



Deploy the ASA Virtual on GCP

You can deploy the ASA Virtual on the Google Cloud Platform (GCP).

- [Overview, on page 1](#)
- [Prerequisites, on page 4](#)
- [Guidelines and Limitations, on page 5](#)
- [Sample Network Topology, on page 6](#)
- [Deploy the ASA Virtual on GCP, on page 6](#)
- [Access the ASA Virtual Instance on GCP, on page 9](#)
- [CPU Usage and Reporting, on page 11](#)

Overview

GCP lets you build, deploy, and scale applications, websites, and services on the same infrastructure as Google.

The ASA Virtual runs the same software as physical ASAs to deliver proven security functionality in a virtual form factor. The ASA Virtual can be deployed in the public GCP. It can then be configured to protect virtual and physical data center workloads that expand, contract, or shift their location over time.

GCP Machine Types and ASA Virtual Versions

Select the Google virtual machine type and size to meet your ASA Virtual needs.

The ASA Virtual supports the following *General-purpose N1, N2* and *Compute-optimized C2* GCP machine types:

Table 1: GCP Supported Compute-Optimized Machine Types

Compute-Optimized Machine types	Attributes		vNICs	ASA Virtual Version
	vCPUs	Memory (GB)		
c2-standard-4	4	16	3	9.15.1 or above
c2-standard-8	8	32	3	9.15.1 or above
c2-standard-16	16	64	3	9.15.1 or above

GCP supports the ASA Virtual on the following GCP machine types:

- c2-standard-4 (ASAv5, ASAv10, and ASAv30)
- c2-standard-8 (ASAv50)
- c2-standard-16 (ASAv100)

Table 2: GCP Supported General Purpose Machine Types

Machine type	Attributes		vNICs	ASA Virtual Version
	vCPUs	Memory (GB)		
n1-standard-4	4	15	3	9.15.1 or above
n1-standard-8	8	30	3	9.15.1 or above
n1-standard-16	16	60	3	9.15.1 or above
n2-standard-4	4	16	3	9.15.1 or above
n2-standard-8	8	32	3	9.15.1 or above
n2-standard-16	16	64	3	9.15.1 or above
n2-highmem-4	4	32	3	9.15.1 or above
n2-highmem-8	8	64	3	9.15.1 or above
e2-standard-4	4	8	3	9.15.1 or above
e2-standard-8	8	16	3	9.15.1 or above
e2-standard-16	16	32	3	9.15.1 or above
e2-highmem-4	4	8	3	9.15.1 or above
e2-highmem-8	8	16	3	9.15.1 or above
e2-highmem-16	16	32	3	9.15.1 or above
e2-highcpu-4	4	8	3	9.15.1 or above
e2-highcpu-8	8	16	3	9.15.1 or above
e2-highcpu-16	16	32	3	9.15.1 or above
n1-highmem-4	4	8	3	9.15.1 or above
n1-highmem-8	8	16	3	9.15.1 or above
n1-highmem-16	16	32	3	9.15.1 or above
n2d-standard-4	4	8	3	9.15.1 or above
n2d-standard-8	8	16	3	9.15.1 or above
n2d-standard-16	16	32	3	9.15.1 or above

Machine type	Attributes		vNICs	ASA Virtual Version
	vCPUs	Memory (GB)		
c2d-standard-4	4	8	3	9.15.1 or above
c2d-standard-8	8	16	3	9.15.1 or above
c2d-standard-16	16	32	3	9.15.1 or above
t2d-standard-4	4	16	4	9.24.10 or above
t2d-standard-8	8	32	8	9.24.10 or above
t2d-standard-16	16	64	8	9.24.10 or above

- The ASA Virtual requires a minimum of 3 interfaces.
- The maximum supported vCPUs is 16.
- The Memory-Optimized machine type is not supported.

You create an account on GCP, launch an ASA Virtual instance using the ASA virtual firewall (ASA Virtual) offering on the GCP Marketplace, and choose a GCP machine type.



Note

- ASA virtual maximum vCPU limit: ASA virtual supports a maximum of 16 vCPUs. Instances exceeding this capacity are not supported. ASAvU (ASAv-Unlimited) is supported only on KVM, VMware, and AWS.
- ASA virtual minimum vCPU requirement: ASA virtual requires a minimum of 4 vCPUs. Instances with fewer than 4 vCPUs are not supported.
- By default, ASA Virtual instances are deployed using UEFI mode.
- User can enable Secure Boot while deploying the ASA Virtual instance.
- ASA virtual maximum vCPU limit: ASA virtual supports a maximum of 16 vCPUs. Instances exceeding this capacity are not supported. ASAvU (ASAv-Unlimited) is supported only on KVM, VMware, and AWS.
- ASA virtual minimum vCPU requirement: ASA virtual requires a minimum of 4 vCPUs. Instances with fewer than 4 vCPUs are not supported.

