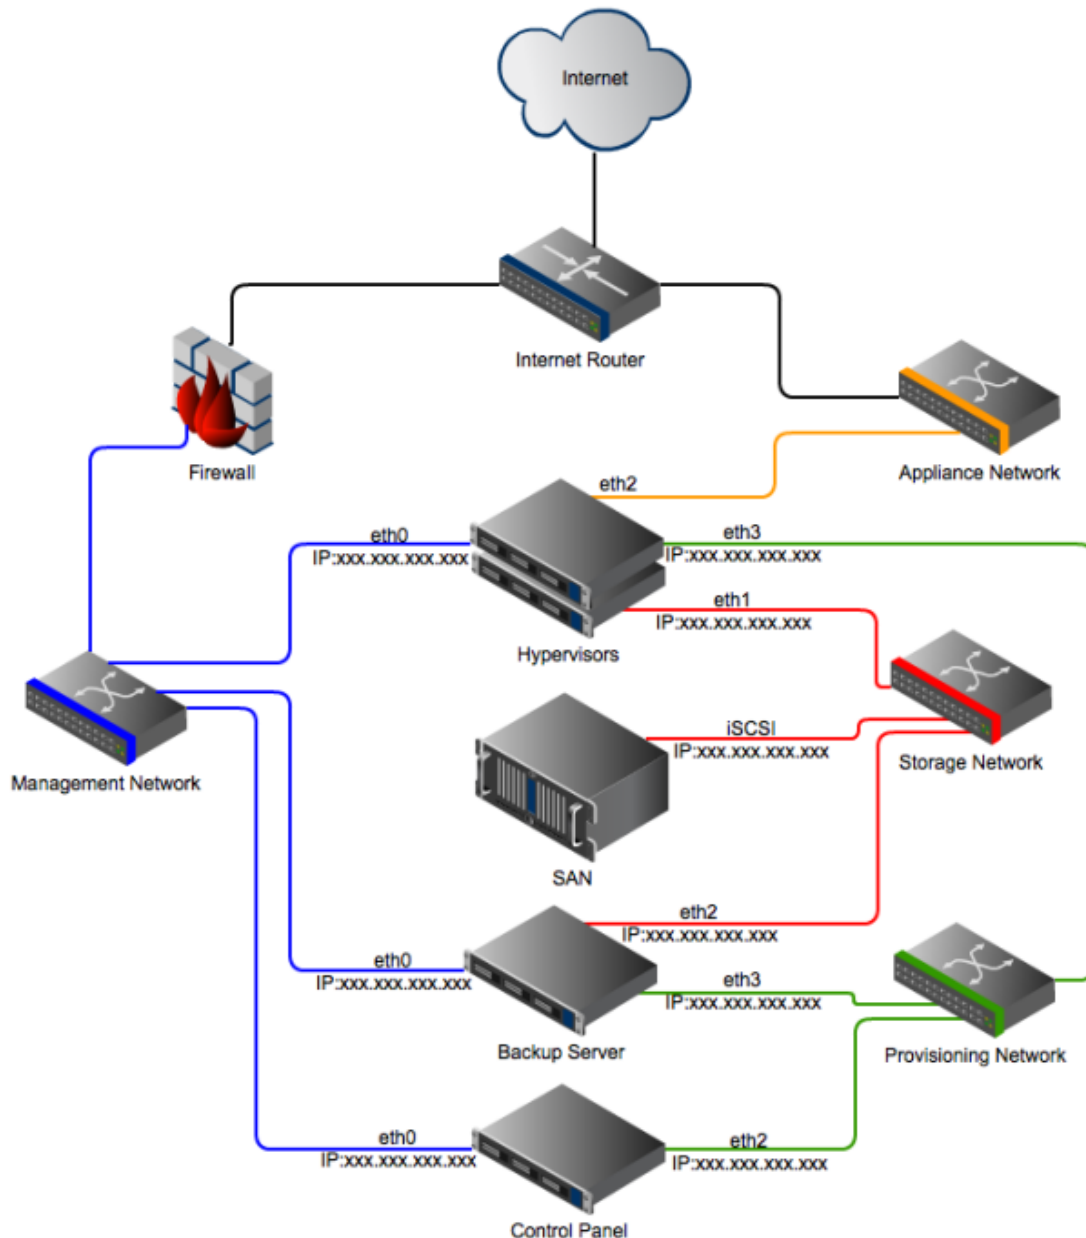
- The *open ssh server* should be configured with an ability for users to access and log into the box.
- The **root** user should be available on the box and configured as *root account/ root user/ superuser* with an access to all files, commands/tools, and services on the system. Installers should be run under the **root**.
- The *curl*, *rpm*, *yum*, and *grub* packages must be installed on the system. The *grub* is a mandatory boot loader only for static compute resources.
- Avoid using additional (not native) repositories for CentOS, such as **Extra Packages for Enterprise Linux** (*epel*) and others.

