Deploy the ASAv On the Microsoft Azure Cloud

You can deploy the ASAv on the Microsoft Azure cloud.

Overview

Microsoft Azure is a public cloud environment that uses a private Microsoft Hyper V Hypervisor. The ASAv runs as a guest in the Microsoft Azure environment of the Hyper V Hypervisor. The ASAv on Microsoft Azure supports the Standard D3 and Standard D3_v2 instances, which supports four vCPUs, 14 GB, and four interfaces.

Table 1: ASAv Licensed Feature Limits Based on Entitlement

Performance Tier	Instance Type (Core/RAM)	Rate Limit	RA VPN Session Limit
ASAv5	D3_v2 4 core/14 GB	100 Mbps	50
ASAv10	D3_v2 4 core/14 GB	1 Gbps	250
ASAv30	D3_v2 4 core/14 GB	2 Gbps	750
ASAv50	D4_v2 8 core/28 GB	5.5 Gbps	10,000

Performance Tier	Instance Type (Core/RAM)	Rate Limit	RA VPN Session Limit
ASAv100	D5_v2 16 core/56 GB	11 Gbps	20,000

You can deploy the ASAv on Microsoft Azure:

- As a stand-alone firewall using the Azure Resource Manager on the standard Azure public cloud and the Azure Government environments
- As an integrated partner solution using the Azure Security Center
- As a high availability (HA) pair using the Azure Resource Manager on the standard Azure public cloud and the Azure Government environments

See Deploy the ASAv from Azure Resource Manager, on page 8. Note that you can deploy the ASAv HA configuration on the standard Azure public cloud and the Azure Government environments.

Prerequisites

- Create an account on Azure.com.

  After you create an account on Microsoft Azure, you can log in, choose the ASAv in the Microsoft Azure Marketplace, and deploy the ASAv.

- License the ASAv.

  Until you license the ASAv, it will run in degraded mode, which allows only 100 connections and throughput of 100 Kbps. See Smart Software Licensing for the ASAv.

Note The ASAv defaults to the ASAv30 entitlement when deployed on Azure. The use of the ASAv5, ASAv10, ASAv30, ASAv50, and ASAv100 entitlement is allowed. However, the throughput level must be explicitly configured to use the ASAv5, ASAv10, ASAv30, ASAv50, and ASAv100 entitlement.

- Interface requirements:

  You must deploy the ASAv with four interfaces on four networks. You can assign a public IP address to any interface; see Public IP addresses for Azure's guidelines regarding public IPs, including how to create, change, or delete a public IP address.

- Management interface:

  In Azure, the first defined interface is always the Management interface.

- Communications paths:
  - Management interface—Used for SSH access and to connect the ASAv to the ASDM.
  - 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 when using the Standard_D3 interface.

- For ASAv hypervisor and virtual platform support information, see Cisco ASA Compatibility.

Guidelines and Limitations

Supported Features

- Deployment from Microsoft Azure Cloud
- Maximum of 16 vCPUs, based on the selected instance type

Note Azure does not provide configurable L2 vSwitch capability.

- Public IP address on any interface

  You can assign a public IP address to any interface; see Public IP addresses for Azure's guidelines regarding public IPs, including how to create, change, or delete a public IP address.

- Routed firewall mode (default)

Note In routed firewall mode the ASAv is a traditional Layer 3 boundary in the network. This mode requires an IP address for each interface. Because Azure does not support VLAN tagged interfaces, the IP addresses must be configured on non-tagged, non-trunk interfaces.

Azure DDoS Protection Feature

Azure DDoS Protection in Microsoft Azure is an additional feature implemented at the forefront of ASAv. In a virtual network, when this feature is enabled it helps to defend applications against common network layer attacks depending on the packet per second of a network's expected traffic. You can customize this feature based on the network traffic pattern.

For more information about the Azure DDoS Protection feature, see Azure DDoS Protection Standard overview.

Password Setup

Ensure that the password you set complies with the guidelines given below. The password must:

- Be an alphanumeric string with a minimum of 12 characters and a maximum of 72 characters
- Comprise of lowercase and uppercase characters, numbers, and special characters that are not '\' or '-'
- Have no more than 2 repeating or sequential ASCII characters
- Not be a word that can be found in the dictionary

If you see an error such as the one given below or any other password-related errors in the boot logs, check if the password that has been set up satisfies the password complexity guidelines.

```
OS Provisioning failed for VM 'TEST-FW-NCSA-QC' due to an internal error. (Code:
OSProvisioningInternal Error)
```

Known Issues

Idle Timeout

The ASAv on Azure has a configurable *idle timeout* on the VM. The minimum setting is 4 minutes and the maximum setting is 30 minutes. However, for SSH sessions the minimum setting is 5 minutes and the maximum setting is 60 minutes.

