



Deploy the ASAv On the Rackspace Cloud

You can deploy the ASAv on the Rackspace cloud.



Important Beginning with 9.13(1), any ASAv license now can be used on any supported ASAv vCPU/memory configuration. This allows ASAv customers to run on a wide variety of VM resource footprints.

- [Overview, on page 1](#)
- [Prerequisites, on page 2](#)
- [Rackspace Cloud Network, on page 3](#)
- [Rackspace Day 0 Configuration, on page 4](#)
- [Deploy the ASAv , on page 6](#)
- [CPU Usage and Reporting, on page 7](#)

Overview

Rackspace is a leading provider of expertise and managed services across all the major public and private cloud technologies. The Rackspace Cloud is a set of cloud computing products and services billed on a utility computing basis.

You can deploy the ASAv for Rackspace as a virtual appliance in the Rackspace cloud. This chapter explains how to install and configure a single instance ASAv appliance.

Instance types in the Rackspace Cloud are referred to as *flavors*. The term flavor refers to a server's combination of RAM size, vCPUs, network throughput (RXTX factor), and disk space. The following table lists Rackspace flavors suitable for ASAv deployment.

Table 1: Rackspace Supported Flavors

Flavor	Attributes		Aggregate Bandwidth
	vCPUs	Memory (GB)	
general 1-2	2	2	400 Mbps
general 1-4	4	4	800 Mbps
general 1-8	8	8	1.6 Gbps

Flavor	Attributes		Aggregate Bandwidth
	vCPUs	Memory (GB)	
compute 1-4	2	3.75	312.5 Mbps
compute 1-8	4	7.5	625 Mbps
compute 1-15	8	15	1.3 Gbps
memory 1-15	2	15	625 Mbps
memory 1-15	4	30	1.3 Gbps
memory 1-15	8	60	2.5 Gbps

About Rackspace Flavors

Rackspace Virtual Cloud Server Flavors fall into the following classes:

- **General Purpose v1**

- Useful for a range of use cases, from general-purpose workloads to high performance websites.
- The vCPUs are oversubscribed and “burstable”; in other words, there are more vCPUs allocated to the Cloud Servers on a physical host than there are physical CPU threads.

- **Compute v1**

- Optimized for web servers, application servers, and other CPU-intensive workloads.
- The vCPUs are “reserved”; in other words, there are never more vCPUs allocated to the Cloud Servers on a physical host than there are physical CPU threads on that host.

- **Memory v1**

- Recommended for memory-intensive workloads.

- **I/O v1**

- Ideal for high performance applications and databases that benefit from fast disk I/O.

Prerequisites

- Create a [Rackspace](#) account.

All Rackspace Public Cloud accounts are set to the Managed Infrastructure service level by default. You can upgrade to the Managed Operations service level inside the Cloud Control Panel. At the top of the Cloud Control Panel, click your account username and then select Upgrade Service Level.

- License the ASA. Until you license the ASA, it will run in degraded mode, which allows only 100 connections and throughput of 100 Kbps. See [Licensing for the ASA](#).
- Interface requirements:

- Management interface
- Inside and outside interfaces
- (Optional) Additional subnet (DMZ)
- Communications paths:
 - Management interface—Used to connect the ASAv to the ASDM; can't be used for through traffic.
 - Inside interface (required)—Used to connect the ASAv to inside hosts.
 - Outside interface (required)—Used to connect the ASAv to the public network.
 - DMZ interface (optional)—Used to connect the ASAv to the DMZ network.
- For ASA and ASAv system compatibility and requirements, see [Cisco ASA Compatibility](#).

Rackspace Cloud Network

Your cloud configuration can include several kinds of networks, connected as appropriate for your needs. You can manage the networking capabilities of your cloud servers in many of the same ways you manage your other networks. Your ASAv deployment will interact primarily with three types of virtual networks in the Rackspace Cloud:

- **PublicNet**—Connects cloud infrastructure components such as cloud servers, cloud load balancers, and network appliances to the Internet.
 - Use PublicNet to connect the ASAv to the Internet.
 - The ASAv attaches to this network via the Management0/0 interface.
 - PublicNet is dual-stacked for IPv4 and IPv6. When you create a server with PublicNet, your server receives an IPv4 address and an IPv6 address by default.
- **ServiceNet**—An internal, IPv4-only multi-tenant network within each Rackspace cloud region.
 - ServiceNet is optimized to carry traffic across servers within your configuration (east-west traffic).
 - It provides servers with no-cost access to regionalized services such as Cloud Files, Cloud Load Balancers, Cloud Databases, and Cloud Backup.
 - The networks 10.176.0.0/12 and 10.208.0.0/12 are reserved for ServiceNet. Any servers that have ServiceNet connectivity will be provisioned with an IP address from one of these networks.
 - The ASAv attaches to this network via the Gigabit0/0 interface.
- **Private Cloud Networks**—Cloud Networks lets you create and manage secure, isolated networks in the cloud.
 - These networks are fully single tenant, and you have complete control over the network topology, IP addressing (IPv4 or IPv6), and which Cloud Servers are attached.
 - Cloud Networks are regional in scope, and you can attach them to any of your Cloud Servers in a given region.