7 Networking

An OnApp Cloud installation requires four separate networks that are Management, Appliance, Storage, and Provisioning network. If you plan to use OnApp Integrated Storage, you don't need to set up the Provisioning network.

The networks should be separated either physically, using different switches, or with VLANs. If you experience MAC address flapping across ports because a switch does not support the balance-rr mode, you can set up separated VLANs per each bond pair for the switch.

The following scheme illustrates how networks handle your cloud environment. For more details on the role and configuration of each network, see the sections below.



Some datacenter vendors may provide limited access to the network configuration, ability to create VLANs, and assign IP addresses. These limitations can prevent such vendors from working with OnApp. To ensure

that your datacenter vendor complies with the requirements, please contact OnApp architecture team.

On this page:

[Management Network](#)

[Appliance Network](#)

[Storage Network](#)

[Provisioning Network](#)

[External Network Connectivity](#)

See also:

[Required Ports](#)

[Storage](#)

[Servers](#)

[ISOs](#)

[OVAs](#)

7.1 Management Network

The Management network serves as a route for communication between the Control Panel server, compute resources and backup servers. The management network should always be a default gateway. If you deploy CloudBoot compute resources, the Control Panel server automatically assigns management network IP addresses via the internal DHCP server to compute resources and backup servers.

The Management network should be a local network behind a gateway device that is capable of bridging traffic to the Internet to allow servers to perform tasks such as DNS resolution, NTP and operating system updates. If the gateway has a DHCP service for allocating private IP addresses, this service must be disabled.

You have to open the 443 port for outgoing connections to the OnApp Licensing Server. The Control Panel server needs to have incoming traffic allowed to ports 80/443 & 30000->40000. This should be configured at the gateway with incoming NAT. If your gateway device doesn't support it, a network can also be an external network. However, you should always use firewall at the gateway to block all incoming traffic except for the ports listed above.

Since the Management network serves as a route for communication between the Control Panel server, compute resources and backup servers, the stability of this network is critical. You should always consider bonding to introduce the network level redundancy and the network bandwidth should be at least 1 Gbit.

If your management network is behind a firewall, please make sure that ports 22/80/443/30000-40000 are open to the Control Panel server and port 22 for all other servers. The 30000-40000 ports are not required if you are going to use HTML5 console as the console proxies over port 80 or 443.

OnApp and vCloud Director Integration Requirements

OnApp and vCloud connection is supported by RabbitMQ. OnApp Control Panel connects to vCloud Director using REST API and requires an outgoing connection to vCloud API interface via ports 80 and 443. By default, the RabbitMQ server is installed by OnApp and the Management network requires an incoming connection to port 5672. The port 15672 is optional for the RabbitMQ server management. If external AMQP server is used, an outgoing connection to the RabbitMQ default port 5672 is required.

7.2 Appliance Network

The Appliance network is used for all virtual servers network traffic. OnApp bridges the appliance NIC and assigns virtual interfaces to it when virtual servers are provisioned, or when additional network interfaces are added to virtual servers via OnApp UI or API. Since the public interface is managed by OnApp, the public NIC requires a blank config.

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

You should configure your network interface file accordingly. You don't need to add any configurations to this NIC, such as subnet, gateway or IP address details. 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). Take it into consideration when you design your compute resources to make sure you have a physical NIC available. The Appliance network bandwidth should be at least 1 Gbit. You should also consider bonding on the Appliance network to introduce redundancy at the network level.