Note Be aware that the ASAv's idle timeout always overrides the SSH timeout and disconnects the session. You can choose to match the VM's idle timeout to the SSH timeout so that the session does not timeout from either side.

Failover from Primary ASAv to Standby ASAv

When an Azure upgrade occurs on an ASAv HA in Azure deployment, a failover may occur from the primary ASAv to the standby ASAv. An Azure upgrade causes the primary ASAv to enter a pause state. The standby ASAv does not receive any hello packets when the primary ASAv is paused. If the standby ASAv does not receive any hello packets beyond the failover hold time, a failover to the standby ASAv occurs.

There is also the possibility of a failover occurring even if the failover hold time has not been exceeded. Consider a scenario in which the primary ASAv resumes 19 seconds after entering the pause state. The failover hold time is 30 seconds. But, the standby ASAv does not receive hello packets with the right timestamp because the clock is synchronized every ~2 minutes. This causes a failover from the primary ASAv to the standby ASAv.

Note This feature supports IPv4 only, ASA Virtual HA is not supported for IPv6 configuration.

Unsupported Features

- Console access (management is performed using SSH or ASDM over network interfaces)
- VLAN tagging on user instance interfaces
- Jumbo frames
- Proxy ARP for an IP address that the device does not own from an Azure perspective
- Promiscuous mode (no sniffing or transparent mode firewall support)

Note Azure policy prevents the ASAv from operating in transparent firewall mode because it doesn't allow interfaces to operate in promiscuous mode.

- Multi-context mode

- Clustering
- ASAv native HA.

Note You can deploy ASAv on Azure in a stateless Active/Backup high availability (HA) configuration.

- VM import/export
- By default, FIPS mode is not enabled on the ASAv running in the Azure cloud.

Note If you enable FIPS mode, you must change the Diffie-Helman key exchange group to a stronger key by using the **ssh key-exchange group dh-group14-sha1** command. If you don't change the Diffie-Helman group, you will no longer be able to SSH to the ASAv, and that is the only way to initially manage the ASAv.

- IPv6
- Gen 2 VM generation on Azure
- Re-sizing the VM after deployment
- Migration or update of the Azure Storage SKU for the OS Disk of the VM from premium to standard SKU and vice versa

Resources Created During Deployment

When you deploy the ASAv in Azure the following resources are created:

- The ASAv machine
- A resource group (unless you chose an existing resource group)

  The ASAv resource group must be the same resource group used by the Virtual Network and the Storage Account.

- Four NICS named vm name-Nic0, vm name-Nic1, vm name-Nic2, vm name-Nic3

  These NICs map to the ASAv interfaces Management 0/0, GigabitEthernet 0/0, GigabitEthernet 0/1, and GigabitEthernet 0/2 respectively.

Note Based on the requirement, you can create Vnet with IPv4 only .

- A security group named vm name-SSH-SecurityGroup

  The security group will be attached to the VM's Nic0, which maps to ASAv Management 0/0.

  The security group includes rules to allow SSH and UDP ports 500 and UDP 4500 for VPN purposes. You can modify these values after deployment.

- Public IP addresses (named according to the value you chose during deployment)

  You can assign a public IP address (IPv4 only).

  to any interface; see Public IP addresses for Azure's guidelines regarding public IPs, including how to create, change, or delete a public IP address.

- A Virtual Network with four subnets (unless you chose an existing network)
- A Routing Table for each subnet (updated if it already exists)

  The tables are named subnet name-ASAv-RouteTable.

  Each routing table includes routes to the other three subnets with the ASAv IP address as the next hop. You may chose to add a default route if traffic needs to reach other subnets or the Internet.

- A boot diagnostics file in the selected storage account

  The boot diagnostics file will be in Blobs (binary large objects).

- Two files in the selected storage account under Blobs and container VHDs named *vm name*-disk.vhd and *vm name*-<uuid>.status
- A Storage account (unless you chose an existing storage account)

Note When you delete a VM, you must delete each of these resources individually, except for any resources you want to keep.

Azure Routing

Routing in an Azure Virtual Network is determined by the Virtual Network's Effective Routing Table. The Effective Routing Table is a combination of an existing System Routing Table and the User Defined Routing Table.