C2 Compute-Optimized Machine Type Limitations

The Compute-Optimized C2 machine types have the following restrictions:

- You cannot use regional persistent disks with compute-optimized machine types. For more information, see the Google documentation [Adding or resizing regional persistent disks](#).
- Subject to different disk limits than general-purpose and memory-optimized machine types. For more information, see the Google documentation [Block storage performance](#).

- Available only in select zones and regions. For more information, see the Google documentation [Available regions and zones](#).
- Available only on select CPU platforms. For more information, see the Google documentation [CPU platforms](#).

Performance Tiers for ASA Virtual

The ASA Virtual supports performance-tiered licensing that provides different throughput levels and VPN connection limits based on deployment requirements.

Performance Tier	Machine type (Core/RAM)	Rate Limit	RA VPN Session Limit
ASAv5	c2-standard-4 4 core/16 GB	100 Mbps	50
ASAv10	c2-standard-4 4 core/16 GB	1 Gbps	250
ASAv30	c2-standard-4 4 core/16 GB	2 Gbps	750
ASAv50	c2-standard-8 8 core/32 GB	7.6 Gbps	10,000
ASAv100	c2-standard-16 16 core/64 GB	16 Gbps	20,000

Prerequisites

- Create a GCP account at <https://cloud.google.com>.
- Create your GCP project. See the Google documentation, [Creating Your Project](#).
- License the ASA Virtual. Until you license the ASA Virtual, it will run in degraded mode, which allows only 100 connections and throughput of 100 Kbps. See [Licenses: Smart Software Licensing](#).
- Interface requirements:
 - Management interface—Used to connect the ASA Virtual to the ASDM; can't be used for through traffic.
 - Inside interface—Used to connect the ASA Virtual to inside hosts.
 - Outside interface—Used to connect the ASA Virtual to the public network.
- Communications paths:
 - Public IPs for access into the ASA Virtual.
- For ASA Virtual system requirements, see [Cisco Secure Firewall ASA Compatibility](#).

Guidelines and Limitations

Supported Features

The ASA Virtual on GCP supports the following features:

- Deployment in the GCP Virtual Private Cloud (VPC)
- Maximum of 16 vCPUs per instance
- Routed mode (default)
- Licensing – Only BYOL is supported

Unsupported Features

The ASA Virtual on GCP does not support the following:

- IPv6
 - Instance-level IPv6 setting is not supported on GCP
 - Only the load balancer can accept IPv6 connections, and proxy them over IPv4 to GCP Instances
- Jumbo Frames
- ASA Virtual native HA
- Autoscale
- Transparent/inline/passive modes

UEFI and Secure Boot Limitations

On GCP, UEFI firmware is supported only for greenfield (fresh) deployments and must be enabled during initial deployment.

Existing brownfield deployments can be upgraded to Version 9.24 without impact. Changing the boot mode or enabling UEFI Secure Boot after deployment is not supported.

Upgrade Restrictions and Limitations

Revert upgrade restrictions




Caution Reverting an upgrade is not supported.

Ensure that you take a backup before upgrading. After upgrading to ASA Virtual 9.24, you cannot roll back to a previous software version. To return to an earlier version, you must redeploy the Management Center.

Sample Network Topology

The following figure shows the recommended network topology for the ASA Virtual in Routed Firewall Mode with 3 subnets configured in GCP for the ASA Virtual (management, inside, and outside).

Figure 1: Sample ASA Virtual on GCP Deployment



Deploy the ASA Virtual on GCP

You can deploy the ASA Virtual on the Google Cloud Platform (GCP).

Create VPC Networks

Before you begin

The ASA Virtual deployment requires three networks which you must create prior to deploying the ASA Virtual. The networks are as follows:

- Management VPC for the management subnet.
- Inside VPC for the inside subnet.
- Outside VPC for the outside subnet.

Additionally, you set up the route tables and GCP firewall rules to allow traffic flow through the ASA Virtual. The route tables and firewall rules are separate from those that are configured on the ASA Virtual itself. Name the GCP route tables and firewall rules according to associated network and functionality. See [Sample Network Topology](#), on page 6.