Configuring a switch trunk port is a preferred method because it gives you additional flexibility and security. Alternatively, you can configure a switch access port. In the latter case, you don't need to specify a VLAN when adding a range to OnApp. To be able to use multiple appliance VLANs, connect your Appliance network to switch ports that are configured in a VLAN trunk mode. This provides your cloud with flexibility to offer private VLANs to users in future. If you choose multiple appliance VLANs, you need to associate your VLAN with a subnet when you add a range to OnApp.

SDN Network

Software-defined networking (SDN) provides the ability to manage networks using VXLAN technology across OnApp cloud compute resources. Thus, you receive a tool to build a level-two network infrastructure with OnApp on top of the existing IP (level three) network. SDN networks belong to appliance networks.

You may consider the following points before the creation of SDN networks:

- OnApp requires Carbon 0.6.2 ODL controller version.
- ODL controller should be accessible from Control Panel with SDN manager host:port and from compute resources with selected connection options (tcp:ip_address:port)
- The SDN network creation is currently supported only on KVM compute resources.
- You will need to ensure the [connection via IPs](#) between the compute resources: configure an IP address to use for tunneling traffic over VXLAN to the other nodes.
- Ensure there is a [connection option](#) to connect a compute resource to OpenDayLight Controller. You may refer to the [Install OpenDayLight Controller](#) guide to check the existing hardware requirements to install OpenDayLight server.

- Control Panel should be able to connect to the ODL controller using host, port, login, and password (it is possible with OnApp Cloud).
- Since VXLAN adds 50 to 54 bytes of additional header information to the original Ethernet frame, you might want to increase the maximum transmission unit (MTU) of the underlying NIC on the corresponding compute resource.

7.3 Storage Network

The Storage network enables a connection between storage devices (e.g. SANs) and compute resources. The type of a network depends on which kind of connectivity your primary storage requires. For example, if you use iSCSI or ATAoE, you need to set up an Ethernet network. If your SAN has fibre connectivity, then the storage network is a fiber network.

The stability of the Storage network is absolutely critical. You should always make redundancy your primary concern when designing this network. See the [Centralized Storage \(SAN\)](#) section for more details.

There are the following requirements to the Storage network that you need to follow:

- The Storage network must be a local network.
- The Storage network should run at least 10 Gbit FibreChannel or InfiniBand to achieve the best 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.
- To achieve better performance and redundancy over 1 Gbit, 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 a switch connecting compute resources to SAN supports jumbo frames: the Storage network on compute resources and SAN(s) must have MTU set to 9000 to optimize performance.

7.4 Provisioning Network

The Provisioning network is used to transfer backup and template data between the provisioning server and the primary storage volumes. The network is used to transfer large sets of data, so we recommend that it runs at least 1 Gbit. It is more preferable for you to consider 10Gbit, FibreChannel, InfiniBand or aggregated 1 Gbit links for maximum throughput. The Provisioning network is not required for clouds that run OnApp Integrated Storage with dedicated backup servers. However, without a provisioning network, all migration traffic will use the management network.

7.5 External Network Connectivity

The following table provides an overview of communications between the OnApp Control Panel server and external networks.

Source	Destination	Port	Description
OnApp CP	licensing.onapp.com	443	For the Control Panel server to communicate with the OnApp licensing dashboard.
OnApp CP	Public Internet	25	For email notifications that are sent outbound from the OnApp Control Panel server.
End Users	OnApp CP	80/443	For users to access the OnApp web interface over the HTTP or HTTPS protocol.

8 Storage

The installation requirements to storage vary depending on a storage strategy that you plan to deploy. In this document, you can find information on how to implement the following storage solutions:

- [Centralized Storage \(SAN\)](#)
- [Integrated Storage \(OnApp Storage\)](#)
- [SolidFire Storage](#)
- [StorPool Storage](#)

On this page:

[Centralized Storage \(SAN\)](#)
[Integrated Storage \(OnApp Storage\)](#)
[SolidFire Integration](#)
[StorPool Integration](#)

See also:

[Integrated Storage Servers](#)

8.1 Centralized Storage (SAN)

The primary storage is critical to your cloud and your SAN has a huge impact on the performance of the whole platform. OnApp offers you a lot of flexibility in your primary storage technology. OnApp supports each solution that is capable of presenting a block device to compute resources. This could be, for example, FiberChannel, SCSI or SAS HBA, iSCSI or ATAoE, or a InfiniBand HCA controller, since all of them present the block device directly. OnApp does not support services such as NFS for primary storage, because they preset a filesystem and not the block device.