Note The ASAv cannot use dynamic interior routing protocols like EIGRP and OSPF due to the nature of Azure cloud routing. The Effective Routing Table determines the next hop, regardless of whether a virtual client has any static/dynamic route configured.

Currently you cannot view either the Effective Routing Table or the System Routing Table.

You can view and edit the User Defined Routing table. When the System table and the User Defined tables are combined to form the Effective Routing Table, the most specific route wins and ties go to the User Defined Routing table. The System Routing Table includes a default route (0.0.0.0/0) pointing to Azure's Virtual Network Internet Gateway. The System Routing Table also includes specific routes to the other defined subnets with the next-hop pointing to Azure's Virtual Network infrastructure gateway.

To route traffic through the ASAv, the ASAv deployment process adds routes on each subnet to the other three subnets using the ASAv as the next hop. You may also want to add a default route (0.0.0.0/0) that points to the ASAv interface on the subnet. This will send all traffic from the subnet through the ASAv, which may require that ASAv policies be configured in advance to handle that traffic (perhaps using NAT/PAT).

Because of the existing specific routes in the System Routing Table, you must add specific routes to the User Defined Routing table to point to the ASAv as the next-hop. Otherwise, a default route in the User Defined table would lose to the more specific route in the System Routing Table and traffic would bypass the ASAv.

Routing Configuration for VMs in the Virtual Network

Routing in the Azure Virtual Network depends on the Effective Routing Table and not the particular gateway settings on the clients. Clients running in a Virtual Network may be given routes by DHCP that are the .1 address on their respective subnets. This is a place holder and serves only to get the packet to the Virtual Network's infrastructure virtual gateway. Once a packet leaves the VM it is routed according to the Effective Routing Table (as modified by the User Defined Table). The Effective Routing Table determines the next hop regardless of whether a client has a gateway configured as .1 or as the ASAv address.

Azure VM ARP tables will show the same MAC address (1234.5678.9abc) for all known hosts. This ensures that all packets leaving an Azure VM will reach the Azure gateway where the Effective Routing Table will be used to determine the path of the packet.

Note The ASAv cannot use dynamic interior routing protocols like EIGRP and OSPF due to the nature of Azure cloud routing. The Effective Routing Table determines the next hop, regardless of whether a virtual client has any static/dynamic route configured.

IP Addresses

The following information applies to IP addresses in Azure:

- You should use DHCP to set the IP addresses of ASAv interfaces.

  The Azure infrastructure ensures that the ASAv interfaces are assigned the IP addresses set in Azure.

- Management 0/0 is given a private IP address in the subnet to which it is attached.

  A public IP address may be associated with this private IP address and the Azure Internet gateway will handle the NAT translations.

- You can assign a public IP address to any interface.
- Public IP addresses that are dynamic may change during an Azure stop/start cycle. However, they are persistent during Azure restart and during ASAv reload.
- Public IP addresses that are static won't change until you change them in Azure.

DNS

All Azure virtual networks have access to a built-in DNS server at 168.63.129.16 that you can use as follows:

```
configure terminal
dns domain-lookup management
dns server-group DefaultDNS
```

```
 name-server 168.63.129.16
end
```

You can use this configuration when you configure Smart Licensing and you don't have your own DNS Server set up.

Deploy the ASAv

You can deploy the ASAv on Microsoft Azure.

- Deploy the ASAv as a stand-alone firewall using the Azure Resource Manager on the standard Azure public cloud and the Azure Government environments. See Deploy the ASAv from Azure Resource Manager.
- Deploy the ASAv as an integrated partner solution within Azure using the Azure Security Center. Security-conscious customers are offered the ASAv as a firewall option to protect Azure workloads. Security and health events are monitored from a single integrated dashboard. See Deploy the ASAv from Azure Security Center.
- Deploy an ASAv High Availablity pair using the Azure Resource Manager. To ensure redundancy, you can deploy the ASAv in an Active/Backup high availability (HA) configuration. HA in the public cloud implements a stateless Active/Backup solution that allows for a failure of the active ASAv to trigger an automatic failover of the system to the backup ASAv. See Deploy the ASAv for High Availability from Azure Resource Manager, on page 11.

Deploy the ASAv from Azure Resource Manager

The following procedure is a top-level list of steps to set up Microsoft Azure on the ASAv. For detailed steps for Azure setup, see Getting Started with Azure.

