

Configure Servers

This section is the part of the OnApp preparation guide.

[Configure Networks](#) > [Configure Storage](#) > **Configure Servers**

Once you have configured networks and storage, proceed to setting up the Control Panel, Backup, and Compute resource servers.

See also:

[Configure Networks](#)

[Configure Storage](#)

[Installation Guide](#)

[Technical Details](#)

On this page:

- [Server Installation Requirements](#)
- [Control Panel Server](#)
- [Backup Server](#)
- [Compute Resource Servers](#)
- [CloudBoot Compute Resource Servers](#)

Server Installation Requirements

This section lists the server installation requirements needed for an OnApp Cloud installation. For minimum hardware specs, see [Technical Details](#). OnApp primarily runs on CentOS but the version depends on what virtualization method you are running.

- We recommend installing CentOS from the minimal CentOS ISO for Control Panel servers, static backup servers and static compute resources.
- CloudBoot is not compatible with CentOS 5 KVM compute resources.
- Full root access: please do not create the user 'onapp' since this is created as part of the RPM installation.
- Currently Emulex hardware does not support 3.x Linux kernels, so it is only compatible with CentOS 5.x.
- When installing CentOS, do not use a partition scheme that will allocate the majority of disk space to a dedicated `/home` partition leaving the root partition a small amount of space. Instead, the majority of disk space should be allocated to the root partition or a dedicated `/onapp` partition.



We strongly recommend that you avoid creating mixed compute zones:

- do not add CloudBoot and static boot compute resources to one compute zone
- do not add both XEN and KVM compute resources to one zone

The reason is that XEN VSs cannot migrate/failover to a KVM compute resource and KVM VSs cannot migrate/failover to a XEN compute resource.

Supported server configuration

- **XEN Compute resources** CentOS 5.x x64 or CentOS 6.x x64
- **KVM Compute resources** 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 (only if running on Xen 3.x) or CentOS 6.x x64

Recommended server configuration

We highly recommend using the following server configuration:

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

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 compute resources 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 compute resources, 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.

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 – 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.

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

In the OnApp cloud with CloudBoot 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 functionality. The backup scheme remains unchanged.

Compute Resource Servers

Compute resources are where virtual servers live 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 are using a centralized SAN, then the virtual servers' disks will live on that SAN, and, the compute resource's own disk will simply be 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 are using OnApp Storage (our integrated SAN), you should obviously 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 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 virtual servers you can fit on a single compute resource is limited by the amount of RAM installed in that compute resource.

Another limitation to consider is that the compute resource's CPU is a shared resource: the physical cores are shared among the VEs running on a compute resource. Do not overload the compute resource with too many virtual servers, as this will 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 will also introduce instability and performance loss to virtual servers (see the Host Components - Compute Resource Connectivity to the Storage Network section for more details).

In the [Recommended Network Configurations](#) chapter, 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 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 compute resources to enable you to provision Windows-based virtual servers on your OnApp cloud!

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, being automatically discovered by the OnApp Control Panel Server, and installed over the network so it boots 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. Compute resource images come pre-installed, 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, 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 compute resources 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 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.

This section is the part of the OnApp preparation guide.

[Configure Networks](#) > [Configure Storage](#) > **Configure servers**