Beyond a type of the 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 designing your SAN.

8.1.1 Fabric Components - Network Fabric Between Compute Resources and SANs

You 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 compute resources and SANs to the cloud, you need to be able to grow the physical connectivity without downtime on the Storage Network.

8.1.2 Host Components - Compute Resource Connectivity to Storage Network

You 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 also need to think about the throughput and whether the connectivity on compute resources will be suitable for the virtual servers they will be running. A bottleneck here will cause major performance issues. Consider adding multiple HBAs or NICs if you plan to run a redundant switching mesh (see the Fabric section) as bonding or multipath will be required, unless the redundancy is built into the physical switch chassis (for example, failover backplanes).

8.1.3 Storage Components - SAN Chassis, Controllers, and Disk Trays

You need to take into consideration physical capacity that is required to build a storage of a necessary size. This gives 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 an increasing use of the storage in future. You should also think about the physical storage hardware and whether you will use SATA, SAS or SSD. 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.

8.2 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 compute resources. With OnApp Storage, you create a virtual data store that spans multiple physical drives in compute resources, with RAID-like replication and striping across drives. The SAN is fully integrated into the compute resource platform and the platform is completely decentralized. There is no single point of failure: for example, if a compute resource fails, the SAN reorganizes itself and automatically recovers the data.

The following requirements are recommended for Integrated Storage implementation:

- Integrated Storage can group together any number of drives across any compute resource. We strongly recommend at least two drives per compute resource to enable redundant datastore configurations.
 - SSD drives are recommended for best performance.
 - At least one dedicated NIC assigned per compute resource for the storage network (SAN).
 - Multiple NICs bonded or 10Gbit/s Ethernet are recommended).
 - MTU on storage NIC: 9000 is recommended.
 - IGMP snooping must be disabled on a storage switch for a storage network.
- Enabling jumbo frames MTU > 1500, up to a maximum of 9000, requires NIC and switch support. Ensure that your network infrastructure has jumbo frame support and that jumbo frames are enabled in any switches, otherwise, leave MTU as default 1500 for storage NICs. Additionally, MTU must be equal for all storage NICs for compute resources, including backup servers.

- To start using Integrated Storage, you must enable it in the system configuration first (**Admin > Settings > Configuration > System Configuration > OnApp Storage**).
- Integrated Storage uses a certain RAM amount on each compute resource but the exact RAM amount depends on the number of drives and controllers which will be configured.
- Note that advanced disk sector format is not supported for Integrated Storage disks. Ensure that your disk drives support the 512-byte sector alignment before installing and using them with Integrated Storage.
- The Bonded NICs for the management/boot interface are not yet available (they will be introduced in future releases).

8.3 SolidFire Integration

OnApp is integrated with the SolidFire storage management system. With the SolidFire integration, it is possible to use the SF SAN directly within the OnApp cloud and manage the SolidFire cluster via the SolidFire API. To be able to use SolidFire in the cloud, you need to install the SolidFire storage system first.

You can perform the following options with SolidFire:

- Use SolidFire SAN in the OnApp cloud.
- Allocate dedicated LUNs from the SF cluster per virtual server disk, when creating a virtual server.
LUN is created per each virtual server disk, with a separate LUN per swap disk.
- Manage SolidFire LUNs automatically via API.
- Create virtual servers without a 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 SolidFire side, using the CreateVolume API call. The LUNs on the SolidFire are managed automatically via SolidFire API. Inasmuch a 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 SolidFire volume, you have to enable export to this device (a compute resource or 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 for SolidFire:

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

8.4 StorPool Integration

OnApp provides and maintains integration with StorPool, whereas each VS disk becomes a separate volume in the StorPool storage system.

OnApp with StorPool can be deployed in two modes - hyperconverged and standalone.

- In 'hyperconverged' mode, StorPool storage servers are running on the OnApp hypervisors (Compute Resources) to pool the disks together to make a datastore that will be managed in its entirety by StorPool.
- In 'standalone' mode, StorPool storage servers are deployed on dedicated servers. In this case, StorPool block device drivers still need to be installed on the OnApp hypervisors and backup nodes, so they are able to access the StorPool storage and consume volumes.