When you deploy the ASAv in Azure it automatically generates various configurations, such as resources, public IP addresses, and route tables. You can further manage these configurations after deployment. For example, you may want to change the Idle Timeout value from the default, which is a low timeout.

Step 1 Log into the Azure Resource Manager (ARM) portal.

The Azure portal shows virtual elements associated with the current account and subscription regardless of data center location.

Step 2 Search Marketplace for Cisco ASAv, and then click on the ASAv you would like to deploy.

Step 3 Configure the basic settings.

a) Enter a name for the virtual machine. This name should be unique within your Azure subscription.

Important If your name is not unique and you reuse an existing name, the deployment will fail.

b) Enter your username.

c) Choose an authentication type, either **Password** or **SSH public key**.

If you choose **Password**, enter a password and confirm. See Password Setup for guidelines on password complexity.

d) Choose your subscription type.

e) Choose a **Resource group**.

The resource group should be the same as the virtual network's resource group.

f) Choose your location.

The location should be the same as for your network and resource group.

g) Click **OK**.

Step 4 Configure the ASAv settings.

a) Choose the virtual machine size.

The ASAv supports Standard D3 and Standard D3_v2.

b) Choose a storage account.

You can use an existing storage account or create a new one. The location of the storage account should be the same as for the network and virtual machine.

c) Request a public IP address by entering a label for the IP address in the Name field, and then click **OK**.

Azure creates a dynamic public IP by default, which may change when the VM is stopped and restarted. If you prefer a fixed IP address, you can open the public-ip in the portal and change it from a dynamic to a static address.

d) Add a DNS label if desired.

The fully qualified domain name will be your DNS label plus the Azure URL: *<dnslabel>.<location>.cloupapp.azure.com*

e) Choose an existing virtual network or create a new one.

f) Configure the four subnets that the ASAv will deploy to, and then click **OK**.

Important Each interface must be attached to a unique subnet.

g) Click **OK**.

Step 5 View the configuration summary, and then click **OK**.

Step 6 View the terms of use and then click **Create**.

What to do next

- Continue configuration using CLI commands available for input via SSH or use ASDM. See Start ASDM for instructions for accessing the ASDM.

Deploy the ASAv from Azure Security Center

The Microsoft Azure Security Center is a security solution for Azure that enables customers to protect, detect, and mitigate security risks for their cloud deployments. From the Security Center dashboard, customers can set security policies, monitor security configurations, and view security alerts.

Security Center analyzes the security state of Azure resources to identify potential security vulnerabilities. A list of recommendations guides customers through the process of configuring needed controls, which can include deployment of the ASAv as a firewall solution to Azure customers.

As an integrated solution in Security Center, you can rapidly deploy the ASAv in just a few clicks and then monitor security and health events from a single dashboard. The following procedure is a top-level list of steps to deploy the ASAv from Security Center. For more detailed information, see Azure Security Center.

Step 1 Log into the Azure portal.

The Azure portal shows virtual elements associated with the current account and subscription regardless of data center location.

Step 2 From the Microsoft Azure menu, choose **Security Center**.

If you are accessing Security Center for the first time, the **Welcome** blade opens. Choose **Yes! I want to Launch Azure Security Center** to open the **Security Center** blade and to enable data collection.

Step 3 On the **Security Center** blade, choose the **Policy** tile.

Step 4 On the **Security policy** blade, choose **Prevention policy**.

Step 5 On the **Prevention policy** blade, turn on the recommendations that you want to see as part of your security policy.

a) Set **Next generation firewall** to **On**. This ensures that the ASAv is a recommended solution in Security Center.
b) Set any other recommendations as needed.

Step 6 Return to the **Security Center** blade and the **Recommendations** tile.

Security Center periodically analyzes the security state of your Azure resources. When Security Center identifies potential security vulnerabilities, it shows recommendations on the **Recommendations** blade.

Step 7 Select the **Add a Next Generation Firewall** recommendation on the **Recommendations** blade to view more information and/or to take action to resolve the issue.

Step 8 Choose **Create New** or **Use existing solution**, and then click on the ASAv you would like to deploy.

Step 9 Configure the basic settings.

a) Enter a name for the virtual machine. This name should be unique within your Azure subscription.

Important If your name is not unique and you reuse an existing name, the deployment will fail.

b) Enter your username.
c) Choose an authorization type, either password or SSH key.

If you choose password, enter a password and confirm. See Password Setup for guidelines on password complexity.

d) Choose your subscription type.
e) Choose a resource group.