Procedure

-
- Step 1** In the GCP console, choose **Networking > VPC network > VPC networks**, then click **Create VPC Network**.
 - Step 2** In the **Name** field, enter the descriptive name for your VPC network, for example, *vpc-asiasouth-mgmt*.
 - Step 3** From the **Subnet creation mode**, click **Custom**.
 - Step 4** In the **Name** field under **New subnet**, enter the desired name, for example, *vpc-asiasouth-mgmt*.
 - Step 5** From the **Region** drop-down list, select the region appropriate for your deployment. All three networks must be in the same region.
 - Step 6** In the **IP address range** field, enter the first network's subnet in CIDR format, such as 10.10.0.0/24.
 - Step 7** Accept the defaults for all other settings, then click **Create**.
 - Step 8** Repeat steps 1-7 to create the remaining two networks in your VPC.
-

Create the Firewall Rules

You apply the firewall rules for the management interface (to allow SSH and HTTPS connections) while deploying the ASA Virtual instance, see [Create the ASA Virtual Instance on GCP, on page 7](#). According to your requirements, you can also create firewall rules for the inside and outside interfaces.

Procedure

-
- Step 1** In the GCP console, choose **Networking > VPC network > Firewall**, then click **Create Firewall Rule**.
 - Step 2** In the **Name** field, enter a descriptive name for your firewall rule, for example, *vpc-asiasouth-inside-fwrule*.
 - Step 3** From the **Network** drop-down list, select the name of the VPC network for which you are creating the firewall rule, for example, *asav-south-inside*.
 - Step 4** From the **Targets** drop-down list, select the option applicable for your firewall rule, for example, **All instances in the network**.
 - Step 5** In the **Source IP ranges** field, enter the source IP address ranges in CIDR format, for example, 0.0.0.0/0.
Traffic is only allowed from sources within these IP address ranges.
 - Step 6** Under **Protocols and ports**, select **Specified protocols and ports**.
 - Step 7** Add your security rules.
 - Step 8** Click **Create**.
-

Create the ASA Virtual Instance on GCP

Complete the following steps to deploy an ASA Virtual instance using the Cisco ASA virtual firewall (ASA Virtual) offering from the GCP Marketplace.

Procedure

- Step 1** Log into to the [GCP Console](#).
- Step 2** Click **Navigation menu** > **Marketplace**.
- Step 3** Search the Marketplace for “Cisco ASA virtual firewall (ASAv)” and choose the offering.
- Step 4** Click **Launch**.
- Step 5** Add a unique **Deployment name** for the instance.
- Step 6** Select the **Zone** where you want to deploy the ASA Virtual.
- Step 7** Select the appropriate **Machine type**. For a list of supported machine types, see [Overview, on page 1](#).
- Step 8** (Optional) Paste the public key from the SSH key pair under **SSH key (optional)**.
- The key pair consists of a public key that GCP stores and a private key file that the user stores. Together they allow you to connect to your instance securely. Be sure to save the key pair to a known location, as it will be required to connect to the instance.
- Step 9** Choose whether to allow or block the project-wide SSH keys to access this instance. See the Google documentation [Allowing or blocking project-wide public SSH keys from a Linux instance](#).
- Step 10** (Optional) Under **Startup script**, provide the day0 configuration for your ASA Virtual. The day0 configuration is applied during the firstboot of the ASA Virtual.
- The following example shows a sample day0 configuration you can copy and paste in the **Startup script** field:
- See the [ASA Configuration Guides](#) and the [ASA Command Reference](#) for complete information on the ASA commands.
- Important**
- When you copy text from this example, you should validate the script in a third-party text editor or validation engine to prevent format errors and remove invalid Unicode characters.

```
!ASA Version 9.15.1

interface management0/0

management-only
nameif management
security-level 100
ip address dhcp setroute
no shut
!
same-security-traffic permit inter-interface
same-security-traffic permit intra-interface
!
crypto key generate rsa modulus 2048
ssh 0 0 management
ssh timeout 60
ssh version 2
username admin password cisco123 privilege 15
username admin attributes
service-type admin
! required config end
dns domain-lookup management
dns server-group DefaultDNS
name-server 8.8.8.8
```

Step 11 Keep the default **Boot disk type** and **Boot disk size in GB** for the provisioned disk space.

Step 12 Configure the interfaces under **Network interfaces**.

- management
- inside
- outside

Note

You cannot add interfaces to an instance after you create it. If you create the instance with an improper interface configuration, you must delete the instance and recreate it with the proper interface configuration.

- From the **Network** drop-down list, select a VPC network, for example, *vpc-asiasouth-mgmt*.
- From the **External IP** drop-down list, select the appropriate option.