With the StorPool integration, it is possible to get control over each individual volume, in terms of monitoring and data placement. To be able to use StorPool in the cloud, you need to install the StorPool storage system first. For details, refer to [StorPool section](#) of our documentation.

9 Servers

When you are finished with networks and storage, you can proceed to setting up the following servers:

- [Control Panel](#)
- [Backup](#)
- [Static Compute Resources](#)
- [CloudBoot Compute Resources](#)

9.1 Installation Requirements

There are some requirements to server installation that you need to follow. OnApp runs on CentOS but a CentOS version depends on a virtualization you are running.

- We recommend installing CentOS from a minimal CentOS ISO for Control Panel servers, static backup servers, and compute resources.
- Full root access: please do not create the user 'onapp' since it is created as a part of the RPM installation.
- When installing CentOS, do not use a partition scheme that allocates the majority of disk space to a dedicated `/home` partition, leaving the `root` partition a small amount of space. Instead, allocate the majority of disk space to the `root` partition or a dedicated `/onapp` partition.

For the list of all requirements, see [Software Specifications](#).

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

On this page:

[Control Panel Server](#)

[Backup Server](#)

[Compute Resource Servers](#)

[CloudBoot Compute Resource Servers](#)

See also:

[Hardware Specifications](#)

[Software Specifications](#)

[Storage](#)

9.2 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 selecting hardware for this server. When your production workloads grow, you need to add more compute resources and SANs, which puts more load on you Control Panel. Selecting the right hardware at the beginning is important and helps to avoid downtime during upgrades later down the line.

The Control Panel server may experience a high load on MySQL as you add more compute resources, so a fast disk array and lots of memory is recommended. For more details, see the [Hardware Specifications](#) document. If you have the Control Panel server specifications in mind, you can send them to your OnApp integrations specialist for a review.

9.3 Backup Server

The backup server stores virtual server backups and templates. It is also responsible for processing any disk transactions running in your cloud, such as provisioning virtual servers, 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 compute resources.

Unlike primary storage, performance is not so essential here so 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 RAID6 is usually ideal for the backup volume. When configuring SAN, take into consideration that a larger block size is recommended owing to the nature of the data being stored on this array.

Backup storage is 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. In the traditional/centralized SAN configuration, you have to bind all your data stores to a backup server. The volume groups of each data store based on SAN must be shared with the backup server.

In a cloud where CloudBoot is enabled, you have to use CloudBoot backup servers instead of dedicated backup servers. To do so, you have to create a CloudBoot compute resource to be used as a backup server. You can set up CloudBoot backup servers and virtual dedicated backup servers to be used with the Integrated Storage. The backup scheme remains unchanged.

9.4 Compute Resource Servers

Compute resources are where virtual servers run in your cloud. A small amount of compute resource CPU, memory, and disk resource is reserved for the OnApp engine: the remainder is available as virtual resources to allocate to virtual servers.

If you use a centralized SAN, then the virtual server disks run on that SAN, and the compute resource own disk is used to boot the compute resource and run the OnApp engine. Performance here is not critical but we recommend introducing some redundancy: RAID1 SATA/SAS would be perfect. If you use OnApp Storage (our integrated SAN), you should factor

more disks into your compute resource 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 compute resources, though you lose the ability to failover from compute resource to compute resource: this is not recommended for an optimal cloud setup.

When you build your hardware, it's important to take into consideration the specifications of the primary components that will be virtualized: 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 virtual servers you can fit on a single compute resource is limited by the amount of RAM installed on that compute resource. Another limitation to consider is that the compute resource CPU is a shared resource: the physical cores are shared among virtual servers running on a compute resource. Do not overload the compute resource with too many virtual servers as this can stretch the available CPU time and degrade the performance of all servers on that compute resource.

It's also important to note that too many virtual servers could potentially saturate the SAN NICs on the compute resource, which can introduce instability and performance loss to virtual servers (see the *Host Components - Compute Resource Connectivity to Storage Network* section for more details).

In the [Networking](#) document, you can see that OnApp requires at least 4 NICs on the compute resources. 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 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 compute resources to enable you to provision Windows-based virtual servers on your OnApp cloud.