The resource group should be the same as the virtual network's resource group.

f) Choose your location.

The location should be the same as for your network and resource group.

g) Click **OK**.

Step 10 Configure the ASAv settings.

a) Choose the virtual machine size.

The ASAv supports Standard D3 and Standard D3_v2.

b) Choose a storage account.

You can use an existing storage account or create a new one. The location of the storage account should be the same as for the network and virtual machine.

c) Request a public IP address by entering a label for the IP address in the Name field, and then click **OK**.

Azure creates a dynamic public IP by default, which may change when the VM is stopped and restarted. If you prefer a fixed IP address, you can open the public-ip in the portal and change it from a dynamic to a static address.

d) Add a DNS label if desired.

The fully qualified domain name will be your DNS label plus the Azure URL: *<dnslabel>.<location>.cloupapp.azure.com*

e) Choose an existing virtual network or create a new one.

f) Configure the four subnets that the ASAv will deploy to, and then click **OK**.

Important Each interface must be attached to a unique subnet.

g) Click **OK**.

Step 11 View the configuration summary, and then click **OK**.

Step 12 View the terms of use and then click **Create**.

What to do next

- Continue configuration using CLI commands available for input via SSH or use ASDM. See Start ASDM for instructions for accessing the ASDM.
- If you need more information on how the recommendations in Security Center help you protect your Azure resources, see the documentation available from Security Center.

Deploy the ASAv for High Availability from Azure Resource Manager

The following procedure is a top-level list of steps to set up a High Availability (HA) ASAv pair on Microsoft Azure. For detailed steps for Azure setup, see Getting Started with Azure.

ASAv HA in Azure deploys two ASAvs into an Availability Set, and automatically generates various configurations, such as resources, public IP addresses, and route tables. You can further manage these configurations after deployment.

Step 1 Log into the Azure portal.

The Azure portal shows virtual elements associated with the current account and subscription regardless of data center location.

Step 2 Search Marketplace for **Cisco ASAv**, and then click on the **ASAv 4 NIC HA** to deploy a failover ASAv configuration.

Step 3 Configure the **Basics** settings.

a) Enter a prefix for the ASAv machine names. The ASAv names will be ‘prefix’-A and ‘prefix’-B.

Important Make sure you do not use an existing prefix or the deployment will fail.

b) Enter a **Username**.

This will be the administrative username for both Virtual Machines.

Important The username **admin** is not allowed in Azure.

c) Choose an authentication type for both Virtual Machines, either **Password** or **SSH public key**.

If you choose **Password**, enter a password and confirm. See Password Setup for guidelines on password complexity.

d) Choose your subscription type.

e) Choose a **Resource group**.

Choose **Create new** to create a new resource group, or **Use existing** to select an existing resource group. If you use an existing resource group, it must be empty. Otherwise you should create a new resource group.

f) Choose your **Location**.

The location should be the same as for your network and resource group.

g) Click **OK**.

Step 4 Configure the **Cisco ASAv settings**.

a) Choose the Virtual Machine size.

The ASAv supports Standard D3 and Standard D3_v2.

b) Choose **Managed** or **Unmanaged OS disk** storage.

Important ASA HA mode always uses **Managed**.

Step 5 Configure the **ASAv-A** settings.

a) (Optional) Choose **Create new** to request a public IP address by entering a label for the IP address in the Name field, and then click **OK**. Choose **None** if you do not want a public IP address.

Note Azure creates a dynamic public IP by default, which may change when the VM is stopped and restarted. If you prefer a fixed IP address, you can open the public-ip in the portal and change it from a dynamic to a static address.

b) Add a DNS label if desired.

The fully qualified domain name will be your DNS label plus the Azure URL: *<dnslabel>.<location>.cloupapp.azure.com*

c) Configure the required settings for the storage account for the ASAv-A boot diagnostics.

Step 6 Repeat the previous steps for the **ASAv-B** settings.

Step 7 Choose an existing virtual network or create a new one.

a) Configure the four subnets that the ASAv will deploy to, and then click **OK**.

Important Each interface must be attached to a unique subnet.

b) Click **OK**.

Step 8 View the **Summary** configuration, and then click **OK**.

Step 9 View the terms of use and then click **Create**.

What to do next

- Continue configuration using CLI commands available for input via SSH or use ASDM. See Start ASDM for instructions for accessing the ASDM.
- See the 'Failover for High Availability in the Public Cloud' chapter in the ASA Series General Operations Configuration Guide for more information about ASAv HA configuration in Azure.