- You can create and manage Cloud Networks via an API or by using the Rackspace Cloud Control Panel.

The ASAv attaches to these networks via Gigabit0/1 through Gigabit0/8 interfaces.

Rackspace Day 0 Configuration

When a VM is deployed in the Rackspace Cloud, a CD-ROM device containing files with Rackspace provisioning information is attached to the VM. The provisioning information includes:

- The hostname
- IP addresses for required interfaces
- Static IP routes
- Username and password (Optional SSH public key)
- DNS servers
- NTP servers

These files are read during the initial deployment and ASA configuration is generated.

ASAv Hostname

By default, the ASAv hostname is the name you assign to your cloud server when you begin to build your ASAv.

```
hostname rackspace-asav
```

The ASA hostname configuration will only accept a hostname that complies with RFCs 1034 and 1101:

- Must start and end with a letter or digit.
- Interior characters must be a letter, a digit or a hyphen.



Note The ASAv will alter the cloud server name to comply with these rules while making it as close as possible to the original cloud server name. It will drop special characters from the beginning and end of the cloud server name, and replace non-compliant interior characters with a hyphen.

For example, a cloud server named **ASAv-9.13.1.200** will have hostname **ASAv-9-13-1-200**.

Interfaces

Interfaces are configured in the following manner:

- Management0/0
 - Named 'outside' because it is connected to the PublicNet.
 - Rackspace assigns both IPv4 and IPv6 public addresses to the PublicNet interface.

- Gigabit0/0
 - Named 'management' since it is connected to the ServiceNet.
 - Rackspace assigns an IPv4 address from the ServiceNet subnet for the Rackspace region.
- Gigabit0/1 through Gigabit0/8
 - Named 'inside', 'inside02', 'inside03', etc. because they are connected to private Cloud Networks.
 - Rackspace assigns an IP address from the Cloud Network subnet.

The interface configuration for an ASA with 3 interfaces would look something like this:

```
interface GigabitEthernet0/0
  nameif management
  security-level 0
  ip address 10.176.5.71 255.255.192.0
!
interface GigabitEthernet0/1
  nameif inside
  security-level 100
  ip address 172.19.219.7 255.255.255.0
!
interface Management0/0
  nameif outside
  security-level 0
  ip address 162.209.103.109 255.255.255.0
  ipv6 address 2001:4802:7800:1:be76:4eff:fe20:1763/64
```

Static Routes

Rackspace provisions the following static IP routes:

- Default IPv4 route via PublicNet interface (**outside**).
- Default IPv6 route via PublicNet interface.
- Infrastructure subnet routes on ServiceNet interface (**management**).

```
route outside 0.0.0.0 0.0.0.0 104.130.24.1 1
ipv6 route outside ::/0 fe80::def
route management 10.176.0.0 255.240.0.0 10.176.0.1 1
route management 10.208.0.0 255.240.0.0 10.176.0.1 1
```

Login Credentials

A username 'admin' is created with a password created by Rackspace. A public key for user 'admin' is created if the cloud server is deployed with a Rackspace Public Key.

```
username admin password <admin_password> privilege 15
username admin attributes
  ssh authentication publickey <public_key>
```

The Day0 SSH configuration:

- SSH via PublicNet interface (**outside**) is enabled for IPv4 and IPv6.
- SSH via ServiceNet interface (**management**) is enabled for IPv4 .
- Configure stronger key exchange group at request of Rackspace.

```
aaa authentication ssh console LOCAL
ssh 0 0 management
ssh 0 0 outside
ssh ::0/0 outside
ssh version 2
ssh key-exchange group dh-group14-sha1
```

DNS and NTP

Rackspace provides two IPv4 service addresses to be used for DNS and NTP.

```
dns domain-lookup outside
dns server-group DefaultDNS
name-server 69.20.0.164
name-server 69.20.0.196

ntp server 69.20.0.164
ntp server 69.20.0.196
```

Deploy the ASAv

You can deploy the ASAv as a virtual appliance in the Rackspace Cloud. This procedure shows you how to install a single instance ASAv appliance.

Before you begin

Review the [Rackspace Day 0 Configuration, on page 4](#) topic for a description of the configuration parameters that the Rackspace Cloud enables for a successful ASAv deployment, including hostname requirement, interface provisioning, and networking information.