9.5 CloudBoot Compute Resource Servers

CloudBoot is a feature that enables fast provisioning of Xen and KVM compute resources without any pre-installation requirements. Using network/PXE boot methods, a new server can be plugged in and powered on. This server is automatically discovered by the OnApp Control Panel Server and installed over the network, so it is booted as a fully configured compute resource ready to host virtual servers.

The Control Panel Server manages IP address to hardware MAC assignment and the booting of a Xen or KVM image on demand. The compute resource images come preinstalled with all the SSH keys and any other settings specific to the node to enable compute resources to come online instantly. Images are booted as a standalone RAM disk. After images are bootstrapped, they operate independently from other servers but without any persistent installation dependency.

This enables booting of *diskless* blades, as well as booting compute resources with Integrated Storage 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 compute resource servers.
- A secondary NIC is recommended for the Control Panel Server to provide a fully isolated network for the compute resource management subnet, including PXE boot and DHCP support for the compute resources.

For resilience, a secondary static TFTP server target can be configured to handle Control Panel server failure and ensure hardware boot consistency in the event of such a failure.

10 ISOs

You can enable users to build and boot virtual servers from ISO images. To enable the usage of ISO, you mount a location where ISO images are stored on a Control Panel server to compute resource servers.

10.1 Mount ISO Locations

When a virtual server is booted from an ISO image, the ISO image is taken from a compute resource server. To mount and share a location where ISO images are stored on a Control Panel server with the compute resource servers, you edit the `onapp.yml` file as follows:

- `iso_path_on_cp` - specifies the location where ISO images are stored on the Control Panel server. By default, the location is `/data`. You can change it to any other suitable location. Make sure that this location is shared with the specified `iso_path_on_hv` location.
- `iso_path_on_hv` - specifies the location where ISO images are located on the compute resource servers. By default, the location is `/data`. You can change it to any other suitable location with the `onappowner` and read/write access. Make sure that this location is mounted to the specified `iso_path_on_cp` location.

CloudBoot compute resources mount the `/data` location at boot to the `/onapp/tools/recovery` and create the symlink at `default/data` automatically on the CentOS 7 compute resource. If you are using CentOS 6 compute resources, you will have to create a symlink manually. For that, create symbolic links in `/onapp`:

```
# unlink /onapp/templates
```

You can store ISO images on a dedicated server at any location with an arbitrary name. In this case, it is necessary to mount the ISO images location on this server to the `iso_path_on_cp` directory on Control Panel and all `iso_path_on_hv` locations on compute resources. This can be a backup server to avoid the usage of the Control Panel resources.

10.2 Enable Permissions in Control Panel UI

Make sure to enable the following permissions for your Admin and other roles in the Control Panel user interface:

- *Any action on ISOs* - the user can take any action on ISOs
- *Create a new ISO* - the user can create a new ISO
- *Destroy any ISO* - the user can delete any ISO (own, user, and public)
- *Destroy own ISO* - the user can delete only own ISO
- *Destroy user ISO* - the user can delete ISOs created by any user, but not public ISOs
- *Make any ISO public* - the user can make public any ISO available to all users
- *Make own ISO public* - the user can make public only own ISOs
- *Make user ISO public* - the user can make public ISOs created by any user
- *Create and manage own ISOs* - the user can create and edit/delete/view own ISOs

- *Manage all ISOs* - the user can manage own/user/public ISOs
- *Create and manage user ISOs* - the user can view/create/edit/delete ISOs created by any user
- *See all ISOs* - the user can view all ISOs in the cloud
- *See own ISOs* - the user can only view the ISOs created by themselves
- *See all public ISOs* - the user can view all public ISOs
- *See user ISOs* - the user can view the ISOs created by any user in the cloud
- *Update any ISO* - the user can edit any ISO in the cloud
- *Update own ISO* - the user can edit only own ISO
- *Update user ISO* - the user can edit ISOs created by any user in the cloud

See also:

[ISOs](#)

[Permissions](#)

11 OVAs

You can enable users to build and boot virtual servers from OVA images. To enable the usage of OVA, you mount a source location to which OVA is uploaded on a Control Panel server with a destination location where OVA is stored on a backup server

11.1 Mount OVA Locations

The can mount OVA locations in the `on_app.yml` file:

- **Source Path on Control Panel Server** (`ova_path`) - specifies the location where OVAs are downloaded/uploaded on the Control Panel server. By default, the location is `/data`. You can change it to any other suitable location. Make sure that this location is mounted to the specified destination location.
- **Destination Path on Backup Server** (`ova_path`) - specifies the location where OVAs are stored on the backup server. By default, the location is `/data`. You can change it to any other suitable location with the `onappowner` and read/write access. Make sure that this location is shared with the specified source location.

