



Deploy the Management Center Virtual On the AWS Cloud

Amazon Virtual Private Cloud (Amazon VPC) enables you to launch Amazon Web Services (AWS) resources into a virtual network that you define. This virtual network closely resembles a traditional network that might operate in your own data center, with the benefits of using the scalable infrastructure of AWS.

You can deploy the management center virtual on the AWS Cloud.

- [Overview, on page 1](#)
- [Guidelines and Limitations, on page 3](#)
- [Configure the AWS Environment, on page 5](#)
- [Deploy the Management Center Virtual, on page 9](#)

Overview

Management Center Virtual Requires 28 GB RAM for Upgrade (6.6.0+)

The management center virtual platform has introduced a new memory check during upgrade. The management center virtual upgrades to Version 6.6.0+ will fail if you allocate less than 28 GB RAM to the virtual appliance.



Important

As of the Version 6.6.0 release, lower-memory instance types for cloud-based management center virtual deployments (AWS, Azure) are fully deprecated. You cannot create the new management center virtual instances using them, even for earlier versions. You can continue running existing instances. See [Table 1: AWS Supported Instances for the Management Center Virtual, on page 2](#).

As a result of this memory check, we will not be able to support lower memory instances on supported platforms.

The following table summarizes the AWS instances types that the management center virtual supports; those that Versions 6.5.x and earlier support, and those that Version 6.6.0+ support.



Note

Version 6.6 adds support for the C5 instance types shown in the following table. Larger instance types provide more CPU resources to your AWS VMs for increased performance, and some allow for more network interfaces.

Table 1: AWS Supported Instances for the Management Center Virtual

Platform	Version 6.6.0+	vCPUs	Memory (GB)	Maximum Number of Interfaces	Version 6.5.x and earlier	vCPUs	Memory (GB)	Maximum Number of Interfaces
Management Center Virtual	c3.4xlarge	16	30	8	c3.xlarge*	4	7.5	4
	c4.4xlarge	16	30	8	c3.2xlarge*	8	15	4
	c5.4xlarge	16	32	8	c3.4xlarge	16	30	8
	—	—	—	—	c4.xlarge*	4	7.5	4
	—	—	—	—	c4.2xlarge*	8	15	4
	—	—	—	—	c4.4xlarge	16	30	8
	*Note that the management center virtual will not support these instance types on Version 6.6.0 and above. Beginning with Version 6.6.0, you must deploy the management center virtual (any version) using an instance with at least 28 GB RAM. See Deprecated Instances and Resizing Instances, on page 2 for more information.							

Table 2: AWS Supported Instances for the Management Center Virtual 300

Platform	Version 7.1.0+
Management Center Virtual 300 (FMCv300)	c5.9xlarge: 36 vCPUs, 72 GB SSD storage: 2000 GB

Deprecated Instances

You can continue running your current Version 6.5.x and earlier management center virtual deployments, but you will not be able to launch the new management center virtual deployments (any version) using these instances:

- c3.xlarge—4 vCPUs, 7.5 GB (DISABLED for the management center virtual after Version 6.6.0+)
- c3.2xlarge—8 vCPUs, 15 GB (DISABLED for the management center virtual after Version 6.6.0+)
- c4.xlarge—4 vCPUs, 7.5 GB (DISABLED for the management center virtual after Version 6.6.0+)
- c4.2xlarge—8 vCPUs, 15 GB (DISABLED for the management center virtual after Version 6.6.0+)

Resizing Instances

Because the upgrade path from any earlier version of management center virtual (6.2.x, 6.3.x, 6.4.x, and 6.5.x) to Version 6.6.0 includes the 28 GB RAM memory check, you need to resize your current instance type to one that supports Version 6.6.0 (see [Table 1: AWS Supported Instances for the Management Center Virtual, on page 2](#)).

You can resize an instance if the current instance type and the new instance type that you want are compatible. For the management center virtual deployments:

- Resize any c3.xlarge or c3.2xlarge to the c3.4xlarge instance type.
- Resize any c4.xlarge or c4.2xlarge to the c4.4xlarge instance type.