For the management interface, select the **External IP** to **Ephemeral**. This is optional for inside and outside interfaces.

- Click **Done**.

Step 13 Apply the firewall rules under **Firewall**.

- Check the **Allow TCP port 22 traffic from the Internet (SSH access)** check box to allow SSH.
- Check the **Allow HTTPS traffic from the Internet (ASDM access)** check box to allow HTTPS connections.

Step 14 Click **More** to expand the view and make sure that **IP Forwarding** is set to **On**.

Step 15 Click **Deploy**.

View the instance details from the VM instance page of the GCP console. You'll find the internal IP address, external IP address, and controls to stop and start the instance. You need to stop the instance if you need to edit it.

Access the ASA Virtual Instance on GCP

Make sure that you have already enabled a firewall rule to allow SSH (TCP connections through port 22) during deployment. See [Create the ASA Virtual Instance on GCP, on page 7](#) for more information.

This firewall rule enables access to the ASA Virtual instance and allows you to connect to the instance using the following methods.

- External IP
 - Any other SSH client or third-party tools
- Serial console
- Gcloud command line

See the Google documentation, [Connecting to instances](#) for more information.



Note You can log in to the ASA Virtual instance using the credentials specified in the day0 configuration, or by using the SSH key pair you created during the instance launch.

Connect to the ASA Virtual Instance Using an External IP

The ASA Virtual instance is assigned with an internal IP and an external IP. You can use the external IP to access the ASA Virtual instance.

Procedure

- Step 1** In the GCP console, choose **Compute Engine** > **VM instances**.
- Step 2** Click the ASA Virtual instance name to open the **VM instance details** page.
- Step 3** Under the **Details** tab, click the drop-down menu for the **SSH** field.
- Step 4** Select the desired option from the **SSH** drop-down menu.

You can connect to the ASA Virtual instance using the following method.

- Any other SSH client or third-party tools—See the Google documentation, [Connecting using third-party tools](#) for more information.

Note

You can log in to the ASA Virtual instance using the credentials specified in the day0 configuration, or by using the SSH key pair you created during the instance launch.

Connect to the ASA Virtual Instance Using SSH

To connect to the ASA Virtual instance from a Unix-style system, log in to the instance using SSH.

Procedure

- Step 1** Use the following command to set the file permissions so that only you can read the file:

```
$ chmod 400 <private_key>
```

Where:

<private_key> is the full path and name of the file that contains the private key associated with the instance you want to access.

- Step 2** Use the following SSH command to access the instance.

```
$ ssh -i <private_key> <username>@<public-ip-address>
```

Where:

<private_key> is the full path and name of the file that contains the private key associated with the instance you want to access.

<username> is the username for the ASA Virtual instance.

<public-ip-address> is your instance IP address that you retrieved from the Console.

<ipv6-address> is your instance management interface IPv6 address.

Connect to the ASA Virtual Instance Using the Serial Console

Procedure

- Step 1** In the GCP console, choose **Compute Engine > VM instances**.
 - Step 2** Click the ASA Virtual instance name to open the **VM instance details** page.
 - Step 3** Under the **Details** tab, click **Connect to serial console**.
See the Google documentation, [Interacting with the serial console](#) for more information.
-

Connect to the ASA Virtual Instance Using Gcloud

Procedure

- Step 1** In the GCP console, choose **Compute Engine > VM instances**.
 - Step 2** Click the ASA Virtual instance name to open the **VM instance details** page.
 - Step 3** Under the **Details** tab, click the drop-down menu for the **SSH** field.
 - Step 4** Click **View gcloud command > Run in Cloud Shell**.
The Cloud Shell terminal window opens. See the Google documentation, [gcloud command-line tool overview](#), and [gcloud compute ssh](#) for more information.
-

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 GCP reported vCPU usage includes the ASA virtual usage as described:

- 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.

GCP CPU Usage Reporting

Click the instance name on GCP Console and then click on the tab **Monitoring**. You will be able to see the CPU usage percentage.

Compute Engine lets you export detailed reports of your Compute Engine usage to a [Google Cloud Storage](#) bucket using the usage export feature. Usage reports provide information about the lifetime of your resources. For example, you can see how many VM instances in your project are running an `n2-standard-4` machine type and how long each instance has been running. You can also review the storage space of a persistent disk, and information about other Compute Engine features.

ASA Virtual and GCP Graphs

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

- The GCP graph numbers are always higher than the ASA Virtual numbers.
- GCP 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.

GCP 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 $\text{Usage in MHz} / \text{number of virtual CPUs} \times \text{core frequency}$