Step 1 On the Rackspace mycloud portal, go to **SERVERS > CREATE RESOURCES > Cloud Server**.

Step 2 On the **Create Server** page, enter your **Server Details**:

- Enter the name for your ASAv machine in the **Server Name** field.
- Choose your region from the **Region** drop-down list.

Step 3 Under **Image**, choose **Linux/Appliances > ASAv > Version**.

Note You would typically choose the most recent supported version when deploying a new ASAv.

Step 4 Under **Flavor**, choose a **Flavor Class** that fits your resource needs; see [Table 1: Rackspace Supported Flavors, on page 1](#) for a list of suitable VMs.

Important Beginning with 9.13(1), the minimum memory requirement for the ASAv is 2GB. When deploying an ASAv with more than 1 vCPU, the minimum memory requirement for the ASAv is 4GB.

- Step 5** (Optional) Under **Advanced Options**, configure an SSH key.
See [Managing access with SSH keys](#) for complete information on SSH keys in the Rackspace Cloud.
- Step 6** Review any applicable **Recommended Installs** and **Itemized Charges** for your ASA, then click **Create Server**.
The root admin password displays. Copy the password, then dismiss the dialog.
- Step 7** After you create the server, the server details page displays. Wait for the server to show an active status. This usually takes a few minutes.
-

What to do next

- Connect to the ASA.
- Continue configuration using CLI commands available for input via SSH or use ASDM. See [Start ASDM](#) for instructions for accessing the ASDM.

CPU Usage and Reporting

The CPU Utilization report summarizes the percentage of the CPU used within the time specified. Typically, the Core operates on approximately 30 to 40 percent of total CPU capacity during nonpeak hours and approximately 60 to 70 percent capacity during peak hours.

vCPU Usage in the ASA Virtual

The ASA virtual vCPU usage shows the amount of vCPUs used for the data path, control point, and external processes.

The Rackspace reported vCPU usage includes the ASA virtual usage as described plus:

- ASA virtual idle time
- %SYS overhead used for the ASA virtual machine
- Overhead of moving packets between vSwitches, vNICs, and pNICs. This overhead can be quite significant.

CPU Usage Example

The `show cpu usage` command can be used to display CPU utilization statistics.

Example

```
Ciscoasa#show cpu usage
```

```
CPU utilization for 5 seconds = 1%; 1 minute: 2%; 5 minutes: 1%
```

The following is an example in which the reported vCPU usage is substantially different:

- ASA Virtual reports: 40%

- DP: 35%
- External Processes: 5%
- ASA (as ASA Virtual reports): 40%
- ASA idle polling: 10%
- Overhead: 45%

The overhead is used to perform hypervisor functions and to move packets between NICs and vNICs using the vSwitch.

Rackspace CPU Usage Reporting

In addition to viewing CPU, RAM, and disk space configuration information for available Cloud Servers, you can also view disk, I/O, and networking information. Use this information to help you decide which Cloud Server is right for your needs. You can view the available servers through either the command-line nova client or the [Cloud Control Panel](#) interface.

On the command line, run the following command:

```
nova flavor-list
```

All available server configurations are displayed. The list contains the following information:

- ID - The server configuration ID
- Name - The configuration name, labeled by RAM size and performance type
- Memory_MB - The amount of RAM for the configuration
- Disk - The size of the disk in GB (for general purpose Cloud Servers, the size of the system disk)
- Ephemeral - The size of the data disk
- Swap - The size of the swap space
- VCPUs - The number of virtual CPUs associated with the configuration
- RXTX_Factor - The amount of bandwidth, in Mbps, allocated to the PublicNet ports, ServiceNet ports, and isolated networks (cloud networks) attached to a server
- Is_Public - Not used

ASA Virtual and Rackspace Graphs

There are differences in the CPU % numbers between the ASA Virtual and Rackspace:

- The Rackspace graph numbers are always higher than the ASA Virtual numbers.
- Rackspace calls it %CPU usage; the ASA Virtual calls it %CPU utilization.

The terms “%CPU utilization” and “%CPU usage” mean different things:

- CPU utilization provides statistics for physical CPUs.

- CPU usage provides statistics for logical CPUs, which is based on CPU hyperthreading. But because only one vCPU is used, hyperthreading is not turned on.

Rackspace calculates the CPU % usage as follows:

Amount of actively used virtual CPUs, specified as a percentage of the total available CPUs

This calculation is the host view of the CPU usage, not the guest operating system view, and is the average CPU utilization over all available virtual CPUs in the virtual machine.

For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely. The virtual CPU usage calculation is Usage in MHz / number of virtual CPUs x core frequency