Be aware of the following before resizing your instance:

- You must stop your instance before you change instance types.
- Verify that your current instance type is compatible with the new instance type that you choose.
- If this instance has an instance store volume, any data on it is lost when the instance is stopped. Migrate your instance store-backed instance before you resize.
- If you're not using an Elastic IP address, the public IP address is released when you stop the instance.

For instructions on how to resize your instance, see the AWS documentation “Changing the Instance Type” (<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html>).

AWS Solution Overview

AWS is a collection of remote computing services offered by Amazon.com, also called web services, that make up a cloud-computing platform. These services operate from 11 geographical regions across the world. In general, you should become familiar with the following AWS services when deploying the management center virtual:

- Amazon Elastic Compute Cloud (EC2)—a web service that enables you to rent virtual computers to launch and manage your own applications and service, such as a firewall, in Amazon's data centers.
- Amazon Virtual Private Cloud (VPC)—a web service that enables you to configure an isolated private network that exists within the Amazon public cloud. You run your EC2 instances within a VPC.
- Amazon Simple Storage Service (S3)—a web service that provides you with a data storage infrastructure.

You create an account on AWS, set up the VPC and EC2 components (using either the AWS Wizards or manual configuration), and chose an Amazon Machine Image (AMI) instance. The AMI is a template that contains the software configuration needed to launch your instance.



Note The AMI images are not available for download outside of the AWS environment.

Guidelines and Limitations

Supported Features (7.1.0+)

- **Management Center Virtual 300 (FMCv300) for AWS**—A new scaled management center virtual image is available on the AWS platform that supports managing up to 300 devices and has higher disk capacity.
- Management Center Virtual high availability (HA) is supported.

Prerequisites

The following prerequisites pertain to the management center virtual on AWS:

- An Amazon account. You can create one at aws.amazon.com.
- A Cisco Smart Account. You can create one at Cisco Software Central (<https://software.cisco.com/>).
- License the management center virtual. See [Management Center Virtual Licenses](#) for general guidelines about virtual platform licenses; see “Licensing the System” in the *Firepower Management Center Configuration Guide* for more detailed information about how to manage licenses.
- The management center virtual interface requirements:
 - Management interface.
- Communication Paths:
 - Public/elastic IPs for access into the management center virtual.
- For the management center virtual and System compatibility, see the [Cisco Firepower Compatibility Guide](#).

Guidelines

The following guidelines pertain to the management center virtual on AWS:

- Deployment in the Virtual Private Cloud (VPC).
- Enhanced networking (SR-IOV) where available.
- Deployment from Amazon Marketplace.
- Maximum of four vCPUs per instance.
- User deployment of L3 networks.
- IPv6 is supported.

Limitations

The following limitations pertain to the management center virtual on AWS:

- The management center virtual appliances do not have serial numbers. The **System > Configuration** page will show either **None** or **Not Specified** depending on the virtual platform.
- Any IP address configuration (either from CLI or management center) must match what is created in the AWS console; you should note your configurations during deployment.
- You cannot add interfaces after boot.
- Cloning/snapshots are currently not supported.
- Transport Layer Security (TLS) Server Identity Discovery is not supported with Geneve single-arm setup on AWS.

Configure the AWS Environment

To deploy the management center virtual on AWS you need to configure an Amazon VPC with your deployment-specific requirements and settings. In most situations a setup wizard can guide you through your setup. AWS provides online documentation where you can find useful information about the services ranging from introductions to advanced features. See [Getting Started with AWS](#) for more information.

For greater control over your AWS setup, the following sections offer a guide to your VPC and EC2 configurations prior to launching instances of the management center virtual:

- [Create the VPC, on page 5](#)
- [Add the Internet Gateway, on page 6](#)
- [Add Subnets, on page 6](#)
- [Add a Route Table, on page 7](#)
- [Create a Security Group, on page 7](#)
- [Create Network Interfaces, on page 8](#)
- [Create Elastic IPs, on page 9](#)

Create the VPC

A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as management center virtual instances, into your VPC. You can configure your VPC; you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings.

Before you begin

- Create your AWS account.
- Confirm that AMIs are available for the management center virtual instances.