CloudBoot compute resources mount the `/data` location automatically at boot to the `/onapp/tools/recovery` on a backup server. For more information on using OVAs on CloudBoot backup servers, refer to the [Administration Guide](#).

11.2 Enable Permissions in Control Panel UI

Make sure to enable the following permissions for your Admin and other roles in the Control Panel user interface:

- *Any action on OVAs* - the user can take any action on OVAs
- *Create a new OVA* - the user can create a new OVA file
- *Destroy any OVA* - the user can delete any OVA (own, user, and public)
- *Destroy own OVA* - the user can delete only own OVA
- *Destroy user OVA* - the user can delete OVAs created by any user but not public OVAs
- *Make any OVA public* - the user can make public any OVA available to all users
- *Make own OVA public* - the user can make public only own OVAs
- *Create and manage OVAs* - the user can create and edit/delete/view OVAs
- *Manage public OVAs* - the user can manage public OVAs
- *Create and manage user OVAs* - the user can view/create/edit/delete OVAs created by any user
- *See all OVAs* - the user can view all OVAs in the cloud
- *See own OVAs* - the user can only view OVAs created by themselves
- *Read all public OVAs* - the user can view all public OVAs
- *See user OVAs* - the user can view OVAs created by any user in the cloud
- *Unlock any OVA* - the user can unlock any OVA that is currently being converted
- *Update any OVA* - the user can edit any OVA in the cloud
- *Update own OVA* - the user can edit only own OVA
- *Update user OVA* - the user can edit OVAs created by any user in the cloud

11.3 Use OVA on CloudBoot

You can enable OVA on Static and CloudBoot resources. To use OVA on a CloudBoot backup server, some additional steps are required. Follow the next procedure to enable OVA for CloudBoot backup servers:

1. Connect to the Control Panel server via SSH.
2. Edit the `/etc/exports` file on the line `/data X.X.X.X/YY(ro,no_root_squash,` where `X.X.X.X/YY` is your network/subnet. Change the permissions from `ro` to `rw` and save the file.
3. Restart the NFS service as follows:

```
/etc/init.d/nfs restart
```

It is not recommended to restart the NFS service at the same time when files are in use from the NFS share.

4. Share the `/data` directory as described in the **Mount OVA Locations** section.
5. Go to **Control Panel > Settings > Compute Resources > Label** of the required CloudBoot compute resource.
6. Click **Actions > Edit** next to the CloudBoot backup server.
7. Add the following into the *Custom Config* box:

```
cp /etc/lvm/lvm.conf /etc/lvm/lvm.conf.orig
sed -i 's/^[[:space:]]*filter = .*/filter = \[ "r|\|/dev|\nbd\|" \]/g'
/etc/lvm/lvm.conf
```

8. Click **Save**.
9. Reboot your CloudBoot backup server.

You can also execute the custom config command directly on a backup server to apply it without the reboot.

On this page:

- [Mount OVA Locations](#)
- [Enable Permissions in Control Panel UI](#)
- [Use OVA on CloudBoot](#)

See also:

- [2020-04-06 21-29-16 OVAs](#)
- [Install Backup Server](#)
- [Permissions](#)

12 Support

12.1 24/7 Support

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

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

12.2 What Does OnApp Support in my Cloud?

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

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

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

12.3 Cloud Troubleshooting

In case of any issues, you may [lock down your cloud](#) and allow the OnApp support team to assist you with troubleshooting.

12.4 Professional Services

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

12.5 OnApp Community

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

12.6 Knowledge Base

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

12.7 Documentation

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

13 What's Next

When you are ready with your infrastructure, you can proceed to installing the Control Panel, compute resources, and backup servers. For instructions on how to install each server instance, see the following documents:

Install Control Panel Server

Install Compute Resources

Install Data Stores

Install Backup Server

You can access all the documents related to the installation procedures at [OnApp Installation Guide](#).