-
- Step 1** Log into aws.amazon.com and choose your region.
- AWS is divided into multiple regions that are isolated from each other. The region is displayed in the upper right corner of your screen. Resources in one region do not appear in another region. Check periodically to make sure you are in the intended region.
- Step 2** Click **Services > VPC**.
- Step 3** Click **VPC Dashboard > Your VPCs**.
- Step 4** Click **Create VPC**.
- Step 5** Enter the following in the **Create VPC** dialog box:
- a) A user-defined **Name tag** to identify the VPC.
 - b) A **CIDR block** of IP addresses. CIDR (Classless Inter-Domain Routing) notation is a compact representation of an IP address and its associated routing prefix. For example, 10.0.0.0/24.

- c) A **Tenancy** setting of Default to ensure that instances launched in this VPC use the tenancy attribute specified at launch.

Step 6 Click **Yes, Create** to create your VPC.

What to do next

Add an Internet gateway to your VPC as described in the next section.

Add the Internet Gateway

You can add an Internet gateway to connect your VPC to the Internet. You can route traffic for IP addresses outside your VPC to the Internet gateway.

Before you begin

- Create a VPC for your management center virtual instances.
-

Step 1 Click **Services > VPC**.

Step 2 Click **VPC Dashboard > Internet Gateways**, and then click **Create Internet Gateway**.

Step 3 Enter a user-defined Name tag to identify the gateway and click **Yes, Create** to create the gateway.

Step 4 Select the gateway created in the previous step.

Step 5 Click **Attach to VPC** and select the VPC you created previously.

Step 6 Click **Yes, Attach** to attach the gateway to your VPC.

By default, the instances launched on the VPC cannot communicate with the Internet until a gateway is created and attached to the VPC.

What to do next

Add subnets to your VPC as described in the next section.

Add Subnets

You can segment the IP address range of your VPC that the management center virtual instances can be attached to. You can create subnets to group instances according to security and operational needs. For the threat defense virtual you need to create a subnet for management as well as subnets for traffic.

Step 1 Click **Services > VPC**.

Step 2 Click **VPC Dashboard > Subnets**, and then click **Create Subnet**.

Step 3 Enter the following in the **Create Subnet** dialog box:

- A user-defined **Name tag** to identify the subnet.
- A **VPC** to use for this subnet.

- c) The **Availability Zone** where this subnet will reside. Select No Preference to let Amazon select the zone.
- d) A **CIDR block** of IP addresses. The range of IP addresses in the subnet must be a subset of the range of IP addresses in the VPC. Block sizes must be between a /16 network mask and a /28 network mask. The size of the subnet can equal the size of the VPC.

Step 4 Click **Yes, Create** to create your subnet.

Step 5 Repeat for as many subnets required. Create a separate subnet for management traffic and create as many subnets as needed for data traffic.

What to do next

Add a route table to your VPC as described in the next section.

Add a Route Table

You can attach a route table to the gateway you configured for your VPC. You can also associate multiple subnets with a single route table, but a subnet can be associated with only one route table at a time.

Step 1 Click **Services > VPC**.

Step 2 Click **VPC Dashboard > Route Tables**, and then click **Create Route Table**.

Step 3 Enter a user-defined **Name tag** to identify the route table.

Step 4 Select the **VPC** from the drop-down list that will use this route table.

Step 5 Click **Yes, Create** to create your route table.

Step 6 Select the route table that you just created.

Step 7 Click the **Routes** tab to display the route information in the details pane.

Step 8 Click **Edit**, then click **Add another route**.

a) In the **Destination** column, enter **0.0.0.0/0**.

b) In the **Target** column, select the Internet Gateway you created above.

Step 9 Click **Save**.

Step 10 Click the **Subnet Associations** tab and click **Edit**.

Step 11 Check the box next to the subnet to be used for the management center virtual's management interface and click **Save**.

What to do next

Create a security group as described in the next section.

Create a Security Group

You can create a security group with rules specifying allowed protocols, ports and source IP ranges. Multiple security groups can be created with different rules which you can assign to each instance. AWS has detailed documentation on Security Groups if you are not familiar with this feature.

-
- Step 1** Click **Services** > **EC2**.
- Step 2** Click **EC2 Dashboard** > **Security Groups**.
- Step 3** Click **Create Security Group**.
- Step 4** Enter the following in the **Create Security Group** dialog box:
- A user-defined **Security group name** to identify the security group.
 - A **Description** for this security group.
 - The **VPC** associated with this security group.
- Step 5** Configure **Security group rules**:
- Click the **Inbound** tab, then click **Add Rule**.

Note HTTPS and SSH access is required to manage the management center virtual from outside AWS. You should specify the Source IP addresses accordingly. Also, if you are configuring both the management center virtual and threat defense virtual within the AWS VPC, you should allow the private IP management subnet access.
 - Click the **Outbound** tab, then click **Add Rule** to add a rule for outbound traffic, or leave the defaults of **All traffic** (for **Type**) and **Anywhere** (for **Destination**).
- Step 6** Click **Create** to create your security group.
-

What to do next

Create network interfaces as described in the next section.

Create Network Interfaces

You can create network interfaces for the management center virtual using static IP addresses. Create network interfaces (external and internal) as needed for your particular deployment.

- Step 1** Click **Services** > **EC2**.
- Step 2** Click **EC2 Dashboard** > **Network Interfaces**.
- Step 3** Click **Create Network Interface**.
- Step 4** Enter the following in the **Create Network Interface** dialog box:
- A optional user-defined **Description** for the network interface.
 - Select a **Subnet** from the drop-down list. Make sure to select the subnet of the VPC where you want to create the instance.
 - Enter a **Private IP** address. It is recommended to use a static IP address rather than **auto-assign**.
 - Select one or more **Security groups**. Make sure the security group has all the required ports open.
- Step 5** Click **Yes, Create** to create your network interface.
- Step 6** Select the network interface that you just created.
- Step 7** Right-click and select **Change Source/Dest. Check**.
- Step 8** Choose **Disabled**, then click **Save**.

Repeat this for any network interfaces you create.

What to do next

Create elastic IP addresses as described in the next section.

Create Elastic IPs

When an instance is created, a public IP address is associated with the instance. That public IP address changes automatically when you STOP and START the instance. To resolve this issue, assign a persistent public IP address to the instance using Elastic IP addressing. Elastic IPs are reserved public IPs that are used for remote access to the management center virtual as well as other instances. AWS has detailed documentation on Elastic IPs if you are not familiar with this feature.



Note At a minimum, you want to create one elastic IP addresses for the management center virtual and two elastic IP addresses for the threat defense virtual management and diagnostic interfaces.

-
- Step 1** Click **Services > EC2**.
- Step 2** Click **EC2 Dashboard > Elastic IPs**.
- Step 3** Click **Allocate New Address**.
- Repeat this step for as many elastic/public IPs that you require.
- Step 4** Click **Yes, Allocate** to create your elastic IP.
- Step 5** Repeat for as many elastic IPs required for your deployment.
-

What to do next

Deploy the management center virtual as described in the next section.

Deploy the Management Center Virtual

Before you begin

- Configure AWS VPC and EC2 elements as described in [Configure the AWS Environment](#).
- Confirm that an AMI is available for the management center virtual instances.



Note The default admin password is the AWS Instance ID, unless you define a default password with user data (**Advanced Details > User Data**) during the initial deployment.

- Step 1** Go to <https://aws.amazon.com/marketplace> (Amazon Marketplace) and sign in.
- Step 2** After you are logged in to the Amazon Marketplace, click the link provided for the management center virtual.
- Note** If you were previously in AWS, you may need to sign out and then sign back in for the link to work.
- Step 3** Click **Continue**, then click the **Manual Launch** tab.
- Step 4** Click **Accept Terms**.
- Step 5** Click **Launch with EC2 Console** in your desired region
- Step 6** Choose an **Instance Type** supported by the management center virtual; see [Overview](#) for the supported instance types.
- Step 7** Click the **Next: Configure Instance Details** button at the bottom of the screen:
- Change the **Network** to match your previously created VPC.
 - Change the **Subnet** to match your previously created management subnet. You can specify an IP address or use auto-generate.
 - Under **Advanced Details > User Data**, add the default login information.
Modify the example below to match your requirements for device name and password.
Sample login configuration:
- ```
#FMC
{
 "AdminPassword": "<enter_your_password>",
 "Hostname": "<Hostname-vFMC>"
}
```
- Caution** Use only plain text when entering data in the **Advanced Details** field. If you copy this information from a text editor, make sure you copy only as plain text. If you copy any Unicode data into the **Advanced Details** field, including white space, the instance may be corrupted and you will have to terminate the instance and re-create it.
- In Version 7.0 and greater, the default admin password is the AWS Instance ID, unless you define a default password with user data (**Advanced Details > User Data**) during the initial deployment.
- In earlier releases the default admin password was **Admin123**.
- Step 8** Click **Next: Add Storage** to configure your storage device settings.
- Edit the settings of the root volume so the volume Size (GiB) is 250 GiB. Less than 250 GiB will limit event storage and is not supported.
- Step 9** Click **Next: Tag Instance**.
- A tag consists of a case-sensitive key-value pair. For example, you could define a tag with **Key** = Name and **Value** = Management.
- Step 10** Select **Next: Configure Security Group**.
- Step 11** Click **Select an existing Security Group** and choose the previously configured Security Group, or create a new Security Group; see AWS documentation for more information on creating Security Groups.
- Step 12** Click **Review and Launch**.
- Step 13** Click **Launch**.
- Step 14** Select an existing key pair or create a new key pair.

**Note** You can select an existing key pair, or create a new key pair. The key pair consists of a public key that AWS 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 may be required to connect to the instance.

**Step 15** Click **Launch Instances**.

**Step 16** Click **EC2 Dashboard > Elastic IPs** and find a previously allocated IP, or allocate a new one.

**Step 17** Select the elastic IP, right-click and select **Associate Address**.

Locate the Instance or Network Interface to select, then click Associate.

**Step 18** Click **EC2 Dashboard > Instances**.

**Step 19** The management center virtual Instance state will show “running” and Status checks will show pass for “2/2 checks” after only a few minutes. However, deployment and initial setup processes will take approximately 30 to 40 minutes to complete. To view the status, right-click the Instance, then select Instance **Settings > Get Instance Screenshot**.

When setup is complete (after approximately 30 to 40 minutes), the **Instance Screenshot** should show a message similar to “Cisco Firepower Management Center for AWS vW.X.Y (build ZZ)” and possibly followed by some additional lines of output.

You should then be able to log in to the newly created the management center virtual using SSH or HTTPS. Actual deployment times may vary depending on the AWS load in your region.

You can access the management center virtual using SSH:

```
ssh -i <key_pair>.pem admin@<Public_Elastic_IP>
```

SSH authentication is handled by a key pair. No password is required. If you are prompted for a password then setup is still running.

You can also access the management center virtual using HTTPS:

```
https://<Public_Elastic_IP>
```

**Note** If you see a “system startup processes are still running” then setup is not yet complete.

If you get no response from SSH or HTTPS, double check these items:

- Make sure deployment is complete. The management center virtual VM Instance Screenshot should show a message similar to “Cisco Firepower Management Center for AWS vW.X.Y (build ZZ)” and possibly followed by some additional lines of output.
- Make sure you have an Elastic IP and that it is associated with the management center's management network interface (eni) and that you are connecting to that IP address.
- Make sure there is an Internet Gateway (igw) associated with your VPC.
- Make sure your management subnet has a route table associated with it.
- Make sure the route table associated with your Management subnet has a route for “0.0.0.0/0” that points to your Internet gateway (igw).
- Make sure your Security Group allows incoming SSH and/or HTTPS from the IP addresses you are connecting from.

**What to do next**

## Configuring Policies and Device Settings

After you install the threat defense virtual and add the device to the Management Center, you can use the management center user interface to configure device management settings for the threat defense virtual running on AWS and to configure and apply access control policies and other related policies to manage traffic using your threat defense virtual device. The security policy controls the services provided by the threat defense virtual, such as Next Generation IPS filtering and application filtering. You configure the security policy on the threat defense virtual using the management center. For information about how to configure the security policy, see the Configuration Guide or the online help in Management Center.

-